Selecting Transmit Powers and Carrier Sense Thresholds for CSMA Protocols

Technical Report (October 2004)*

Jason A. Fuemmeler
fuemmele@uiuc.edu

Nitin H. Vaidya
nhv@uiuc.edu

Venugopal V. Veeravalli
vvv@uiuc.edu

University of Illinois at Urbana-Champaign

1 Introduction

Spatial reuse is an important aspect of wireless ad hoc network design. However, because simultaneous transmissions can interfere with one another and prevent correct reception, it is important to properly control channel access. In recent years, the 802.11 MAC protocol has become the de facto standard for wireless ad hoc networks. This protocol employs two types of carrier sensing to prevent collisions. The first type is physical carrier sensing, where nodes wait until the power on the channel is below a threshold before transmitting. The second type is virtual carrier sensing, in which RTS/CTS control frames are exchanged to inform surrounding nodes that the channel is being reserved for a packet transmission.

In this paper, we focus on the physical carrier sensing mechanism. We start from the premise that collisions should be prevented. We then develop collision prevention conditions on the transmit powers and carrier sense thresholds used in the network. The main result of this analysis is that senders should keep the product of the transmit power and the carrier sense threshold equal to a fixed constant. We then describe how to apply this analysis to a practical protocol. Results are then presented that illustrate the potential benefits of this scheme. Before concluding, we place our work in the context of previous work in the area of MAC protocol design for spatial reuse.

2 Analysis

2.1 Assumptions

We will be defining certain functions that map some subset of $[0, \infty]$ to some other subset of $[0, \infty]$. Note that a value of ∞ may be included in either the range or domain. For a continuous function f, we have the following conventions:

- $f(\infty) = y$ is equivalent to $\lim_{x\to\infty} f(x) = y$. Note that it is possible to have $y = \infty$.
- $f(0) = \infty$ is equivalent to $\lim_{x\to 0^+} f(x) = \infty$.

*This work was funded in part by NSF and a NSF Graduate Research Fellowship.

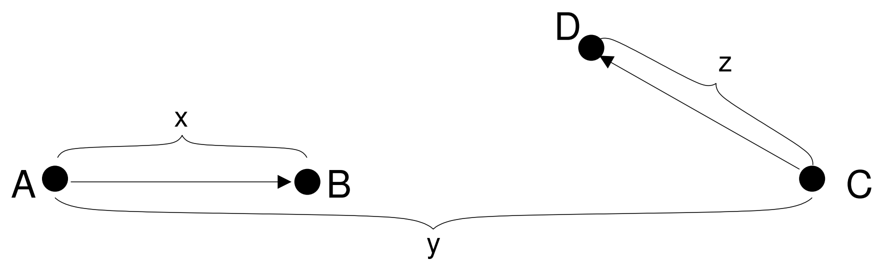

Figure 1: Wireless Network with 4 Nodes

These conventions allow us to always work with closed intervals and are for convenience of notation only. We will not be focused on performing any calculations involving ∞.

The physical layer assumed has the following characteristics:

- A node's transmission is successful if the SINR at the intended destination is at or above the threshold γ.
- The thermal noise at each node is η.
- The power gain between two nodes separated by a distance x is given by the function $g(x)$. The function $g(x)$ is a continuous, nonnegative, strictly decreasing function of x defined on the interval $[0, \infty]$. It is assumed that $g(0) = g_{max}$ and that $g(\infty) = 0$. In a physical network, we have that g_{max} is finite since infinite gain cannot be achieved. However, sometimes for the purposes of analysis a function that has $g_{max} = \infty$ may be used. An example of such a function is $g(x) = g_0 x^{-\alpha}$.
- Because of the above properties of $g(x)$, the inverse gain function, $g^{-1}(x)$, exists and is a continuous, nonnegative, strictly decreasing function of x defined on the interval $[0, g_{max}]$. We then have $g^{-1}(0) = \infty$ and $g^{-1}(g_{max}) = 0$.

We will consider a CSMA protocol that uses only physical carrier sensing. A node which wishes to send a packet selects two parameters: a transmit power and a carrier sense threshold. We specify that these parameters will be functions of x, the distance to the intended receiver. Thus the transmit power is given as $p_t(x)$ and a the carrier sense threshold is given as $p_{cs}(x)$. Note that x can be computed as $g^{-1}(g(x))$ if $g(x)$ has been estimated. The sender is allowed to begin transmitting only if the total received power from ongoing transmissions (not including thermal noise) is less than or equal to $p_{cs}(x)$. Note that we assume that this received power can be easily measured; this assumption may not be realistic (see [1]).

2.2 Conditions for Collision Prevention

Consider the arrangement of nodes shown in Figure 1. Node A is attempting to communicate with node B at a distance x. Node C is an interferer at a distance y from the sender. To maximize the interference, C is placed in the same direction as the receiver B. We therefore call C a worst-case interferer. Node C is attempting to send to its own destination, node D, at a distance z. The direction of C's transmission is immaterial and the value of z is unknown to node A. We wish to establish conditions on the transmit power used by node A ($p_t(x)$) such that A's transmission will be successful. In this section, we assume that C has not yet begun transmitting when A's transmission begins. This is valid since any rules obtained will also be followed by C, causing C's transmission to also be successful.

The fact that C is indeed allowed to interfere implies that its carrier sense threshold has not been exceeded. Mathematically, we have

$$p_t(x)g(y) \le p_{cs}(z) \Longrightarrow y \ge g^{-1}\left(\frac{p_{cs}(z)}{p_t(x)}\right) \tag{1}$$

For ease of exposition, we have assumed that $p_{cs}(z)/p_t(x) \le g_{max}$. Actually, we will assume that $p_{cs}(z)/p_t(x) \le g(x)$ to ensure that $y \ge x$. This assumption will be discussed further in Section 2.4. Since $y \ge x$, the worst-case interference occurs when y is minimized — i.e. when the above inequality is replaced with equality.

We can express the SINR requirement on node A's transmission as

$$\frac{p_t(x)g(x)}{kp_t(z)g(y-x)+\eta} \ge \gamma \tag{2}$$

Since there may be multiple interferers, we have introduced a factor k into this equation which is the number of worst-case interferers assumed. By substituting the minimum value of y from (1) into (2), rearranging terms, and maximizing over z we get the inequality

$$p_t(x)g(x) \ge k\gamma \max_z \left\{p_t(z)g\left(g^{-1}\left(\frac{p_{cs}(z)}{p_t(x)}\right)-x\right)\right\}+\gamma\eta \tag{3}$$

The inequality in (3) is somewhat problematic since the function p_t is defined in terms of itself and the function p_{cs}. To simplify matters, an approximation is in order; namely, we choose to eliminate the x term subtracted inside the gain function. We will refer to this as the colocation approximation. This is because it is equivalent to assuming that node A is moved to become colocated with B while the gain between A and B remains at $g(x)$. Note that a similar approximation is made in [2]. The inherent advantage of this arrangement is that the sender A can perfectly estimate the interference seen at the receiver B. A fringe benefit is that since g and g^{-1} cancel on the right-hand side of (3), it is no longer necessary for the sender to know x and the function g. Instead, it is sufficient to know the gain to the receiver $g(x)$.

The problem with the colocation approximation is that it is *less conservative* than the original statement. However, we can partially explain this away by realizing that if there are indeed multiple interferers, they will be spaced all around nodes A and B instead of being concentrated in the worst-case position. We discuss this approximation further in Section 2.4.

Making our approximation, (3) becomes

$$p_t(x)g(x) \ge k\gamma\frac{\max_z\{p_t(z)p_{cs}(z)\}}{p_t(x)}+\gamma\eta \tag{4}$$

We can see from this equation that the maximization over z determines a key parameter for our analysis. By using the notation $\beta = \max_z\{p_t(z)p_{cs}(z)\}$, (4) reduces to a quadratic equation (with an inequality) whose solution is

$$p_t(x) \ge \frac{\gamma\eta+\sqrt{\gamma^2\eta^2+4k\gamma\beta g(x)}}{2g(x)} \tag{5}$$

Note that we have now implicitly imposed the constraint that $p_t(x)p_{cs}(x) \le \beta$ for all values of x and must meet this constraint in future analysis.

2.3 Transmit Power and Carrier Sense Functions

For maximum spatial reuse, it seems reasonable to use low transmit power and high carrier sense thresholds. However, we must meet the constraints in the previous section. Using these guidelines, we arrive at the following formulae for the transmit power and carrier sense threshold functions:

$$p_t(x) = \frac{\gamma\eta + \sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x)}}{2g(x)} \tag{6}$$

$$p_{cs}(x) = \frac{\beta}{p_t(x)} \tag{7}$$

Note that for all x, $p_t(x)p_{cs}(x) = \beta$.

It is worthwhile to rearrange (6) to give a formula for the power seen at the receiver, as

$$p_t(x)g(x) = \frac{\gamma\eta + \sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x)}}{2} \tag{8}$$

From this expression we can see that if $\beta = 0$, there is no interference margin at the receiver. We also know from (7) that if $\beta = 0$, the carrier sense threshold is always 0. Thus if $\beta = 0$, energy usage is minimized but there is no spatial reuse. On the other hand, if β is large, spatial reuse increases at the expense of energy consumption. Note that as $\beta \to \infty$, the noise becomes inconsequential and the amount of spatial reuse approaches some (unquantified) limit.

Through algebraic manipulation, we can find an interesting alternate interpretation of (6) and (7). Start with (7) and substitute (6) as

$$p_{cs}(x) = \frac{\beta}{p_t(x)} = \frac{2\beta g(x)}{\gamma\eta + \sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x)}} \tag{9}$$

Multiplying the numerator and denominator of this expression by a specific quantity yields (after simplification)

$$p_{cs}(x) = p_{cs}(x)\frac{\sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x)} - \gamma\eta}{\sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x)} - \gamma\eta} = \frac{\sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x)} - \gamma\eta}{2k\gamma} \tag{10}$$

From (8) we see that we can rewrite this as

$$p_{cs}(x) = \frac{p_t(x)g(x) - \gamma\eta}{k\gamma} = \frac{1}{k}\left(\frac{p_t(x)g(x)}{\gamma} - \eta\right) \tag{11}$$

Thus we see that the carrier sense threshold being used is the interference margin at the receiver divided by k. The close relationship to the interference margin is not surprising since under the colocation approximation the sender can perfectly estimate the current interference at the receiver to determine whether a potential transmission will be successful.

2.4 Bounding the Approximation Error

In this section, we will see that it is possible to compensate for the colocation approximation by choosing appropriately high values of k. We will then show that these k satisfy the assumptions made in Section 2.2. This discussion is not meant to completely characterize the effects of the

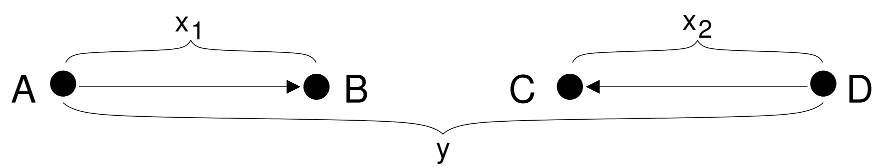

Figure 2: Another Wireless Network with 4 Nodes

colocation approximation; instead, the intent is to illustrate that the preceding analysis was indeed reasonable.

Consider the setup of Figure 2, which is very similar to that in Figure 1. In this setup, there are two links whose receivers are both in worst-case positions. Link 1 has distance x_1 and uses a k of k_1. Link 2 has distance x_2 and uses a k of k_2. We seek conditions on k_1 and k_2 such that collisions are prevented on both links. If we allow the transmit power and carrier sense functions to depend on the values of k chosen, we have the following condition for collision prevention on link 1:

$$p_t(x_1,k_1)g(x_1) \geq \gamma p_t(x_2,k_2) g\left(g^{-1}\left(\frac{p_{cs}(x_2,k_2)}{p_t(x_1,k_1)}\right) - x_1\right) + \gamma\eta \tag{12}$$

A similar condition applies to link 2. For a general gain function, it is difficult to completely characterize the set of (k_1, k_2) such that both conditions are satisfied. In order to gain some insight, we set $k_1 = k_2 = k$ and try to find conditions on k. This allows us to drop the dependence on k_1 and k_2 and rewrite (12) as

$$\frac{1}{\gamma}\frac{p_t(x_1)g(x_1) - \gamma\eta}{p_t(x_2)} \geq g\left(g^{-1}\left(\frac{p_{cs}(x_2)}{p_t(x_1)}\right) - x_1\right) \tag{13}$$

Again, there is a corresponding condition for link 2.

We first consider the case where $\eta = 0$. In this case, the transmit power and carrier sense functions ((6) and (7)) have particularly simple forms. Substituting for these functions in (13) and rearranging terms, we have

$$g^{-1}\left(\frac{1}{k\gamma}\sqrt{g(x_1)g(x_2)}\right) - g^{-1}\left(\frac{1}{\gamma}\sqrt{g(x_1)g(x_2)}\right) \geq x_1 \tag{14}$$

If we assume $\gamma \geq 1$ (which is reasonable) and $k \geq 1$ (which will be verified shortly), the left-hand side of this inequality is guaranteed to be defined. Note that the first term in this inequality is $g^{-1}(p_{cs}(x_2)/p_t(x_1))$, which is just the minimum value of y as in (1). We can combine (14) with the equivalent statement for link 2 to yield

$$g^{-1}\left(\frac{1}{k\gamma}\sqrt{g(x_1)g(x_2)}\right) - g^{-1}\left(\frac{1}{\gamma}\sqrt{g(x_1)g(x_2)}\right) \geq \max\{x_1, x_2\} \tag{15}$$

This expression can then be easily solved to give a condition on k as

$$k \geq \frac{\frac{1}{\gamma}\sqrt{g(x_1)g(x_2)}}{g\left(\max\{x_1,x_2\} + g^{-1}\left(\frac{1}{\gamma}\sqrt{g(x_1)g(x_2)}\right)\right)} \tag{16}$$

As expected, we have that $k \geq 1$. For the canonical case of $g(x) = g_0 x^{-\alpha}$, (16) reduces to

$$k \geq \frac{\left(\gamma^{\frac{1}{\alpha}} + \frac{\max\{x_1,x_2\}}{\sqrt{x_1 x_2}}\right)^{\alpha}}{\gamma} = \frac{\left(\gamma^{\frac{1}{\alpha}} + \sqrt{\frac{\max\{x_1,x_2\}}{\min\{x_1,x_2\}}}\right)^{\alpha}}{\gamma} \tag{17}$$

Note that if the ratio of the two link lengths becomes very large, k must also be very large. Suppose $\gamma = 10$, $x_1/x_2 = 10$, and $\alpha = 2$. Applying our formula, we see that we need $k \geq 4$. If instead $\alpha = 4$, we then need $k \geq 59.58$.

Now suppose $\eta \neq 0$. As in the previous case, by substituting for the transmit power and carrier sense functions ((6) and (7)) in (13) and rearranging terms we obtain an inequality similar to (14). Because η is nonzero, (6) and (7) do not simplify and the resultant condition is somewhat more complicated. However, we can choose to write it as

$$g^{-1}\left(\frac{c}{k\gamma}\sqrt{g(x_1)g(x_2)}\right) - g^{-1}\left(\frac{c}{\gamma}\sqrt{g(x_1)g(x_2)}\right) \geq x_1 \tag{18}$$

where c is not a constant but is the k-dependent quantity

$$c = \frac{4k\gamma\beta\sqrt{g(x_1)g(x_2)}}{\left(\gamma\eta + \sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x_1)}\right)\left(\gamma\eta + \sqrt{\gamma^2\eta^2 + 4k\gamma\beta g(x_2)}\right)} \leq 1 \tag{19}$$

Since c and the left-hand side of (18) are symmetric in x_1 and x_2, it is easy to combine (18) with the equivalent condition for link 2 to yield

$$g^{-1}\left(\frac{c}{k\gamma}\sqrt{g(x_1)g(x_2)}\right) - g^{-1}\left(\frac{c}{\gamma}\sqrt{g(x_1)g(x_2)}\right) \geq \max\{x_1, x_2\} \tag{20}$$

Note that (15) and (20) are identical except for the presence of c in (20). Although it is not possible to easily solve for a condition on k, we can see that for k sufficiently large, $c \to 1$ and (20) will be satisfied.

Interestingly, for $g(x) = g_0 x^{-\alpha}$, values of k which satisfy (17) (and thus (15)) also satisfy (20). To see this, first note that in this special case $g^{-1}(cx) = c^{-\frac{1}{\alpha}} g^{-1}(x)$. Since $c \leq 1$ for any k, the left-hand side of (20) is greater than the left-hand side of (15), which is greater than the right-hand side of (15), which is the same as the right-hand side of (20). Thus (20) is satisfied.

Recall that in Section 2.2 we assumed that $p_{cs}(z)/p_t(x) \leq g(x)$ (referring to Figure 1). If this assumption did not hold, some of the arguments to the gain and inverse gain functions would not be defined. The worst-case interferer position would then always be at $y = x$ and our formulas would not apply. In the present case, we need two assumptions to hold: $p_{cs}(x_2)/p_t(x_1) \leq g(x_1)$ and $p_{cs}(x_1)/p_t(x_2) \leq g(x_2)$. Fortunately, it is easy to show that (16) implies both of these assumptions. We show this only for the first assumption; the other case is derived similarly. We first write

$$g(y) \leq p_{cs}(x_2)/p_t(x_1) = \frac{c}{k\gamma}\sqrt{g(x_1)g(x_2)} \leq \frac{1}{k\gamma}\sqrt{g(x_1)g(x_2)} \tag{21}$$

In these relations, the first inequality is a consequence of (1). The middle equality can be seen by substitution or by recognizing that (just as in the $\eta = 0$ case) the first term in (18) is equal to the minimum value of y in (1). The final inequality is simply because $c \leq 1$. Using (16) and the fact that g is a decreasing function, we can continue (21) as

$$\frac{1}{k\gamma}\sqrt{g(x_1)g(x_2)} \leq g\left(\max\{x_1, x_2\} + g^{-1}\left(\frac{1}{\gamma}\sqrt{g(x_1)g(x_2)}\right)\right) \leq g(\max\{x_1, x_2\}) \tag{22}$$

Since $g(\max\{x_1, x_2\}) \leq g(x_1)$, the assumption is verified. Thus, our analysis remains valid provided that k is appropriately chosen.

3 Toward a Protocol

3.1 General Principles

The preceding analysis was based primarily on a worst-case analysis. In an actual network, the interference around a link may be significantly less than the worst case. In addition, we have also not specified how to determine k, the number of assumed worst-case interferers. For these reasons, we propose that each link should adjust its transmit power dynamically to use the minimum power necessary to protect its transmissions. In keeping with the above analysis, this implies that the carrier sense threshold on each link should also be adjusted so that $p_t(x)p_{cs}(x) = \beta$.

An advantage of this scheme should be emphasized. As can be seen mathematically in (4), the relation $p_t(x)p_{cs}(x) = \beta$ effectively bounds the amount of interference a single node can pose to a nearby link. This can be understood intuitively by realizing that if a node chooses a high transmit power, its low carrier sense threshold causes it to defer to transmissions that are farther away. As a result of this bounded interference, the prevention of collisions on a given link is primarily (but not entirely) a function of the transmit power chosen for that link. It is therefore reasonable to assume that a collision on a given link is due to insufficient transmit power being used on that link.

This property is quite useful since it allows links to have access to the required feedback. Through the use of MAC acknowledgements, a link is able to determine whether a packet was successfully received. The link can then adjust its transmit power based on this feedback. In contrast, consider a scheme in which a link chooses some fixed transmit power and then adjusts its carrier sense threshold to prevent collisions. In this case, if the carrier sense threshold is selected too high, a collision could be caused at a nearby link. However, the link using the improper carrier sense threshold would not know that this collision has occurred. This of course makes adaptation difficult.

A potential problem with our scheme is that fairness may be adversely affected. Of course, this depends on the definition of fairness being used. The source of the problem is that different senders in the same region may be using widely different carrier sense thresholds. As a result, those with low carrier sense thresholds are at a certain disadvantage since they need to wait for quieter conditions before transmitting. In the meantime, senders with high carrier sense thresholds may access the channel, further postponing the transmissions of the disadvantaged transmitters. Note that short links will generally have an advantage over long links. This is a price paid for aggressively insisting on many simultaneous transmissions.

3.2 Further Details

We have assumed that on every link, the sender knows the gain to the receiver. In practice, this gain will need to be estimated. One way this can be accomplished is by including the transmit power in RTS and/or data frames sent to the receiver. The receiver can then estimate the gain and relay it back to the sender in CTS and/or ACK frames. The receiver's transmit power when replying should be identical to that used originally by the sender. If the sender and receiver use different transmit powers, it becomes unclear to each who was at fault when a transmission is unsuccessful.

Unfortunately, not all packet losses at the MAC layer are due to improper transmit power or carrier sense thresholds. Collisions can also occur when two nodes decide to transmit at the same time. The probability of this type of collision increases with the transmit power used by neighboring nodes. Thus, the transmit power control loops should be biased toward low powers to prevent this effect.

Instead of varying the transmit power directly, we have chosen to adjust the effective number of worst-case interferers k. One reason for doing so is that we can reduce k to zero without falling below the minimum necessary received power $\gamma\eta$. Another reason is that we can easily compare the values of k across links without having to normalize by the individual link gains. A final reason for varying k is that there are some values of k which seem reasonable — perhaps on the order of 1 to 6. Of course, we have seen in Section 2.4 that this is not always the case.

We now give some design suggestions. We have already stated that β should be large for maximum spatial reuse. However, if there is a maximum power constraint using large β severely limits the maximum possible transmit distance. More specifically, we have the following constraint on β:

$$\beta \leq \frac{p_t^{max}}{k_{max}} \left(\frac{p_t^{max} g(x_{max})}{\gamma} - \eta \right) \tag{23}$$

where x_{max} is the desired maximum transmit distance, p_t^{max} is the maximum power constraint, and k_{max} is the maximum anticipated value of k. The value of k_{max} to use may need to be determined from experience since the worst-case values obtained in Section 2.4 occur with small probability.

Another design suggestion is to set a minimum transmit distance x_{min} since from Section 2.4 we know that the worst-case value of k can depend on the ratio of the lengths of two adjacent links. Thus if a gain to a receiver is found to be larger than $g(x_{min})$, then $g(x_{min})$ should be used instead to compute the transmit power. This gives the minimum transmit power to be used as

$$p_t^{min} = \frac{\gamma\eta}{g(x_{min})} \tag{24}$$

4 Results

4.1 General Simulation Setup

The ns-2 network simulation package was used for our simulations; however, some significant changes were needed. In the default wireless physical layer code, the carrier sense threshold parameter is used for two somewhat different functions. First, any single packet received with power above the threshold prohibits the transmitter from transmitting. Second, any single packet received above the threshold causes the receiver to begin tracking that packet for reception. Because we wanted to vary the carrier sense threshold without changing the behavior of the receiver, this second function was undesired. Moreover, we needed to keep track of the *cumulative* interference at the receiver (not just look at a single packet) to compute an accurate SINR.

Table 1: Simulation Parameters

Parameter	Value
SINR Threshold (γ)	10 dB
Thermal Noise (η)	0 W
Link Rate	1 Mbps
Packet Size	512 Bytes
Slot Time	20 μs
Fixed CW Size	31 slots
RTS/CTS	Disabled
Effective Link Rate	767 kbps
Offered Load Per Node	100 pkt/s

In our modified code, a list of received powers for all active packets is kept at each node. A thermal noise parameter is also present. Thus an SINR can be computed. Whether a packet is tracked for reception is then based on whether the SINR is at or above the threshold γ. Thus, the carrier sense threshold is used only by the transmitter when deciding whether to access the channel. In line with our previous discussion, the modified code has the receiver reply to the sender using the same power originally used by the sender. It is assumed that each sender and receiver has perfect knowledge of gain on the link connecting them; thus, there is no increase in overhead for gain estimation.

The gain function used in simulations was the default one in ns-2. It is given as

$$x = \begin{cases} \left(\frac{75}{914\pi}\right)^2 x^{-2} & \text{if } x < 27.42\pi \\ \left(\frac{3}{2}\right)^4 x^{-4} & \text{if } x \geq 27.42\pi \end{cases} \tag{25}$$

where x is measured in meters. Many of the other simulation parameters were the defaults. A table of simulation parameters can be found in Table 1. The thermal noise was set to 0 to maximize spatial reuse and to make the choice of β irrelevant. Also, there was no minimum or maximum power. For simplicity, the contention window size was fixed and the use of RTS/CTS frames was disabled. Thus, virtual carrier sensing was effectively disabled. The effective link rate listed in Table 1 is the throughput of a network containing a single link and is less than the link rate due to overhead. In all simulations, the network was very heavily loaded.

4.2 Ring Topology

In this set of simulations, nodes were placed equally spaced on a circle of radius 250 m. Each of the nodes transmitted UDP packets to its nearest neighbor in a counterclockwise direction. The total network throughput was recorded for various values of k. The results are shown in Figure 3 for various numbers of nodes in the ring. The results indicate that the optimal throughput occurs for k between 1 and 3. This seems reasonable since in a large ring we would expect to need to account for 2 interfering links — one link in the clockwise direction and one in the counterclockwise direction.

Because of the symmetry of this topology and the absence of noise, each value of k in Figure 3 can be mapped to an equivalent ratio of transmit power to carrier sense threshold (this was also

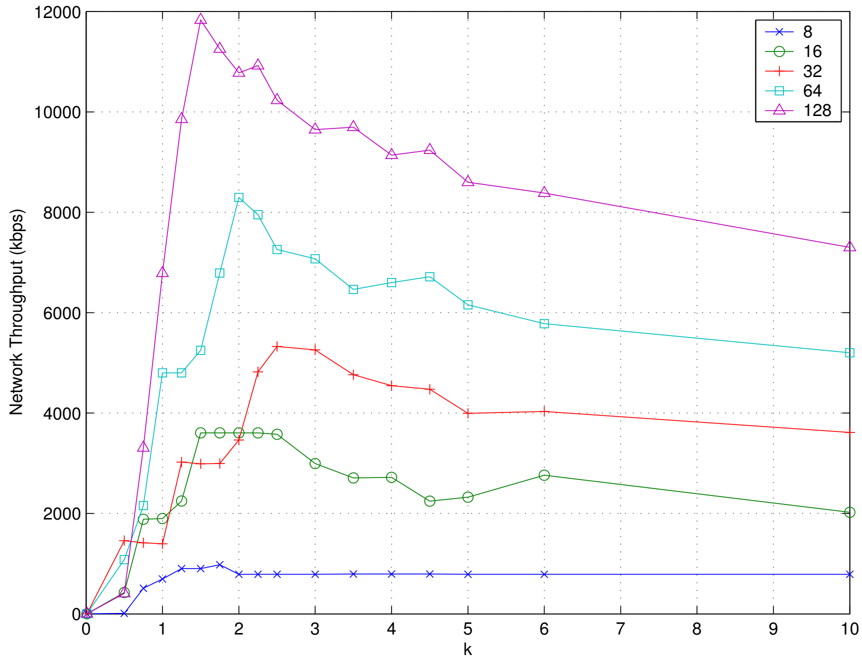

Figure 3: Ring Topology Throughput vs. k

verified through simulation). Thus we can transform the data to generate Figure 4, which plots the network throughput for a fixed transmit power (0.282 W) and varying carrier sense threshold. Note that in this figure, the optimal throughputs are obtained at widely different carrier sense thresholds. This is only a small inconvenience, however.

There is no performance advantage to our scheme for a ring topology. However, if the parameters were to be adjusted dynamically then our scheme may have advantages for reasons already mentioned in Section 3.1.

4.3 Random Topology

To generate a random topology, 30 links were placed within a circle of radius 250 m. The link lengths were Rayleigh distributed with mean length $250/\sqrt{30} = 45.64$ m. The Rayleigh distribution was selected because in a network of uniform density, the distance to a nearest neighbor is Rayleigh distributed. The mean link length was selected so that the area of 30 circles of this radius would equal the area of one circle with radius 250 m. This was deemed to yield a reasonable density of links. A sample random topology is shown in Figure 5.

We simulated two different random topologies under the following three scenarios:

1. The receive power at each receiver was kept fixed at 3.652×10^{-10} W while the carrier sense threshold was varied uniformly across the network. The network throughput curves are shown in Figure 6.

2. The transmit power was kept fixed at 0.282 W while the carrier sense threshold was varied uniformly across the network. The network throughput curves are shown in Figure 7.

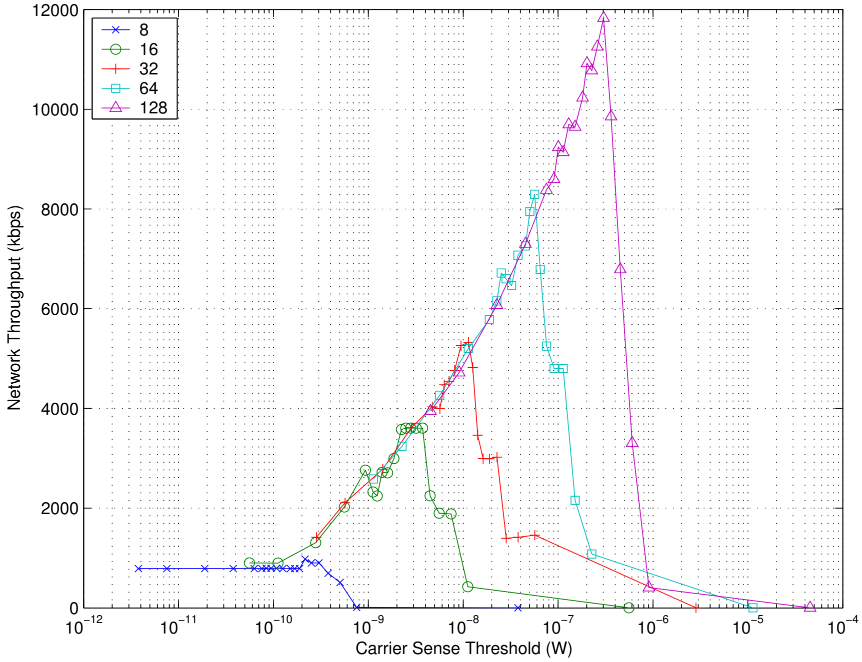

Figure 4: Ring Topology Throughput vs. CS Threshold

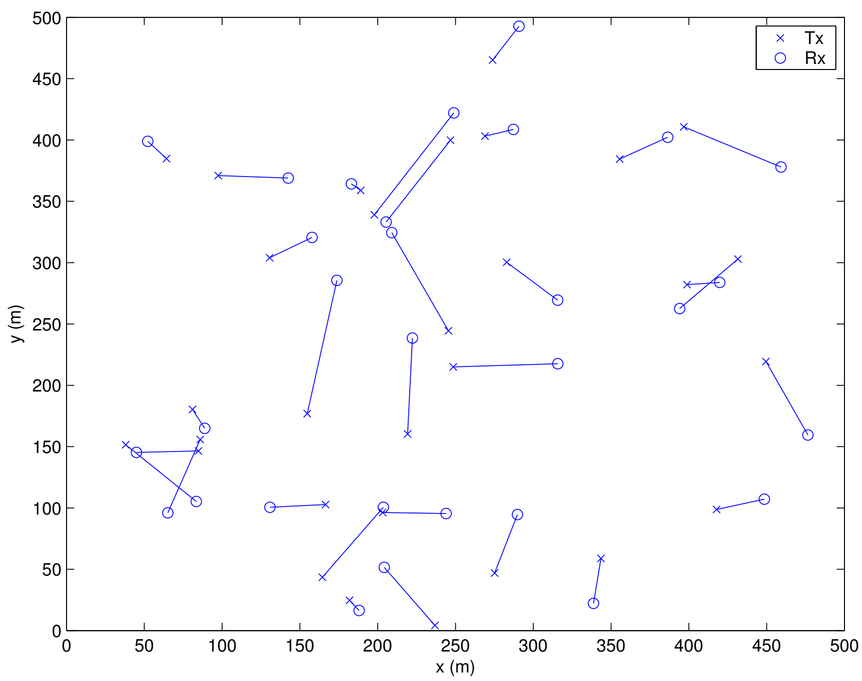

Figure 5: Sample Random Topology

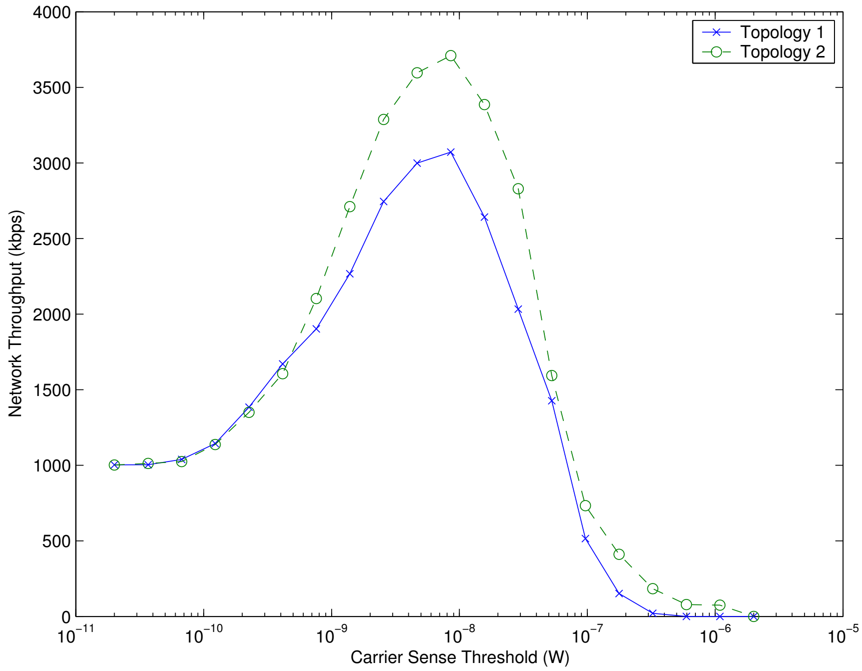

Figure 6: Throughput with Constant Rx Power

Table 2: Peak Throughputs (in kbps) for Various Scenarios

Scenario	Topology 1	Topology 2
Fixed Rx Power	3071	3710
Fixed Tx Power	3627	3944
Static k	4780	5218
Dynamic k	4955	5365

3. Our scheme was used, varying k uniformly across the network. The value of β used was 1 W^2 but this value was immaterial since there was no thermal noise. The network throughput curves are shown in Figure 8.

The peak throughputs for each scenario are given in the first three rows of Table 2

In examining these data, it can be seen that our scheme outperforms both other scenarios. More specifically, our scheme achieves 40% to 50% higher throughput than the constant receive power scenario and approximately 30% higher throughput than the constant transmit power scenario. Moreover, once the optimal throughput is reached the throughput decreases rather slowly with increasing k. This indicates a certain robustness for this scheme.

4.4 Dynamic Adjustment of k

Suppose a node wants to send a packet. The node waits until its carrier sense threshold is not exceeded and then attempts to send the packet. For each such attempt, the packet transmission

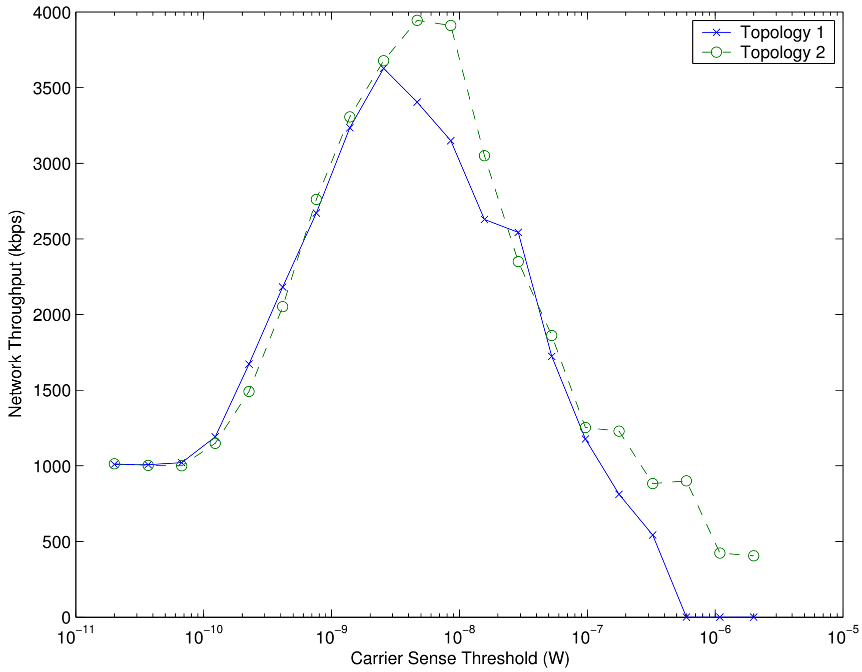

Figure 7: Throughput with Constant Tx Power

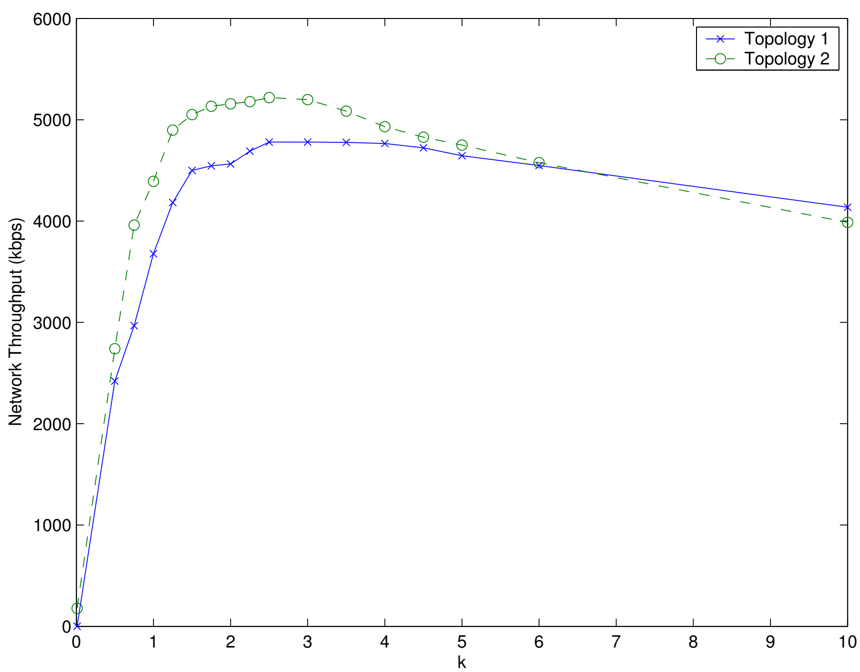

Figure 8: Throughput as a Function of k

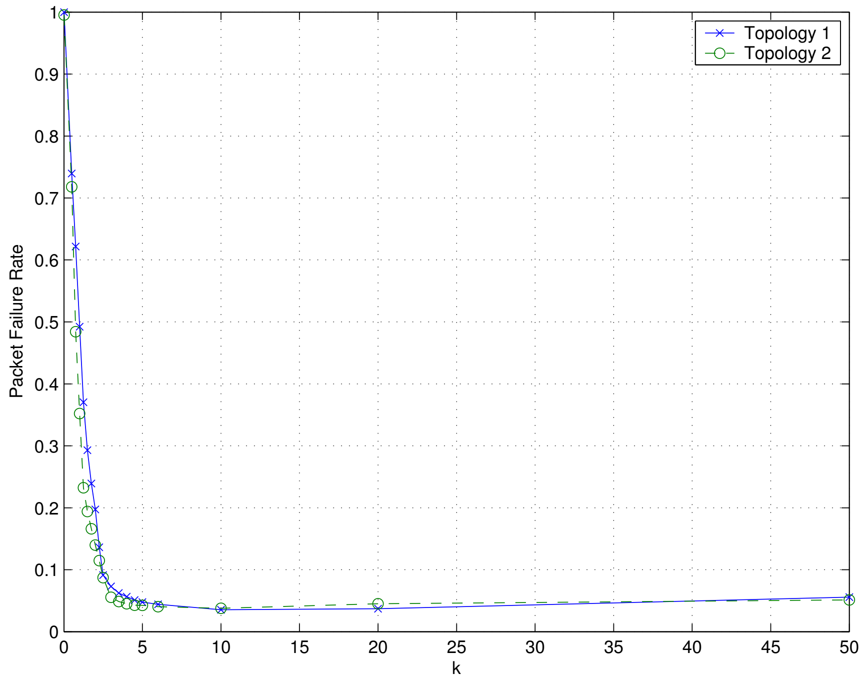

Figure 9: Failure Rate as a Function of k

either succeeds or fails. For the random networks of the previous section, we can plot the failure rate as a function of k as in Figure 9. Comparing these data with those of Figure 8 reveals that a reasonable control algorithm (run independently for each link) might try to approximately minimize the failure rate. At the same time, the control loop should be biased toward low values of k to prevent nodes from becoming too cautious.

We now describe such a control algorithm. This algorithm is the result of trial and error and should not be considered optimal in any sense. The algorithm has two phases. In the first phase, the link has an initial k of 0. On a packet failure k is increased by 0.1 and on a packet success k is decreased by 0.1. This continues until at least 5 transmission attempts have been made and the failure rate averaged over all attempts is less than 0.75. This should happen eventually since the failure rate should converge in the mean to 0.5. The purpose of this phase is to reach a state where both successes and failures are occurring so that the next phase will be able to determine how to minimize the failure rate.

The second phase is based on a gradient-descent algorithm. Suppose k_i is the current value of k. First, 20 packet attempts are performed using a k of k_i. The number of failures that occur in these 20 attempts is n_1. Next, 20 packet attempts are performed using a k of $k_i + 0.5$. The number of failures is n_2. If $f(k)$ is the failure rate as a function of k, we can estimate the gradient at k_i as

$$\nabla f(k_i) = \frac{n_2 - n_1}{(0.5)(20)} \tag{26}$$

The next value of k to use, k_{i+1}, is determined through the update relation

$$k_{i+1} = k_i - 0.1(\nabla f(k_i) + 0.1) \tag{27}$$

Note that instead of trying to achieve $\nabla f = 0$, the algorithm tries to achieve $\nabla f = -0.1$. Since $f(k)$ seems to decrease and then increase with increasing k (at least judging from Figure 9), we see that this implies the desired bias toward low values of k.

Simulations were run using the above control algorithm for each of the topologies considered in the previous section. The length of each simulation was 120 seconds. The resultant throughputs are given in the last row of Table 2. These values are approximately 3% higher than the peak values for static k. This slight gain probably stems from allowing different links to use differing values of k. Regardless, it seems clear that the control algorithm is working well since the network throughput is not decreased from the optimal for static k.

5 Comparison with Other Work

5.1 PCMA

The PCMA protocol is described in [3]. It is not a CSMA protocol but is worth considering because it is concerned with collision prevention. In PCMA, any nodes that are currently receiving a packet periodically transmit busy tones on an out-of-band channel. The power at which these busy tones are transmitted is selected as a function of the amount of additional interference each receiver can tolerate. As a result, nodes hearing the busy tones can compute the maximum transmit power they can use without disturbing any ongoing transmissions. Thus, collisions are prevented.

The feedback used for collision prevention in PCMA is receiver-based since it is the receivers that send the busy tones. In contrast, our scheme uses transmitter-based feedback — the transmissions themselves. Receiver-based feedback should be preferred over transmitter-based feedback since it is only the SINR at the receiver that impacts reception. In fact, the root cause of the colocation approximation error is that our feedback is transmitter-based. In this respect, PCMA is superior to our scheme. However, because our scheme uses the transmissions themselves to prevent future interference, we do not require the use of a separate channel or for nodes to have multiple transceivers. We can effectively say that the transmissions themselves are our busy tones.

Because both our scheme and PCMA attempt to prevent collisions, an important constant that appears in [3] is very similar to our β. Using their notation, the transmit power for a busy tone transmitted by receiver j is computed as

$$Pt\_BT_j = \frac{Pt\_Max \cdot CS\_Thresh}{E_j} = \frac{C}{E_j}$$

where Pt_Max is the maximum transmit power, CS_Thresh is the minimum received power at which a busy tone can be detected (this is *not* the same as the carrier sense threshold in a CSMA protocol), C is equal to $Pt\_Max \cdot CS\_Thresh$, and E_j is the additional interference the receiver j can tolerate. Thus, a node transmitting at power Pt_Max that is too far away to detect the busy tone will not cause a collision at receiver j. Thus, the interference posed by a single node becomes bounded. This is precisely the role of β in our analysis. Moreover, β is the product of a transmit power and a carrier sense threshold — essentially the same form as C in the above equation.

In PCMA, transmissions are initiated using a modified RTS/CTS exchange. One purpose of this exchange is to select the transmit power so that there is some margin for future interference at the receiver. The authors of [3] choose to have an initial SIR at the receiver that is 4 dB higher than the minimum necessary for packet reception. However, there is no justification provided for this transmit power selection strategy. In contrast, our analysis provides for transmit power selection and collision prevention simultaneously. In other words, our analysis considers the entirety of the transmission process — from initiation to completion. This would seem to be advantageous

5.2 PCDC and POWMAC

The PCDC and POWMAC protocols (described in [4] and [5] respectively) use very similar techniques and are thus discussed together. They are not CSMA protocols but are worth considering since they are concerned with collision prevention. In both protocols, control frames are sent at relatively high powers so that neighboring nodes are aware of ongoing transmissions. In PCDC, control frames are sent in an out-of-band channel. In POWMAC, a single channel is used but the RTS/CTS exchanges are augmented to include a third control frame. In addition, an access window is used so that multiple links can send control frames before any data frames are sent. Thus, both PCDC or POWMAC require significant additional overhead. In contrast, our scheme does not require such significant changes in the 802.11 protocol.

The control frames in PCDC and POWMAC are used to select the transmit power on a given link. The two protocols use somewhat different rules for transmit power selection and these rules are somewhat complex. In both cases, the interference margin at each receiver is maximized under certain constraints (generally related to the maximum transmit power). However, increasing the interference margin at one link causes increased interference in the rest of the network. Thus it is unclear whether performing such a maximization at each link leads to maximum spatial reuse. It is possible that our scheme, which attempts to use the minimum power necessary for collision prevention, may provide better results.

5.3 Carrier Sense Selection Research

In [2], analysis is performed with the goal of computing an optimal carrier sense threshold for the 802.11 protocol. Interestingly, some of their results represent specific cases of our analysis. For example, one major conclusion of their work is that the ratio of the carrier sense threshold to the received power at the receiver should satisfy (using our notation)

$$\frac{p_{cs}(x)}{p_t(x)g(x)} \leq \left(1+\gamma^{\frac{1}{\alpha}}\right)^{-\alpha} \tag{28}$$

Correspondingly, if in (13) we set $x_1 = x_2 = x$ and rearrange terms we have

$$\frac{p_{cs}(x)}{p_t(x)} \leq g\left(x + g^{-1}\left(\frac{p_t(x)g(x)-\eta}{\gamma p_t(x)}\right)\right) \tag{29}$$

Setting $\eta = 0$ and dividing both sides by $g(x)$ yields

$$\frac{p_{cs}(x)}{p_t(x)g(x)} \leq g\left(x + g^{-1}\left(\frac{g(x)}{\gamma}\right)\right) \Big/ g(x) \tag{30}$$

For the canonical gain function $g(x) = g_0 x^{-\alpha}$, this inequality becomes identical to (28). Thus, our analysis arrives at the same conclusion. The authors of [2] also note that where $\gamma^{\frac{1}{\alpha}} \gg 1$, the transmitter and receiver are essentially colocated and (28) can be approximated as

$$\frac{p_{cs}(x)}{p_t(x)g(x)} \leq \frac{1}{\gamma} \tag{31}$$

We can reach this same result by eliminating the x term inside the gain function in (30) (this is indeed the colocation approximation) and again using $g(x) = g_0 x^{-\alpha}$.

Despite these similarities, the analysis of [2] does not deal with the selection of the transmit power. In fact, it is assumed that all links use the same power and that all links have the same link lengths. Our analysis is therefore more general. More specifically, our analysis not only specifies the ratio of the carrier sense threshold to the transmit power (as in (28)), but also the product of these two quantities (which we denote as β).

In [6], a distributed algorithm is proposed to adapt the carrier sense threshold to be used by all nodes in the network. The idea in this algorithm is to select the largest carrier sense threshold that allows collisions to be prevented throughout the network. The results for this algorithm approach the optimal performance for a statically selected carrier sense threshold. However, our results indicate that by allowing nodes to use different carrier sense thresholds, network throughput can be enhanced in topologies with heterogeneous links.

6 Conclusions

By analyzing conditions for collision prevention we have argued that in CSMA protocols, the product of the transmit power and the carrier sense threshold should remain constant. We have discussed the practical implementation of such a scheme and provided some preliminary results that show that spatial reuse can indeed be enhanced. A comparison of our scheme to previous work in this area has led to two main conclusions. The first conclusion is that in contrast with some previous work, our scheme allows benefits to be obtained without radically altering the 802.11 protocol. The second conclusion is that although previous analyses have obtained results similar to ours, our analysis considers the entirety of the transmission process and is thus a more comprehensive solution.

Further research in this area is needed. As noted in Section 2.1, the estimation of signal powers in the presence of noise may not always be possible. This could render our scheme (and other proposed schemes) impractical. For simplicity, some of the features of the 802.11 protocol such as RTS/CTS exchanges and exponential backoff were disabled in our simulations. These features would need to be included for more accurate results. Also, the various schemes proposed to increase spatial reuse need to be directly compared. One key question to answer is whether the increase in spatial reuse that arises from using receiver-based feedback instead of transmitter-based is worth the increase in overhead needed to implement such a scheme. It is also important to learn whether aggressively increasing spatial reuse at the MAC layer can adversely impact higher-layer metrics such as fairness, end-to-end throughput, and end-to-end delay. If this is in fact the case, MAC layer protocols will need to be modified to make the necessary tradeoffs.

References

[1] A. Sahai, N. Hoven, and R. Tandra, "Some fundamental limits on cognitive radio," in *Forty-Second Allerton Conference on Communication, Control, and Computing*, Sept. 2004.

[2] J. Zhu, X. Guo, L. L. Yang, and W. S. Conner, "Leveraging spatial reuse in 802.11 mesh networks with enhanced physical carrier sensing," in *IEEE International Conference on Communicatons, 2004*, vol. 7, pp. 4004–4011, June 2004.

[3] J. P. Monks, V. Bharghavan, and W. W. Hwu, "A power controlled multiple access protocol for wireless packet networks," in *INFOCOM 2001*, vol. 1, pp. 219–228, Apr. 2001.

[4] A. Muqattash and M. Krunz, "Power controlled dual channel (PCDC) medium access protocol for wireless ad hoc networks," in *INFOCOM 2003*, vol. 1, pp. 470–480, Mar. 2003.

[5] A. Muqattash and M. Krunz, "A single-channel solution for transmission power control in wireless ad hoc networks," in *MobiHoc 2004*, pp. 210–221, May 2004.

[6] J. Zhu, X. Guo, L. L. Yang, W. S. Conner, S. Roy, and M. M. Hazra, "Adapting physical carrier sensing to maximize spatial reuse in 802.11 mesh networks," *to appear in Special Issue on Emerging WLAN Technologies and Applications, Wiley Journal of Wireless Communications and Mobile Computing (WCMC)*, Nov. 2004.