Optimizing 802.11 Wireless Mesh Network Performance Using Physical Carrier Sensing

*Sumit Roy**, *Hui Ma, Rajiv Vijayakumar*
Dept. of Electrical Engineering
U. Washington Box 352500
Seattle, WA 98195

Jing Zhu
Communications Technology Lab, Intel Corpn.
2111 NE 25th St. , JF3-206
Hillsboro, OR 97124

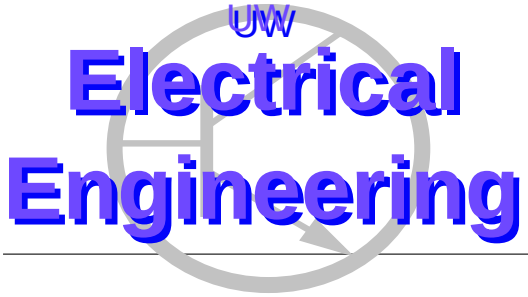

UWEE Technical Report
Number UWEETR-2006-0005
April 2006

Department of Electrical Engineering
University of Washington
Box 352500
Seattle, Washington 98195-2500
PHN: (206) 543-2150
FAX: (206) 543-3842
URL: http://www.ee.washington.edu

Optimizing 802.11 Wireless Mesh Network Performance Using Physical Carrier Sensing

Sumit Roy*, Hui Ma, Rajiv Vijayakumar
Dept. of Electrical Engineering
U. Washington Box 352500
Seattle, WA 98195

Jing Zhu
Communications Technology Lab, Intel Corpn.
2111 NE 25th St. , JF3-206
Hillsboro, OR 97124

University of Washington, Dept. of EE, UWEETR-2006-0005
April 2006

Abstract

Multi-hop ad-hoc networks suffer from the presence of hidden and exposed nodes that diminish aggregate network throughput. While there are various approaches to mitigating these problems, in this work we focus exclusively on the role of physical carrier sensing (PCS). Specifically, tuning the PCS threshold leads to a trade-off between the 'exposed' and 'hidden' areas. The contributions of this work are (at-least) two-fold:
i. we develop an analytical formulation (assuming constant rate nodes) to determine the optimal PCS threshold. It is shown that choosing *the carrier sensing range equal to the interference range* is the appropriate optimality principle for PCS based network optimization. An important component of this analysis is an *improved* link layer model in contrast to much of the existing literature, that leads to significantly *new and different* insights into enhancing 802.11 network performance. Further,
ii. in attempting to support the analysis via simulations with OPNET, several problems with various aspects of OPNET's implementations of the 802.11 link and MAC layers were exposed. Our description of these issues as well as fixes are likely to be of interest to the broader research community interested in using OPNET for their own investigations in 802.11 network performance[1].

Physical Carrier Sensing, Throughput Optimization, Hidden and Exposed Terminals, 802.11 WLAN, OPNET Simulations

1 Introduction

Physical carrier sensing (PCS) is one of the two main interference mitigation (contention resolution) mechanisms defined in the PHY/MAC layers of 802.11 WLANs - the other being RTS/CTS or Virtual Carrier Sensing (VCS). A node that intends to transmit first assesses the current channel state (this is generically termed as Clear Channel Assessment or CCA) by comparing the measured on-air received energy against a predefined PCS threshold to determine if it should contend for channel access as per the CSMA/CA protocol. Each node samples the *net* energy level on-air and initiates channel access only if the detected value is below the threshold, indicating that the channel is 'free' of significant ongoing transmissions.

*Author for All Correspondence: roy@ee.washington.edu, Ph: 206 221 5261
[1]The revised OPNET modules are available by request from the authors; see http://www.ee.washington.edu/research/funlab/pcs_model.htm

The PCS threshold effectively defines a carrier sensing *range* that denotes an area wherein a secondary transmitter is prevented from contending for access so as to not disrupt the reference transmission.

However (as is well known), sensing the channel at the transmitter's location is not always an accurate predictor of the channel state at the receiver, leading to the well known *hidden and exposed terminals*, both of which degrade aggregate throughput. By accepted definition, **hidden nodes** occur when *two transmitters* that are outside the mutual *carrier sensing range* attempt to transmit; if the secondary source lies *within* the interference range of a common (reference) receiver, it results in a packet loss (or 'collision'). Conversely, **exposed terminals** occur when a source lies *within* the sensing range of the reference transmitter but *outside* the interference range of the reference receiver. Hidden and exposed nodes lower aggregate network throughput in different ways: while, hidden nodes *disrupt* the reference transmission (lead to incorrect decoding or loss of the reference packet), exposed nodes represent *lost* throughput by needlessly suppressing simultaneous transmissions that could have occurred without disruption to the reference transmission. Restated, exposed nodes represent a loss of spatial reuse that is fundamental to improving aggregate network throughput[2].

The preferred approach to collision avoidance in the .11 literature has leaned towards use of RTS/CTS, i.e. virtual carrier sensing (VCS) whereby the reference sender-destination pair engage in a handshake to 'reserve' the channel prior to actual data transmission. All nodes within a *transmission range R_{tr} of both the sender and the receiver defer transmission* by this mechanism for the duration of the subsequent DATA and ACK; however it is easy to see that whenever the *interference range* $R_I > R_{tr}$, a secondary source within the interference range but outside the transmission range of the destination receiver will cause loss of reference packet even if RTS/CTS were employed. This point was investigated in depth by [1] who argued that VCS is optimal when $R_{tr} = R_I$, i.e., the transmission range equals the interference range. However, we underscore the difficulty of achieving this in practice due to the fact that R_I is implicitly a function of R_{tr} and cannot be independently tuned to achieve the desired equality. Further note that in the 802.11 standard, the use of RTS/CTS has been provided as an *option* for longer payloads only (i.e., is not recommended for short packets) so as to amortize the additional overhead due to the RTS/CTS exchange. In fact, RTS/CTS is triggered based on packet size and is not, generally speaking, under user control and thus cannot be relied on for network optimization. We refer the reader to [4] which shows several examples where use of RTS/CTS actually leads to lower throughput than if it were disabled.

For all these reasons, we have argued in prior work [16] for a deeper investigation of PCS as a *pragmatic preferred alternative* to VCS for optimizing aggregate network throughput. The fundamental premise is that PCS allows better (more fine grained) control and tuning of network performance than is possible through VCS for two reasons: i. PCS threshold is a parameter *independent* of transmit/interference range, and ii. PCS is anyway mandatory in .11. Several current .11 WLAN hardware/firmware support one or more parameters for PCS control; some of these are available at run-time for user definition via open source linux drivers such as IPW2200 [5] for Intel's Pro/Wireless chipsets.

In this paper, we develop a simple (yet effective) geometric model for the effect of hidden and exposed terminals in a spatially homogeneous (uniform node density) network that yields an optimization problem in terms of the PCS threshold. We show that *under suitable conditions*, the optimal PCS threshold is obtained when the *CS range equals the interference range* [12]. OPNET simulation results provide values of optimal PCS threshold setting in more general scenarios than is accurately predicted by our analytical model.

[2]In this work, we exclusively focus on aggregate network throughput as the metric; we defer considerations of important topics such as per node throughput fairness to later work.

2 Link Layer Modelling: Literature Review and Commentary

The above discussion identifies the four parameters central to CCA or PCS in an ad-hoc network: i) Node separation distance D, ii) Transmission Range R_{tr}, iii) Carrier Sensing Range R_{cs} and iv) Interference Range R_I. While the above quadruplet of parameters essentially characterize the conditions when a *transmit* node seeks channel access, we also introduce v) Reception Range R_{rx} to highlight its role in *packet reception* of a receive node [3].

A survey of the literature on simulation studies of .11 ad-hoc networks reveals a wide range of choices of the above parameters often without adequate justification which sometimes leads to misconceptions or unexamined assumptions. A further difficulty is the lack of *consistency* in definitions of the same terms across various authors (notably for R_{tr}, R_{cs}, R_I); hence we first establish definitional clarity and consistency in this regard.

2.1 CCA in 802.11 Systems

The contention resolution mechanism within DCF in the 802.11 standard states that a station may begin access only when the channel is 'sensed idle'. In the language of the 802.11 standard [19], the PHY layer performs a "Clear Channel Assessment" (CCA) and indicates channel status to the MAC layer via the CCA indicator which takes on the values BUSY or IDLE.

It is worth noting that the standard does *not* specify the algorithmic details of how the CCA is to be performed; this is left for vendor specific implementation and warrants scrutiny due to possible variations between different hardware. Nonetheless, two classes of CCA mechanisms are broadly accepted: a) Energy Detection (ED) or b) PHY Preamble Detect (PD). As the labels indicate, energy detect implies use of a (simple) non-coherent detector that estimates the *net* energy on the channel from *all* ongoing transmissions at that instant, and uses a threshold to set the CCA flag; preamble detect is typically implemented using a (sliding) correlator that is matched to the known training symbols/preamble within the PHY packet header, whose output at the appropriate sampling instant is then thresholded to set the CCA flag. There is a subtle but significant difference between these two mechanisms: ED is a universal, low-cost (in terms of power dissipation) and robust approach that does *not* require any knowledge of how many transmissions are ongoing (or indeed, the identity of the sources) [4]. PD on the other hand is typically accomplished using the PHY header and potentially allows separation of the various active sources since their respective PHY preambles arrive at the receiver at different time instants.

In this paper, we will use CCA/ED as equivalent to (physical) carrier sensing. Further, in several places in the 802.11 standards document [19], the term *carrier sensing* is actually used to refer to PD contributing to further potential confusion - we trust that our use of PCS is unambiguous. Since we rely on supporting OPNET simulations in this work, we note that OPNET CCA implementations exclusively rely on ED.

2.2 Link Layer Model Parameters and OPNET Verification

In view of the discussion in Sec.2, we first provide unambiguous definitions of key link layer parameters pertinent to PCS. Further, these models are all verified via simulations conducted with OPNET v.11 that provides an implementation of PHY/MAC layers for both .11b/g (2.4 GHz) and .11a (5 GHz). Our experiments were conducted only with .11a models by defining IP flows between any Source(S)-Destination (D) pair, i.e. the traffic generator produces Layer 3 packets (IP

[3] The role of R_{rx} will become clearer subsequently, particularly in relation to implementation issues in OPNET.

[4] Discovering the source identity following ED based on the MAC address can only occur upon subsequent decoding of the PHY payload.

Table 1: Rate-Range for .11a in OPNET (PER=10%)

Link Rate (Mbps)	Transmission Range (m)
6	304
12	216
24	90
54	39

datagrams of size 1500 bytes) that are input to the Layer 2 buffer at the source node. However all measurements reported - e.g. send rate (measured at the sender), receive rate (measured at receiver) etc. are obtained at the MAC layer (Layer 2), i.e. all OPNET simulation results for network throughput quoted are in terms of Layer 2 packets.

Defn. 1: The *transmission range* R_{tr} is defined by a sole Tx-Rx pair in the presence of noise only (i.e. no other concurrent transmission) such that the received power exceeds a SNR threshold S_0 given by

$$S_0 = \frac{P_{ref}/R_{tr}^{\gamma}}{P_n} \tag{1}$$

where P_{ref} is the transmit power at a reference distance (from Tx.), P_n is the additive background noise power and γ is the path loss exponent. It is worth (re) emphasizing that R_{tr} denotes the maximum link distance at which *packet decoding* is possible with high probability *at a given link rate*. Thus VCS is only meaningful within R_{tr} since RTS/CTS packets need to be successfully decoded.

The rate dependency of R_{tr} is obvious from Fig. 1 which shows the various transmission ranges as a function of link rate obtained from OPNET simulations conducted with a set-up similar to that shown in Fig 2 with a single reference (S1-D1) pair with variable separation D. Recall that the variable rate vs. range property is based on adaptive link layer modulation, i.e. the number of bits per symbol (equivalently, the link rate) conveyed over the link changes depending on the received SNR required to decode the packet. While conceptually link layer rate adaptation may be deemed continuous, in practice the link capacity is varied only over a discrete set (e.g. 6, 12, 24, 54 Mbps among others in 802.11a as shown in Fig. 1).

The range-rate curves in Fig. 1 were obtained as follows: the maximum throughput over the link is first obtained for a sufficiently small D by measuring the *receive rate* at the receiver (this should correspond to the chosen link capacity discounted by the MAC header overhead). Then D is gradually increased till 10% packet error rate (PER) relative to the maximum throughput is observed; the distance at which this occurs is defined as the transmission range for that link capacity. The results are listed below in Table 1 and effectively represents the hull of the various rate-range curves in Fig. 1. Since our subsequent simulations are conducted for link rate of $12\,Mbps$, we note the corresponding $R_{tr} = 216\,m$.

Defn. 2: *Carrier Sensing Range* R_{cs} is defined by

$$\gamma_{cs} = \frac{P_{ref}}{R_{cs}^{\gamma}} \tag{2}$$

i.e. the distance at which the receiver signal power from the source measured by ED equals the sensing threshold value γ_{cs}. The implication is that only one among all contending nodes within an area defined by R_{cs} may transmit, and the others defer transmission via the CSMA mechanism in .11 DCF whenever the received power as measured by ED mechanism exceeds γ_{cs}.

The OPNET set-up for estimating R_{cs} is shown in Fig.2 where the separation DIS between the two sources S1 and

S2 is varied to determine R_{cs} (the value of $D < R_{tr}$ is immaterial to determining R_{cs} and was chosen as D = 100 m in our experiment). When DIS $< R_{cs}$, the transmit nodes contend for the channel via CSMA/CA. Since each of S1, S2 transmits at the same rate, they get an equal time share of the channel. However when DIS $> R_{cs}$, both flows can exist simultaneously although the receive throughput on each is less than the maximum achievable for a solitary flow due to the mutual interference as seen at the respective receivers. Thus, the value of DIS at which this transition occurs allows us to identify R_{cs}. Our experiment in OPNET for 12 Mbps link rate with the threshold $\gamma_{cs} = -95\, dBm$ yielded $R_{cs} = 260\, m$.

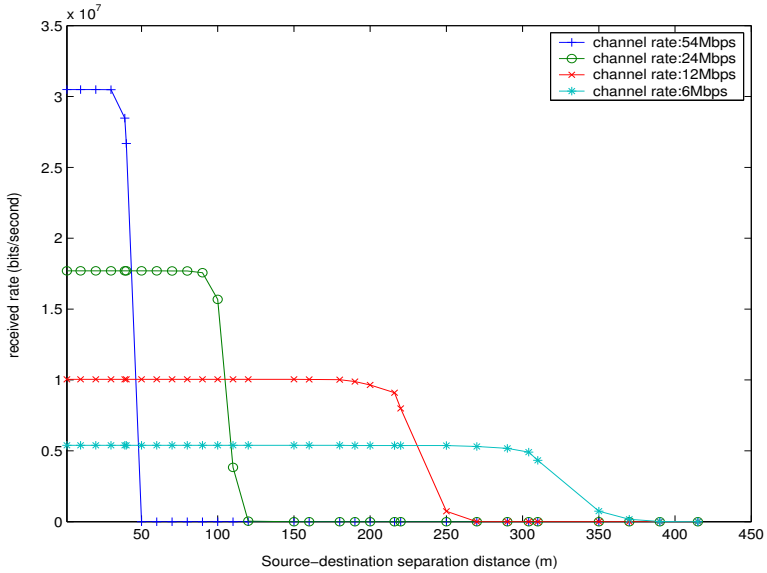

Figure 1: Transmission Range in OPNET (11a)

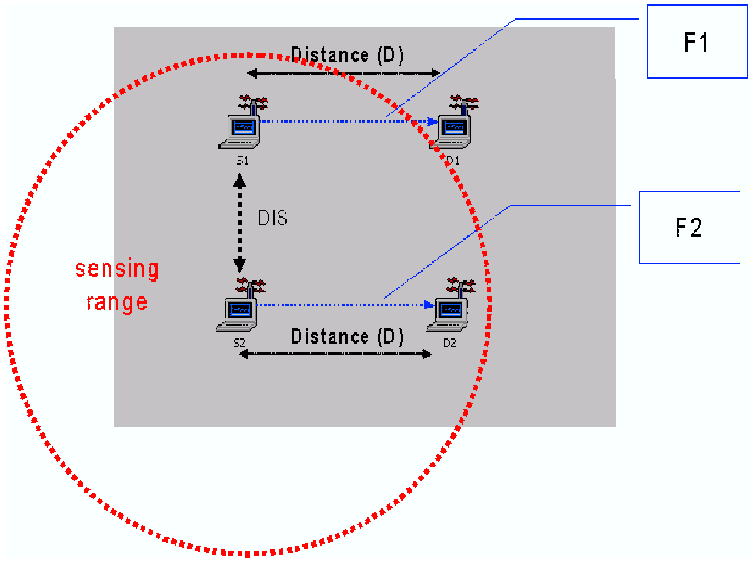

Figure 2: Carrier Sense Range setup in OPNET

Defn. 3: *Interference Range*

For nominal Tx (S1)-Rx(D1) separation D (where $D < R_{tr}$) and ONE interfering concurrent source S2 at a distance DIS from the receiver R1 as shown in Fig.3, the signal-to-interference cum noise ratio SINR at R1 is given by

$$SINR(R1) = \frac{P_{ref}/D^{\gamma}}{P_n + P_{ref}/\mathrm{DIS}^{\gamma}} \tag{3}$$

The interference range R_I is the value of DIS whereby the $SINR(R1)$ from Eq.1 equals S_0 in Eq.3. Equating Eq.1,3 yields after some simple manipulations

$$R_I = \left(\frac{P_{ref}/P_n}{(R_{tr}/D)^{\gamma} - 1}\right)^{\frac{1}{\gamma}} = (S_0)^{\frac{1}{\gamma}} \frac{R_{tr}}{((R_{tr}/D)^{\gamma} - 1)^{\frac{1}{\gamma}}} \tag{4}$$

Eq.4 provides a *theoretical* expression for R_I based on the amount of link margin available when $D < R_{tr}$; i.e., only a secondary transmitter within a radius R_I of the intended *receiver* will disrupt the reference transmission (i.e. lead to loss of reference packet). It can be shown by taking derivative of Eq.4 that R_I is monotonic increasing with D, implying that increasing $D \rightarrow R_{tr}$ results in the link becoming more vulnerable to interference as can be expected. Specifically, as $D \rightarrow R_{tr}$, $R_I \rightarrow \infty$ implying the loss of all link margin in the limit and hence ANY concurrent transmitter (at whatever distance from the reference receiver) causes the reference packet to be dropped as the receiver SINR drops below the threshold needed for packet decoding.

R_I is estimated in OPNET via the experimental set-up depicted in Fig.3. A reference flow F1 (S1-D1) is established for a given separation D and the throughput at D1 measured. A second flow F2 (S2-D2) is then introduced as shown where the distance S1-S2 exceeds the carrier sensing range. The minimum value of DIS at which the throughput of F1 shows a 10% PER relative to the measured throughput in the absence of F2 is determined to be the interference range. The results are shown in Fig.4 with Eq.4 included for comparison; the functional dependance of R_I on D as predicted by Eq.4 is confirmed.

Of particular interest is the fact that for $D << R_{tr} = 216\,m$, R_I is approximately proportional to D as shown Fig.5 (which is merely an expanded version of Fig.4 for lower range of D). Note also that in this region, the interference range $R_I < R_{tr}$ which belies the folk theorem that R_I always exceeds R_{tr}; the above in fact is a necessary condition for the phenomena of 'stronger last collisions' [15] that has been observed in practice.

This is also a good occasion to review and comment on the values of R_{tr}, R_{cs}, R_I that have been adopted to date in the literature for performance evaluation of MESH networks. It is likely that the choice of parameter values used by various authors have been influenced by the default settings in the simulation tools of their choice, e.g. ns-2 and Qualnet. For example, the values used in [9], [4] for experimentation ($R_{tr} = 250\,m, R_{cs} = R_I = 550\,m$) are the default in ns-2 802.11 simulator and may suggest (by incorrect generalization) that $R_{tr} \leq R_I \leq R_{cs}$. Alternately, [2] states that ".. for open space environment, the interference range is 1.78 times the transmitter-receiver distance" (and as a corollary, RTS/CTS does not work when $D > \frac{R_{tr}}{1.78}$). This fact has been used by other authors subsequently such as [12] for performance studies using the Qualnet simulator.

Both the above assumptions are worthy of closer scrutiny since in general, these have significant impact on network performance. While it is certainly reasonable that $R_{tr} < R_{cs}$ in general, an inequality such as $R_{tr} < R_I$ is *not* true under all conditions, as we show conclusively. Further, statements such as $R_I \approx 1.78 R_{tr}$ leave (perhaps unwittingly) an impression that R_I is typically a constant (2-3) times the transmit range R_{tr}, and further, there exists a linear relationship

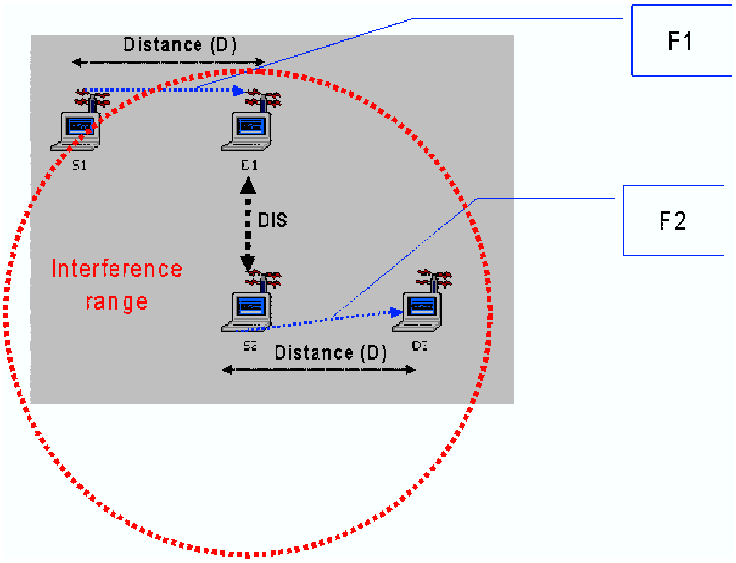

Figure 3: Interference Range setup in OPNET

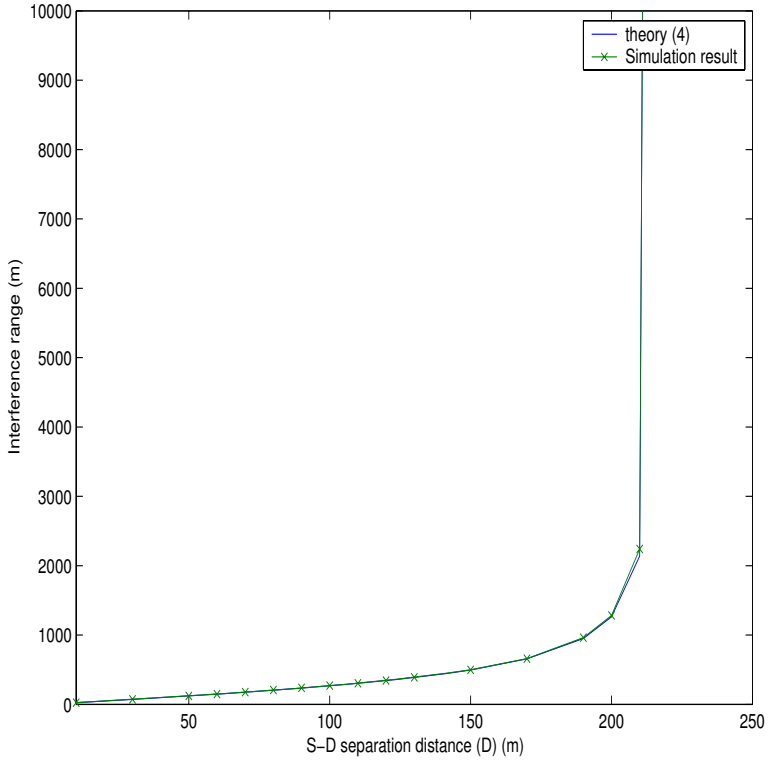

Figure 4: Estimating Interference Range: $R_{tr} = 216\,m$

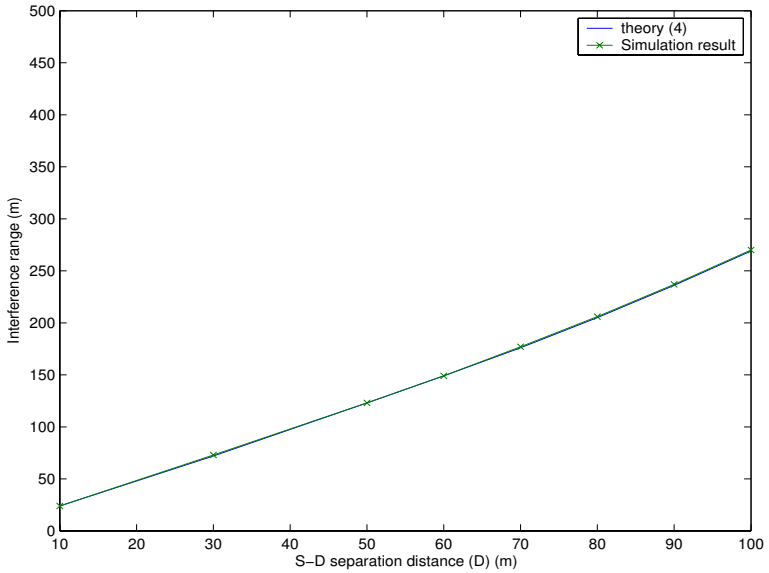

Figure 5: Linearity of R_I vs. D, $D << R_{tr}$

between the two variables. This is certainly untrue; as we show, when $D \rightarrow R_{tr}$, R_I increases rapidly to infinity ! A look at the derivation of the above result in [2] reveals the root of the problem; they ignore the presence of background noise and base the interference range on a threshold of purely the received SIR and not received SINR, since it is implicitly assumed that the noise power is negligible compared to the interference power. Such an approximation is only supported when the interference range is small, or equivalently when $D << R_{tr}$ (i.e. the S-D separation is only a fraction of the transmission range) and not otherwise.

3 Default OPNET .11 Link Layer Implementation: A Critical Review

The commentary in this section is aimed at educating users who are contemplating use of 802.11 PHY/MAC as implemented in OPNET v.11. This is timely and useful because of some important shortcomings of OPNET 802.11 code which was discovered during our investigations that are pertinent to a) physical carrier sensing and b) packet decoding. Since these are potential contributors to (significant) deviations between OPNET simulation results from those based on analytical models, we have had to identify these 'bugs' and provide fixes in the new OPNET codes.

At the outset, it is useful to reiterate that OPNET represents the PHY as pipelined 'black-box' stages, whose input-output characteristics are based using analytical (equivalently, table lookup) functions. We state a few basic facts next about OPNET link layer implementations.

1. In the absence of any ongoing transmission, the receiver at any node observes thermal noise; the thermal noise power level is constant and set at $P_n = -101$ dBm for the 20 MHz bandwidth of an 802.11a channel.

2. If a source node i is transmitting, OPNET computes the *average* power at the reference receiver due to this transmission as a constant $P_r(i)$ based on a) the transmit power of the source node i (which is configurable) and b) the distance between the source i and the reference node. The relation used for this computation in default OPNET is

the standard Friis path loss model for RF propagation with a path loss exponent of $\gamma = 2$, i.e. free space path loss (which is clearly inappropriate for typical indoor scenarios). *The default OPNET model also does NOT include any short-term (multipath) fading*.

3. Although CCA and packet decoding at a node are two distinct functions (the first relates to a transmit node seeking channel access, the later relates to PHY layer reception at a receive node) that should logically be triggered by *two* distinct thresholds (i.e. the PCS threshold γ_{cs} and the receiver sensitivity P_{rx}), OPNET defines *only a single threshold* (called 'packet reception power threshold' in OPNET) that is used for both purposes. As we explain below, this (and other OPNET implementation artifacts) leads to significant performance deviations and necessitated modifications to the relevant OPNET PHY modules for the work reported here. All results presented here are thus based on modified OPNET code [5]. This is consistent with several 802.11 hardware designs such as Wavelan-II [8] and Intersil for client-side adapter cards and the Alcatel OmniAccess APs that use two thresholds for transmit and receive.

3.1 CCA in Default OPNET

OPNET is a discrete-time event-driven simulator and whenever a node transmits a packet it generates an interrupt at every other node in the network. This behavior allows every node to *know exactly when a new transmission begins and the associated source ID (along with the received power which is computed as described above), which is not possible in physical systems*. Thus, any OPNET node seeking channel access uses the configurable parameter 'Packet Reception Power Threshold' P_R to set the CCA flag busy whenever the received power from any other single transmit node i exceeds P_R, i.e.,

$$P_r(i) > P_R \tag{5}$$

3.2 Packet Reception in Default OPNET

The following are *necessary* (but not sufficient) conditions for packet detection at a receive node from a source s in default OPNET v.11 implementation -

1. The CCA flag must be IDLE initially(i.e., the node cannot be contending for channel access);

2. The flag remains idle till a packet is received from a node s for which the received power $P_r(s) > P_R$; in that case, the CCA flag is set busy and packet detection is triggered [6]; the receive node is then 'locked' into decoding this first arrival;

3. If a subsequent packet (from a node i) arrives within the duration of the reference packet from s, it is treated as interference; a) should $P_r(i) > P_R, i \neq s$, this leads to an automatic collision and the reference packet is lost; b) if $P_r(i) < P_R$, the node conducts packet decoding as described next;

4. The decoding of the reference packet is conducted by several pipelined stages in OPNET. One of the pipeline stages divides the reference packet duration into several sub-segments with constant SINR; for each such segment, the

[5] A catalog of differences in results from identical simulations performed with default OPNET v.11 and UW-modified OPNET v.11 is available at http://www.ee.washington.edu/research/funlab.

[6] After the busy timer of a receiver expires, the CCA flag will revert back to idle.

effective SINR is computed and used in a table lookup for the given uncoded modulation scheme to determine the number of bit errors in that segment. This process is continued over all the pipeline stages to find the cumulative number of bit errors in the entire packet; if this exceeds a threshold value, the packet is deemed to be lost[7].

3.3 Comments on Some Aspects of Default OPNET v.11 Implementations and Resulting Modifications

The above discussion, although brief, suffices to highlight multiple problems with default OPNET 802.11 implementations; these, along with our modifications, are listed below in no particular order.

I. CCA

1. The default CCA/ED module Eq (5) does *not* aggregate the power from (multiple) ongoing transmissions for ED [8]. The corrected rule is now

$$\sum_i P_r(i) + P_n > \gamma_{cs} \tag{6}$$

When the threshold γ_{cs} is exceeded, the CCA flag is set busy. The l.h.s is computed at each instant when a (new) packet transmission begins and ends, and the flag remains busy as long as Eq. 6 is satisfied; it reverts back to idle when the sum power drops below γ_{cs}.

2. Along with the above, *two* distinct thresholds γ_{cs} and P_R are now introduced in the modified PCS module respectively - the second threshold P_R is used to set the PHY RX flag described next.

3. The fact that packet detection is triggered using the *same* threshold as CCA in default OPNET has implications for network performance evaluation. Consider the following: if the CCA is already busy (i.e. the receiver is already attempting to decode a reference packet), any subsequent new packet arrival whose received power exceeds P_R is deemed to automatically result in a collision, i.e. the reference packet will be dropped. Clearly, this is unduly conservative; the condition for decoding a packet should be based on a received SINR threshold, and not on a threshold on the interference power. Re-stated, it is possible in principle that the receiver with CCA busy can still detect an incoming packet, but this was not possible in default OPNET v.11.

II. Packet Reception

The following are revised *necessary* (but not sufficient) conditions for the packet to be received successfully at a receive node from a source s:

1. The PHY RX flag at the receive node is IDLE initially (i.e., the received power of no single packet exceeds P_R);
2. The PHY RX flag remains idle till a packet is received from a node s for which the received power $P_r(s) > P_R$; in that case, the PHY RX flag is set to busy and packet detection is triggered; the receive node is then 'locked' into decoding this first arrival;
3. If a subsequent packet (from a node i) arrives within the duration of the reference packet from s, it is treated as interference; but irrespective of the received power of the interference packet, the packet decoding pipelined chain is always triggered.

[7] The default error correction code threshold value in OPNET v.11 for the number of code errors is 0, implying no error correction coding is accounted for in the link layer model.

[8] To the best of our knowledge, Intel 2915 cards implement CCA per Eq.(6) and thus OPNET based simulation results can be expected to differ from those in a test-bed using the above hardware.

We reiterate two important changes between the modified and default OPNET implementation for packet detection. First, packet detection is triggered by PHY RX flag set to busy which is *different* from the CCA busy condition in the modified module. Second, when the receiver is attempting to decode a reference packet, the arrival of a subsequent packet whose received power $P_r(i) > P_R$ does not automatically imply a collision.

III. Pipelined Packet Decoding

In the pipelined packet decoder in default OPNET v.11, the decision to drop a packet is based on a table-lookup BER. Unfortunately, this is incorrectly done - it uses the symbol SNR whereas $\frac{E_b}{N_0}$ should be properly used (in other words, the symbol SNR should be converted to bit SNR taking into account the modulation scheme). This bug has also been corrected in our modified modules. Further, as has already been remarked, the link layer is presumed to be uncoded; future link models should account for the coding gains due to the error control coding schemes implemented in 802.11 standard.

3.4 Receiver Sensitivity and 'Stronger Last' Collisions

As has already been described, our modified OPNET PCS module introduces a separate (from the threshold γ_{cs} for CCA) receiver sensitivity threshold P_R that is applied to the received *power* from any source - the receiver will attempt to decode the packet from that source only when the received power exceeds P_R. This defines a receiving range R_{rx} given by

$$R_{rx} = (\frac{P_{ref}}{P_R})^{1/\gamma} \tag{7}$$

Obviously, it follows that $D < \min.(R_{rx}, R_{tr})$ is necessary and typically P_R is set such that $R_{rx} < R_{tr}$.

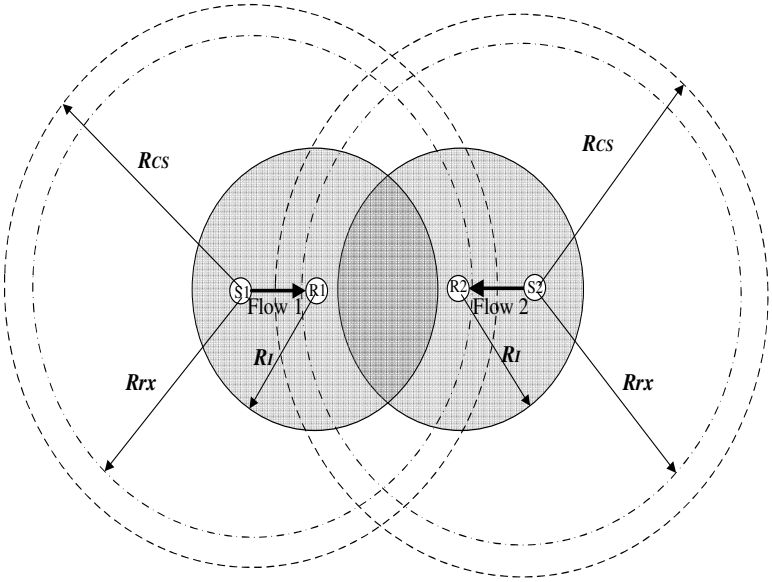

Figure 6: Stronger Last Collisions

We now demonstrate that this added degree of freedom via P_R can be used to counteract the phenomena called "stronger-last" collision [15]. This is a variation of the hidden node problem where two sources that are outside R_{cs} distance of each other simultaneously transmit. However, it differs from the conventional hidden terminal problem that assumes two sources transmitting to a *common* receiver; as shown in Fig. 6, the two sources S1 and S2 transmit to their

respective receivers R1, R2. The reference packet from S1 to R1 may be lost even if S2 lies *outside* the interference range of R1, provided the transmission from S2 *precedes* that from S1 to R1. The nomenclature indicates the fact that the weaker (from more distant source) and undesired packet arrives first at R[9] and the stronger packet comes second. The receiver is locked initially to whichever packet arrives first, even if it is *undesired* and thus loses out on receiving the desired packet that arrives later.

The necessary condition for such collisions to occur is :

$$R_I < min.(R_{tr}, R_{rx}) = R_{rx} \tag{8}$$

The above condition is not commonly valid since typically $R_I > max.(R_{tr}, R_{rx})$ whenever D is a significant fraction of (R_{rx}, R_{tr}) [10]. In a grid network where all link distances equal D, 'Stronger Last' Collisions do not occur whenever $R_{cs} \approx R_I$—which we will show to be the optimal carrier sensing range—since with $R_{rx} \leq R_{cs} \approx R_I$, the necessary condition for Stronger Last Collisions in Eq. 8 will not hold. However, when $R_{cs} > R_I$, then $R_{rx} > R_I$ is possible, leading to potential for Stronger Last Collisions.

Clearly, setting the reception sensitivity threshold P_R sufficiently high such that $R_{rx} < R_I$ avoids stronger last collision; hence the signal from S2 will not trigger packet reception at R1 since R1 is now outside R_{rx} range of S2 and a subsequent transmission from S1 to R1 can be correctly decoded by R1. In all the simulation results reported later, the above was implemented for better match with our analytical estimates, as our analytical model does not account for stronger last collision phenomena.

4 Optimizing Physical Carrier Sensing for Homogeneous Networks: Analytical Model

We make the following assumptions for our analytical model in this paper :

1. All sources have identical, fixed transmission power 0 dBm (1 mW)
2. The channel between any two nodes is identical and non-fading. The mean received signal power is related to the transmit power by a standard propagation power law characteristic (path loss exponent $\gamma = 2$).
3. The link capacity (W) is identical for each link. All simulations are thus conducted on a uniform 2-D grid topology with identical 1-hop separation D between any S-D pair.
4. Network is large enough to ignore edge effects and is homogeneous, i.e., it is spatially uniform with a fixed density. Thus the net interference environment as seen at any station is the same, on average.
5. The presence of a hidden terminal impacts aggregate network throughput the same as an exposed node; the latter prevents a rate W transfer while the former wastes a successful transmission opportunity due to a collision.

We now proceed to illustrate the trade-off between hidden and exposed terminals via a concrete formulation. In Figure 7, A_H indicates "hidden terminal" area, which lies within the interference range of R, but outside the carrier sensing range of S; A_E indicates "exposed terminal" area, which is outside the interference range of R, but within the carrier sensing range of S.

[9] Note that the received signal power at R from both the packets exceeds P_R as is necessary.
[10] However, the possibility of stronger last collision is increased in a *dense* network where $D << R_{tr}$ on average.

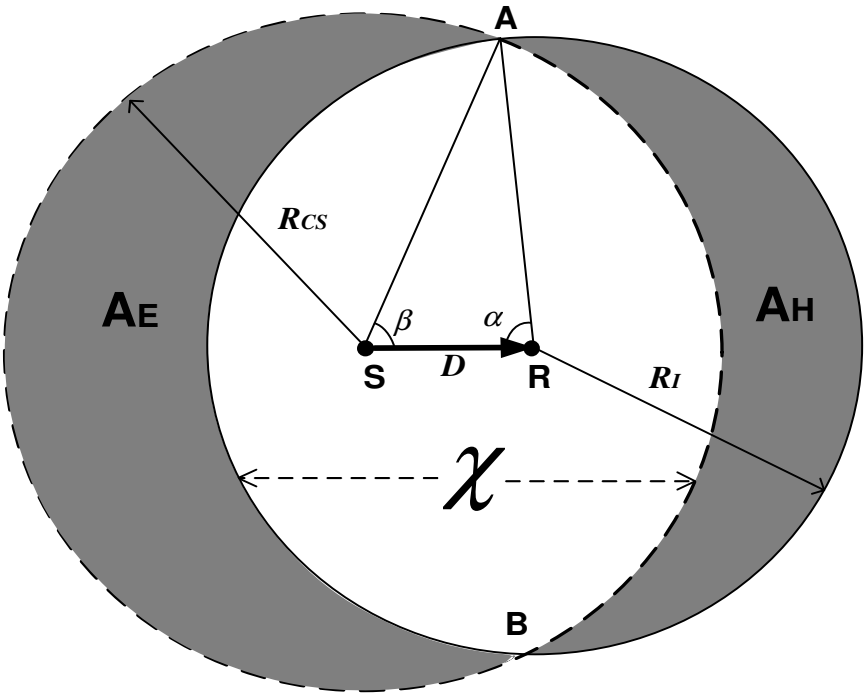

Figure 7: Hidden and Exposed Area Geometry

Using the definition of the area χ (area defined by the intersections of the two circles) as shown in Fig. 1, it follows that

$$\begin{aligned} \mathcal{A}_E + \chi = &\quad \pi\, R_{cs}^2 \\ \mathcal{A}_H + \chi = &\quad \pi\, R_I^2 \end{aligned} \tag{9}$$

from which it readily follows that $A_H - A_E = \pi\, \left(R_{cs}^2 - R_I^2\right)$. Thus, it is clear that $A_H \gtreqless A_E \;\leftrightarrow\; R_{cs} \gtreqless R_I$.

The area χ is divided into two regions by the segments AR and RB; let χ_1 denote the (area of) the region which includes S, and let χ_2 denote the (area of) the region that does not include S. We have $\chi = \chi_1 + \chi_2$. Then, from elementary geometry

$$\begin{aligned} \chi_1 = &\quad\quad \alpha\, R_I^2 \\ \chi_2 = &\quad \beta\, R_{cs}^2 - 2\Delta SAR \end{aligned} \tag{10}$$

where the angles α, β are defined by the triangle SAR and ΔSAR denotes the area of the triangle. Using the law of cosines,

$$\begin{aligned} R_{cs}^2 &= D^2 + R_I^2 - 2D\, R_I\, cos\, \alpha \\ R_I^2 &= R_{cs}^2 + D^2 - 2R_{cs}\, D\, cos\, \beta \end{aligned} \tag{11}$$

and $2\Delta SAR = D\,R_I sin\alpha$. Hence

$$\chi = \alpha\,R_I^2 + \beta\,R_{cs}^2 - D\,R_I\,sin\alpha \tag{12}$$

From Fig.7, the trade-off between the hidden A_H and exposed A_E areas is intuitively obvious for $R_I - D < R_{cs} < R_I + D$ such that either of the two areas is not a proper subset of the other. For a given R_I, as R_{cs} is increased, the hidden area reduces while the exposed area increases.

Assumption (5) suggests a cost function using the total area of hidden and exposed terminal area:

$$F = A_H + A_E = \pi R_{cs}^2 + \pi R_I^2 - 2\,\chi \tag{13}$$

The carrier sensing range R_{cs} is the design parameter. Hence, the problem can be simply formulated as:

$$\text{Minimize} \quad F, \quad \text{subject to} \quad R_{cs} \geq 0 \tag{14}$$

Below we will find analytical solution of Eq. 14 for the case $D << R_{cs}, R_I$.

Substituting χ and R_{cs} in Eq. 13 with Eq.12 and Eq. 11, we have

$$F = (\pi - 2\beta)(D^2 + R_I^2 - 2D\,R_I\,cos\,\alpha) + (\pi - 2\alpha)R_I^2 + 2D\,R_I\,sin\alpha \tag{15}$$

which can be rewritten as

$$\begin{aligned} F &= -(\pi - 2\alpha - 2\theta)(D^2 + R_I^2 - 2D\,R_I\,cos\,\alpha) + (\pi - 2\alpha)R_I^2 + 2D\,R_I\,sin\alpha \\ &= -(\pi - 2\alpha - 2\theta)(D^2 - 2D\,R_I\,cos\,\alpha) + 2\theta R_I^2 + 2D\,R_I\,sin\alpha \\ &\approx -(\pi - 2\alpha)(D^2 - 2D\,R_I\,cos\,\alpha) + 2\theta R_I^2 + 2D\,R_I\,sin\alpha \end{aligned}$$

where θ is the angle of SAR, which can be ignored as being order of magnitude smaller than α.

Because $D << R_{cs}, R_I$ and $|R_{cs} - R_I| \leq D$, we have

$$2\theta R_I^2 \approx 2sin\theta\,R_I R_{cs} = 2D\,R_I\,sin\alpha \tag{16}$$

Thus, Eq. 14 can be rewritten as

$$\text{Minimize} \quad F(\alpha), \quad \text{subject to} \quad 0 \leq \alpha \leq \pi \tag{17}$$

where

$$F(\alpha) = [-(\pi - 2\alpha)(D - 2R_I\,cos\,\alpha) + 4R_I\,sin\alpha]D \tag{18}$$

Taking 1st and 2nd derivative of the above w.r.t α, we have

$$\frac{d}{d\alpha}F(\alpha) = 2D[D - 2R_I sin\alpha(\pi/2 - \alpha)] \tag{19}$$

$$\frac{d^2}{d\alpha^2}F(\alpha) = 4DR_I\,[sin\alpha - (\pi/2 - \alpha)cos\alpha] \tag{20}$$

With $D << R_I$, solving $\frac{d}{d\alpha} F(\alpha) = 0$, we have

$$\alpha_1 = cos^{-1}\left(\frac{D}{2R_I}\right) \tag{21}$$

$$\alpha_2 = sin^{-1}\left(\frac{D}{\pi R_I}\right) \tag{22}$$

Because $\frac{d^2}{d\alpha^2} F(\alpha_1) > 0$ and $\frac{d^2}{d\alpha^2} F(\alpha_2) < 0$, the global minimum of F is achieved when $cos\alpha = cos\alpha_1 = \frac{D}{2R_I}$, which corresponds to $R_{cs} = R_I$. The global minimum of F is given by

$$F_{min} = 4R_I\, sin\alpha_1 D = 8\Delta SAR = 2D\sqrt{2}\,4R_I^2 - D^2 \approx 4D\, R_I \tag{23}$$

Thus,when $R_{cs}, R_I >> D$, the optimal PCS threshold should be set such as $R_I = R_{cs}$, or equivalently $A_E = A_H$. Note that, remarkably, this equality condition is the default setting in ns-2; however, the conditions for which this is appropriate is now exposed by this analysis.

4.1 Numerical results

Our analytical solution for optimal PCS threshold in previous section assumes that $R_{cs}, R_I >> D$. Here, we turn to numerical computations using MATLAB to determine the optimum PCS threshold in more general conditions. We search for the minimum of F for a range of values of S_0 and γ with $D = 1$.

Fig. 8 shows the effect of varying sensing range with the path loss exponent γ between 1.5- 5.0 and fixed $S_0 = 6dB$. The resulting F is normalized by F_{min}. $R_{cs} = R_I + D$ represents the case of exposed terminals only (no hidden terminals) and $R_{cs} = R_I - D$ represents hidden terminal only (no exposed terminals), corresponding to the cost function F given by A_H and A_E as two extreme cases, respectively. Note that the optimal point is achieved in between these two extreme cases ($D = 1$, and $D = -1$), close to $R_{cs} - R_I = 0$ as suggested by our analysis (i.e. balance between exposed and hidden areas).

To see how S_0 and γ impact the optimal R_{cs} (denoted as R_{csopt}), we use global search; the results are shown in Fig.9. Clearly, higher the SIR threshold S_0 or smaller the path loss exponent γ, the greater the interference range and the closer the optimal sensing range R_{csopt} is to the interference range R_I.

In the next experiment, we investigate the sensitivity of using $R_{cs} = R_I$ on network performance even when it is not optimal. Fig.10 shows $(F - F_{min})/F_{min}$ as the function of S_0 for various values of γ. The additional hidden+exposed terminal area introduced by using $R_{cs} = R_I$ is smaller than 10% when $S_0 \geq 5$dB; this drops to 2.5% when S_0 increases to 20dB.

In Fig. 11, the sensitivity to separation D is explored for various S_0, γ. While it increases with γ and lower threshold S_0, it is nonetheless within 5% of the optimal. In summary, the condition $R_{cs} = R_I$ appears to be a robust initialization point for adapting to network conditions at run-time.

4.2 OPNET Simulation Experiments

Our analysis in the preceding sections aims to maximize throughput indirectly by minimizing the sum of the hidden and exposed terminal areas which in turn minimizes losses or wasted opportunities. In this section we use OPNET simulations to directly study the effect of modifying the carrier sense range on the network throughput.

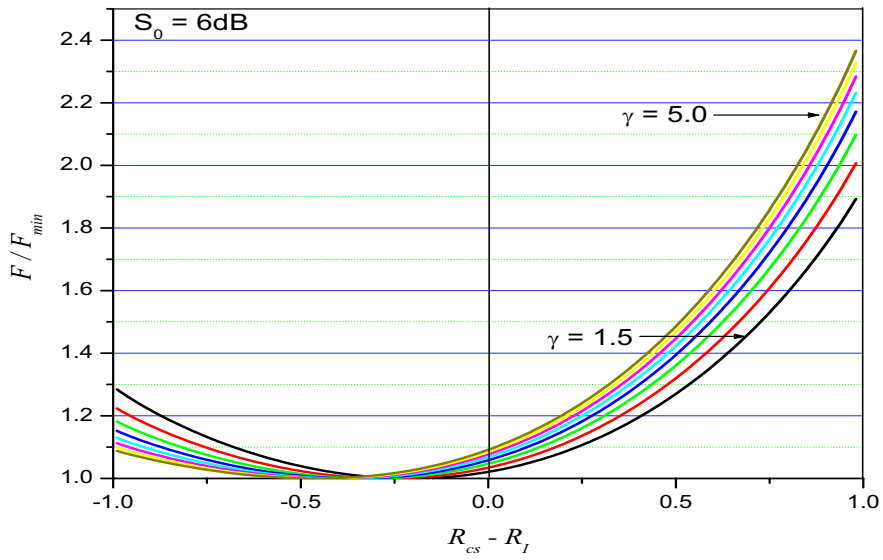

Figure 8: The effect of sensing range R_{cs}

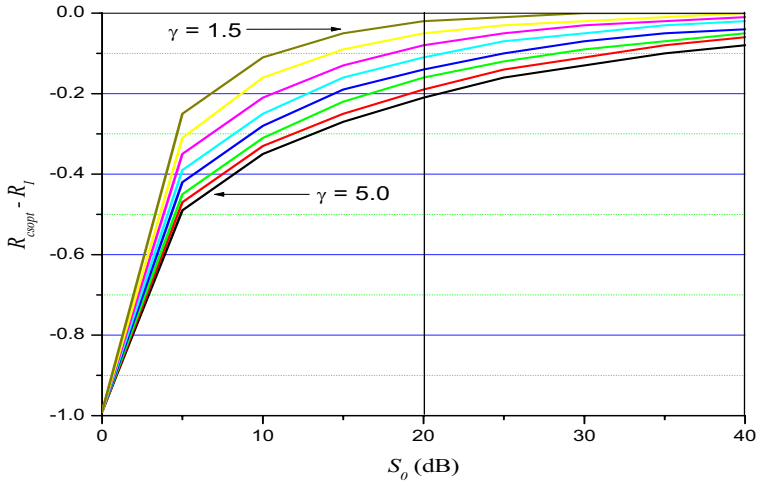

Figure 9: Optimal sensing range R_{cs}

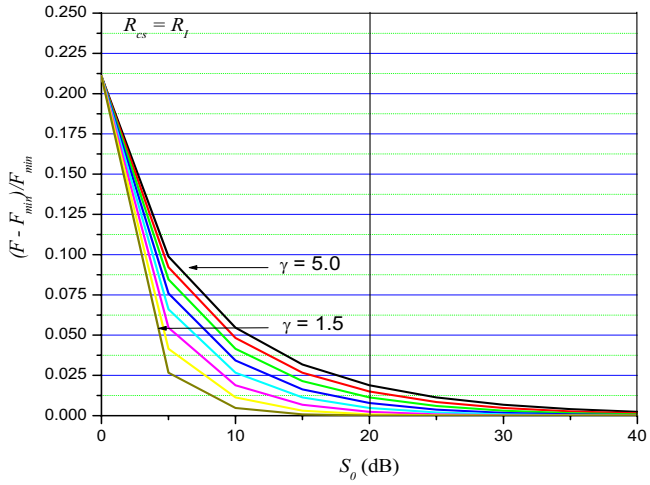

Figure 10: Relative Sensitivity to PCS Threshold for $R_{cs} = R_I$

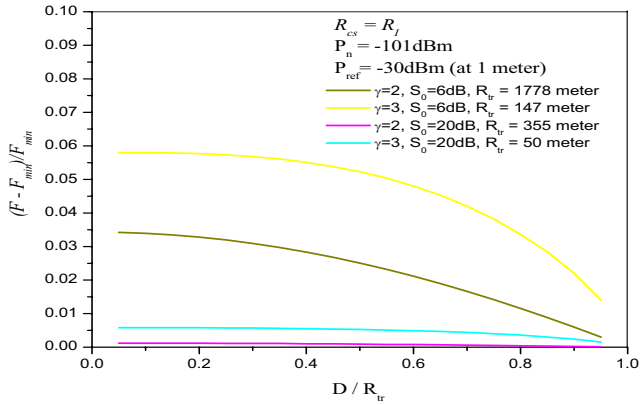

Figure 11: Relative Sensitivity to separation D for $R_{cs} = R_I$

4.2.1 Setup

OPNET simulations were run on a 10x10 square grid of nodes with a grid spacing of 10m; see Figure 12. The reception sensitivity was set such that the reception range was 10m; thus a node can only receive packets from its "one-hop" neighbors, i.e. those neighbors that are 10m away. The physical layer used was 802.11a at 12 Mbps. The carrier sense range R_{cs} was varied from 10m to 128m; $R_{cs} = 128$ ensures that all nodes in the simulation are within carrier sense range of each other.

A one-hop traffic flow was set up on each edge of the grid in both directions, for a total of 360 flows. Each flow consisted of a Poisson stream of packets generated directly at the IP layer, i.e. no transport protocol was used. The packet size (including IP headers) was a constant 1500 bytes.

For each carrier sense range, the offered rate of all flows was simultaneously increased until the fraction of offered packets network-wide which got dropped reached 10%. Note that a packet can be dropped for two reasons: either because the MAC layer buffer (which has a capacity of 21 packets) is full and cannot accept another packet from the IP layer, or because the number of retransmission attempts for the packet exceeds the retransmission limit of 7. We denote the highest offered rate (per flow) for which the packet drop rate stayed below 10% by T_{max}. Since T_{max} is the per-flow throughput and there are 360 flows in all, this means that when the offered traffic per flow is T_{max}, the carried traffic network-wide is $0.9 \times 360 \times T_{max}$.

T_{max} is an upper bound to the maximum traffic that can be carried simultaneously on each link while maintaining a low packet loss rate. We use T_{max} as our metric for comparison of network performance across various values of R_{cs}. Since

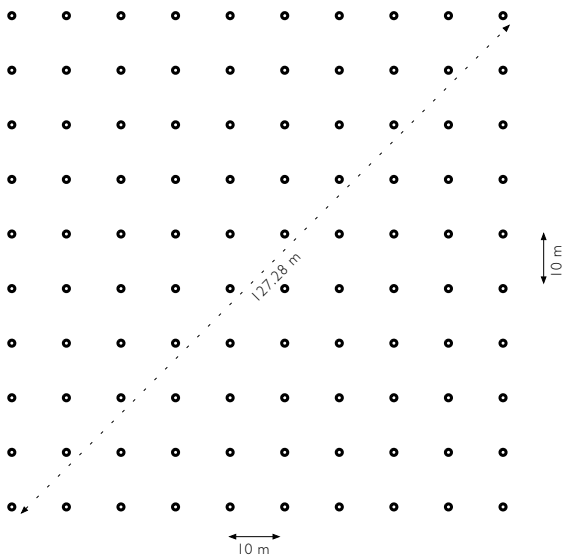

Figure 12: The 10x10 grid topology used for OPNET simulations.

we increase rates on all links simultaneously, we are essentially considering a "fair" scenario in which all links are used equally. Our interest is in the variation of the throughput metric with R_{cs}, and as such, other choices of metric could also have been be used. In particular, the choice of a packet loss rate of 10% is somewhat arbitrary; it was chosen to provide a metric that can be measured reliably without excessively long simulation runs. We also chose *not* to use "saturated" sources, since saturated sources are known to cause unfairness in topologies where not all nodes can sense each other.

4.2.2 Results

T_{max} is plotted as a function of R_{cs} in Figure 13.

The maximum value that T_{max} attains in Figure 13 is 104 kbps when $R_{cs} = 29$m. This optimal carrier sensing range is slightly higher than the interference range, which is 24m from the results in Figure 4. For $R_{cs} = 11$ m, T_{max} drops significantly to 24kbps. At the other end of the scale, when all nodes can sense each other (i.e. R_{cs}= 128 m), T_{max} drops to 50 kbps, which is about 50% of the maximum achievable rate.

We also noted in the simulations that, for large R_{cs} values, packet losses are almost entirely due to MAC layer buffer overflow. Conversely for low R_{cs} values, packet losses are almost entirely due to the number of retransmissions exceeding the retry limit, and not due to MAC layer buffer overflow.

We discuss these results in more detail below.

4.2.3 Discussion

In Figure 14, we show a source-destination (S-D) pair in the grid, along with all the nodes within interference range of D; there are 19 such nodes. The nodes are labelled according to their distance from the *sender*; the three nodes labelled "1" are closest to S (at a distance of 10m) and form the first "tier" of potential interferers, while the two nodes labelled "7" are furthest from S (at a distance of 31.62m) and are the last tier of interferers. As R_{cs} is increased starting from 10m,

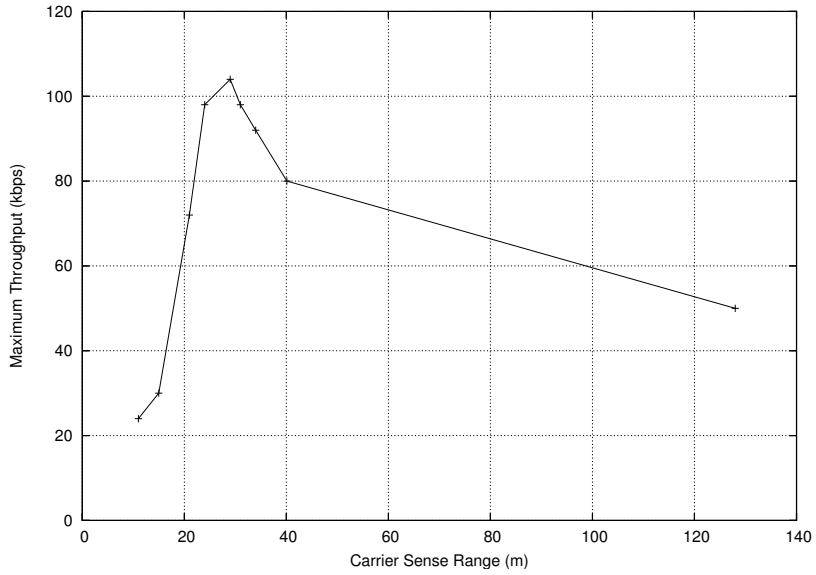

Figure 13: T_{max} as a function of R_{cs} from OPNET simulations on a 10x10 grid

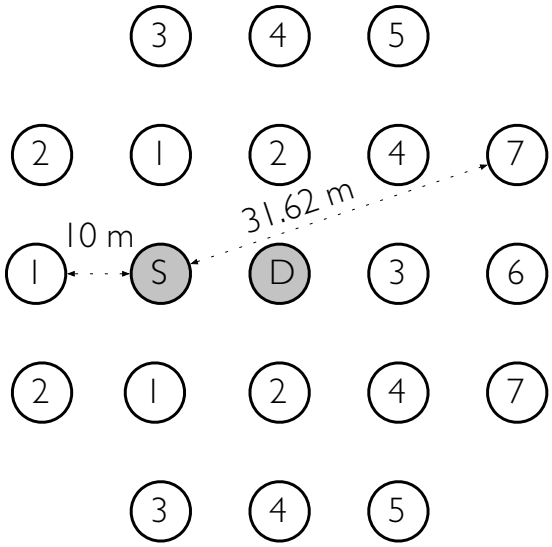

Figure 14: A source-destination pair and the set of nodes with interference range of D. The nodes are numbered in increasing order of their distance from S. The distances to the closest and farthest tier are shown.

additional tiers of potential interferers come within carrier sense range of S, thus reducing the number of hidden nodes (while increasing the number of exposed nodes). For $R_{cs} > 31.62$, all potential interferers are within CS range of S, thus eliminating hidden nodes entirely.

In Table 2, we list the distance of each "tier" of interferers from the sender, the number of hidden and exposed nodes when R_{cs} is set to include that tier but not higher tiers, and the corresponding value of T_{max}. We see that the throughput is maximized when interferers up to the 5th tier are contained within sensing range, leaving 3 hidden terminals. For the corresponding R_{cs} of 29m, the sum of the number of hidden and exposed terminals is 10, which is the minimum possible. However, the number of hidden and exposed terminals are not equal, possibly reflecting a greater penalty for hidden terminals than for exposed terminals in this particular scenario. Figure 15 illustrates the hidden and exposed node areas for the optimal R_{cs}.

Table 2: Distance, number of hidden and exposed nodes, and T_{max} for each tier.

Tier	Distance	# Hidden	# Exposed	T_{max}
1	10.00	16	0	24
2	14.14	12	0	30
3	20.00	9	1	72
4	22.36	5	5	98
5	28.28	3	7	104
6	30.00	2	10	98
7	31.62	0	16	92

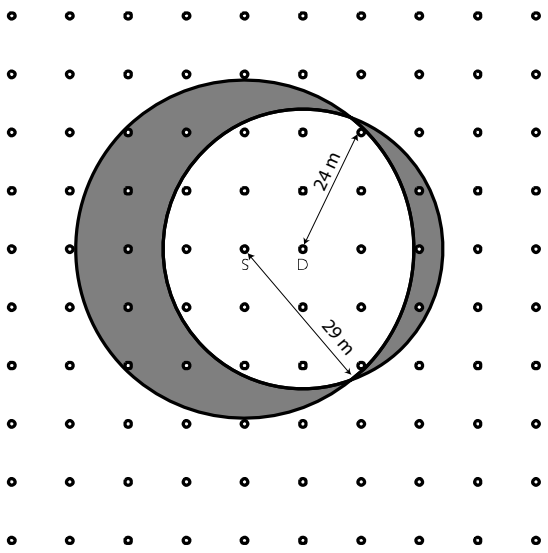

Figure 15: Hidden and exposed areas for the optimal carrier sensing range in the OPNET simulations. There are 3 hidden terminals and 7 exposed terminals when R_{cs} is set to its optimum value of 29 meters.

Note that the loss in throughput from setting $R_{cs} = R_I$ is about 5.7%, which again supports the conclusion that setting $R_{cs} = R_I$ provides a robust initialization point for adapting to network conditions at run-time.

The value of $T_{max} = 50$ for R_{cs}= 128 m also reveals an interesting consequence of our modified OPNET model. With 360 flows in all, $T_{max} = 50$ corresponds to a network wide throughput of 18 Mbps, which is higher than the channel rate of 12 Mbps even though all nodes are within sensing range of each other and hence ostensibly share the available 12 Mbps. This is because—in our revised OPNET model—when two senders pick the same backoff value and hence transmit simultaneously, both transmissions can be successful if the transmitters are spaced sufficiently far apart. This is in contrast to the default OPNET model where two simultaneous transmissions always produce a collision that causes both transmissions to fail. The revised model, in essence, allows for some spatial reuse to occur even when all nodes can sense each other.

5 Conclusion

We have demonstrated via analytical modelling that the optimal PCS threshold in 802.11 should be set such that the carrier sensing range equals the interference range. The validity of this result was investigated by OPNET simulations, that necessitated significant changes to OPNET code for the link layer decode chain as well as the CCA mechanism. The results underscore the improvements to aggregate throughput achievable via suitable tuning of PCS threshold.

References

[1] F. Ye, S. Yi and B. Sikdar, "Improving Spatial Reuse of IEEE802.11 Based Ad Hoc Networks," *Proc. IEEE Globecom*, San Fransisco, CA, Nov. 2003.

[2] K. Xu, M. Gerla and S. Bae, "How Effective is the IEEE 802.11 RTS/CTS Handshake in Ad-Hoc Networks," *Proc. IEEE Globecom*, Taipei, Nov. 2002.

[3] L-J. Chen, T. Sun, G. Yang, M. Y. Sanadidi, and M. Gerla, "AdHoc Probe: Path Capacity Probing inWireless Ad Hoc Networks," *1st IEEE Int. Conference on Wireless Internet*, Budapest, Hungary, Aug. 2005.

[4] G. Anastasi, E. Borgia, M. Conti and E. Gregori, "WiFi in Ad-Hoc Mode: A Measurement Study," *Proc. 2nd IEEE Conf. Pervasive Computing & Commn.*, 2004.

[5] "Intel Pro/Wireless 2200BG Driver for Linux," http://ipw2200.sourceforge.net

[6] Y. Zhou and S. Nettles, "Balancing the Hidden and Exposed Node Problems with Power Control in CSMA/CA Based Wireless Networks," *Proc. IEEE WCNC*, New Orleans, 2005.

[7] J.Li, C. Blake, S. J. De Couto, H. Lee and R. Morris "Capacity of Ad Hoc Wireless Networks," *Proc. ACM Mobicom* July 2001, pp. 61-69.

[8] A. Kamerman and L. Montaban, "Wavelan-II: A High-Performance Wireless LAN for the Unlicensed Band,"*Bell Systems Tech. J.*, Summer 1997.

[9] P. C. Ng and S. C. Liew, "Offered Load Control in IEEE802.11 Multi-hop Ad-hoc Networks,"*1st IEEE Conf. on Mobile Ad-Hoc and Sensor Systems*, Ft. Lauderdale, FL, 2004.

[10] X. Yang and N. H. Vaidya, “On the Physical Carrier Sense in Wireless Ad Hoc Networks,” *Proc. IEEE Infocom*, Mar.2005, Miami, FL.

[11] K. Jamieson, B. Hull, A. Miu and H. Balakrishnan, “Understanding the Real-World Performance of Carrier Sense, ”*ACM Sigcomm Workshop on Experimental Approaches to Wireless Network Design and Analysis*, Philadelphia, PA, Aug. 2005.

[12] A. Vasan, R. Ramjee and T. Woo, “Enhanced Capacity 802.11 Hotspots,” *Proc. IEEE Infocom*, Mar. 2005, Miami, FL.

[13] G. Anastasi, M. Conti and E. Gregori, “IEEE 802.11 Ad Hoc Networks: Protocols, Performance and Open Issues,” in *Mobile Ad-Hoc Networking*, ed. S. Basagni et al., IEEE Press, 2004, pp.117-138.

[14] J. Deng, B. Liang and P. K. Varshney, “Tuning the Carrier Sense Range of IEEE 802.11 MAC,” *Proc. IEEE Globecom*, Nov. 2004.

[15] Kamin Whitehouse, Alec Woo, Fred Jiang, Joseph Polastre, David Culler, “Exploiting the Capture Effect for Collision Detection and Recovery,” *The Second IEEE Workshop on Embedded Networked Sensors (EmNetS-II)*, Sydney, Australia, May 2005.

[16] J. Zhu, X. Guo, L. L. Yang, W. S. Conner, and S. Roy and M. M. Hazra, “Adapting Physical Carrier Sensing to Maximize Spatial Reuse in 802.11 Mesh Networks”, *Wiley J. Wireless Commun. Mob. Comput.* 2004, vol. 4, pp. 933-946.

[17] J. Zhu, B. Metzler, X. Guo and Y. Liu, “High Density WLAN: A Hardware Prototyping Approach,” *Proc. IEEE Infocom*, Barcelona, Spain, 2006.

[18] S. DeSilva and R. V. Boppana, “On the Impact of Noise Sensitivity in 802.11 based Ad-Hoc Networks,” *Proc. IEEE ICC*, June 2004, Paris, France.

[19] IEEE Standard for Wireless LAN Medium Access Control: MAC and PHY Specifications, ISO/IEC 8802-11, 1999(E), Aug. 1999.