Optimal Design of High Density 802.11 WLANs

Vivek P. Mhatre*
Thomson Research Lab, Paris, France
mhatre@gmail.com

Konstantina Papagiannaki
Intel Research Cambridge, UK
dina.papagiannaki@intel.com

Abstract: The provisioning of high throughput performance infrastructure wireless networks necessitates the deployment of a high density of Access Points. While the latter improves wireless link quality to the clients, it can also introduce additional interference unless the network is carefully planned and tuned. It has been shown that the reaction of CSMA/CA protocol to interference is unnecessarily conservative in high density environments. In this work, we study the problem of infrastructure wireless network design, and the interaction between high density and MAC parameter tuning. Through analysis and numerical results, we provide recommendations on (i) optimum dimensioning of high density networks, and (ii) optimum tuning of their MAC parameters. We demonstrate that 802.11a networks are inherently noise-dominated, while 802.11g networks are interference-dominated, thus requiring different network design approaches. In sharp contrast to previous work, we establish that MAC parameter tuning has limited benefit in properly planned 802.11a networks. On the other hand, analytical results on the optimal tuning of MAC parameters in interference-dominated 802.11g deployments show substantial throughput improvements. Using the insight gained through our analysis, we propose an algorithm for the optimal tuning of MAC parameters in unstructured high density environments. Opnet simulations show that the proposed algorithm results in up to 260% improvement in network throughput.

1. INTRODUCTION

The wide acceptance of the IEEE 802.11 protocol for wireless access to the Internet has already led to high density unstructured networks in urban areas, as well as high density planned networks in enterprise environments for extended coverage and higher throughput support. While high density is already present in certain areas, it is still unclear how these networks need to be planned and tuned to optimally address the interplay between increased interference due to the closer proximity of APs, and the gains of improved link quality to the clients [1].

In 802.11 MAC, each terminal (client as well as AP) senses the wireless channel before attempting a transmission. In this mechanism, termed "physical carrier sensing", if the received power during channel sensing is greater than a certain threshold (referred to as Clear Channel Assessment Threshold, or CCA threshold), then the terminal infers that the wireless channel is currently busy, and therefore defers its transmission. By using a small CCA threshold, the amount of interference can be reduced by suppressing concurrent transmissions in the network, and thus the data rates can be improved. However suppression of concurrent transmissions also results in low network throughput due to reduced spatial reuse. Interestingly enough, the aforementioned problem only attracted the interest of the research community in the recent past [2], [3], [4], [5].

In this work, we study the problem of optimal design and tuning of high density wireless networks to cover a particular geographical area at the minimum cost (number of APs) while offering the best throughput performance to users. Interference is expected to play a key role in such networks. Unlike cellular networks, where interference mitigation is done through per-user power control, WLANs rely on the MAC protocol for a similar task. Consequently, we introduce an analytical framework that borrows design elements from cellular network design, but significantly expands to accurately incorporate the MAC layer behavior. In contrast to previous works [2], [3], [4], [5], our framework takes into account factors such as *AP density, available number of orthogonal channels, different modulation and coding schemes, and noise power*. We demonstrate that all these factors have a significant impact on the problem.

Using the aforementioned framework, we identify regimes under which a WLAN is dominated by noise or interference. Identifying the operating-regime is of prime importance, because *in a noise-dominated network, the carrier sensing parameter has almost no role to play*. On the other hand, in an interference-dominated network, the carrier sensing parameter has a pivotal role to play in determining the network throughput. In this context, the key insight that we provide is that *typical 802.11g networks are interference-dominated due to a shortage of orthogonal channels (3 orthogonal channels), while typical 802.11a networks are noise-dominated due to the relatively high number of orthogonal channels (12 orthogonal channels)*. We also show that while in high density WLANs, high data rates could be possible between a

---

[*] This work was done when the first author was at Intel Research Cambridge, UK.

ACM CoNEXT 2006 Lisboa, Portugal.

client and its AP due to close proximity, it turns out that *using the highest data rates may not be optimum from the overall network capacity perspective*.

Our analytical results are instrumental in guiding us in the design of a CCA adaptation algorithm that can be used by a wireless sender to mitigate the interference it experiences in its environment. Our algorithm can work in arbitrary network topologies, in 2D as well as 3D environments, irrespective of the physical layer employed, irrespective of the efficiency of the frequency selection mechanism, and without making any assumptions about the operating regime of the network. Driven by actual channel measurements, our algorithm is capable of balancing throughput gains and the amount of interference it accepts from the network. Using Opnet simulations it is further shown to lead to up to 260% throughput improvement compared to today's default MAC settings, while outperforming a recently proposed counterpart algorithm, ECHOS [18], by up to 30%.

The rest of the paper is organized as follows. In Section 2 we describe the problems faced in high density wireless networks, and list the network design and MAC layer parameters determining user throughput. The interplay between network design and MAC layer tuning is studied analytically in Section 3. Relaxing the assumption on flexibility in the deployment of the dense WLAN, we derive appropriate formulations for CCA adaptation in Section 4. Our analytical findings are used to formulate a novel CCA adaptation algorithm in Section 5 that can work in arbitrary networks. Its performance is evaluated using Opnet simulations in Section 6. We discuss some of the related work in this area in Section 7, and conclude in Section 8.

2. WLAN DESIGN AND CSMA/CA

The primary objective of network operators deploying wireless networks is to provide geographical coverage of a certain region (such as an enterprise environment, or a university campus) at the lowest network cost (AP density), while delivering the best possible performance to the end user. A single 802.11a/g AP can typically provide coverage up to a distance of about 100 m at a data rate of 6 Mbps [13], [11]. However, in the presence of multiple users and time-varying wireless conditions, the actual throughput on the cell boundary may be much lower than 6 Mbps. Hence for improved coverage and higher throughput, it is necessary to deploy multiple APs over the region. In such high density wireless networks (cell radius of 20 to 50 m), interference is expected to play a key role [2, 3, 4, 5, 8].

2.1 The impact of interference

When two nodes communicate wirelessly, the maximum rate at which this communication can be sustained is determined by the Signal to Interference and Noise Ratio (SINR) at the receiver which is given by:

$$\text{SINR} = \frac{S_R}{S_N + S_I}, \tag{1}$$

where S_R is the power at which the transmitted signal is received at the receiver, S_N is the power of noise in the receiver circuitry (explained in more details in Section 2.3), and S_I is the total interference power. The higher the SINR, the higher the data rate that can be sustained.

Given that clients associate with their closest AP, increased AP density implies a reduction in the distance between a client and its closest AP. This in turn implies that the signal strength received by a client from its associated AP (S_R in equation (1)) can be increased, and this results in an improvement in the SINR as well as the data rate. For example, although the communication range of 802.11a is 100m, the rate it can support at that distance is just 6 Mbps [11]. If a user is within 20m of the AP, then a rate of up to 36 Mbps can be supported. Therefore high density networks aim at shrinking the cell size, so that higher data rates can be supported on the cell boundary.

However, note that the available number of orthogonal channels in 802.11 is limited (3 in 802.11g, and 12 in 802.11a [11]). As the AP density increases, the distance between co-channel APs decreases, and therefore the interference in the network (S_I in equation (1)) increases. When the interference around a transmitter is significant, the 802.11 carrier sensing prevents the transmitter from accessing the medium so as not to corrupt the ongoing communication. The Clear Channel Assessment (CCA) Threshold corresponds to that amount of S_I that determines whether a transmitter is allowed to access the medium or not. If the CCA threshold of 802.11 MAC is set to a small value, then concurrent transmissions on the same channel are suppressed, and in spite of having a high density of APs, not all the APs are able to transmit simultaneously. Thus, although the data rates could be very high, the fraction of time for which an AP transmits can be very small, having an adverse effect on the throughput. On the other hand, if the CCA threshold of 802.11 MAC is set to a large value, then multiple co-channel APs may transmit concurrently. The increased interference can, however, nullify the benefit of improved signal strength due to significant degradation of the receiver SINR.

The above description clearly outlines the complex interdependence, between the achievable throughput, the AP density, the available number of orthogonal channels, and MAC layer parameters, such as the CCA threshold in high density 802.11 networks. Consequently, it is imperative that a single unified framework be used to study this problem. Previous work has typically looked at MAC layer parameter tuning in isolation [2], [3], [4] and [5].

2.2 Network Design Choices

Based on the above, there are only a few tools at the disposal of the network operator to avoid the adverse effect of interference in high density deployments: (i) identification of the appropriate physical layer technology (802.11g or 802.11a) for the density of the planned network, since different technologies offer different coverage and spatial reuse factors, (ii) choice of AP density so that the desired throughput performance can be achieved, (iii) careful selection of the channels assigned to APs, so as to minimize contention among nearby APs, and (iv) careful tuning of MAC parameters to mitigate interference in the network. We call the assembly of the four aforementioned tasks as "network design", and provide recommendations on optimum network design of high density 802.11 networks in Section 3.

Notice that the first three tasks listed above are similar to the network design and planning tasks carried out in cellular networks. One important difference, however, is that cellular networks mitigate interference by using *per-user* power control [6]. However in 802.11, the time granularity over which a user is served is much smaller than in a cellular voice network, and hence per-user power control is difficult.

Thus, power control in 802.11 is an altogether independent problem, and not the focus of this work. In this work, we only focus on the tuning of 802.11 MAC for interference mitigation, and propose an algorithm that can enable today's 802.11 networks to make optimal use of their high density. Note that even though both cellular and 802.11 networks may rely on micro-cell deployments for high throughput, as we show later, the theoretical framework borrowed from cellular network design needs to be adjusted to accurately incorporate the impact of the MAC layer.

2.3 High Density MAC parameters

Even though appropriate choices in the initial design of a high density wireless network may be sufficient in achieving high performance (as will be shown in the next section), this may not always be the case either due to incremental deployment, or less flexibility in terms of the AP placement and equipment capabilities (directional antennas could for instance mitigate interference). In settings where interference cannot be avoided by design, the behavior of a wireless transmitter and receiver depends on a number of parameters which control their reaction to noise and interfering signals, as well as their aggressiveness to seize the channel for transmission.

Wireless Transmitter: The behavior of a wireless transmitter is primarily determined by the **CCA Threshold**. In 802.11 MAC, a transceiver decides if the wireless channel is currently busy based on the relative value of interference as compared to the CCA Threshold (Clear Channel Assessment Threshold). Before attempting to transmit a frame, the transceiver measures the strength of the received interference on the wireless channel. If the measured interference is higher than the CCA Threshold, the channel is assumed to be busy, and the transmission is deferred.

Wireless Receiver: The behavior of the wireless receiver further depends on a series of parameters capturing the quality of the receiver's circuitry, the power of the signal needed for successful decoding, as well as the selectivity of the receiver in terms of the transmissions that it attempts to decode.

Receiver Noise Power refers to the amount of noise generated in the receiver circuitry. Let N_0 denote this quantity in dB, defined as $N_0 = 10 \log S_N$ where S_N is measured in milli Watts. For typical state-of-the-art 802.11 APs such as [11], N_0 is in the range of -96 to -91 dB.

The receiver noise power directly impacts what is termed as **Receiver Sensitivity**, which refers to the minimum power of the desired signal at the receiver, that is required for the successful decoding of the signal *in the presence of noise alone.* The Receiver Sensitivity depends on the modulation and coding scheme. If the minimum required SINR for successful decoding of a signal for a given modulation and coding scheme is β (in dB), then the Receiver Sensitivity (in dB) is given by $N_0 + \beta$.

Whenever a wireless transmission is received by a particular client at a power greater than the Receiver Sensitivity, the client circuitry attempts to decode the signal, even if that signal may not be addressed to the client in question. Any subsequent transmission targeted to the client will not be decodable for the duration of the initial transmission. This phenomenon is termed as "strongest last-collision", and may occur in high density environments [4, 17]. The **Receiver Threshold** defines that value of received power below which the receiver circuitry will not attempt to decode the received signal. It should be further set to be greater than the Receiver Sensitivity and approximately equal to the power the client should expect from its associated AP. By default, the Receiver threshold and the CCA threshold are set to be approximately equal to, and slightly higher than, the noise power [21].

Notice that the Receiver Sensitivity and the Receiver Threshold are altogether different quantities. While the former is determined and *fixed* by the hardware capabilities, the latter is *configurable.* In Section 3.5, we will show that the Receiver Threshold can be easily determined based on the CCA Threshold.

3. OPTIMUM NETWORK DESIGN

In Section 2, we identified the key factors impacting the performance of high density wireless networks. In this section, we study the interplay between optimal network design and high density MAC parameter tuning, deriving recommendations on the optimal deployment and configuration of such environments. Optimal design within such a context refers to the provisioning of a wireless network that covers the desired geographical area at the smallest cost, while offering the best possible throughput to the users. Consequently, the metric of interest is what we call Throughput-Coverage.

Definition 1. *A network provides a Throughput-Coverage of $\mathcal{C}$ bits per second per square meter when the total throughput delivered to all the clients in a unit area is at least $\mathcal{C}$ bits per second (even if the clients are located on the edge of the cell).*

The above definition gives a *lower bound* on throughput performance. Users closer to the AP can get higher throughput through Auto Rate Fallback, but we focus on the worst case performance when all the users are assumed to be on the cell boundary. In this section, our objective is to analytically derive the optimum AP density, and the optimum CCA Threshold values that guarantee a desired Throughput-Coverage performance. We borrow elements from cellular network design and provide a framework for the analysis of the performance of dense WLANs that incorporates the effects of CSMA/CA.

A common practice in cellular network design is the modeling of cells using hexagons. The use of a hexagon pattern for a cell can be shown to be optimal in the sense that one can cover a certain geographic region with the smallest number of cells, while a hexagon closely approximates a circular radiation pattern [9]. We define a WLAN cell to be the geographic region around an AP in which the received signal strength from the given AP is stronger than the received signal strength from any other AP. Assuming that the APs have a regular hexagonal lattice placement over a two-dimensional region, we denote the number of APs per unit area by ρ, and the number of orthogonal channels available for assignment by K. Note that hexagonal lattice deployment is used for clarity of insights. All results can be easily extended to the case of random uniform AP deployment by replacing the summation used for computing the total interference, by an appropriate integral in the following arguments.

802.11g has three orthogonal channels, while in 802.11a, up to 12 orthogonal channels are available in US and Europe

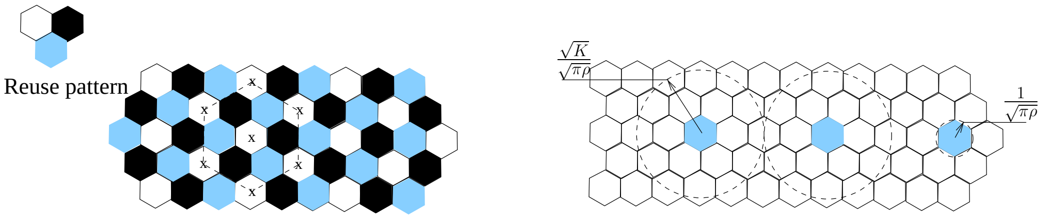

(a) Reuse pattern (b) Co-channel APs

Figure 1: Reuse pattern and channel allocation.

Index	1	2	3	4	5	6	7	8
Rate (Mbps)	54	48	36	24	18	12	9	6
SINR (dB)	24.6	24	18.8	17	10.8	9	7.8	6

Table 1: SINR requirements for different data rates for 802.11a/g

[11]. For analytical tractability, we assume that APs operating on a common channel also have a regular hexagonal lattice placement. In such regular environments, typically found in enterprises, appropriate reuse patterns can be designed as long as there are at least three orthogonal channels [9]. For example, Fig. 1(a) shows a channel allocation scheme with three orthogonal channels. We can see that the co-channel APs also have a regular lattice pattern.

Given an AP density of ρ, and K orthogonal channels, the AP density for each channel type is $\frac{\rho}{K}$. The average cell radius, R, and the average separation between co-channel APs, D, are given as follows:

$$R = \frac{1}{\sqrt{\pi\rho}}, \quad D = 2\sqrt{\frac{K}{\pi\rho}} \tag{2}$$

These distances are depicted in Fig. 1(a). The three shaded hexagons represent co-channel cells. The higher the AP density, the smaller the cell-radius, and hence the stronger the received signal strength at the boundary of the cell. The higher the number of channels, K, the larger the separation between the co-channel APs, and hence the lower the interference.

The client throughput depends on (i) the client-AP separation, (ii) the amount of interference, and (iii) the transmission rates supported by the AP. The largest client-AP separation is determined by the AP density, while the CCA threshold determines the highest co-channel interference received by the client. These quantities together determine the worst case SINR in a cell (via equation (1)). We only consider 802.11a/g transmission rates, since these physical layers are guaranteed to lead to higher throughput regimes than 802.11b. In 802.11a/g, as Table 1 shows, there are 8 discrete data rates that can be supported for different values of client SINR [13].

Our objective in this section is to derive the smallest AP density that can guarantee a desired throughput-coverage performance. For this, we take the following approach. *For each data rate i in Table 1 (i from 1 to 8), we determine the optimum AP density (or equivalently, the optimum cell radius R_i), and the optimum CCA threshold, CCA$_i$, that can provide the desired throughput-coverage of $\mathcal{C}$. We then choose that data rate, k, for dimensioning and configuring our network, which requires the smallest AP density.* Together with Definition 1, this ensures that data rate k can be supported on the cell boundary at the smallest network cost. Also note that schemes corresponding to data rates lower than that of k can also be supported for the *same* choice of AP density and CCA threshold, due to their lower SINR requirements.

In Sections 3.1 to 3.3, we formulate the constraints of our optimization problem, and then solve the problem in Section 3.4. The only inputs to our optimization problem are hardware and system parameters, capturing: (i) the number of channels, K, (ii) the AP transmit power, P, (iii) the receiver noise power N_0, and (iv) a propagation loss model $P \cdot ax^{-\alpha}$, where a is a fixed constant, and α is the path loss exponent.

3.1 SINR Constraint

Consider the ith modulation and coding scheme. Let C_i be its data rate, and let β_i be the corresponding minimum required SINR for this rate. Let R_i be the cell radius, and CCA$_i$ be the CCA Threshold required for guaranteeing the desired throughput-coverage of $\mathcal{C}$ *with this scheme*. Assume that the carrier sensing distance is d_i, i.e., two co-channel APs can transmit concurrently only if they are separated by a distance of at least d_i. Note that this *does not* mean that CCA$_i = P \cdot ad_i^{-\alpha}$. This is because there are multiple APs that are located at a distance greater than d_i from the given AP, and it is the *cummulative* interference from all these APs that determines CCA$_i$. CCA$_i$ also represents the maximum allowable interference power for scheme i. The worst case client SINR is at the boundary of the cell, and is given by

$$\text{SINR} = \frac{P \cdot aR_i^{-\alpha}}{N_0 + \text{CCA}_i}, \Rightarrow \frac{1}{\text{SINR}} = \frac{N_0}{P \cdot aR_i^{-\alpha}} + \frac{\text{CCA}_i}{PaR_i^{-\alpha}} \tag{3}$$

It is well-known that due to the strong propagation fall-off, the total interference in a cell is dominated by the closest tier of interfering APs [20, 3, 5]. Hence, we only consider the first tier of co-channel APs in computing the total interference (see Fig. 1(a)). By using simple geometric arguments, we can obtain this interference power as a function of P, a, R_i, and d_i. Furthermore, if we define an auxiliary variable X_i as follows,

$$X_i = \frac{d_i}{R_i}, \tag{4}$$

then we can show that (see [5])

$$\frac{\mathrm{CCA}_i}{PaR_i^{-\alpha}} \triangleq I(X_i) = \frac{1}{(X_i+1)^\alpha} + \frac{1}{(X_i-1)^\alpha} + \frac{2}{\left(\left(\frac{X_i}{2}+1\right)^2 + \left(\frac{\sqrt{3}X_i}{2}\right)^2\right)^{\frac{\alpha}{2}}} + \frac{2}{\left(\left(\frac{X_i}{2}-1\right)^2 + \left(\frac{\sqrt{3}X_i}{2}\right)^2\right)^{\frac{\alpha}{2}}} \quad (5)$$

In (4) and (5), X_i is the ratio of the cell radius and the carrier sensing range, while $I(.)$ is a measure of the cummulative interference and the received signal strength. Since we want to ensure that data rate C_i can be supported on the cell boundary, using (3), we get the following SINR constraint.

$$I(X_i) + \frac{N_0 R_i^\alpha}{Pa} \leq \frac{1}{\beta_i} \quad (6)$$

3.2 Throughput-Coverage Constraint

Since we only allow concurrent transmissions between co-channel APs separated by at least d_i, circles of radius $\frac{d_i}{2}$ around concurrently transmitting co-channel APs are disjoint. Hence maximum spatial reuse is attained when such disjoint circles are minimally packed over the entire region. If the entire region to be covered has area A, then the maximum number of APs transmitting concurrently on a given channel is $\frac{A}{\pi\left(\frac{d_i}{2}\right)^2} = \frac{4A}{\pi d_i^2}$. If we assume that the cell radius R_i is chosen such that data rate of C_i is supportable on the cell boundary, then the collective data rate of all the APs on a given channel is at least $\frac{4C_iA}{\pi d_i^2}$ (even when all the users are on the cell boundary). Since there are K channels, the collective rate at which data is transferred over the entire network is $\frac{4KC_iA}{\pi d_i^2}$. Hence throughput per unit area is $\frac{4KC_i}{\pi d_i^2}$. Note that although the data rate is C_i, the actual useful throughput at the application layer is lower due to MAC-layer and protocol-specific overheads. For simplicity, we do not include these overheads. Since we require the throughput per unit area to be at least $\mathcal{C}$, using (4), we have the following constraint for throughput-coverage.

$$\frac{KC_i}{\pi\left(\frac{X_iR_i}{2}\right)^2} \geq \mathcal{C} \quad \Rightarrow \quad X_iR_i - 2\sqrt{\frac{KC_i}{\pi\mathcal{C}}} \leq 0 \quad (7)$$

3.3 Constraint on the number of Channels

Note that the number of available channels determines the physical separation of co-channel APs through (2). The carrier sensing distance d_i cannot be smaller than D (given by (2)), since there are no co-channel APs at a distance smaller than D from an AP. Hence,

$$d_i \geq \frac{2\sqrt{K}}{\sqrt{\pi\rho}} \quad (8)$$

Combining (2), (4), and (8), we get the following constraint.

$$X_i \geq 2\sqrt{K} \quad (9)$$

Note that no such constraint was considered in [2], [3], [4] and [5]. The above constraint shows that the number of channels has a significant role to play in this problem. For example, with 12 channels, $X_i \geq 6.9$. A choice of X_i for *every* modulation and coding scheme must be at least as high as 6.9. "Optimum" values of X_i computed for several scenarios in [3] and [5] are much lower than this value. This means that those values are not even feasible when we have a typical 802.11a setting with at least 12 channels. In essence, a large number of channels results in a larger separation between the co-channel APs, and can thus reduce interference. Hence this constraint has to be taken into account in determining the optimum amount of allowable interference, i.e., the optimum carrier sensing threshold.

3.4 Optimization Problem and its Solution

Combining (7), (9) and (6), the problem of optimum network dimensioning (finding the minimum AP density), and optimum tuning of the carrier sensing threshold CCA_i, so that a throughput coverage of $\mathcal{C}$ is guaranteed, can be formulated as follows. Since minimizing AP density is equivalent to maximizing cell radius R_i, i.e., minimizing $-R_i$, we get:

$$\text{Minimize:} \quad f(X_i, R_i) = -R_i \quad (10)$$

$$\text{Subject to:} \quad g_1(X_i, R_i) = X_iR_i - 2\sqrt{\frac{KC_i}{\pi\mathcal{C}}} \leq 0 \quad (11)$$

$$g_2(X_i, R_i) = 2\sqrt{K} - X_i \leq 0 \quad (12)$$

$$g_3(X_i, R_i) = I(X_i) + \frac{N_0R_i^\alpha}{Pa} - \frac{1}{\beta_i} \leq 0 \quad (13)$$

where $I(X_i)$ is given by (5). Using the Karush-Kuhn-Tucker Theorem [10], we can show that the solution to the above optimization problem is as follows (we do not include the proof due to space constraints).

$$R_i^* = \min\left\{\sqrt{\frac{C_i}{\pi\mathcal{C}}}, \hat{R}_i\right\} \text{ where } \hat{R}_i \text{ is given by} \quad (14)$$

$$\frac{1}{\beta_i} = \frac{N_0\hat{R}_i^\alpha}{Pa} + I\left(\frac{1}{\hat{R}_i}\sqrt{\frac{4KC_i}{\pi\mathcal{C}}}\right) \quad (15)$$

$$X_i^* = \frac{1}{R_i^*}\sqrt{\frac{4KC_i}{\pi\mathcal{C}}} \quad (16)$$

We solve the above optimization problem for each of the 8 modulation-coding schemes, and choose that modulation-coding scheme which results in the largest cell radius while providing the desired throughput-coverage. We then choose the corresponding AP density for network dimensioning, and set the carrier sensing threshold of the entire network, $\mathrm{CCA}_{\mathrm{thr}}$, with respect to this modulation scheme using (5).

$$j = \mathrm{argmax}_i\,(R_i), \quad \rho = \frac{1}{\pi R_j^2}, \quad \mathrm{CCA}_{\mathrm{thr}} = P \cdot aR_j^{-\alpha} I(X_j).$$

Notice that we use a network-wide common CCA threshold. This is because if the network is interference dominated, nearby co-channel APs can hear each other. If the two nearby APs choose different CCA Thresholds, then the AP with a higher CCA Threshold (the AP that can accept more interference) will acquire the channel more aggressively than the AP with a lower CCA Threshold. This may lead to starvation of the AP that uses a lower CCA threshold. The approach of using a common network-wide CCA Threshold is adopted in [4, 5] as well. However, the ECHOS algorithm proposed in [18], does not necessarily result in a common network-wide CCA threshold.

3.5 Numerical Results

In this section, we obtain numerical results for the optimum network dimensioning and carrier sensing parameters as a function of the desired throughput-coverage. In Table 2, we have listed the values of the system parameters that we use in obtaining numerical results in this section. Propagation loss parameters are as per ITU recommendations for 802.11a/g in indoor office environment [14]. Transmit power of the APs, noise power and number of channels in Table 2 are typical values for Cisco [11]. Note that since 802.11a uses lower transmit power, and has a stronger propagation fall-off, 802.11a cells usually have smaller coverage than 802.11g cells. In Table 1, we have listed the minimum required SINR for supporting different modulation-coding schemes in 802.11a/g [13]. In all the plots, the x-axis is the target throughput-coverage multiplied by $\pi(100)^2$. This normalization constant is included so that the x-axis can be interpreted as the average throughput delivered to a circular region of radius 100 m. We obtain numerical results for the cases when the noise power is -91 dB or -96 dB, and the underlying network employees 802.11a or 802.11g. The reason for *discontinuities* in all the plots is that there are 8 *discrete* modulation and coding schemes in 802.11a/g.

	802.11a	802.11g
Transmit Power, P	+17 dBm	+20 dBm
Number of Channels, K	12	3
Path loss exponent, α	3.1	3.0
Path loss constant, a	-46.5 dB	-40 dB

Table 2: System Parameters for Numerical Results

The optimum value of cell radius and the corresponding AP density as a function of target throughput-coverage are plotted in Fig. 2(a) and Fig. 2(b). The plots show that to achieve higher throughput-coverage, we must reduce the cell radius, i.e., the network architecture must comprise of a high density of microcells. Also, the higher the noise power, the smaller the required cell radius, i.e., the higher the required AP density.

In Fig. 3, we plot the data rate which is used for system dimensioning and CCA Threshold tuning. Due to the SINR constraint in Section 3.1, this is also the data rate that can be sustained on the edge of every cell. We can see from Fig. 3 that as the target throughput-coverage increases, we must employ higher modulation and coding schemes to satisfy the throughput-coverage requirements. We note from Fig. 3 that for 802.11a, despite the fact that a data rate of 54 Mbps can be sustained at a distance of 20 m [11], to provide a throughput-coverage of 500, a cell radius of 20 m (Fig. 2(b)), and corresponding data rate of 36 Mbps are chosen (Fig. 3). This is because in order to support a data rate of 54 Mbps, a substantially large number of co-channel APs need to be suppressed, and this results in poor spatial reuse. If instead, a data rate of 36 Mbps is used, then fewer co-channel APs need to be suppressed. Thus our analytical and numerical results prove that in high density networks the benefits of improved spatial reuse may more than offset what we lose by using a lower data rate. Hence, *in high density 802.11 networks, for the overall network good, using the highest data rates is not necessarily optimum.*

The optimum values of CCA Threshold as a function of desired throughput-coverage are plotted in Fig. 4(a,b,c,d). We also plot the Receiver Sensitivity for the modulation and coding scheme that can be supported on the cell boundary. There are three distinct operational regions in terms of the impact of interference on MAC layer behavior. If the power of the interfering signal is below the noise power, CCA adaptation has no impact in that environment. If the power of the interfering signal is above the noise power but under the receiver sensitivity, the signal can be received but cannot be decoded. Consequently, media access is governed by physical carrier sensing, and depending on the value of the CCA threshold the transmitter will suppress or continue with a transmission. Lastly, if the interfering signal is stronger than the Receiver Sensitivity, it can be successfully decoded. In this case, the terminal must respect the Network Allocation Vector field (NAV field) in the interfering signal (if RTS/CTS is enabled), and must defer its own transmission for the duration of the interfering transmission. This is the virtual carrier sensing mechanism of 802.11.

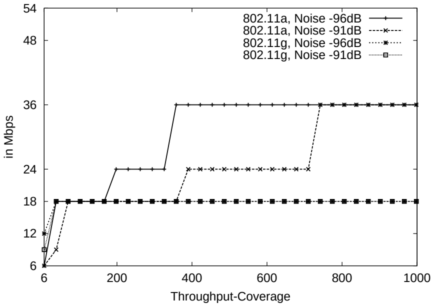

Figure 3: Optimum Rate on cell boundary as a function of $\mathcal{C} \cdot \pi(100)^2$ Mbps/m^2.

According to the above, we note that the optimum CCA threshold for 802.11a (Fig. 4(a),(b)) is comparable to the noise level even in highly dense scenarios. In fact, for $N_0 = -96$ dB, the CCA threshold exceeds the noise power by 7 dB at most, and for $N_0 = -91$ dB, this difference is no more than 2 dB. Note that typical shadow-fading variations are on the order of 5-10 dB [14], and hence CCA threshold has a limited role to play for the above settings. We also note from Fig. 4(a),(b) that the CCA Threshold is never above the Receiver Sensitivity. This means that an AP cannot successfully decode the interfering signals. Thus, we conclude that *in regular high density 802.11a networks, the extent of co-channel interference is almost negligible, and the network is noise-dominated.* Also note that increased throughput-coverage *cannot* be provided by tweaking the CCA threshold, an observation that verifies the experimental results in [4], where the optimum CCA threshold for 10-20% raw frame error rate was found to be very low (-90 dB). The reason for 802.11a being noise-dominated, and not interference-dominated, is the availability of a large number of orthogonal channels which results in large inter-AP separation for co-channel APs, and thereby substantially lower co-channel interference. *Thus, in 802.11a networks, the key design parameter is AP density and appropriate frequency selection for optimal spatial reuse, and not the CCA threshold.*

On the other hand, in 802.11g networks, due to shortage of orthogonal channels, interference plays an important role in determining the optimum system parameters. From Fig. 4(c),(d), we can see that the CCA threshold is substantially higher than the noise power for most settings. In fact, the optimum CCA Threshold is higher than the Receiver Sensitivity. Our results indicate that in such scenarios, 802.11g APs *should not defer their transmission even if they can successfully decode the received signal from an interfering AP*. This can be achieved as follows. The Receiver Threshold (described in Section 2) can be configured to be

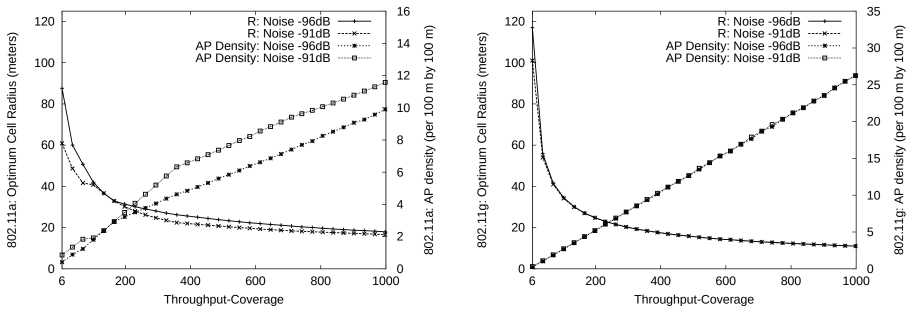

(a) 802.11a (b) 802.11g

Figure 2: Optimum Cell Radius, R, as a function of $\mathcal{C} \cdot \pi(100)^2$ Mbps/m^2.

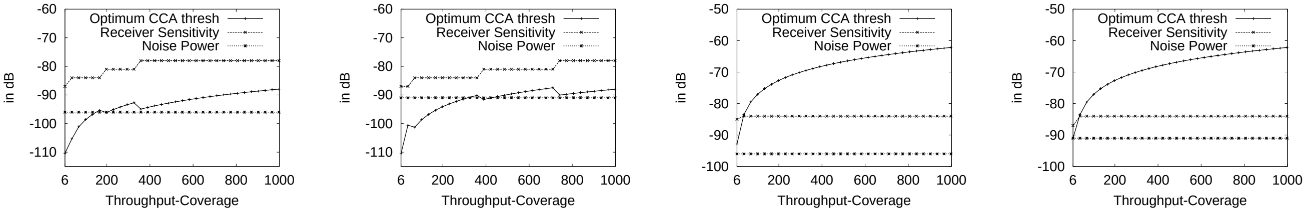

(a) 802.11a, $N_0 = -96$ dB (b) 802.11a, $N_0 = -91$ dB (c) 802.11g, $N_0 = -96$ dB (d) 802.11g, $N_0 = -91$ dB

Figure 4: Optimum CCA as a function of $\mathcal{C} \cdot \pi(100)^2$ Mbps/m^2. Receiver Sensitivity is equal to $N_0 + \beta$, where N_0 is the noise power, and β is the minimum required SINR for a given modulation and coding scheme.

equal to the optimum CCA Threshold, so that the receiver does not attempt to decode interfering frames. This will ensure that the virtual carrier sensing mechanism of 802.11 is over-ridden by the physical carrier sensing mechanism to avoid the problem of strongest-last collision. Such a recommendation is also made in [16] and [4]. Also note from Fig. 2(b) that the optimum cell radius is almost independent of the noise power for 802.11g; the plots for thedifferent noise powers are overlapping. Thus, in sharp contrast to 802.11a networks, we note that *regular 802.11g networks are interference dominated, and tuning of MAC parameters plays an important role in ensuring high network throughput.*

In summary, in this section we showed that due to differences in the available number of orthogonal channels, and due to different path loss models, the design of 802.11a and 802.11g networks requires substantially different approaches. Optimal AP density and proper frequency selection can mitigate interference to a large extent in 802.11a networks. On the other hand, the design of 802.11g networks needs to address AP density as well as MAC layer tuning under a unified framework to optimize network performance.

4. OPTIMUM MAC LAYER TUNING

In Section 3, we established design recommendations for regular dense wireless networks when one controls both, the density of the APs, as well as the tuning of the MAC parameters. In this section, we relax the first assumption and study the tuning of MAC parameters in regular networks with a predefined AP density. The insights gained through our analysis are then employed in Section 5, where our solutions are generalized to arbitrary network topologies. Another way to perceive the topic in question is that we attempt to identify the regimes where CCA adaptation can result in throughput improvement in existing networks. As in Section 3, we assume a regular lattice deployment of APs in our analysis. We also assume that channel allocation has been performed in such a way that even the co-channel APs form a regular pattern. As noted in Section 3, as long as there are at least 3 orthogonal channels, such channel allocation is always possible in regular lattice structures.

Our objective in this section is to select the optimum CCA threshold, given a fixed cell radius R, so as to maximize the throughput-coverage of the network. As in Section 3, *we first determine the optimum CCA Threshold for each of the 8 modulation schemes, and then choose that scheme for setting the network-wide CCA Threshold which results in the highest throughput-coverage.* Note that although we use a similar approach to Section 3, the constraints and the objective function of the optimization problem considered in this section are different. From Section 3.2, the throughput-coverage of the network under the ith modulation-coding scheme is $\frac{KC_i}{\pi\left(\frac{X_iR}{2}\right)^2} = \frac{C_i}{X_i^2}\frac{K}{\pi\left(\frac{R}{2}\right)^2}$. Since R is fixed, we have the following optimization problem for the ith modulation-coding scheme.

$$\text{Maximize: } f(X_i) = \frac{C_i}{X_i^2} \tag{17}$$

$$\text{Subject to: } X_i \geq 2\sqrt{K} \tag{18}$$

$$g_2(X_i) = I(X_i) + \frac{N_0 R^\alpha}{Pa} - \frac{1}{\beta_i} \leq 0 \tag{19}$$

In the above, note that (19) *may not* be feasible in general for all modulation-coding schemes. In other words, it may not be possible to support all the data rates on the boundary of a cell due to the non-zero contribution of noise.

The solution to the optimization problem in (17)-(19) is as follows. For a given network (given R), we check if each of the modulation and coding schemes can be supported on the cell boundary, i.e., we check if for each i, (19) holds for some X_i. If so, given the fact that $I(X_i)$ is strictly decreasing, (19) can be rewritten as

$$X_i \geq \eta_i \tag{20}$$

where η_i is a constant that depends only on R, β_i, N_0, a, P and α. Then, the solution to the optimization problem in (17)-(19) is

$$\hat{X}_i = \min\left\{2\sqrt{K}, \eta_i\right\} \tag{21}$$

We then choose that modulation and coding scheme for setting CCA Threshold which results in the highest value of throughput-coverage, i.e.,

$$j = \text{argmax}_i \left(\frac{C_i}{\hat{X}_i^2}\right) \tag{22}$$

where argmax is taken over those modulation-coding schemes i, that satisfy (18) and (19). The CCA Threshold is then set using (5) as follows.

$$\text{CCA}_{\text{thr}} = P \cdot aR^{-\alpha} I(X_j) \tag{23}$$

4.1 Numerical Results

In this sub-section, we obtain numerical results for setting the optimum CCA Threshold in regular networks with fixed AP density using the above problem formulation. In Fig. 5(a),(b) we plot the optimum CCA threshold as a function of R for an 802.11a network. The optimum CCA Threshold is below the noise power for a cell radius larger than 35 m. Even when the CCA Threshold exceeds the noise power, it is substantially smaller than the Receiver Sensitivity, and therefore interfering signals cannot be decoded. These results confirm that interference in planned 802.11a high density networks (where the AP deployment pattern and channel allocation avoids reuse of the same frequency in neighboring APs) has limited impact, and the system is noise dominated. Setting the CCA threshold optimally is critical only for very small cell sizes (20 to 30 m).

In Fig. 5(c),(d), we plot the optimum CCA threshold for an 802.11g network for different noise powers. When $N_0 = -96$ dB (respectively -91 dB), the interference is comparable to the noise power when the cell radius is smaller than 170 m (respectively 120 m). For a cell radius larger than these values, the system is noise dominated, and hence the CCA threshold can be set to the noise power. For small cell sizes (20 to 60 m), interference is higher than the noise power as well as the Receiver Sensitivity. Hence, as in Section 3.5, we set the Receiver Threshold to be equal to the optimum CCA Threshold so that the physical carrier sensing mechanism overrides the virtual carrier sensing mechanism.

If, instead of setting the CCA threshold optimally, an 802.11g network were to operate with the fixed default CCA threshold, then physical, as well as virtual carrier sensing will suppress concurrent transmissions in a large neighborhood, and will bring down the throughput substantially. In Fig. 6, we plot the percentage improvement in throughput-coverage when the CCA threshold is set optimally, compared to the use of the default CCA threshold that is equal to the noise power. We note that there are substantial gains in throughput-coverage. For 802.11g, Fig. 6 shows that even for a cell radius of 40 m, *the improvement in throughput-coverage is as high as 175 to 300%.* For 802.11a, CCA Adaptation is beneficial only for very small cell sizes (20 to 35 m). For example, for a cell radius of 20 m, the throughput can be improved by 50 to 300 % by choosing the CCA Threshold optimally.

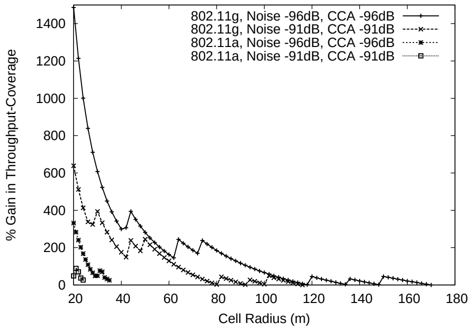

Figure 6: Percentage improvement in throughput-coverage of optimum CCA selection over the default CCA threshold.

Note that these results were obtained for a regular topology with optimal channel allocation. However, the results strongly suggest that CCA adaptation is beneficial. Later on, in Section 6, we show that even for random deployment, CCA adaptation results in substantial throughput improvement.

In Fig. 7(a), we plot the supportable rate on the cell boundary for 802.11a and 802.11g for two different receiver noise powers. Note that Fig. 7(a) shows rate as a function of distance in the presence of noise *as well as interference*, and hence should not be compared to the range that is typically advertised by the vendors [11]. The latter is computed for an interference-free scenario. We can clearly see that the supportable rate decreases as the cell size increases. Also note that the maximum cell size for 802.11a is lower than the maximum cell size for 802.11g. This is because 802.11a uses lower transmit power and has a stronger path loss (see Table 2). More importantly, we note that high transmission rates, such as 48 or 54 Mbps, are never optimal from the perspective of overall network throughput due to their strict requirements in terms of SINR (as in Fig. 3).

In Fig. 7(b), we plot the throughput-coverage (on log scale) as a function of the cell radius. We note that 802.11a offers better throughput-coverage for smaller cell sizes, however the throughput-coverage drops sharply to zero. On the other hand, 802.11g provides lower throughput-coverage for smaller cell sizes, while providing extended coverage.

5. ORCCA: OPTIMAL-RATE CCA ADAPTATION

The analysis presented in Section 4, even though applicable to environments where co-channel APs are deployed in a regular lattice, offers the following insights. First, it shows that the potential gains of using optimum CCA Threshold in high density networks are substantial (see Fig. 6). Second, it provides insight into designing an algorithm for optimally setting the CCA threshold in random topologies. In this section, we present ORCCA, an algorithm that sets the CCA Threshold of a given network *purely based on channel measurements.* ORCCA exploits the results from Section 4 and *does not make any assumptions about the placement of*

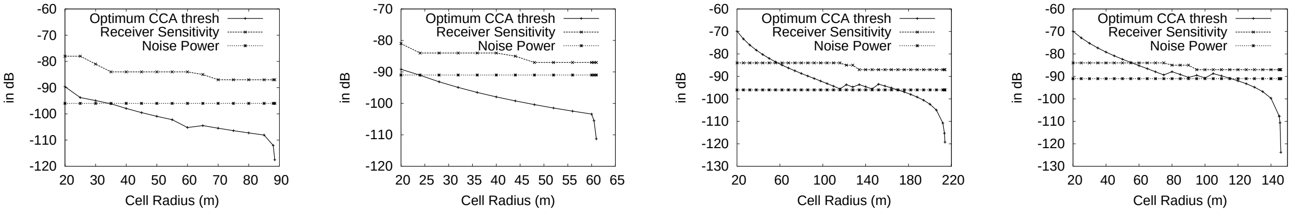

(a) 802.11a, $N_0 = -96$ dB (b) 802.11a, $N_0 = -91$ dB (c) 802.11g, $N_0 = -96$ dB (d) 802.11g, $N_0 = -91$ dB

Figure 5: Optimum CCA Threshold as a function of cell radius R. Receiver Sensitivity is equal to $N_0 + \beta$, where N_0 is the noise power, and β is the minimum required SINR for a given modulation and coding scheme.

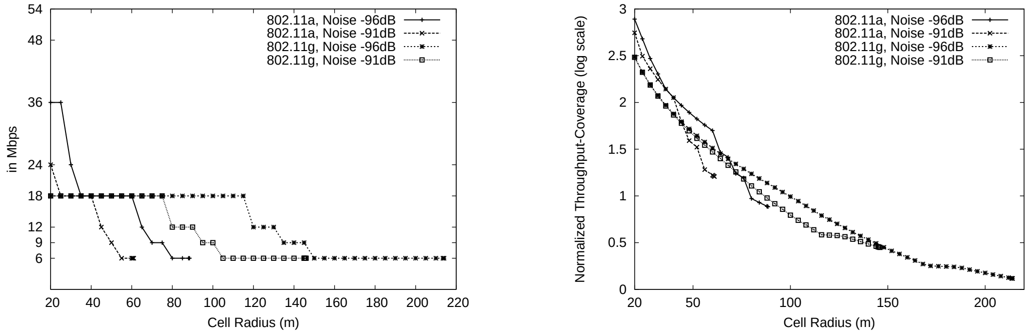

(a) Optimum Rate (b) Throughput-coverage on log-scale $\left(\log(\mathcal{C} \cdot \pi(100)^2)\right)$

Figure 7: Optimal rate and throughput-coverage for different cell sizes.

clients, APs, cell shapes, or channel assignment.

Consider an 802.11 a/g network whose CCA Threshold is to be optimally configured. We make the following assumptions in our algorithm:

1. Each AP keeps track of the signal strengths of all the co-channel APs that are operating in its neighborhood. This is required as part of the Radio Resource Management framework of 802.11k [15], and can be easily implemented by capturing all the co-channel beacons, and passing them up to the driver for monitoring purposes.

2. We assume that the interference received by a client is approximately equal to the interference received at the serving AP. While this assumption might seem restrictive at first, in high density scenarios where the cell sizes are very small (20-50 meters) this assumption is not too restrictive. To relax this assumption, we could instead assume that the client reports the received signal strength of all the interfering APs to its serving AP, and ORCCA will still work under the latter assumption. However, as of now, there is no mechanism in the 802.11 standard for the client to convey a detailed interference report. Subsequent revisions of the 802.11 standard may support this feature.

3. We assume that the APs communicate with a central controller, and report their measured quantities (described later in this section). The central controller uses these measurements to compute the optimum CCA Threshold, and communicates it back to the APs. Several commercial products use a central controller for network parameter tuning [19].

Consider all the APs operating on a certain fixed channel, say channel 1. Assume that there are M APs in the network operating on this channel (not all of them may interfere with each other). Let variable l index these APs. The proposed algorithm is as follows.

At an AP: Each AP, l, performs the following steps:

Step 1: Measure the beacon strength of all the co-channel APs, sort them in decreasing order, and save them in $P_I^{(l)}(k)$, where k varies from 1 to M. Note that AP l may not receive measurable signal from all the APs, in which case the corresponding $P_I^{(l)}(k)$ are set to 0.

Step 2: Measure received signal strength of all the clients associated with itself (AP l), and find the client whose signal is received at the lowest power. Let this power be $P_R^{(l)}$. Assuming a symmetric channel, $P_R^{(l)}$ is also the power received by the client from AP l. Thus, this client is located at the edge of the cell, and hence is used for calibrating the CCA threshold of AP l.

Step 3: Let i be the index of supportable modulation-coding schemes. Let C_i be the supportable data rate, and β_i be the required SINR for scheme i. For each scheme i, AP l does the following.

1. Determine

$$I_i^{(l)} = \left(\frac{P_R^{(l)}}{\beta_i} - N_0 \right) \tag{24}$$

This is the maximum allowable interference while still supporting rate i for the user with the weakest received signal strength. Note that if $I_i^{(l)} < 0$, then this user cannot be served at rate i. If $I_i^{(l)} > 0$, then the CCA Threshold required for supporting the ith data rate in cell l is:

$$\text{CCA}_i^{(l)} = \left(\frac{P_R^{(l)}}{\beta_i} - N_0 \right) = I_i^{(l)} \tag{25}$$

2. Find the smallest index $m_i^{(l)}$ such that

$$\sum_{k=1}^{M} P_I^{(l)}(k) = \sum_{k=1}^{m_i^{(l)}} P_I^{(l)}(k) + \sum_{k=m_i^{(l)}+1}^{M} P_I^{(l)}(k), \tag{26}$$

and

$$\sum_{k=m_i^{(l)}+1}^{M} P_I^{(l)}(k) < I_i^{(l)} \tag{27}$$

The above implies that if the $m_i^{(l)}$ strongest interfering APs are suppressed, then the interference from the remaining APs can be accommodated while still supporting rate i for the worst case user associated with AP l.

3. The achievable throughput for scheme i in cell l, denoted by $\gamma_i^{(l)}$, is given by

$$\gamma_i^{(l)} = \begin{cases} \frac{C_i}{\left(m_i^{(l)}+1\right)} & \text{if } I_i^{(l)} > 0. \\ 0 & \text{otherwise} \end{cases} \tag{28}$$

Although the data rate C_i is supportable when $I_i^{(l)} > 0$, there is a reduction in throughput due to time sharing of the wireless channel by virtue of the CCA Threshold. When the CCA Threshold is chosen such that (26) and (27) are satisfied, we note that AP l time shares its channel with $m_i^{(l)}$ other APs. Hence the term $(m_i^{(l)}+1)$ in the denominator of (28). Note that in analysis, this was taken into account by the term X_i^2 in the denominator of throughput-coverage in (17). After each AP l determines its expected throughput $\gamma_i^{(l)}$, for ith modulation-coding scheme, it sends all the 8 triplets $\left(i, \gamma_i^{(l)}, \text{CCA}_i^{(l)}\right)$ to the central controller.

At the Central Controller: The central controller determines the network-wide common CCA Threshold based on the information gathered from all the APs. Assume that the central controller chooses a network-wide CCA Threshold CCA*. Let the *highest* supportable modulation-coding scheme for AP l with this CCA Threshold be $i(l)$. Then,

$$\text{CCA}_{i(l)}^{(l)} \leq \text{CCA}^* \tag{29}$$

Using (28), the overall network throughput is given by,

$$\text{Network Throughput} = \sum_{1}^{M} \gamma_{i(l)}^{(l)} \tag{30}$$

The central controller determines that value of CCA* which maximizes the above network throughput. This amounts to performing an exhaustive search for $\text{CCA}^* = \text{CCA}_i^{(l)}$ over all i and l (with M APs, at most $8M$ combinations). Only those (i, l) triplets are considered for which $\gamma_i^{(l)} > 0$. This ensures that the maximization of network throughput does not starve any user.

Note that the task of the central controller is considerably simplified since each AP l computes its 8 triplets $\left(i, \gamma_i^{(l)}, \text{CCA}_i^{(l)}\right)$. Even in a network with M APs where M is large, the central controller just needs to perform $O(M)$ additions to determine the optimum CCA Threshold. Also note that ORCCA relies on measurements, and hence is not restricted to the case of regular topologies. In fact, ORCCA can function in both 2D as well as 3D environments. ORCCA does not assume any specific path-loss or shadow-fading model. It can also be easily modified to account for a heterogeneous collection of APs in which 802.11g and 802.11b APs co-exist on the same channel. This would require minor modifications to ORCCA to take into account the modulation and coding schemes that 802.11b supports in addition to taking into account the 802.11g modulation and coding schemes. This requires performing optimization over 12 schemes (8 of 802.11g and 4 of 802.11b) in Step 3 of the algorithm.

6. SIMULATION RESULTS

In this section, we present simulation results on ORCCA's performance. We use Opnet Modeler version 11.0 for our simulations. Opnet has an excellent physical layer model for simulating wireless links. It has support for all the 8 modulation-coding schemes of 802.11a/g, and it simulates the transmission of individual bits over the wireless physical layer. Unlike the Network Simulator ns-2, it incorporates the capture phenomenon which is a crucial aspect of wireless communication in the presence of interference. We simulate three schemes for two different topologies (six scenarios in total). The three schemes are (i) the base scheme of CSMA/CA MAC which uses the default CCA Threshold (equal to the noise power), (ii) ORCCA, and (iii) the ECHOS algorithm which is a recently proposed algorithm for CCA adaptation [18]. To the best of our knowledge, ECHOS and the CCA adaptation algorithm in [4] are the only two other algorithms which attempt to configure the CCA threshold of a network. We do not simulate the latter algorithm, since it relies on measurement of packet error rates for tuning the CCA threshold, and the current version of Opnet (version 11.0) does not support shadow-fading which has a significant impact on packet error rates. We simulate the following 2 topologies:

Regular topology: 16 APs are deployed in a hexagonal regular structure (4×4) with each AP having 4 clients associated with it. The APs have an inter-AP separation of 70 m, while the client-AP separation within each cell is 20 m. This is equivalent to having a regular hexagonal cell pattern of 48 APs with 3 channels (an 802.11g network with a cell radius of 20 m). We focus on one of these three channels, and study the subset of the APs operating on that channel. The area of the total region covered is about 300m by 300m.

Random topology: 48 APs and 292 clients (about 4 clients per AP) are deployed randomly and uniformly over a region of an area 300m by 300m. The users associate with their closest AP. We assume that there are 3 orthogonal channels available (802.11g), and the channels are allocated randomly to the APs[1]. We consider one of the three channels, and study the subset of the network operating on this channel.

Link level throughput performance is measured using saturated UDP traffic from the APs to the clients for a duration of 10 minutes. Transmit power of 20 dBm, noise power of -91 dBm, and path loss exponent of 3.0 were used in the simulations, emulating 802.11g links. The RTS/CTS handshake was disabled while rate adaptation, and shadow fading was not supported. The channel model consists of path loss computed as per the ITU Recommendation [14]. Studying the impact of time-varying channel gain on the performance

---

[1] More efficient channel selection algorithms such as [8] could also be incorporated, however channel allocation is not the focus of this work.

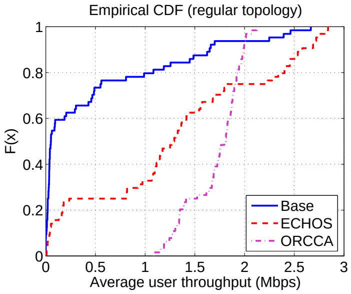

(a) Regular Topology

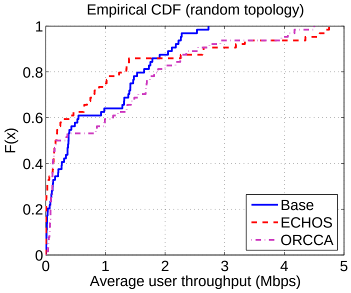

(b) Random Topology

Figure 8: Cumulative Distribution of user throughputs.

of the proposed CCA adaptation algorithm is part of our future work. Notice that our formulation does incorporate such an effect but unfortunately cannot be tested using today's 802.11 simulation tools.

Scenario	Total	Mean	Median	σ
	Regular topology			
Base case	30.3	0.47	0.05	0.72
ECHOS	83.3	1.3	1.3	0.9
ORCCA	109.2	1.7	1.8	0.27
	Random topology			
Base case	50.6	0.78	0.38	0.82
ECHOS	51.2	0.79	0.17	1.3
ORCCA	65.1	1.0	0.26	1.2

Table 3: Throughput results (Mbps).

From Fig. 8(a),(b), we note that the CDF of the throughput per user for ORCCA is found to the right of the plot, thereby showing a higher average throughput. In addition, the minimum throughput achieved by ORCCA in the regular topology is significantly greater than in the other two cases, clearly demonstrating the risk of user starvation under ECHOS, whereby no network-wide CCA is computed. From Table 3 we note that for a regular topology, ORCCA results in 260% (respectively 31%) improvement in the total network throughput as compared to the base case (respectively ECHOS). Even for a random topology the gains are 28% and 27% respectively. ORCCA is shown to feature the smallest standard deviation σ, for a regular topology (Table 3) thereby demonstrating good fairness properties. For a random topology, the standard deviation is bound to be high due to large variations in the channel qualities resulting from the random node distribution. Thus, the simulation results clearly demonstrate that ORCCA succeeds in identifying the best CCA threshold for the network and thus in optimizing overall network throughput.

7. RELATED WORK

In this work, we studied the problem of CCA adaptation and network dimensioning in high density wireless networks, under a unified framework. To the best of our knowledge, prior work has primarily focused on CCA adaptation *in isolation* to network design. In [2, 3, 5], the problem of optimal CCA adaptation is addressed within an analytical framework by assuming a regular hexagonal lattice placement of APs, and a noise-free environment. In [5], the authors bring to light the fact that different modulation-coding schemes have different SINR requirements, and hence different optimum CCA Thresholds. The authors do not take into account the impact of the available number of orthogonal channels on CCA tuning. Thus, the optimum CCA threshold that is analytically computed for several scenarios is not even feasible with 802.11a, where as per our results, 12 orthogonal channels are typically enough to mitigate interference. Besides, the approach adopted by the authors of ignoring noise right at the beginning of the analysis does not let them identify noise-dominated and interference-dominated regimes.

In [18], the authors propose a CCA adaptation algorithm called ECHOS, that uses a measurement-driven approach for tuning the CCA Threshold of APs and clients. The clients report the measured interference levels to their respective APs. Each AP, if possible, tries to set its CCA Threshold so as to support the highest data rate at the farthest client. The algorithm has three key limitations. Firstly, it does not take into account the impact of high SINR requirements on the overall network throughput. As we showed in Section 4, using the highest supportable data rate is not necessarily optimum in high density networks. Secondly, in determining the worst case SINR in a cell, the algorithm takes into account only the strongest interferer heard by the client. We know that interference is cumulative, and hence the *sum* of the powers of all the interfering signals should be considered as in (26) and (27). Lastly, and more importantly, ECHOS cannot preclude user starvation since each node selects its optimal CCA threshold in isolation, something also demonstrated through our results.

In [4], the authors propose an algorithm that measures the packet error rate across the entire network, and adjusts the CCA threshold based on the collected measurements. The CCA threshold is increased (additively) if packet error rates are below the target threshold, and decreased (multiplicatively) if the packet error rates are above the threshold. For reasonable target packet error rates (10-20%), there seems to be no gain from the proposed algorithm. The proposed algorithm only provides gains when the allowable packet error rates are up to 50%. The fact that the proposed algorithm does not provide any gains for 10-20% packet error rates can be easily inferred from the observation that the algorithm does not result in any change in the initial preset CCA threshold.

Cell-planning has been used in the cellular community for several years. The capacity planning phase of such networks consists in modeling the networks as hexagonal lattices, obtaining expressions for worst-case SINR, and then dimensioning the network parameters appropriately [9]. However, as noted earlier, in a cellular voice network, a user has a *dedicated* channel which is allocated to the user for several minutes. Hence, there is no issue of MAC contention. Interference mitigation is achieved by using per user power control [6]. In cellular data networks such as HSDPA, interference mitigation is achieved by using advanced modulation coding techniques such as Early Packet Termination and Hybrid-ARQ [7].

8. CONCLUSION AND FUTURE WORK

We determined optimum settings for the AP density and CCA threshold parameters in dense 802.11a/g networks, both for original deployments, as well as for the fine-tuning of existing installations. We showed that 802.11g networks are interference dominated, and require optimum tuning of the CCA threshold. On the other hand, the impact of CCA tuning in 802.11a networks is limited, and proper AP deployment and channel allocation can mitigate interference in most settings. However, as AP density increases, interference is likely to play an important role even in 802.11a networks. In such scenarios, CCA adaptation is required to counter interference. Interestingly, our analysis revealed that in dense 802.11 networks, for the benefit of the overall network throughput, it may not be optimum to use the highest supportable data rates.

We proposed a measurement-driven algorithm for CCA adaptation of 802.11 networks with arbitrary topologies, which we call ORCCA. Due to its measurement-driven nature our algorithm does not require any assumptions about the dimensionality of the environment (2-D vs. 3-D), the shadowing and fading properties of the space or the interference regime it is meant to address. Simulation results demonstrate that ORCCA results in up to 260% improvement in network-wide throughput for a regular topology, and up to 28% improvement for a random topology as compared to using the default CCA threshold, today's current best practice. It can further lead to throughput improvements of 31% and 27% respectively when contrasted to recent solutions.

Based on the encouraging simulation results, we are currently developing a testbed to experimentally study the performance of ORCCA. ORCCA, currently, requires a central controller for the computation of the optimum CCA threshold. Designing a distributed algorithm that uses local information to solve this problem is part of our future work. The current work also assumes full co-operation between all the APs in the network for CCA adaptaion. However when APs belong to different entities some of the APs may not participate in the CCA adaptation algorithm. Studying scenarios of non-cooperating APs and incremental deployment is a part of our future work. Lastly, in future work we intend to investigate ways in which we can relax our assumption on purely downlink traffic that was used for analytical tractability.

9. REFERENCES

[1] G. Davis, "Using picocells to build high-throughput 802.11 networks," `http://rfdesign.com/mag/407rfdf1.pdf`, Jul. 2004.

[2] J. Zhu, X. Guo, L. L. Yang, and W. S. Conner, "Leveraging Spatial Reuse in 802.11 Mesh Networks with Enhanced Physical Carried Sensing," in Proc. of *IEEE Infocom*, Miami, Florida, USA, 2005.

[3] X. Yang, and N. Vaidya, "On Physical Carrier Sensing in Wireless Ad Hoc Networks," in Proc. of *IEEE Infocom*, Miami, Florida, USA, 2005.

[4] J. Zhu, B. Metzler, X. Guo and Y. Liu, "Adaptive CSMA for Scalable Network Capacity in High-Density WLAN: A Hardware Prototyping Approach," in Proc. of *IEEE Infocom*, Barcelona, Spain, Apr 2006.

[5] H. Zhai, and Y. Fang, "Physical Carrier Sensing and Spatial Reuse in Multirate and Multihop Wireless Ad Hoc Networks," to appear in Proc. of *IEEE Infocom*, Barcelona, Spain, April, 2006.

[6] G. Foschini and Z. Miljanic, "A Simple Distributed Autonomous Power Control Algorithm and its Convergence," *IEEE Transactions on Vehicular Technology*, vol. 42, No. 2, pp. 641-646, Nov 1993.

[7] Q. Wu and E. Esteves, "The CDMA2000 High Rate Packet Data System," *Chapter 4 of Advances in 3G Enhanced Technologies for Wireless Communications*, March 2002, Editors: Jiangzhou Wang, Tung-Sang Ng.

[8] B.-J. Leung and K. K. Kim, "Frequency assignment for IEEE 802.11 wireless networks," in Proc. of *58th IEEE Vehicular Technology Conference (VTC-2003-Fall)*, vol. 3, pp. 1422-1426, Oct 2003.

[9] T. Rappaport, "Wireless Communications: Principles and Practice," second edition, Prentice Hall.

[10] S. Boyd and L. Vandenberghe, "Convex Optimization," Cambridge University Press, 2004.

[11] Data sheet for Cisco Aironet 1230AG series AP. Available at `http://www.cisco.com`.

[12] The Wireless Grid, Whitepaper, Aruba Wireless Networks, `http://www.arubanetworks.com`.

[13] J. Yee and H. Pezeshki-Esfahani, "Understanding Wireless LAN Performance Tradeoffs," `http://www.commsdesign.com`, Nov 2002.

[14] Recommendation ITU-R P.1238-4, "Propagation data and prediction method for the planning of indoor radio communication systems and local area networks in the frequency range of 900 MHz to 100 GHz," 2005.

[15] Working Group for IEEE 802.11k, Radio Resource Management, `http://grouper.ieee.org/groups/802/11/`.

[16] Intersil Direct Sequence Spread Spectrum Baseband Processor, Doc# FN 4816.2, Feb. 2002, `http://www.intersil.com/`.

[17] K. Whitehouse, A. Woo, F. Jiang, J. Polastre and D. Culler, "Exploiting the Capture Effect for Collision Detection and Recovery," under submission to The Second IEEE Workshop on Embedded Networked Sensors (EmNetS-II), Sydney, Australia, May 30-31, 2005.

[18] A. Vasan, R. Ramjee and T. Woo, "ECHOS - Enhanced Capacity 802.11 Hotspots," in Proc. of *IEEE Infocom*, Miami, Florida, USA, 2005.

[19] Meru Networks. `http://www.merunetworks.com`.

[20] B. Hajek, A. Krishna and R. O. Rao, "On the capture probability for a large number of stations," *IEEE Transactions on Communications*, vol. 45, no. 2, pp. 245-260, 1997.

[21] Intel PRO/Wireless 2200 BG driver for linux.