Towards a More Accurate Carrier Sensing Model for CSMA Wireless Networks

Cai Hong Kai, Soung Chang Liew
Department of Information Engineering
The Chinese University of Hong Kong, Hong Kong SAR
{chkai6, soung}@ie.cuhk.edu.hk

Abstract—In the majority of studies on CSMA wireless networks, a contention graph is used to model the carrier sensing relationships among links. This is a 0-1 model in which two links can either sense each other completely or not. In real experiments, we observed that this is generally not the case: the carrier sensing relationship between the links are often probabilistic and can vary dynamically over time. This is the case even if the distance between the links is fixed and there is no drastic change in the environment. Furthermore, this "partial carrier sensing" relationship is prevalent and occurs over a wide range of distances between the links. This observation is not consistent with the 0-1 contention graph and implies that many results and conclusions drawn from previous theoretical studies need to be re-examined. This paper establishes a more accurate carrier sensing model with the objective of laying down a foundation for future theoretical studies that reflect reality. Towards that end, we set up detailed experiments to investigate the partial carrier sensing phenomenon. We discuss the implications and the use of our partial carrier sensing model in network analysis.

Index Terms—Contention graph, partial carrier sensing, CSMA.

I. INTRODUCTION

This paper concerns the carrier sensing (CS) behavior in wireless networks. Carrier sensing is an important feature of CSMA networks (i.e., 802.11 networks) to avoid packet collisions. A node that has packets to send must first sense the channel. If no nearby node is transmitting, it transmits immediately. If a nearby node is transmitting, it defers, waiting until the end of the interfering transmission before attempting to transmit. That is, nodes that can sense each other will not transmit simultaneously to avoid packet collisions.

In the majority of prior studies, the carrier sensing relationships among the links (a link is a transmitter-receiver pair) are modeled by a contention graph. The links are represented by vertices, and an edge joins two vertices if the associated transmitters can sense each other. In other words, the carrier sensing behavior between two links is a 0-1 relationship in that they can either sense or not sense each other. A model in which the CS relationship is a simple function of distance is often assumed. If the distance between two links is shorter than a threshold (called the Carrier Sensing Range, CSRange), they can always sense each other; otherwise the two links will never

hear each other. Although widely adopted in both theoretical and simulation studies, this model has not undergone rigorous verification in practice.

Indeed, our real-network experimental data indicated that this 0-1 CS model is not accurate. Very often a node could only hear another node some but not all of the time. This phenomenon is likely to be caused by the intricate interplay of channel fading and circuit noise, as well as co-channel interference. Surprisingly, theoretical work to-date on CSMA networks has largely shunned the issue of channel fading and circuit noise when it comes to carrier sensing, although their effects on signal interference and correct data reception have been well studied. This calls into question the findings and conclusions reached by the past investigations.

To build a more realistic CS model, we conducted extensive experiments to characterize the partial carrier sensing relationship among links. We show that a probabilistic CS model matches experimental data more closely than the prior absolute 0-1 CS model.

The main contributions of this paper are as follows: 1) we show that there is a long range of distance between two links over which the carrier sensing between them is partial rather than full, and therefore one can expect partial CS to be prevalent in a typical CSMA wireless network; 2) we show that partial CS has a significant effect on link throughputs, and therefore one should not simply approximate the partial CS model with the 0-1 CS model in analysis; 3) we propose a more realistic probabilistic partial CS model that match the experiment results; 4) we discuss the implications and the use of the partial CS model in analytical studies.

Related Work

Ref. [1]–[8] are some recent attempts to study the effects of carrier sensing. Ref. [1], [2] investigated optimal choices of carrier sensing parameters for AWGN channel. Ref. [3], [4] studied the impacts of carrier sensing on system performance under AWGN and slow fading channel. Ref. [1]–[4], however, considered the less interesting all-inclusive carrier-sensing network in which all nodes can hear each other. Ref. [5], [6] are attempts to develop more sophisticated analytical models for packet reception and carrier sensing. However, the intricate details of carrier sensing have been ignored and only energy detection is taken into account (in practical 802.11 devices, preamble detection is often used in conjunction with energy detection for more effective carrier sensing). A gray zone of

3dB where carrier sensing and interference apply partially was observed in [7] and the authors propose to use continuous values to capture the gray zone in [8].

In this paper we show that in reality the CS relationship is probabilistic over a long range of distance. Due to the significant effects of partial CS on network performance, many previous theoretical and simulation frameworks need to be revisited. Some examples are as follows: 1) The NS2 simulator [9] is by far the most popular simulation tool used for the studies of 802.11 networks; however, it uses the unrealistic 0-1 CS model. 2) With partial CS, the throughput distributions among links in a CSMA network may be quite different from those derived under the 0-1 CS model. In particular, many reported problems, such as link starvation and unfairness [10] may be alleviated under a partial CS model, i.e., a link which is expected to be starved in the 0-1 CS model may also obtain some throughput due to not sensing a neighbor link's transmission in practice. 3) Many previous studies based on the 0-1 contention graph and their conclusions will need to be re-examined given the existence of partial CS. For example, the "island states" and the "phase transition" phenomenon as reported in [10] and [11] may not be common in practice.

II. LARGE TRANSITION RANGE OF CARRIER SENSING IN REAL NETWORKS

This section first gives a quick review of the carrier sensing mechanisms as defined in the IEEE 802.11 standard [12]. After that, we show that there is a long range of distance between two links over which partial carrier sensing occurs. In particular, the transition from full carrier sensing to no carrier sensing is a gradual rather than an abrupt process as the distance varies. Also, we show that even if the distance between the two transmitters is not changed, partial CS still exists when the transmit power varies.

A. Carrier sensing in IEEE 802.11 standards

In 802.11 networks, physical carrier sensing (PCS) is performed by the Clear Channel Assessment (CCA) function, which monitors the channel to determine whether it is free. The 802.11 standard defines three CCA operation modes. The channel is declared as busy when (i) the energy detected exceeds a threshold CS_{th} (Energy Detection); (ii) a valid 802.11 signal is detected, even if the power is below CS_{th} (Preamble Detection); (iii) either (i) or (ii) occurs. The research community has largely considered only (i) in the investigations of CSMA networks, although modes (ii) and (iii) are often used in real 802.11 equipments.

Under (ii), in the event that a correct PLCP (Physical Layer convergence Procedure) Header is received, the CCA signal may be held inactive (channel busy) for the full duration of the packet as indicated by the PLCP LENGTH field. Even if a loss of carrier occurs in the middle of reception, the CCA will indicate a busy medium for the intended duration of the transmitted packet.

B. Long transition range of carrier sensing in real networks

In the static 0-1 carrier sensing model, the carrier sensing range (CSRange) is defined to be the distance between two transmitters within which concurrent transmissions are not allowed. Consider the setup as shown in Fig. 1. Suppose that the link lengths of the two links are very small (say, 1m) compared with the separation of the links d (say, $d \geq 10$ m) so that packets, if transmitted, are most surely received successfully. In this case, the link throughput measured at the receiver also corresponds to the packet transmission rate at the transmitter.

For the 0-1 CS model, within the CSRange, the two links can sense each other with probability 1; beyond CSRange, they do not sense each other and can transmit independently. That is, when $d <$ CSRange, the two links will share the channel and each of them will get half of the medium airtime. The normalized aggregate throughput of the two links is thus 1. Once $d \geq$ CSRange, the two links can transmit as if they are isolated links. The normalized aggregate throughput jumps to 2 immediately. As illustrated in Fig. 2(a), the aggregate throughputs of the two links change abruptly at $d =$ CSRange in the 0-1 CS model.

In real networks, as illustrated in Fig.2(b), the normalized aggregate throughput of the two links increase gradually from 1 to 2 as d increases. Fig. 3(a) shows the measured throughput of one of the links in a real 802.11a two-links network in which the transmit power of each link is fixed to 1mW. Instead of an abrupt jump, the measured throughput increases gradually with d. We are interested in the underlying causes of this long transition range. One possibility is signal capture, which has been reported in [13]. Another possibility is partial carrier sensing, which we find to be the dominating factor over the long transition range, as will be explained in Section III.
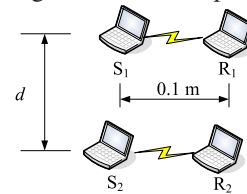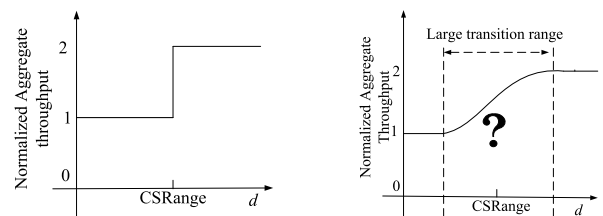

Fig. 1.   Experiment setup of a two-links 802.11a network.

C. Another set of experiments for the existence of partial CS

We conduct another set of experiments in which the link separation is fixed to $d = 20$ m. Instead of moving the nodes, we alter the transmit powers of the two links. Fig. 3(b) shows the measured throughput of one of the links with

(a) under the 0-1 CS model   (b) under a partial CS model
Fig. 2.   Normalized aggregate throughput of the two links versus d under different CS models.
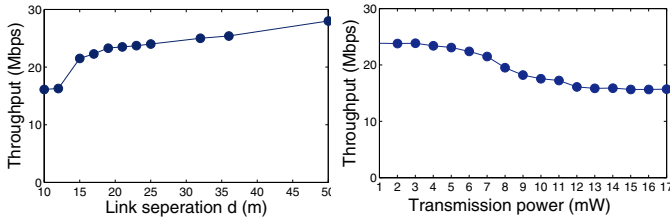
(a) Fix transmit power=1mW and alter d (b) Fix d = 20 m and alter transmit power
Fig. 3.   Measured throughput in two sets of real experiments.

respect to different transmit powers. Note that even when the transmit power is only 1mW, the interaction between the two links cannot be completely isolated. The degree of partial CS varies as the transmit power varies without the channel being changed.

III. EXPERIMENT SETUP AND RESULTS

Our experiments are based on wireless cards that use the Atheros 802.11 chips. We want to measure the CS relationship between two links. However, we do not have direct access to the CCA information. We therefore design an indirect method for measuring CS. By looking at the variations of link throughputs, Packet Loss Ratios and the transmit attempt rates (numbers of transmission attempts per second) with respect to d, we demonstrate the existence of partial carrier sensing.

A. Experiment setup

We set up experiments with two pairs of DELL Latitude D505 laptops with 1.5GHz Celeron Mobile CPU. Each laptop had a NETGEAR WAG511v2 wireless card, and ran Fedora5 with MadWifi driver [14]. All Atheros chipset extensions were disabled. The network setup is shown in Fig. 1. Most of our experiments were conducted outdoor in an empty athletic field during summer vacations to make sure that there was no drastic change in the environment. We used the IEEE 802.11a channel 36, in which there was no other 802.11 sources. The distance between each sender-receiver pair was set to 0.1m to remove hidden-node effects. To make the experiment easier to control, the transmission power of each link was set to the minimum value allowed by the hardware (1mW). OmniPeek, a network analysis software [15], was installed in another laptop to serve as a "sniffer" to collect traffic traces.

Typical 802.11a parameters were used in the experiments: (i) fixed data rate and basic rate of 54 Mbps and 6 Mbps, respectively; (ii) packet payload of 1460 Bytes; (iii) CW_{min} of 15 and mini-timeslot of 9 μs, where CW is the contention window; (iv) basic mode of DCF. IPerf, a network testing tool [16], was used to generate UDP data streams and measure the throughputs. For each UDP session, the date rate was set to 30 Mbps to ensure link saturation. Each experiment lasted for 60 seconds and was repeated three times. Note that although all the data presented in this paper were gathered in outdoor experiments, the long transition range of carrier sensing was also observed in indoor experiments.

B. Experiment results

We define Packet Loss Ratio (PLR) as the ratio between the number of packets lost during transmission and the number of

packets transmitted at the transmitter. That is, PLR = (# of packets transmitted at the transmitter - # of packets received at the receiver)/ # of packets transmitted at the transmitter. The retransmissions of the same packet at MAC layer is counted as multiple transmissions in our measurements. Table I lists the statistics of one of the links versus link separation d.

TABLE I
THROUGHPUT/PLR/TRANSMIT ATTEMPT RATE WITH RESPECT TO LINK SEPARATION d

d	0.2m	1m	3m	6m	12m	15m	20m	50m
Throughput (Mbps)	14.6	15.6	16.1	16.2	16.3	22.7	23.7	27.8
PLR	12.9%	5.0%	1.0%	1.0%	0.7%	1.5%	1.5%	0.4%
# of transmit attempts/sec	1412	1410	1393	1403	1404	1968	2055	2392

C. Signal capture versus Partial carrier sensing

Signal capture refers to the successful reception of the packet with higher signal strength during a collision [13]. In Fig. 1, when $d_{S_2R_1} \gg d_{S_1R_1}$, the power from S_1 to R_1 is much larger than that from S_2; hence, the packet from link 1 can be captured when links 1 and link 2 transmit together. In 802.11a networks the countdown time is uniformly chosen from [0, 15] timeslots, even with full carrier sensing, with probability $1/8.5 = 12\%$ the two links could still transmit together when they count down simultaneously to zero. Signal capture allows link throughputs to be higher since the desired packet can be extracted in a collision.

As d increases, partial carrier sensing may cause the two links to transmit together more often, also giving rise to higher link throughputs. There is the issue of whether signal capture or partial carrier sensing is the dominant effect. Our experimental results indicate that partial carrier sensing may have a more significant effect.

1) Capture effect dominating range (d = 0.2 ∼ 12 m): When d = 0.2 m, the measured PLR is 12.9% (consistent with our computation above). That is, there is no signal capture here. As d increases, the measured PLR decreases, indicting that more packets can be captured due to the stronger power. This factor, however, results in only slightly higher throughput. As shown in Table I, as d increases from 0.2m to 12m, the transmit attempt rate remains more or less constant, indicating that CS is "full" over this range. In particular, the throughput increase comes from the smaller PLR as d increases. We refer to this range of d as the "capture effect dominating range". The throughput increase is not significant over this range (from 14.6 Mbps to 16.3 Mbps).

2) Partial carrier sense dominating range (d = 12 ∼ 50 m) : As shown in Table I, as d increases from 12m to 50m, the throughput increase is much larger (from 16.3 Mbps to 27.8 Mbps). PLR over this range remains more or less constant. The boost in throughput mainly comes from the increase in transmit attempt rates: the ratio of throughput to the number of transmit attempts per second is constant throughout this range. Since the transmitters can make more transmit attempts per second only when they cannot fully sense each other, we conclude that in this range partial carrier sensing kicks in. Also, as d increases, when the two links simultaneously transmit, signal captures have a good chance to occur. However, signal capture alone cannot explain the

large increase in throughput without partial carrier sensing. The experimental data in Table I indicate that there is a large "partial carrier sense dominating range" over which the likelihood of carrier sensing varies from 100% to 0%.

D. Theoretical computation to validate our experiments

Next we conduct some analytical computation to validate the data shown in Table I. For an isolated link, the time consumed by a successful packet transmission consists of (i) PACKET duration consisting of physical-layer preamble/header, MAC Header, and data payload; (ii) SIFS; (iii) ACK; (iv) DIFS; (v) the random number of backoff countdown timeslots. For each packet, the airtime within its carrier-sensing range that must be exclusively dedicated to it is

$$T_{tr} = \text{PACKET} + \text{SIFS} + \text{ACK} + \text{DIFS} \qquad (1)$$

In addition, it also consumes a random backoff countdown time (i.e., component (v) above). Theoretically, in our setup the unshared time needed to transmit a packet is $340\,\mu s$ (about 38 timeslots). So the theoretical throughput of an isolated link is

$$1460 * 8 * 10^6/(340 + 7.5 * 9) = 28.66 \text{ Mbps.}$$

If two links fully hear each other, the average time cycle needed to transmit a packet is the sum of the time consumed by two packets transmission plus a random backoff countdown time (Note that the countdown time is shared by the two links). Hence, the throughput of each link is

$$1460 * 8 * 10^6/(2 * 340 + 7.5 * 9) = 15.63 \text{ Mbps.}$$

With perfect signal capture, the throughput of each link is

$$1460 * 8 * 10^6/((1 + 7.5/8.5) * 340 + 7.5 * 9) = 16.51 \text{ Mbps.}$$

Note that in our experiment setup, when $d = 12 \sim 50$ m, PLR keeps quite low so that its effects (i.e., EIFS and exponential backoff) does not have significant impacts on throughputs. The range of throughput between 16.51Mbps and 28.66Mbps as computed above map roughly to the range of throughputs between $d = 12$ m and $d = 50$ m, the partial carrier sense dominating range. The slightly lower experimental throughputs in Table I are attributed to the fact that our theoretical computation above ignores the periodic beacons sent out by APs and the random packet loss due to noise.

IV. PARTIAL CARRIER SENSING MODELING

This section attempts to build a more realistic carrier sensing model based on experimental results.

A. Statistics of the measured total countdown times

Recall that in 802.11 networks when a link is counting down during its backoff stage, the countdown process may be frozen when it senses a neighbor link transmitting. The total countdown time (active plus frozen countdown time) between two successive transmissions of a link contains important statistics on its partial carrier sensing relationships with other links. The total countdown time can be measured by tracing the inter-packet transmission time (i.e., the time gap between two

transmissions of the same link) at the transmitter (Note that the total countdown time = the inter-packet transmission time - T_{tr} define in (1)). Since all the packets in our experiments are of the same length and are transmitted using the same data rate, the inter-packet arrival time at the sniffer is equal to the inter-packet transmission time at the transmitter plus measurement error. Hence, we collected the inter-packet arrival times at the sniffer using OmniPeek to get the statistics of the total countdown times.

1) An isolated link: For an isolated link, there is no frozen event. The total countdown time is equal to the active countdown time, which is uniformly distributed over [0, 15] timeslots. Fig. 4 plots the probability distributions of the countdown time according to our analysis and according to the experiment. As can be seen, the experimental measurements match well with analysis with only very small deviations from the uniform distribution.
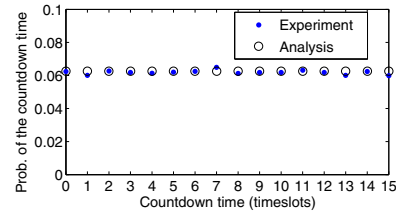

Fig. 4. Countdown time distributions of an isolated link.

2) Two links with full carrier sensing: When the two links fully sense each other, once a link is frozen, it will be frozen for the whole packet transmission time. As a result, the total countdown time (active plus frozen countdown time) of one link will fall into several clusters: [0, 15] timeslots (no frozen events between two successive transmissions of the link); or 38 + [0, 15] timeslots (frozen once between two transmissions); or 76 + [0, 15] timeslots (frozen twice between two transmissions), and so on. In our real-network experiments shown in Section III, when $0\text{m} \leq d \leq 12\text{m}$, the two links can fully sense each other. The total countdown times measured at the sniffer, fall into several clusters. To verify the accuracy of our measurements, we compare the measured distribution of the countdown time with that of theoretical computation in Fig. 5. As can be seen, our measurements fully agree with the theoretical result computed assuming full carrier sensing.
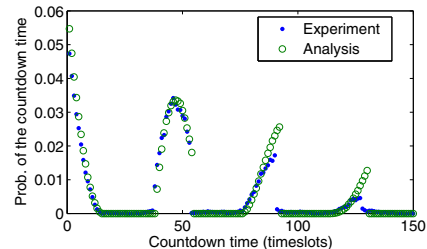

Fig. 5. Countdown time distributions with full carrier sensing.

3) Two links with partial carrier sensing: As argued in Section III, carrier sensing is partial when $12\text{m} < d < 50\text{m}$. Fig. 6 shows the distributions of countdown time (active plus frozen) with respect to d, which we found to be quite stable over different measurement runs. The clusters become less and less distinct as d increases. After $d > 16$ m, all the packets fall into a much wider first cluster of [0, 38] timeslots, indicating

that the link can be frozen *part of the time* when the other link is transmitting. Note from the difference between Fig. 4 (the isolated link) and Fig. 6(d) that the link in Fig. 6(d) did get affected by the other link and did freeze from time to time (just that never more than one packet duration). We see that the carrier-sensing relationship may change in a short-time order (Specifically, from timeslot to timeslots during the backoff stage).
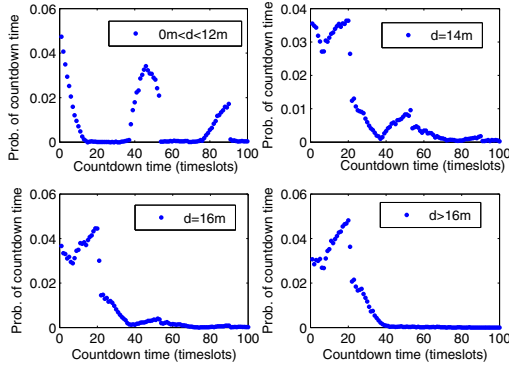

Fig. 6.   Countdown time (active plus frozen) distributions with partial carrier sensing with respect to different d.

B. A more accurate carrier sensing model

We find that a partial carrier sensing model as described in the next paragraph can match the experimental results in Fig. 6 rather well. We emphasize that we are not proposing a new carrier sensing mechanism or design. Rather, we explore how a model that more accurately reflects the operation of the existing carrier sensing mechanism can be constructed. Our model below is consistent with the 802.11 standard specifications shown in Section II-A.

Over each timeslot, a node in idle state tries to detect the presence of a physical preamble with CS threshold set to -82 dBm (receiver sensitivity required for data rate of 6 Mbps in 11a networks). Once a detection event is triggered, CCA will determine whether it is the start of a new 802.11 signal by preamble detection. If not, the node only freezes for this timeslot and gives up receiving immediately. Otherwise, it will spend time trying to track the carrier. It will do so for at least for 4 or 5 timeslots (time needed to receive a complete PHY header). If the PHY header is decoded successfully, the node will reserve the channel for the whole frame transmission time. If it can not decode the PHY header, then it adopts energy detection, with a threshold 20 dB above the minimum 6 Mbps sensitivity. Fig. 7 shows the procedure of physical carrier sensing. Indeed, there are three kinds of frozen events in real networks: 1) it freezes for the whole packet transmission time; 2) it only freezes for one timeslot and 3) it freezes for a period of 4 or 5 timeslots for trying to receive the physical header. As d increases, 2) and 3) become more prevalent.

In our special case of two links, when the transmitter of link 1 finishes a transmission and begins to count down, link 2 is either transmitting or counting down. If link 2 has been transmitting, obviously link 1 will not detect its PHY header because part of the packet of link 2 has already been transmitted. In this case, link 1 will keep sensing with CS
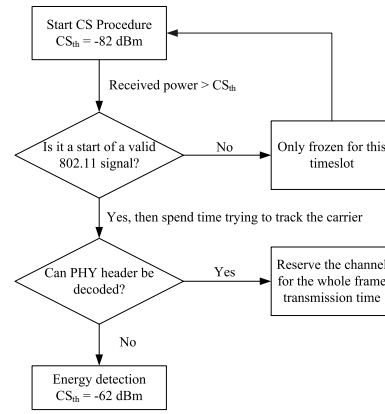

Fig. 7.   Procedure of physical CS.

threshold equal to -82 dBm. Link 1 will be frozen with probability p in each of the subsequent timeslots.

If link 2 is also counting down, over each timeslot with probability q link 2 begins to transmit first and link 1 detects this start of the physical preamble of link 2. Assuming the new generated countdown time of link 1 is k, which is uniformly distributed over [0, 15], then with probability $1-(1-q)^k$, link 1 will detect the beginning of the transmission of link 2. Then link 1 spends 4-5 timeslots trying to track the carrier of link 2. We have the following two possibilities:

i) With probability r, the PHY header of link 2 can be decoded successfully. Once the PHY header is decoded successfully, link 1 will reserve the channel for the whole frame transmission time and its countdown will be frozen (this case does not occur when $d > 16$ m in our experiments).

ii) With probability $1 - r$, link 1 can not decode the PHY header of link 2. After that link 1 adopts energy detection, with a threshold 20 dB above the minimum 6 Mbps sensitivity. Since even the PHY header transmitted at 6 Mbps can not be decoded, when using a 20dB higher threshold, it is rare to find a busy channel to be frozen due to energy detection in the subsequent timeslots.

Fig. 8 shows the comparison between experimental results and countdown time distributions under the partial CS model above when $d = 26$ m. Our model fits the experimental curve very well when the parameters are set as follows: $r = 0$, $p = 0.47$, $q = 0.04$. Note and recall that when $d > 16$ m, no packet falls into the second cluster. The link never reserves the channel for the whole frame transmission (i.e., $r = 0$).
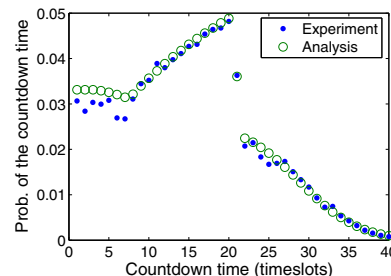

Fig. 8.   Comparison between experimental results and carrier sense model.

V. Implications and Applications of Partial Carrier Sensing

We give three concrete applications to show the necessity of partial carrier sensing modeling.

A. Throughput analysis and the likelihood of starvation

It is known that link throughput distributions of CSMA networks are quite unfair and extreme under the 0-1 CS model [10]. A link suffers from starvation, for example, when it is sandwiched between other links that keep transmitting. Consider the contention graph in Fig. 9. In this network, link 1 and link 3, link 1 and link 4 can transmit together. Together, either link 1 and link 3 or link 1 and link 4 grab the access to the channel most of the time, leaving little chance for link 2 to actively count down and transmit. This results in a normalized throughput distributions of [1, 0, 0.5, 0.5], where link 2 is starved [10]. However, if we assume each pair of links only can hear each other with probability $p = 0.8$, it can be shown that the normalized throughputs are [0.80, 0.25, 0.57, 0.57] (this computation is not presented here due to limited space). Link 2 will not be starved any more and better fairness can be achieved.

In a practical wireless network, due to partial carrier sensing, the starved links under the 0-1 CS model may actually obtain some throughputs. The unfairness and starvation problem may not be as bad as commonly reported in the literature.
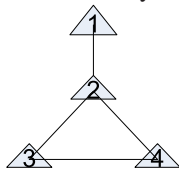

Fig. 9. An example for starvation analysis under the partial CS model.

B. Island States, Phase Transition, and Non-locality Effect

The island-state, phase transition, and non-locality phenomena are corollaries of the static contention graph model [10], [11]. These phenomena may disappear (at least, less distinct) due to partial carrier sensing. Because of partial carrier sensing, the contention graph structure is not static and may change with time. These changes may undo the island-state, phase-transition, and non-locality phenomena, since they are highly dependent on and specific to a static contention-graph structure. Interesting as they are theoretically, it is not clear that these phenomena would occur in practice. Certainly we have not observed these effects in actual operation of Wi-Fi networks today. The issue is an interesting topic left for future investigation.

C. Resource Allocation

Under the 0-1 CS model, the links within a clique cannot transmit together, and therefore their aggregate transmission airtime cannot be larger than 1. This has been the basis for many prior investigations on bandwidth/resource allocation in CSMA networks. Due to partial carrier sensing, however, things become probabilistic and more complicated. It is possible that more than two links transmit together. We believe that in a partial carrier-sensing network with complex clique

formations, the problem of bandwidth allocation will need to be reformulated and re-investigated.

VI. CONCLUSION

In this paper we show that in practice there is a long range of distance where two links partially sense each other (i.e., they sometimes can hear each other, and sometimes cannot). We identify a more accurate *probabilistic* carrier-sensing model based on real network measurements. Although in the paper we only consider the basic two-links case, we believe this is a first important step in exploring the carrier sensing behavior in general wireless CSMA networks.

A goal of this paper is to initiate a new research direction so as to take into account the phenomenon of "partial carrier sensing" in future analytical work. To motivate our call for a re-evaluation and a change in direction, we show some examples on how a partial carrier sensing model can lead to conclusions that are different from those obtained under the 0-1 carrier sensing model. For example, an analytical conclusion from the 0-1 CS model is that link starvation and throughput unfairness can easily arise in many network topologies. We believe that due to partial carrier sensing, starvation and throughput unfairness are actually less common (at least less severe) in real networks.

REFERENCES

[1] B. Zhen, H. Li, S. Hara, and R. Kohno, "Clear Channel Assessment in Integrated Medical Environment," *Hindawi Publishing Corporation EURASIP Journal on Wireless Communications and Networking*, 2008.

[2] I. Ramachandran and S. Roy, "On the impact of clear channel assessment on MAC performance," in *IEEE Globecom*, 2006.

[3] J. Chong, Y. Sung, and D. Sung, "Analysis of CSMA/CA Systems under Carrier Sensing Error: Throughput, Delay and Sensitivity," in *IEEE Globecom*, 2008.

[4] J. Sheng and K. Vastola, "Physical Carrier Sensing Outage in Single Hop IEEE 802.11 Ad Hoc Networks with Slowly Moving Stations," in *IEEE WCNC*, 2008.

[5] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in Static Wireless Networks," in *Proceedings of ACM SIGCOMM*, 2006.

[6] L. Qiu, Y. Zhang, F. Wang, M. Han, and R. Mahajan, "An analysis of short-term fairness in wireless media access protocols," in *Proceedings of ACM MOBICOM*, 2007.

[7] W. Kim, J. Lee, T. Kwon, S. Lee, and Y. Choi, "Quantifying the interference gray zone in wireless networks: A measurement study," in *Proceedings of IEEE ICC*, 2007.

[8] J. Lee, S. Lee, W. Kim, D. Jo, T. Kwon, and Y. Choi, "RSS-based carrier sensing and interference estimation in 802.11 wireless networks," in *IEEE SECON*, 2007.

[9] "NS2," *http://www.isi.edu/nsnam/ns/*.

[10] S. Liew, C. Kai, J. Leung, and B. Wong, "Back-of-the-Envelope Computation of Throughput Distributions in CSMA Wireless Networks," *to appear in IEEE Transactions on Mobile Computing*.

[11] M. Durvy, O. Dousse, and P. Thiran, "Border Effects, Fairness, and Phase Transition in Large Wireless Networks," in *IEEE INFOCOM*, 2008.

[12] IEEE 802.11 Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE 802.11-1997.

[13] A. Kochut, A. Vasan, A. Shankar, and A. Agrawala, "Sniffing out the correct Physical Layer Capture model in 802.11b," in *Proceedings of 12th IEEE ICNP*, 2004.

[14] "MadWifi," *http://madwifi-project.org/*.

[15] "OmniPeek," *http://www.wildpackets.com/*.

[16] "IPerf," *http://www.noc.ucf.edu/Tools/Iperf/*.