

In Defense of Wireless Carrier Sense

Micah Z. Brodsky
MIT CSAIL

micahbro@csail.mit.edu

Robert T. Morris
MIT CSAIL

rtm@csail.mit.edu

ABSTRACT

Carrier sense is often used to regulate concurrency in wireless medium access control (MAC) protocols, balancing interference protection and spatial reuse. Carrier sense is known to be imperfect, and many improved techniques have been proposed. Is the search for a replacement justified? This paper presents a theoretical model for average case two-sender carrier sense based on radio propagation theory and Shannon capacity. Analysis using the model shows that carrier sense performance is surprisingly close to optimal for radios with adaptive bitrate. The model suggests that hidden and exposed terminals usually cause modest reductions in throughput rather than dramatic decreases. Finally, it is possible to choose a fixed sense threshold which performs well across a wide range of scenarios, in large part due to the role of the noise floor. Experimental results from an indoor 802.11 testbed support these claims.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design – *wireless communication*

General Terms

Performance, Measurement, Design

Keywords

Carrier sense, CSMA, medium access control, analytical modeling

1. INTRODUCTION

Carrier sense observes the wireless channel at a sender to decide whether the sender should transmit. In principle this approach is hard to justify, since it is channel conditions at receivers which govern whether they receive transmissions. When all nodes are tightly clustered carrier sense might be expected to work well, because channel conditions at all the nodes are highly correlated. But, for any network large enough to contemplate spatial reuse, the use of carrier sense is suspect. Its failures are usually divided into two categories, exposed terminals and hidden terminals. In the former, nodes that could have transmitted concurrently do not, sacrificing potential concurrency. In the latter, nodes that shouldn't transmit concurrently do, potentially destroying one or both transmissions. These weaknesses have inspired much research (e.g. [2] [14] [24] [10]).

While it is clear that carrier sense *can* misbehave, it is not so clear whether the gap between it and optimum performance is large enough in practice to warrant replacement. When does carrier

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'09, August 17–21, 2009, Barcelona, Spain.

Copyright 2009 ACM 978-1-60558-594-9/09/08...\$10.00.

sense make a poor choice about whether to transmit? How much throughput is sacrificed as a result?

This paper quantitatively analyzes the throughput efficiency of carrier sense. It assumes radios with adaptive bitrate, since that is the MAC's single most powerful tool for improving throughput. The main results come from analytical modeling, based on standard statistical radio propagation models from the EE community, reasonable assumptions about network layout, and Shannon's capacity formula, which gives a rough approximation of throughput for an adaptive bitrate radio. The answer to the questions posed above is that carrier sense sacrifices a surprisingly small amount of throughput.

The key intuition behind this result is that interference is global, affecting to some extent all nodes everywhere. The fact that a sender is communicating with a given receiver bounds how different their channel conditions can be, which limits the risk associated with carrier sense's sender-based decisions. Adaptive bitrate allows a sender to adjust to the level of interference the receiver experiences, helping to compensate for carrier sense errors. Catastrophic mistakes, where interference from concurrent transmissions forces bitrates close to zero, are rare. Conversely, when senders unnecessarily take turns ("multiplex"), they usually don't lose much throughput, because the multiplexing reduces interference and allows a higher bitrate. Carrier sense's decisions are not perfect, but given the flexibility of adaptive bitrate, they are quite reasonable. Further tweaking offers only limited benefits – the theoretical analysis indicates average throughput is typically less than 15% below optimal. Experimental results from an 802.11 testbed support these conclusions.

This work presents three main contributions. First, an analysis showing that carrier sense provides nearly optimal throughput in the common case, and a discussion of the underlying causes for this good behavior. Second, the identification of several distinct behavioral regimes for carrier sense, most of which perform well and only one of which typically encounters bad behavior. Finally, a model for the high-level properties of adaptive bitrate radio throughput, focused not on the worst case but on average-case behavior under realistic radio propagation, whose conclusions in this analysis are confirmed by experiment.

2. LIMITATIONS

While the intent of this paper is to model carrier sense as it would behave on real networks using 802.11 and similar hardware, there are ways in which the model differs from reality, and assumptions that the model makes that are not true of all networks. This section outlines the main limitations.

The model assumes that a sender chooses the modulation that results in the fastest delivery of bits to the receiver; that is, the sender adapts its bitrate to the receiver's SNR, including both noise and interference. This assumption is important to the paper's results. Existing bitrate selection algorithms often react slowly to changing conditions, but this is an area of active research, with encouraging recent work [25].

The paper assumes that radios can achieve throughputs that, as a function of SINR, are proportional to Shannon's capacity bound. Most efficient modulation schemes track the general shape of Shannon capacity even though they lag it by some capacity or SNR penalty, and with the advent of turbo codes, the gap is rapidly closing [4].

The paper assumes that the overall impact of multipath fading is limited to a few dB, which generally holds true for wideband modulations like 802.11's.

The paper considers only situations with two contending senders. For a handful of active senders, the analysis is far more complicated and as yet unfinished, but the conclusions appear to generalize. Many wireless LANs fit this description [4]. The paper's conclusions would likely not hold for large numbers of nearby senders, for example dozens.

Real radios often have carrier sense algorithms more complex than the model's power level thresholding, involving, for example, correlation on preamble patterns. It is not clear whether this is important, however, since correlation strength is strongly a function of signal strength.

The model assumes that competing senders agree on the choice between multiplexing and concurrency, and that when multiplexing, the senders take turns. Real senders will occasionally disagree, for example, if they experience different levels of background noise.

The model assumes that a receiver, when faced with concurrent transmissions, will decode the packet intended for it, if SNR allows, rather than a concurrent packet intended for a different receiver. Some radio hardware is capable of this, such as our AR521x-based 802.11 cards (though the experiments did not use this functionality), but not all.

The model ignores protocol-specific MAC-level mechanisms such as ACK, RTS, CTS, and exponential backoff. These make senders' responses to concurrency more complex than in the model and can potentially complicate or restrict concurrent transmission. Also, the model's throughput calculations do not include overheads caused by control frames, slotting, backoff, and preambles/headers sent at lower bitrates.

The model assumes that all nodes have omni-directional antennas, and the paper assumes all nodes use the same fixed transmit power level. This is typical of most WLANs.

3. THEORETICAL MODEL

3.1 Introduction

A sender's MAC relies on carrier sense to decide whether to transmit at any given time. The decision involves two questions: "Is my intended receiver listening?" and "Is it a good idea to transmit now, given the interference level?" This paper focuses on the latter.

There are several ways to implement carrier sense. The simplest approach is for the sender to monitor the received signal level and to postpone transmission while this power is above some threshold. The MAC combines this with scheduling and back-off mechanisms to approximate fairness among competing senders. In practice, standards (e.g. 802.11 [12]) and implementations are somewhat more complex. The common thread, though, is comparison of a received power metric at the sender against some

fixed threshold, independent of the identities of the sender and receiver and independent of the details of the packet.

For the sake of tractability, this paper considers only two competing sender-receiver pairs. The senders agree to transmit concurrently or to multiplex, depending on the (equal) signal power levels each senses from the other compared to a fixed threshold.

The remainder of this section presents a formal model and uses it to analyze how close carrier sense comes to optimal throughput and how sensitive it is to the choice of threshold.

3.2 Formal Model

3.2.1 Overview

This section presents a model of MAC throughput under different concurrency policies, comparing carrier sense to pure concurrency, pure multiplexing, and optimal selection between concurrency and multiplexing. The model considers expected throughput averaged over all possible configurations of a network, given a set of parameters. Figure 1 shows the basic scenario: A sender communicates with a receiver found at some random location within a maximum radius R_{max} in two-dimensional space. The R_{max} parameter represents the maximum range of interest, beyond which either capacity is too low to be useful or no nodes of interest are found. The two senders themselves are separated by distance D . For a particular D and R_{max} , the model calculates expected total throughput by averaging over all possible receiver positions. The parameters R_{max} and D , for which the analysis does not assume any prior distribution, will turn out to be key in determining network efficiency and carrier sense behavior.




Figure 1 – Generalized model scenario – two senders and their receivers, illustrating the geometry variables used in the model.

Section 2 outlines the main assumptions required by the model. In addition, the model assumes that receivers are distributed independently and uniformly within R_{max} of their senders. This assumption is reasonable for access point networks but unlikely to be true for multi-hop routing. For senders with many receivers scattered around them, the model directly reflects average sender throughput. For senders with only a few possible receivers, such as residential access point networks, the model predicts the average over an ensemble of networks.

3.2.2 Modeling Radio Capacity

The effectiveness of carrier sense ultimately hinges on the properties of radio propagation. The basic path loss – shadowing – fading model (e.g. [1] [22]) breaks radio propagation into three components: "path loss", a deterministic power law decay with distance, "shadowing", a lognormal statistical variation from place to place due to obstacles and reflections, and "fading", a fine-grained Rayleigh or Rician statistical variation in both space

and frequency. The path loss power law exponent α typically varies from 2 to 4 [22] (p166) [13], and shadowing standard deviation σ is typically in the realm of 4-12dB [1] [13] [6] (§4.7.6).¹ Wideband channels, employed by modern packet radios such as 802.11, allow fading to be largely averaged away compared to the variance due to shadowing.

The other main ingredient in a capacity model is a means to estimate average achievable throughput given signal power, interference power, and a reasonable bitrate adaptation algorithm (such as SampleRate [3]). Shannon's capacity formula $Capacity / Bandwidth(Hz) = \log(1 + SNR)$ is an upper bound but in practice can be used as a rough proportional estimate. This paper includes interference in the noise term of Shannon's formula, giving $\log(1 + SINR)$, which is a theoretically reasonable simplification.

To achieve such throughputs, bitrate adaptation must converge well before conditions, including interference, have changed. A sufficient criterion would be to assume long-running, saturating flows. However, any stable state where the preponderance of time is spent under consistent channel conditions would suffice. A robust adaptation algorithm like SampleRate can optimize for the common case, without requiring excessive feedback and without serious penalty to throughput.

3.2.3 Analytical formulation

The model places one sender at the origin and its receiver at any polar coordinates (r, θ) with $r < R_{max}$. The second sender, the "interferer", lies on the x-axis at $(D, 0)$.

The model formulates carrier sense capacity as a piecewise function, switching between concurrent transmission and multiplexing depending on the interferer signal strength. Thus the model starts with expressions for concurrent and multiplexing capacity.

For the cases of concurrent transmission, multiplexing, and carrier sensing, the location of the interferer's intended receiver does not matter to the sender – the allocation of channel resources and distribution of interference are unchanged. These cases are symmetric, allowing the model to consider throughput one sender-receiver pair at a time.

On the other hand, any notion of "optimal" must consider both sender-receiver pairs because there may be a need to trade off a small loss at one for a large gain at the other. The model's optimal throughput is maximum sum of throughputs on the two links, subject to the fairness constraint that the senders get equal time on the air. Thus the model's hypothetical optimal MAC must choose whether to transmit concurrently or to multiplex equally between the two senders, whichever maximizes the resulting sum of throughputs. This choice depends on both receivers' locations, as well as their shadowing and fading environments.

The following formulas give an individual sender-receiver pair's throughput, which the model will average over $r = 0..R_{max}$, $\theta = 0..2\pi$. A single sender without any competition has a throughput proportional to the Shannon capacity:

$$C_{single}(r, \theta) = \log(1 + P_0 \cdot r^{-\alpha} \cdot L_\sigma / N_0)$$

Where P_0 is the signal power at unit distance, L_σ is a random variable drawn from the lognormal shadowing distribution, and N_0 represents the thermal noise floor in the channel. Without loss of generality, P_0 can be factored into the noise term, yielding

$$C_{single}(r, \theta) = \log(1 + r^{-\alpha} \cdot L_\sigma / N)$$

For most calculations, the model will use -65dB for $N = N_0 / P_0$. The precise value of N is not important; changing the power level (or noise floor), technically given in units of $(power) \cdot (distance)^\alpha$, is equivalent to rescaling the distances: scaling P_0 / N_0 by ρ is equivalent to scaling distance by $\rho^{-1/\alpha}$. $N = -65dB$ is convenient for networks like 802.11 (which transmit around 15dBm and have a noise floor around -95dBm), because it scales the units of distance roughly as meters. This can be verified empirically or through appropriate use of the Friis transmission equation.

For this power level, under a typical $\alpha = 3$, $r = 20$ gives roughly 26dB SNR, which is reasonable for 802.11a/g 54Mbps. Most of the paper's analyses will consider r up to 120, corresponding to an SNR just shy of 3dB, about the minimum for one-megabit 802.11b.

When two senders of equal transmit power are present, an ideal TDMA MAC gives each a throughput of

$$C_{multiplexing}(r, \theta) = \log(1 + r^{-\alpha} \cdot L_\sigma / N) / 2$$

If instead, the senders transmit concurrently, each sender contributes to the noise at the others' receiver, giving each receiver a throughput of

$$C_{concurrent}(r, \theta) = \log(1 + r^{-\alpha} \cdot L_\sigma / [N + L'_\sigma \cdot (\Delta r)^{-\alpha}])$$

where $\Delta r = \sqrt{[(r \cdot \cos(\theta) - D)^2 + (r \cdot \sin(\theta))^2]}$ (distance between interferer and receiver) and where L_σ and L'_σ are independent values drawn from the same lognormal shadowing distribution.

For a carrier sensing MAC,

$$C_{cs}(r, \theta) = \begin{cases} C_{multiplexing}(r, \theta) & \left| \begin{array}{l} D^{-\alpha} \cdot L''_\sigma > P_{threshold} \\ D^{-\alpha} \cdot L''_\sigma < P_{threshold} \end{array} \right. \\ C_{concurrent}(r, \theta) \end{cases}$$

for some specified threshold power $P_{threshold}$. The threshold can instead be stated as a distance, $D_{threshold} = P_{threshold}^{1/\alpha}$ which, in the absence of shadowing, is the separation distance at which the two senders begin to transmit concurrently.

An optimal MAC would achieve an average throughput of:

$$C_{max}(r_1, \theta_1, r_2, \theta_2) = \frac{1}{2} \cdot \text{Max} \left[\begin{array}{l} C_{concurrent}(r_1, \theta_1) + C_{concurrent}(r_2, \theta_2), \\ C_{multiplexing}(r_1, \theta_1) + C_{multiplexing}(r_2, \theta_2) \end{array} \right]$$

As explained above, this depends jointly on the positions of both pairs of nodes. The paper will sometimes use an upper bound on optimal capacity that ignores the negative effects of optimizing for one receiver without regard for the other:

$$C_{UBmax}(r, \theta) = \text{Max}[C_{concurrent}(r, \theta), C_{multiplexing}(r, \theta)]$$

In all cases, the expected throughput is found by integrating over the area of the R_{max} -radius circle around the sender:

¹ Applied to our own indoor 802.11 testbed at 2.4GHz, we find $\alpha \approx 3.5$, $\sigma \approx 10dB$ and a reasonably good fit.

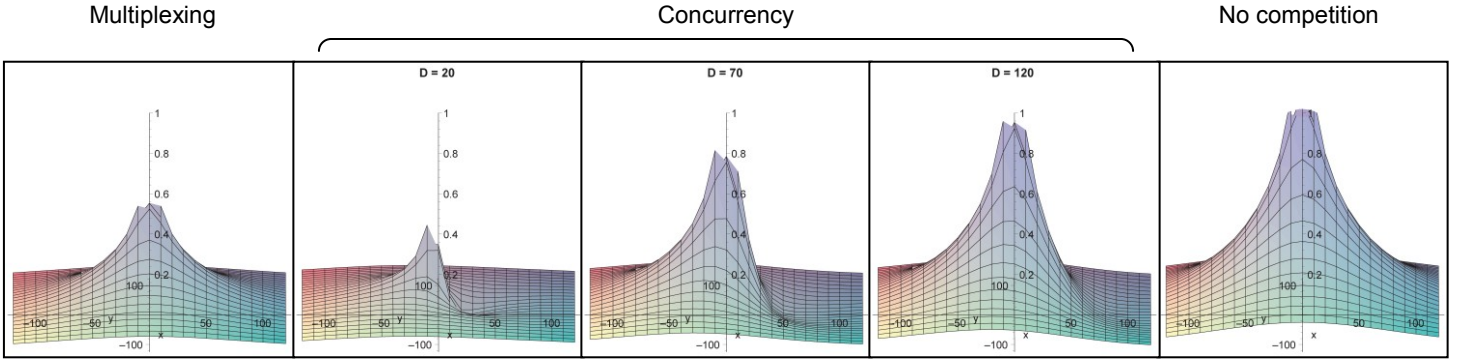


Figure 2 – Capacity “landscape” $C_i(r, \theta)$; $\alpha = 3$, $\sigma = 0$, $P_0 / N_0 = 65\text{dB}$. Depicts capacity as a function of receiver position (x-y plane), with the sender at the origin and an interferer on the x-axis at distance D . D is identified explicitly for the concurrency cases, but not for multiplexing, where capacity is independent of D . The “no competition” case is equivalent to concurrency with $D \rightarrow \infty$. Distance units are in multiples of the $P = 65\text{dB}$ unit distance.

$$\langle C_i \rangle (R_{\max}, D) = 1/\pi R_{\max}^2 \cdot \int_0^{R_{\max}} \int_0^{2\pi} C_i(r, \theta) r \cdot d\theta dr$$

In most cases, this integral cannot be evaluated analytically, but we can investigate it numerically in the physical regimes of interest.

3.2.4 The Capacity Landscape

While most of this paper will present capacity averaged over R_{\max} -radius circles, this section first illustrates how capacity varies within such a circle, before averaging. Figure 2 shows how capacity varies with the positions of the receiver and the interfering sender. Each plot’s x- and y-axes indicate receiver position, with the sender at (0,0). The z-axis indicates $C_i(r, \theta)$. The plots differ in their choice of interferer distance D and concurrency or multiplexing. These plots include no averaging and no shadowing.

All the graphs show high capacity for receivers near the transmitter, rising to an unbounded peak at the transmitter. The multiplexing graph has capacities that are independent of the location of the interferer, sloping down smoothly with the receiver’s distance from the transmitter, everywhere providing half the capacity of the contention-free channel. Capacity under concurrency, however, depends on the interferer’s location. The three concurrency plots illustrate capacity for three values of interferer distance D . In each case, a dip in capacity is visible for receiver positions near the interferer. Capacity throughout the landscape increases as the interferer gets farther from the sender.

3.2.5 Carrier Sense Behavior: A First Look

How well does carrier sense perform with two competing sender-receiver pairs, according to the model? There are two easy limiting cases. First, when the pairs are very far apart ($D \gg R_{\max}$), concurrency is clearly optimal – for all receiver locations within R_{\max} , the pairs interfere insignificantly. Carrier sense handles this properly, since sufficiently distant senders will have mutually sensed power less than any reasonable sense threshold. Second, when the pairs are very close together ($D \ll R_{\max}$), multiplexing is nearly always optimal, since concurrent transmission would lead to an SNR approaching 0dB, regardless of receiver position. For sufficiently close separations, carrier sense also makes the right choice. In both these limiting cases, for a given D and R_{\max} , the same choice is best for essentially all receiver locations: either they all prefer concurrency or they all prefer multiplexing. This property is favorable for carrier sense, since carrier sense chooses without knowing the receiver conditions. So long as all receivers agree, carrier sense can be made optimal with an appropriate choice of threshold.

As results in subsequent sections show (e.g. Figure 4), for reasonable parameter regimes, these “near” and “far” limiting cases or close approximations to them cover much of the relevant range of D . In the “transition region” between the two limits, however, some receiver locations do best with concurrency and others with multiplexing; no one threshold can satisfy them all. Figure 3 illustrates the geographical distribution of receiver preferences, for the same model scenario as Figure 2; white and light grey identify a preference for multiplexing, while dark grey




Figure 3 – Receiver preference regions: a receiver in the dark shaded areas prefers concurrency, in the light shaded areas prefers multiplexing, and in the white areas prefers multiplexing and will be starved ($<10\%$ of $C_{UB\max}$) without it. Circle marks interferer position. $\alpha = 3$, $\sigma = 0$, $P_0 / N_0 = 65\text{dB}$.




Figure 4 – MAC throughput curves for the non-shadowing model, averaged over an R_{\max} -radius network for three different values of R_{\max} . $\alpha = 3$, $P_0 / N_0 = 65\text{dB}$.

marks preference for concurrency. For a nearby interferer at $D = 20$, multiplexing is optimal for all R_{\max} up to about 100. Similarly, for a distant interferer at $D = 120$, concurrency is optimal for all R_{\max} up to about 50. In between, however, at $D = 55$, half the receivers prefer concurrency and half multiplexing. Carrier sense, since it decides without knowledge of receiver conditions, cannot make the best choice for all receivers. This illustration does not quantify the magnitude of the throughput sacrificed, but it does illustrate the basic dilemma.

Besides the distance between transmitters, another crucial factor in MAC behavior is the distance from the sender to the farthest receivers. In “long range” networks (R_{\max} large, hence many receivers are relatively far from their sender), links are weak and network capacity tends to be dominated by the noise floor (thermal noise), even in the presence of interference. Receivers close to the interferer see significant interference, but most receivers are farther away, where the interference is weak and blends into the noise floor; the effects of interference are relatively localized to the vicinity of the interferer. This case is difficult for carrier sense, because receivers close to the interferer prefer multiplexing while the rest prefer concurrency. On the other hand, in “short range” networks (R_{\max} small, so all the receivers are near their sender), an interferer close enough to cause trouble will have roughly the same effect on all receivers. Once interference is strong enough to matter to some receivers, all receivers feel its effects, because it is already too powerful for path loss to drive it into the noise floor within R_{\max} . Carrier sense performs well in this case.

3.2.6 Performance Overview

The tables below summarize the model’s quantitative results. They report carrier sense throughput as a percentage of optimal MAC throughput, computed in Maple with Monte Carlo integration, across a representative set of points in the model’s parameter space: short and long range R_{\max} (within the bounds of typical WLAN capabilities), near, transition, and far interferer distance D , and values of α and σ around typical real-world values.

The tables answer two questions: is carrier sense adequate even when sub-optimal, that is, in the transition region and at long range? And does that performance hold up without specially tuning the sense threshold for each different environment?

The first table shows carrier sense efficiency with a fixed sense threshold $D_{\text{threshold}} = 55$, and with $\alpha = 3$ and $\sigma = 8\text{dB}$:

Table 1 – Carrier sense throughput efficiency as a percentage of optimal throughput; generic threshold $D_{\text{threshold}} = 55$.

$R_{\max} \setminus D$	20	55	120
20	96%	88%	96%
40	96%	87%	96%
120	89%	83%	92%

Carrier sense performs very well: it achieves a high fraction of the optimal MAC throughput. Its weakest areas are in the sender/interferer distance transition region ($D = 55$) and at long range ($R_{\max} = 120$).

For comparison, the next table shows the same configurations as above but with thresholds optimized for each value of R_{\max} by the criteria in Section 3.3.3:

Table 2 – Carrier sense throughput efficiency as a percentage of optimal throughput; environment-optimized thresholds.

$R_{\max} \setminus D$	20	55	120
20 ($D_{\text{thresh}} = 40$)	93%	91%	99%
40 ($D_{\text{thresh}} = 55$)	96%	87%	96%
120 ($D_{\text{thresh}} = 60$)	89%	83%	92%

Note that the performance numbers are not strictly better because shadowing eliminates the notion of a globally unique optimal threshold; a trade-off inevitably exists between conservatively favoring small D and aggressively favoring large D (see Section 3.4). In any case, efficiency with optimized thresholds is only modestly better than without in Table 1. Carrier sense in these configurations provides good average performance that is robust in the face of mis-chosen thresholds.

Within typical propagation parameters, the above conclusions are surprisingly robust. Very little change results from varying α from 2 to 4 and σ from 4dB to 12dB. Smaller α tends to make a network look more short range, and larger α more long range, but the qualitative features are the same and the performance metrics similar. The model on four parameters we argued captured the essential features of carrier sense thus effectively collapses to two parameters: range (short and long) and interferer distance (near, transition, and far). Focusing now on these two parameters, we discuss the underlying causes below.




Figure 5 – Shaded plot of $R_{\max} = 55$, non-shadowing. The vertical line marks the carrier sense threshold. Dark shading is inefficiency under multiplexing and light shading is inefficiency under concurrency, of which the light shaded “triangle” below the multiplexing line represents inefficiency due to poor threshold choice.

3.3 Throughput without Shadowing

Why does carrier sense perform so well, and why does this good performance not require careful threshold selection? This section qualitatively explores these questions using the model simplified by eliminating shadowing ($\sigma = 0$). Section 3.4 reintroduces shadowing.

3.3.1 Carrier Sense vs. Optimal

Figure 4 plots MAC throughput as a function of inter-sender distance D . Each point shows the model’s predicted throughput averaged over all receiver locations within a given R_{\max} . The three curves in each graph show throughput with multiplexing ($C_{\text{multiplexing}}$), concurrent transmission ($C_{\text{concurrent}}$), and the optimal MAC policy (C_{max}). The latter chooses the better of multiplexing and concurrency for each individual receiver, prior to averaging. The horizontal scale is in multiples of the 65dB distance (see Section 3.2.3), while the vertical scale is normalized as a fraction of $R_{\max} = 20$, $D = \infty$ throughput.

Multiplexing throughput in Figure 4 is independent of inter-sender distance. Concurrency throughput varies from low when the senders are coincident towards twice the multiplexing throughput when the senders are widely separated. The graph does not separately plot carrier sense throughput; it is the same as multiplexing to the left of the threshold $D_{\text{threshold}}$ and the same as concurrency to the right. This is, in general, slightly less than the optimal throughput (C_{max}) curve.

Carrier sense’s throughput approaches optimal at both ends of the graph, converging to the concurrency throughput for large D and the multiplexing throughput for small D . In the transition region between the two extremes, receivers disagree on the best choice, those nearer the interferer preferring multiplexing and those farther preferring concurrency. The need to compromise on a single threshold for all receivers results in the gap between carrier sense and optimal throughput.

Although different receiver locations do disagree within the transition region, bitrate adaptation reduces the loss of throughput for the receivers that don’t get their choice. “Hidden terminal” scenarios, where a receiver prefers multiplexing but doesn’t get it, are not usually a disastrous case of “concurrency makes transmissions fail” but rather “a less-than-ideal bitrate is needed to succeed”. Similarly, “exposed terminals” are not so much “concurrent transmissions would succeed” but rather “somewhat better throughput would be achieved with concurrency, though at a lower bitrate”. Given properly chosen bitrates, the right way to view hidden and exposed terminals is that they reduce efficiency. “Hidden terminal inefficiency” is the gap between carrier sense and optimal throughput to the right of the threshold in Figure 5

(shaded lightly), and “exposed terminal inefficiency” is the gap between carrier sense and optimal to the left of the threshold (shaded darkly).

3.3.2 Transition Region Performance

This section discusses why carrier sense performs well in the transition region between small and large D , when some receiver locations benefit most from concurrency but others prefer multiplexing.

One of the most important reasons is the “non-locality” of interference – that its effects extend indefinitely in space, decaying smoothly with distance, rather than being localized to a cookie-cutter region. The significance of this non-locality is that the interference power sensed by the sender is not dramatically different from that seen by receivers within R_{\max} , and different receiver locations see interference levels that do not differ dramatically from each other. Receivers that prefer concurrency see enough interference that multiplexing doesn’t provide too much less throughput; receivers that prefer multiplexing don’t see so much interference that concurrency is unacceptable.

The “landscape” plots of Figure 2 illustrate this non-locality – comparing the concurrency plots against the “no competition” plot shows that interference decreases throughput for all receiver locations, by an amount that decreases smoothly with distance. Throughput decreases gradually for receivers closer and closer to the interferer. A real protocol would implement this by adaptively lowering the bitrate with decreasing SNR, so that even if carrier sense chooses concurrency, receivers near the interferer still get some throughput. In contrast, a fixed-bitrate protocol would generally exhibit a sudden large decrease in throughput for locations close enough to the interferer that the SNR is not able to support the bitrate. Such an SNR boundary would present a difficult situation for carrier sense – no one threshold could satisfy receivers on both sides of the boundary.

Another source of non-locality is the rapid growth of surface area with distance, so that many more receiver locations are far from an interferer than close. In a model that considers only a 1D world, carrier sense consistently averages 2%-4% farther below optimal within the transition region.

The degree of locality depends on the size of the network and the distance to the interferer. Noise-dominated, long-range networks, where interference tends to decay to insignificance before reaching the far side of the network, experience much more strongly localized interference effects than short-range networks. In short range networks, if interference is significant anywhere within R_{\max} it is likely significant everywhere to roughly the same degree. If an interferer is close enough to produce a substantial differential impact across the network, the sender will likely already be using multiplexing. For these reasons carrier sense is significantly more efficient in short range networks than in long.

In short range networks, not only is average throughput good, but every receiver has a reasonable share, because whenever concurrency is employed, interferers are too far from the network to have a localized impact. In long range networks, however, while throughput remains good for most receivers, when an interferer is inside R_{\max} itself, nearby receivers get little throughput. The white regions in Figure 3, showing receivers that receive less than 10% of C_{UBmax} under concurrency, illustrate this effect. In the $D = 120$ and $D = 55$ frames, the white region has the




Figure 6 - Optimal threshold (expressed as the equivalent distance at $\alpha = 3$, for consistency) versus network radius for several values of α , with $\sigma = 8\text{db}$ (included because shadowing has a significant qualitative impact at long range). The erratic ripples on the right are artifacts of the numerical solution method. Dashed lines mark $R_{\max} = R_{\text{thresh}}$ and $R_{\max} = 2 * R_{\text{thresh}}$.

potential to be inside a sufficiently long-range network that is still operating under concurrency; the receivers in this region could reasonably be described as hidden terminals. Overall, while average throughput in long range networks remains robust, fairness can suffer.

3.3.3 Picking a Threshold

What choice of threshold is best for carrier sense efficiency? In the non-shadowing case, for a given R_{\max} , there is a threshold that simultaneously minimizes average inefficiency for all D . Fairness and the distribution of allocations do vary as a function of threshold, but in terms of average throughput alone, the optimal threshold is the D at which the concurrency and multiplexing throughput curves cross, the point where concurrency provides half of the competition-free capacity. Figure 5 illustrates this point, depicting the inefficiencies associated with a suboptimal threshold. Whenever the threshold is leftward of optimal, a “triangle” region of added hidden terminal inefficiency appears (the portion of the light shaded region below the multiplexing line). Similarly, when the threshold is rightward of optimal, a triangle region of added exposed terminal inefficiency appears. These are both eliminated at the concurrency-multiplexing intersection point, which is therefore the best possible choice.

One can thus find the optimal threshold by solving for the intersection of the concurrency and multiplexing curves. However, this relies on knowing the value of R_{\max} and the particular parameters of the propagation environment, unreasonable when a manufacturer must configure a default value. Fortunately, as Section 3.3.4 will show, the precise choice of threshold does not matter too much. As long as the threshold is somewhere in the transition region, the additional losses due to a suboptimal threshold will be small. So a good default threshold lies roughly in the middle of the span of optimal thresholds for the typical operating range of the hardware. For example, for $\alpha = 3$, 802.11g’s bitrate flexibility ranges from around $r = 20$ to $r = 120$ (as explained in Section 3.2.3). $R_{\max} = 20$ corresponds to an optimal threshold about $D_{\text{thresh}} \approx 40$, and $R_{\max} = 120$ corresponds

to $D_{\text{thresh}} \approx 75$. So, a $D_{\text{thresh}} \approx 55$ (equivalent to $P_{\text{thresh}} \approx 13\text{dB}$) would be a reasonable compromise.

Also, observe that, in the $R_{\max} = 20$ example, the optimal threshold is roughly twice R_{\max} , a demonstration of strong non-locality. In general, a short range network has an optimal threshold whose equivalent distance is well outside the network boundaries. On the other hand, a long-range network (e.g. $R_{\max} = 120$) may have an optimal threshold *inside* the network, so that the sender multiplexes only if the interferer is close enough to impact the network in all directions. One can usefully define long range to mean $R_{\text{thresh}} < R_{\max}$ and short range to mean $R_{\text{thresh}} > 2R_{\max}$.

3.3.4 Threshold Robustness

The previous section discussed why carrier sense has good performance with a threshold optimal for the propagation environment and network size. But why is performance good even when the threshold is not optimal, as the Table 1 in 3.2.6 shows?

In order for a small threshold change to produce bad behavior, the gaps between concurrency and multiplexing throughput (as in Figure 4) would need to abruptly grow large on either side of the optimal threshold. However, given adequate bitrate adaptation, the effect of interference varies gradually with interferer distance, changing only on the length scale of the network radius.²

But what about the possibility of large differences from one environment to another? Why don’t variations in α , σ , and R_{\max} require a custom-tuned threshold?

The answer appears to arise in large part from the “sweet spot” operating regime chosen by most data networking hardware. While often capable of ranging a bit wider, most data networking hardware is designed to operate in the regime roughly around 10-25dB SNR; 802.11 [12], Bluetooth, and 802.15.4 [11] all target this region. This is likely a practical consequence of capacity and economic considerations. It’s valuable to have an SNR well above the noise floor, in order to provide a reasonable amount of capacity and a reasonably affordable receiver, but increasing the SNR further leads to diminishing returns due to the logarithm in Shannon capacity, and doing so often faces regulatory and power handling constraints as well.

This operating regime fortuitously turns out to have special significance for MAC behavior: it is the intermediate region between the long range and short range limits. This can be demonstrated using the quantitative criteria proposed at the end of Section 3.3.3, identifying long range networks as those with optimal thresholds inside the network boundary and short range networks as those with optimal thresholds substantially outside the network boundary.

Figure 6 shows the variation of optimal threshold with the size of the network, tabulated at several values of α , representing different propagation environments ($\alpha = 2 - 4$ being typical). On the left lies the short range limiting behavior, thresholds scaling towards zero roughly as the square root of R_{\max} ³ and clustered

² Variations of this statement can be proven mathematically. E.g., for $\alpha = 3$, $\sigma = 0$, the slope of the concurrency curve (in our $R_{\max} = 20$ normalized capacity units) is bounded above by $1.37 / R_{\max}$ for all $D > R_{\max}$.

³ Optimal $D_{\text{thresh}} \approx e^{-1/4} R_{\max}^{1/2} N^{-1/2\alpha}$ (actual distance units, not $\alpha = 3$ equivalents) for very short range networks, which can be derived by taking the limit as $N \rightarrow 0$ and approximating Δr as D_{thresh} .




Figure 8 – MAC average throughput curves for the full model with $\sigma = 8\text{dB}$ shadowing, along with non-shadowing curves for reference. Averaged over an R_{\max} -radius network for three different values of R_{\max} . $\alpha = 3$, $P_0 / N_0 = 65\text{dB}$, as before.




Figure 7 - Conceptual graphic of propagation pathways past a barrier. The dashed arrows show desired transmissions, while the dotted arrow shows a potential source of “hidden terminal” interference. The thick, red arrows show several propagation paths for a carrier sense signal.

closely together in spite of α variation. On the right lies the long range limiting behavior, where threshold growth tapers off in R_{\max} but spreads out in α . In between, centered around $R_{\max} = 40$ or so (corresponding to $\sim 17\text{dB}$ SNR at the edge of the network), the curves show a gradual change in behavior from the one extreme to the other. Neatly enclosing this range are the two dashed lines, representing the criteria of 3.3.3, $R_{\text{thresh}} = R_{\max}$ and $R_{\text{thresh}} = 2 \cdot R_{\max}$. For typical $\alpha \approx 3$, this range is roughly $18 < R_{\max} < 60$, equivalent to $12\text{dB} < \text{SNR} < 27\text{dB}$ at the edge of the network.

Why does the fact that data networking hardware favors the intermediate regime help carrier sense? In the limit of short ranges, optimal threshold scales rapidly with network size; in this limit, while carrier sense performs well for a well-tuned threshold, thresholds are not particularly robust to uncertainty in power or distance. If, for example, transmit power were boosted by 20dB or distances were scaled down by a factor of 5, the differences between the two tables of Section 3.2.6 would be much starker. Transition region losses due to suboptimal thresholds would range around 3%-16% instead of 1%-6%.

On the other hand, in the long range limit, thresholds vary quite slowly. However, as discussed in Section 3.2.5, carrier sense efficiency falls off; optimal thresholds are robust but still perform poorly. In the middle, however, is a compromise that approaches the best of both worlds. Coincidentally, this is the primary operating regime for wireless networking gear.

3.4 Throughput with Shadowing

This section adds shadowing to the model. Real environments have obstacles and reflections, and one might expect these to produce hidden and exposed terminals. This section shows that such irregularities do not ruin the intuitions of the previous section and that carrier sense continues to provide close-to-optimal average performance.

One might imagine, for example, that a hidden terminal configuration could be constructed by inserting a barrier between senders as in Figure 7. Radio propagation, however, is not so easily confined. Most building materials are not particularly opaque to radio; typical attenuation through an interior wall is less than 10dB [6] (§4.6-4.7). However, even if the obstruction were opaque (for example, a metal barrier), reflection from the far wall could still connect the two transmitters. Reflections are difficult to suppress; typical reflection losses are less than 10dB . Yet, even if there were no far wall, only open space, a weak signal would still round the corner, as a result of diffraction. Using the knife-edge approximation and a 5-meter distance to the barrier, the diffraction loss at 2.4GHz would be around 30dB .

Propagation is complex and difficult to constrain – so complex, in fact, that one can reasonably apply the central limit theorem and lump together the effects of all barriers and reflectors into a single Gaussian random variable. This is the origin of lognormal shadowing [1]. Empirically, the standard deviation is typically of the same order of magnitude as the losses quoted above, typically $4\text{-}12\text{dB}$ (see Section 2). Such uncertainty, though significant, is not sufficient to destroy the efficiency of carrier sense.

This lognormal variation affects carrier sense by adding randomness in three ways: to the signal power at the receiver, to the interference power at the receiver, and to the interferer’s power as sensed by the transmitter. Each of these powers now has an independent random component.⁴ To estimate (somewhat pessimistically) the potential effect on a sender’s ability to estimate its receiver’s SNR, the three effects’ variances can be summed, yielding $\sigma_{\text{SNRrest}} = \sigma\sqrt{3} \approx 14\text{dB}$ uncertainty (less if the interferer power is comparable to the noise floor), assuming $\sigma = 8\text{dB}$ shadowing as below.

To help visualize the implications, consider what happens under roughly equivalent distance variations lumped into the interferer distance D .⁵ To take a concrete example, in a short range network of size $R_{\max} = 20$ with threshold $D_{\text{thresh}} = 40$ (close to the $\sigma = 0$ optimum), an interferer that, to the receiver, appeared to be at $D =$

⁴ For this analysis and throughout, we assume that the shadowing distributions are uncorrelated. This is not quite true, but it is good enough for our purposes.

⁵ For receivers approximately equidistant between sender and receiver, around 3/2-fold variation in D (equivalent to 6.5dB) covers 8dB interference power variation. Sense power uncertainty contributes 8dB directly, and effects roughly equivalent to sender power variation require just shy of 4dB . Combining these gives about 11dB , roughly $2.3\times$ distance.

20, would have about a 20% chance of appearing to the sender as beyond D_{thresh} , thereby triggering concurrent transmission. This mistake would leave the receiver with a very low, sub-0dB SNR about 20% of the time (approximately the fraction of the R_{max} disc's area closer to $D = 20$ than to the sender). Combining the probabilities, for the given interferer position, the effects of shadowing on carrier sense would cause very poor SNR in around 4% of configurations but otherwise would behave reasonably most of the time.

Figure 8 shows the results of the shadowing model, using $\sigma = 8\text{dB}$ and the same $\alpha = 3$ path loss as before. The non-shadowing results from Figure 4 are also plotted, for comparison. The results are qualitatively very similar. Short range carrier sense does well, hugging the optimal curve closely, while long range is less efficient. Best performance is found in the limits of near and far interferer distance, while some efficiency is lost in the middle. Shadowing causes carrier sense to make more mistakes, and it introduces more disagreement among receiver locations, but Figure 8 shows that in this scenario shadowing introduces a relatively small gap between optimal and carrier sense.

4. EXPERIMENTAL RESULTS

This section presents experimental results from an indoor testbed, comparing the qualitative predictions of the model with the behavior of real radios. We find mainly broad agreement, as well as a few interesting discrepancies.

The experiments measure average throughput for competing sender-receiver pairs under concurrency, multiplexing, and carrier sense. The tests use an indoor testbed of Atheros AR5212 and AR5213-based 802.11a/b/g wireless adapters each with a “rubber duck”-style antenna, installed in roughly 50 Soekris single-board computers running kernel-mode Click [16], scattered over two closely-coupled floors of a large, modern office building. All experiments run in 11a mode.

To measure throughput, each of the two senders attempts to broadcast 1400-byte packets continuously for 15 seconds, and with each sender's receiver counting the number of packets received. For concurrent throughput, we disable carrier sense so both senders transmit simultaneously. For multiplexing throughput, we run each sender-receiver pair alone, one after another. For carrier sense, we enable default hardware carrier sense. To determine throughput given an optimal bitrate, we repeat every run at each of 6, 9, 12, 18, and 24Mbps, independently identifying the maximum throughput bitrate for each transmitter. Driver problems prevented use of lower and higher bitrates.

Rather than communicating with all receivers within a given R_{max} , each sender communicates only with receivers to which the 6 Mbps bitrate works well. When mimicing short range scenarios, senders communicate only with receivers to which more than 94% of packets are delivered at 6 Mbps; this results in an average SNR of about 27dB, which is roughly similar to that of an $R_{\text{max}} = 30$ model network. Long range scenarios also include receivers with between 80% and 95% delivery probabilities (average SNR about 16dB, roughly similar to $R_{\text{max}} = 70$).

4.1 Short Range

Figure 9 and Figure 10 show the results of the short range (links 94% at 6Mbps) experiment. The first graph compares the throughputs of pure multiplexing and pure concurrency against

hardware carrier sense (CS). Each set of three vertically collinear points represents one pair of transmitter-receiver pairs. The pink square's y value shows total (combined) throughput in packets per second for concurrent transmission, the blue diamond's y value shows total throughput for multiplexing, and the yellow triangle shows carrier sense throughput. The pairs are sorted by CS throughput, and plotted using CS throughput as the x value. One would expect to see the each CS point coinciding with either the corresponding concurrency point or the corresponding multiplexing point. To the extent that carrier sense works well, the other point should be below the CS point. For most pairs this is true: carrier sense is good at making the optimal choice.

The second figure plots the same data with the x axis indicating average RSSI between the two senders (measured separately). RSSI should roughly reflect sender-sender distance, and it is also the principal basis for the hardware's carrier sense decision. The vertical bars highlight points where CS performance is below optimal. At the very right (plotted in the column above zero, spread randomly for visibility) are the points for which no test packets were received, implying a signal strength somewhere below the receive threshold.




Figure 9 – Experimental short range comparison: multiplexing and concurrency throughput vs. carrier sense.




Figure 10 – Experimental short range throughput vs. sender-sender mean RSSI.

Because link quality for all pairs is quite good, the main performance-limiting factor in these graphs is pair-pair interference, with which both x-axes are strongly correlated. The two plots show roughly the same set of features: a close-distance region on the left, where multiplexing and CS performance

coincide and concurrency does quite poorly, a transition region in the middle (roughly 20dB-10dB in Figure 10), where concurrent performance catches up and sometimes exceeds both CS and multiplexing, and a long-distance region on the right, where CS and concurrent performance coincide and multiplexing lags behind by a factor approaching two. This breakdown is as the theory predicts.

An important feature here is that, in this short range data set, carrier sense is quite infrequently bested by multiplexing or concurrency. Even when it is, the gains are not especially compelling, limited to a small set of weakly exposed terminals. Averaging throughput over all runs (that is, over the slightly arbitrary ensemble of pairs and distances present in the testbed) yields an average carrier sense throughput of 1703 packets/second, and an average optimal (maximum of multiplexing and concurrent) throughput of 1753 packets/second. Carrier sense approaches the optimal strategy quite closely, consistent with the model's predictions in the short-range case.

One notable oddity is that, even in cases of very wide pair-pair separation, CS performance often slightly exceeds concurrency. This conceivably might be due to time variation in the channel, properly exploited by carrier sense, causing concurrency to be suboptimal for a small fraction of the time. Alternatively, there may be some subtle experimental bias.

4.2 Long Range

Figure 11 and Figure 12 below show the results of the long range (links 80% to 95% at 6Mbps) experiment. The format is the same as the short range plots above. This dataset is not quite as "long range" as we would like, but pushing farther into the long range regime runs up against the limits of bitrate adaptability in 11a mode, forcing many links to operate at the minimum 6Mbps even without interference. This both hurts concurrent performance and introduces behavior intermediate between variable bitrate and fixed bitrate, which we haven't attempted to model.

Although not as clear-cut as the short-range case, this long-range dataset shows a similar set of features: a close-distance region where multiplexing is preferred, a transition region (roughly 10dB - 5dB in Figure 12), and a long-distance region where concurrency is preferred. In the plot against RSSI, these are cleanly visible from left to right. Note that the transition region mistakes consist mainly of undesirable concurrency - "hidden terminals" - rather than undesirable multiplexing, as the theory predicts given a threshold optimized for the "average case" rather than specifically for long range. Also note that the transition region, located just shy of 10dB, is several dB lower than the roughly 15dB of the short-range case. Again, we would expect this sort of shift from the theory.




Figure 11 – Experimental long range comparison: multiplexing and concurrency throughput vs. carrier sense.




Figure 12 – Experimental long range throughput vs. sender-sender mean RSSI.

In Figure 11, however, the regions are somewhat muddled together. Many of the transition region's cases of undesirable concurrency result in low throughput, and these cases end up on the left of the plot, rather than in the middle.

One oddity about many of these "undesirable concurrency cases" is that they are intermediate in throughput between pure concurrency and pure multiplexing. Many of the corresponding pure concurrency cases have near-zero throughput, possibly due to limited bitrate flexibility or the receive hardware's inability to listen to the relevant concurrent packet, but the CS cases fare somewhat better. One possible explanation is that the CS decision itself is fluttering due to noise, and so these cases aren't actually pure concurrency but instead are a mixture. Another possible explanation is asymmetric carrier sensing, where one node defers while the other does not, leading to a mix of concurrency and unfair multiplexing.

Although carrier sense in the long-range here is not quite as close to optimal as it was in the short-range (as the theory would predict), it is still quite good overall and significantly better than either pure multiplexing or pure concurrency. Averaging throughput over all runs, as before, yields an average carrier sense throughput of 923 packets/second and an average optimal throughput of 1029.

4.3 Summary

These experiments show that the model’s prediction of good average performance for carrier sense is consistent with observed behavior in commodity hardware, within the hardware’s bitrate adaptation range. Short range performance is particularly good and free of starvation, while long range performance is solid on average, with a limited number of cases of low throughput. The experiments show carrier sense behavior splitting up as a function of interferer distance into three distinct regimes, near, intermediate, and far, just as the theory claims, with modest inefficiency in the intermediate regime but nearly optimal performance near and far. Thus, within the parameter range explored, the experimental results concur with the theory, supporting its claim that carrier sense’s average throughput leaves limited room for improvement.

5. DISCUSSION AND FUTURE WORK

Under reasonable assumptions, the preceding analysis and experimental results show that carrier sense has good overall performance, both in theory and in practice, solving the spatial reuse problem with an efficiency approaching optimal. With adaptive bitrate, there is little need for threshold tuning in WLAN-like scenarios. There are also no dramatic losses of efficiency due to differing channel conditions among the nodes, the differences that Karn first warned of [15]. Better exploitation of hidden and exposed terminals could improve behavior in several corner cases, but it would have little effect on average performance. Because the corner cases are infrequent, it is fair to address them with blunt mechanisms that are wasteful of bandwidth, for example, RTS/CTS-like reservations, provided we only engage said mechanisms when necessary. Better treatment of hidden terminals, in particular, could improve fairness and reliability at long range – though long range links are still likely to be failure-prone without much deeper bitrate adaptation than commonly provided today.

Pursuing exposed terminals, however, seems significantly less interesting. Unlike hidden terminals, which, when they do arise (see Sections 3.3.3 and 3.4), can be acute, localized failures, exposed terminals are merely diffuse opportunities for modest improvements in throughput. The primary reason why these results are so much more pessimistic than earlier claims, for example [24], is that they assume adaptive bitrate. Both adaptive bitrate and increased concurrency can exploit excess channel capacity, but unless nodes are widely separated or SNRs are low, adaptive bitrate is strictly more efficient. In informal experiments using the same test set-up as Section 4.1 (short range, where exposed terminal scenarios are found), exploiting even the weak bitrate adaptation the driver supports (6, 9, 12, 18, and 24 Mbps) more than doubles average throughput compared to the base rate. The alternative, perfectly exploiting exposed terminals, provides only 10% increased throughput.⁶ Attempting to combine the two, exploiting exposed terminals on top of bitrate adaptation, yields only about 3% more than bitrate adaptation alone. Worse, injecting new concurrency into the exposed terminal case makes

⁶ The results [24] reported were higher by small constant multipliers. This appears to be because they include results only for high levels of concurrency; as they show, the potential gains of exposed terminal exploitation increase with increasing concurrency. Their best result, 47% average improvement, required *six* concurrent senders – and [4] suggests this level of concurrency is quite a rarity in typical WLAN deployments.

bitrate adaptation there as complicated and interference-dependent as the hidden terminal case; if the paper’s assumptions about bitrate adaptation break down, exploiting exposed terminals might actually cause throughput to fall.

Finally, carrier sense implementations can exhibit several notable pathologies not captured by the theoretical model, such as threshold asymmetry (where only one node defers to the other), slot collisions (where two nodes that normally multiplex transmit concurrently due to randomly picking identical starting times from a limited pool of MAC-defined slots), and chain collisions (where nodes transmit after failing to detect packets whose preamble signatures were hidden beneath another transmission, thereby hiding their own preambles and perpetuating the pathological state). Other interesting issues include carrier sense compatibility across unrelated PHY schemes (generally lacking at present) and challenges associated with convergent flows to a single node, where signals may be weak but the concurrency decision must be based not on carrier sensing but on the contents of the demodulated MAC header (an issue which 802.11 particularly exacerbates by transmitting the MAC header at full data bitrate rather than at base rate like the PLCP header). We have several ideas for possible future research to address these issues, if demonstrated to warrant attention.

6. RELATED WORK

Carrier sense has a long history. It has been in wide use in wired networks since at least [17] and has been applied to wireless networks in various forms, e.g. [20] [2] [7] [24] [19]. Early efforts focused on the hidden terminal problem and the inability of nodes to detect certain nearby transmissions. These difficulties were likely due to the fixed-bitrate, narrowband (and hence fading-prone) modulations that the earlier generations of hardware were limited to. [2], in particular, used fixed-bitrate modulation and an unusual “near-field” radio technology that, although avoiding fading, had an extraordinarily high path loss exponent that made the hidden terminal problem very real.

A large body of prior work explores analytical and simulation-based approaches to modeling and improving carrier sense-based MACs, but these focus mainly on fixed-bitrate systems (e.g. [27] [21]) and fixed-bitrate, multi-hop (e.g. [9] [26]). Bitrate adaptation is a key reason the results presented here are so different, aside from our use of a more realistic radio propagation model with shadowing and our emphasis on average-case, single-hop scenarios at the expense of multi-hop.

A number of these works (e.g. [27] [26] [21]) advocate dynamic threshold adaptation for carrier sense. This paper’s results argue that adaptation on the basis of environment is largely unnecessary in the adaptive bitrate case. On the other hand, we do not exclude adaptation to exploit statistical variation among networks with small numbers of nodes – though we do establish restrictive upper bounds on average potential improvement.

Another line of work augments carrier sense with new mechanisms or replaces it altogether. [14] suggested that bitrate-dependent tuning might improve carrier sense performance, noting the very high exposed terminal concurrencies possible at low bitrates on high SNR links. However, increasing bitrate would improve performance more than increasing concurrency. In [19] the Overlay MAC Layer scheduled transmissions to provide impressive fairness properties but did not improve throughput. [24] showed improved throughput with a contention graph-based

MAC protocol, in the context of fixed low bitrate and high concurrency ($n \geq 3$), on a network capable of supporting much higher bitrates.

For some network structures, however, the assumptions of the model presented in this paper are unjustified, and the conclusions do not necessarily hold. For example, [18] uses highly directional antennas for long, dedicated links spanning tens of kilometers. For such networks, non-carrier-sense MAC protocols may well be appropriate. Saturated multi-hop paths are also out of scope, because they violate the independent receiver distribution assumption.

Information theoretically, it is possible to achieve higher throughputs than our model here allows [8], using physical layer techniques such as joint decoding and interference cancellation [23]. It is not obvious whether meaningful additional capacity is available in practical data networking scenarios – though one might be able use techniques like those of this paper to conduct an analysis.

7. CONCLUSION

This paper argues theoretically and experimentally that carrier sense achieves near-optimal average throughput in common network environments. The main reasons are that interference has global enough reach that senders usually agree with receivers about which of concurrency or multiplexing is best, and adaptive bitrate partially compensates for senders' incorrect choices. The smooth decay of interference with distance and the intermediate-range SNR regime targeted by data networking hardware make carrier sense's performance robust in the face of imperfect choice of sense threshold. Shadowing by environmental obstacles does contribute to reduced performance, but its impact is too weak to make a dramatic difference. Hidden terminals rarely decrease mean throughput significantly. Exposed terminals rarely hurt anyone.

The distinction between short and long range networks is a key contribution of this work. In short range networks, carrier sense performs superbly. Long range networks, owing to the increasing locality of interference, are more challenging and do include corner cases with poor fairness (on top of the unreliability already present at long range), but on average, performance is still good.

This paper's results differ from prior work in part because it relies on radios with adaptive bitrate, in part because it focuses specifically on point-to-point topologies, in part because it considers average conditions rather than worst case, and in part because time has allowed radio hardware to mature, eliminating awkward quirks that plagued earlier research, such as deep fading.

8. REFERENCES

- [1] Akaiwa, Y. *Introduction to Digital Mobile Communication*. Wiley-Interscience, 1997.
- [2] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. "MACAW: A Media Access Protocol for Wireless LANs". ACM SIGCOMM 1994.
- [3] J. Bicket. "Bit-rate Selection in Wireless Networks, MIT Master's Thesis, February 2005.
- [4] A. Burr. "Turbo-codes: the ultimate error control codes?". *Electronics & Communication Engineering Journal* 13 (4), August 2001.
- [5] Y. Cheng et. al. "Jigsaw: solving the puzzle of enterprise 802.11 analysis." ACM SIGCOMM 2006.
- [6] COST Action 231. "Digital Mobile Radio Towards Future Generation Systems, Final Report". European Cooperation in the Field of Scientific and Technical Research, EUR 18957, 1999.
- [7] C. Fullmer and J.J. Garcia-Luna-Aceves. "Floor Acquisition Multiple Access (FAMA) for packet radio networks". ACM SIGCOMM 1995.
- [8] R. Gallager. "A perspective on multiaccess channels". *IEEE Transactions on Information Theory*, vol. 31, no. 2, 1985.
- [9] M. Garetto, J. Shi, and E. Knightly. "Modeling Media Access in Embedded Two-Flow Topologies of Multi-hop Wireless Networks". ACM MobiCom 2005, Cologne, Germany, August 2005
- [10] S. Gollakota and D. Katabi. "ZigZag Decoding: Combating Hidden Terminals in Wireless Networks". ACM SIGCOMM 2008.
- [11] IEEE Standard 802 Part 15.4. "Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specification for Low-Rate Wireless Personal Area Networks (WPANs)". Rev. 2006.
- [12] IEEE Standard 802 Part 11. "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications". Rev. 2007.
- [13] Recommendation ITU-R P.1238-1. "Propagation data and prediction methods for the planning of indoor radiocommunication systems and radio local area networks in the frequency range 900 MHz to 100 GHz." International Telecommunication Union, 1999.
- [14] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. "Understanding the Real-World Performance of Carrier Sense". ACM SIGCOMM E-WIND Workshop, 2005.
- [15] P. Karn. "MACA - A New Channel Access Method for Packet Radio." ARRL/CRRL Amateur Radio 9th Computer Networking Conference, September 1990.
- [16] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. "The Click modular router". *ACM Transactions on Computer Systems*, 18 (4), November 2000.
- [17] R. Metcalfe and D. Boggs. "Ethernet: Distributed Packet Switching for Local Computer Networks". *Communications of the ACM* 19 (5): 395-405, July 1976.
- [18] R. Patra et. al. "WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks". 4th USENIX Symposium on Networked Systems Design & Implementation, 2007.
- [19] A. Rao and I. Stoica. "An Overlay MAC Layer for 802.11 networks". Proceedings of Mobisys 2005, Seattle, April 2005.
- [20] F. Tobagi and L. Kleinrock. "Packet Switching in Radio Channels: Part II – The Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution". *IEEE Transactions on Communications*, vol. 23, no. 12, 1975.
- [21] A. Vasan, R. Ramjee, and T. Woo. "ECHOS - enhanced capacity 802.11 hotspots". IEEE INFOCOM 2005.
- [22] R. Vaughan and J. Andersen. *Channels, Propagation, and Antennas for Mobile Communication*. Institute of Electrical Engineers, 2003.
- [23] S. Verdú. *Multuser Detection*. Cambridge University Press, 1998.
- [24] M. Vutukuru, K. Jamieson, and H. Balakrishnan. "Harnessing Exposed Terminals in Wireless Networks". 5th USENIX Symposium on Networked Systems Design and Implementation, April 2008.
- [25] M. Vutukuru, H. Balakrishnan, and K. Jamieson. "Cross-Layer Wireless Bit Rate Adaptation". ACM SIGCOMM 2009.
- [26] J. Zhu et. al. "Leveraging Spatial Reuse in 802.11 Mesh Networks with Enhanced Physical Carrier Sensing." Proceedings of IEEE ICC, 2004.
- [27] J. Zhu, X. Guo, S. Roy, and K. Papagiannaki. "CSMA Self-Adaptation based on Interference Differentiation". IEEE GlobeCom 2007.