

Partially Overlapped Channels Not Considered Harmful

Arunesh Mishra¹ Vivek Shrivastava² Suman Banerjee^{3*} William Arbaugh⁴

^{1,2,3}University of Wisconsin, Madison, USA.
{arunesh,viveks,suman}@cs.wisc.edu

⁴University of Maryland, College Park, USA.
waa@cs.umd.edu

ABSTRACT

Many wireless channels in different technologies are known to have partial overlap. However, due to the interference effects among such partially overlapped channels, their simultaneous use has typically been avoided. In this paper, we present a first attempt to model partial overlap between channels in a systematic manner. Through the model, we illustrate that the use of partially overlapped channels is not always harmful. In fact, a careful use of some partially overlapped channels can often lead to significant improvements in spectrum utilization and application performance. We demonstrate this through analysis as well as through detailed application-level and MAC-level measurements. Additionally, we illustrate the benefits of our developed model by using it to directly enhance the performance of two previously proposed channel assignment algorithms — one in the context of wireless LANs and the other in the context of multi-hop wireless mesh networks. Through detailed simulations, we show that use of partially overlapped channels in both these cases can improve end-to-end application throughput by factors between 1.6 and 2.7 in different scenarios, depending on wireless node density. We conclude by observing that the notion of partial overlap can be the right model of flexibility to design efficient channel access mechanisms in the emerging software radio platforms.

Categories and Subject Descriptors

C.4 [Computer Systems Organization]: Performance of Systems;
C.2.1 [Computer Communication Networks]: Network Architecture and Design—*Wireless Communication*

General Terms

Measurement, Algorithms, Experimentation, Performance.

Keywords

IEEE 802.11, channel assignment, partially overlapped channels.

1. INTRODUCTION

Wireless communication mostly uses electromagnetic signals to transmit information. While a wireless signal occupies a large

*Supported by NSF Grant CNS-0520152.

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMetrics/Performance'06, June 26–30, 2006, Saint Malo, France.
Copyright 2006 ACM 1-59593-320-4/06/0006 ...\$5.00.




Figure 1: Channels with and without partial overlap.

range of frequencies, the energy is typically concentrated in a relatively narrow range of frequencies. The wireless spectrum is partitioned into ranges of frequencies, referred to as spectral bands, usually by regulatory bodies such as the Federal Communications Commission (FCC) in the US. Since different wireless technologies use different signal modulation and access mechanisms, not all of which are compatible with each other, the FCC and other such regulatory bodies define spectrum usage policies that dictate technology usage constraints on these spectral bands.

In order to resolve rights to transmit on the wireless medium among competing transmitters, many wireless technologies use a two-tier approach. First, they split the spectral band into sub-ranges called “channels” and each transmitter (and its corresponding receivers) are required to operate on one of these channels. Clearly, if there are N potential transmitters and M channels, then the contention problem is reduced by a factor of M . Within a given channel, different technologies use different mechanisms, examples being Time Division Multiple Access (TDMA), Code Division Multiple Access (CDMA), or even random access mechanisms.

We instantiate an example with the IEEE 802.11 (a/b/g)-based wireless technologies. The 802.11b extension operates in the 2.4 GHz spectral band, which is split into 11 channels. The bandwidth of each channel is 44 MHz. When operating in the 802.11 Distributed Coordination Function (DCF) mode, all wireless transmitters assigned to this channel use a random access contention mechanism, such as RTS-CTS handshakes. The wireless signal itself is modulated using a Frequency Hopping Spread Spectrum (FHSS) or Direct Sequence Spread Spectrum (DSSS) technique that allows it to potentially co-exist with transmissions from other wireless technologies and ambient noise.

A transmission on a given channel interferes with any other transmission on the same channel within a certain range, called the interference range. The interference range of a transmission depends on the transmission power used. Therefore, the choice of transmission power determines the amount of spatial re-use of the same channel, i.e., the physical separation required for two simultaneous transmitters to co-exist on the same channel. In order to increase spatial re-use, each wireless technology imposes specific limits on the permissible transmit power on its channels. (FCC in the US also regulates the maximum transmission power permissible in each spectral band.)

Defining channel boundaries: Let us examine the following question: given that each channel in a specific technology occupies a

fixed bandwidth, how far apart should usable channels be spaced in the frequency domain?

To answer this question, we first examine the relationship between the energy of a transmitted signal and the information capacity of a wireless channel. While the exact nature of this relationship depends on the precise choice of physical layer modulation scheme, other technological characteristics, as well as the noise characteristics in the environment, it is interesting to consider the upper bound on information capacity of a wireless channel, as defined by Shannon, as given below:

$$C = B \log_2(1 + SNR)$$

where C is the data capacity, B is the channel bandwidth, and SNR is the signal to noise ratio. As the signal energy increases, the value of SNR increases, and so does the channel capacity. Each wireless technology defines precise limits on the transmitter's output power for each frequency within its channels. Therefore, to maximally utilize the capacity of a given channel within the transmit power bound, a transmitter should emit the maximum permissible power in all frequencies of the channel, e.g., Figure 1(a). The *transmit spectrum mask* which is used by the transmitter to limit the output power on different frequencies will then be required to emulate an *ideal* band-pass filter. Under such a construction it is logical to have neighboring channels (A_1 , A_2 , and A_3) to be non-overlapping, i.e., two neighboring channels do not share any frequency. While such a construction of channels can be efficient in terms of capacity, design of such ideal transmit masks is not possible in practice.

In particular, realistic transmit masks are not ideal band-pass filters and instead are closer to what is shown in Figure 1(b), which implies that the capacities of such channels are lower than the ideal ones shown in Figure 1(a).

We now answer our question on how neighboring channels are constructed in wireless technologies such as 802.11. Neighboring channels in the 802.11 standards are constructed in two different ways. Channels in 802.11a standards are constructed similar to what is shown in Figure 1(b) with no overlap between neighboring channels¹. Hence, channels B_1 and B_2 can be used simultaneously for transmissions in the same physical vicinity.

In the 802.11b standards, channels are constructed similar to that shown in Figure 1(c), where neighboring channels (e.g., C_1 and C_2) partially overlap in the frequency domain. The implication of such a construction is that simultaneous transmissions on channels C_1 and C_2 within close physical proximity will cause interference. We call such channels, *partially overlapped channels*. Hence, in many situations such partially-overlapped channels cannot be used simultaneously. In the example in Figure 1(c) we can see that only channels C_1 , C_4 , and C_7 have no overlap in the frequency domain. We call them *non-overlapped channels*. When two in-range transmitters operate on the same channel, they interfere with each other. Such interference is known as *co-channel interference*. When two transmitters operate on adjacent channels that partially overlap, they cause lesser degree of interference, which is referred to as *adjacent channel interference*. Finally, two transmitters operating on non-overlapped channels will not interfere with each other.

Can partially-overlapped channels be used? Assignment of channels to communicating wireless nodes is an important problem in any wireless environment. The 802.11b standards define 11 channels that are operational in the US, of which only three are non-

¹In practice, there is some overlap between neighboring channels in the IEEE 802.11a standards. However, the energy in the overlapped part is quite low that we can ignore it for practical purposes.


overlapped channels, namely 1, 6, and 11. However, under current best practices, most users and wireless LAN administrators configure their wireless interfaces to use one of these three non-overlapped channels only. This is true even under many dense deployment scenarios, where limiting the choice of channels to only three alternatives imply that two nearby (and potentially interfering) nodes are actually assigned to the same channel. Such an approach is adopted due to the following reason.

In most typical scenarios, interference on the same channel, i.e., co-channel interference, can be directly detected and can be explicitly handled through contention resolution mechanisms, e.g., the RTS-CTS handshake in 802.11 networks. In contrast, adjacent channel interference often contributes to background noise and cannot be handled in an explicit manner by channel contention techniques. Hence, systematic approaches to handle adjacent channel interference is usually considered difficult. Due to the detrimental effects of adjacent channel interference, all prior wireless channel assignment approaches in diverse wireless technological scenarios (e.g., cellular networks, 802.11 WLANs, etc.) have made use of non-overlapped channels alone. In fact, in an overwhelming number of recent as well as classic papers, the notion of channel has come to be defined (and almost rightly so) as a path of information flow which is perfectly isolated from other paths of information flow, i.e., other channels. For examples in context of multi-hop wireless networks, see [1, 2, 3, 4, 5, 6, 7], in the context of single-hop wireless LANs, see [8, 9, 10, 11] and in the context of cellular networks, see [12, 13, 14, 15, 16, 17] and the references therein.

In this paper, we visit the following questions: (i) is spectrum used efficiently when only non-overlapped channels are used?, and (ii) if it is not, how can spectrum utilization be improved?

The answer to the first question is negative. It is fairly easy to see that the use of only non-overlapped channels (Figure 1(b)) leads to *wastage of wireless spectrum capacity*. This answer intuitively follows from Shannon's channel capacity observations. Non-overlapped channels and the practical limits on the shape of transmit spectrum masks imply that there are many frequencies in which the transmitted power is lower than the maximum permissible limit, which degrades the SNR, and hence the maximum achievable channel capacity. It also follows that use of partially-overlapped channels can lead to better utilization of the spectrum. However, an ad-hoc use of partially overlapped channels can actually degrade performance. Therefore, the focus of this paper is to examine systematic approaches to exploit partially overlapped channels efficiently to improve spectrum utilization. In particular, we first describe a model that captures interference effects of partially-overlapped channels, then illustrate how such a model can be effectively used in improving design of channel assignment algorithms, and finally conduct a detailed evaluation study to demonstrate how such models and algorithmic approaches lead to improved utilization of the wireless spectrum. Our examples are drawn from two application scenarios — channel assignment necessary in wireless LANs (WLANs) and multi-hop wireless mesh networks.

Relationship to physical layer coding techniques: At a first glance it might seem that better physical layer modulation techniques can utilize an entire range of spectrum while also allowing for different transmissions to 'co-exist'. For example, in the Frequency Hopping Spread Spectrum method, a single transmission is encoded over different frequencies at different times. The sequence of frequencies selected for transmission is determined a-priori using a 'hopping' pattern. Thus, nodes with different hopping patterns can co-exist in the same frequency domain. However, all nodes using the same physical layer modulation technique have to share the capacity of the wireless medium as determined by the modulation



(a) UDP throughput at 11Mbps data rate (b) TCP throughput at 11Mbps data rate

Figure 2: TCP/UDP throughputs versus physical distance.

method used. Thus, as long as the power dynamics of two transmissions cause them to interfere, the respective communicating nodes experience a reduction in capacity. The focus of this paper is to manage simultaneous transmissions carefully within the frequency domain in order to improve spectrum utilization — a mechanism which is complementary to physical layer methods.

Key contributions: The following are the main contributions of this work:

- We present a systematic and detailed model of partially-overlapped channels in wireless communication that is general in nature and applies to a wide range of communication technologies. The model is motivated through detailed experimentation.
- We use the model to modify two existing algorithms for channel assignment and management in different wireless scenarios and show how the new model allows significant enhancement in utilization of the wireless spectrum.

We believe that the work presented in this paper is a first step and can lead to interesting future directions in more online construction of partially overlapped channels. In particular this line of research can lead to agile, channel-access mechanisms that complement the ongoing developments of software radio platforms.

Roadmap: The rest of this paper is structured as follows: In Section 2, we first present a detailed measurement study that quantifies possible benefits of using partially overlapped channels. Next in Section 3, we discuss our proposed partially-overlapped channel model and analysis. In Section 4, we discuss the implications of using such channels in improving overall utilization of spectral resources. In Section 5, we discuss the applications of partially-overlapped channels in two different environments, namely WLANs and wireless mesh networks. In Section 6, we discuss related work, and finally we conclude in Section 7.

2. MEASURING PARTIAL OVERLAP

In this section, we measure the benefits of using partially overlapped channels through careful experimentation. In particular, we demonstrate how using such channels can yield improvement in throughput and spatial re-use of spectrum. These observations lay the motivation for building an analytical model (Section 3) that captures partial overlap among adjacent channels.

In our prior work [18], we performed an extensive set of experiments to carefully study the impact of using partially overlapped channels. Here, we summarize results from this evaluation. We also supplement these results with additional experiments which used different physical layer modulation methods. Below, we discuss the salient points of these experiments:

Our experiments in [18] measured the amount of partial overlap between channels by studying the impact of such overlap on




Figure 3: The measurement setup.

MAC-level and application-level metrics. These experiments were performed using the IEEE 802.11b standards that operate in the 2.4 GHz band. As discussed before, the 802.11b standards define 11 channels, each with a bandwidth of 44 MHz, while the center frequency of neighboring channels are placed 5 MHz apart. As a consequence, only three of these channels are non-overlapping, namely 1, 6, and 11. To quantify the degree of interference between various channels, we present additional results from experiments which use the following setup:

Setup: Two pairs of communicating wireless nodes built of commodity 802.11 hardware (IBM Thinkpad laptops with 802.11 a/b/g wireless interfaces) were placed as shown in Figure 3. In each pair, the two communicating nodes were placed in close proximity of each other. The lower node in each pair sent a flow of traffic to the upper node. In order to communicate, in all experiments both nodes in each pair were configured to use the same wireless channel. The physical separation and the channel separation between the two pairs of nodes was varied. Experiments in [18] used 1 and 2 Mbps data-rate for the physical layer which uses the Binary Phase Shift Keying (BPSK) modulation. Here, we report results using the Complementary Code Keying (CCK) modulation specified by the 802.11b standard which provides for 11 Mbps of data-rate. Figure 2 shows these results for TCP and UDP traffic at the application-level. We note the following salient points of such results:

(i) As the physical separation increased, the amount of interference decreased and this led to increase in throughputs. This is evident from results in Figure 2.

(ii) However, the same level of throughputs could be achieved at much lower physical separation by increasing the channel separation between the two pairs of nodes. For example, as shown in Figure 2 a channel separation of three (say channel 1 and 4) was enough for both nodes to reach maximum possible throughput with a physical separation of about 10 m. However, operating on the same channel required a physical separation of about 60 m for both links to operate without interference.

Thus, partially overlapped channels can provide much greater spatial re-use if used carefully. In the next section, we model this behaviour analytically by studying its impact on the signal-to-noise ratio and the bit-error probabilities of packet reception. Later, we build mechanisms to utilize partially overlapped channels in wireless LANs and mesh networks.

3. MODEL FOR PARTIAL OVERLAP

A wireless signal has a certain *finite* bandwidth which corresponds to the range of frequencies in which most of its energy is concentrated. When a transmitting node emits a wireless signal in a specific wireless channel, it uses a transmit spectrum mask. The transmit spectrum mask specifies the upper limit of power that is permissible for each frequency of the transmitted signal. Figure 4 illustrates the transmit spectrum mask for IEEE 802.11 standards using DSSS modulation. The channel bandwidth is 44 MHz. At the center frequency, F_c , the mask limits output power to 0 dB^2 — the output power is equal to the input power and the signal is passed unaffected. At frequencies beyond $F_c + 11 \text{ MHz}$ and $F_c - 11 \text{ MHz}$,

²Relative power in $\text{dB} = 10 \log(P_{out}/P_{in})$.




Figure 4: The Transmit Spectrum Mask for IEEE 802.11 DSSS modulation.

Rel. Freq (MHz)	13	14	14.4	14.8	22.4	28	56
Rel. Power (dB)	0	-15	-20	-28	-34	-42	-52

Table 1: The transmit spectrum mask for a 28 MHz channel of the WirelessMAN physical layer in the IEEE 802.16 standard.

the power is attenuated down by -30 dB and further to -50 dB at $F_c \pm 22$ MHz, where F_c is the center frequency for the channel c . Similarly, Table 1 tabulates the transmit spectrum mask for the IEEE 802.16 standards (WiMAX), which has a similar structure.

Note that this transmit spectrum mask is ideal and in reality only some continuous approximation is achieved.

To receive a given signal, a receiver uses another band-pass filter to selectively receive a certain frequency band. The band-pass filter ‘allows’ a certain band around the center frequency and eliminates all other frequencies to pass through the radio circuitry at the receiver. The power with which a signal is received depends on the amount of *overlap* between the receiver’s band-pass filter and the transmitter’s signal distribution (usually limited by the transmit spectrum mask) in the frequency domain. Therefore, if the center frequency of the receiver’s band-pass filter is not aligned with the center frequency of the transmitter’s signal distribution, the received power of the signal reduces. In particular, if the center frequency of the receiver’s filter does not overlap with the transmitter’s spectrum mask, the signal is not perceptible at the receiver.

Based on these observations, it is possible to quantify the notion of partial overlap between two wireless channels. We illustrate this with an example shown in Figure 5. Shown is the distribution of a transmitter’s output power made on a certain channel with center frequency F_c . For example, if $F_c = 2.437$ GHz, this corresponds to channel 6 of IEEE 802.11b standards. The signal occupies a bandwidth of 44 MHz around this center frequency. A receiver with an ideal band-pass filter is positioned at $F_c + 10$ MHz. Since 802.11b channels are separated by 5 MHz, in the example this implies that the receiver is tuned to channel 8. The signal transmitted on channel 6 is received with a lower received power on channel 8, and this power is given by the energy in the shaded region in the figure. In addition, if the receiver’s filter is translated continuously to the right, the received power will decrease in a corresponding *continuous* manner.

Developing a model: The first step to developing mechanisms which take advantage of partial overlap is to build a model that captures such an overlap in a quantitative fashion. We introduce the notion of an *interference factor* (I-factor for short) that captures the *amount of overlap* between a transmission on a certain frequency F_T and reception on a certain frequency F_R . The amount of overlap is captured quantitatively by calculating the area of intersection between a signal’s spectrum and a receiver’s band-pass filter. We first provide a general definition of I-factor as a continuous function and later derive a discrete version which applies to 802.11 based transmissions which have a ‘discrete’ concept of channels.

Let a transmitter T be stationed on frequency F_T and correspondingly a receiver R is positioned to receive on frequency F_R . Let $S_T(f)$ denote the signal’s power distribution across the fre-




Figure 5: Reduction in power due to decreased overlap between the receiver’s and the transmitter’s channels.

quency spectrum. Essentially $S_T(f)$ is computed by taking the Fourier Transform [19] of the output signal. Let $B_R(f)$ denote the band-pass filter’s *frequency response*. That is, the amount of power received at a certain frequency f is given by multiplying the frequency response of the band-pass filter to the incoming signal. Let $\tau = F_T - F_R$. Based on these notations, we can define I-factor function that applies to the transmission of any band-limited signal regardless of their modulation method (such as OFDM, DSSS, etc.), as:

$$IF_{(T,R)}(\tau) = \int_{-\infty}^{+\infty} S_T(F)B_R(F - \tau) df \quad (1)$$

The parameter τ represents the amount of overlap as a continuous variable. $\tau = 0$ indicates that both signals have the same center frequency and a increasing value of $|\tau|$ indicates reduction in the overlap.

The above definition of I-factor captures the overlap between a transmission and reception on any two frequencies. However, the current wireless communication standards (802.11, 802.16 etc) define a set of discrete channels for radios to operate on. We devise a discrete version of the I-factor which calculates the overlap between two discrete channels assuming similar transmission and reception characteristics (i.e. same modulation technique).

Specifically for transmissions based on IEEE 802.11 standard, we define an idealized discrete model of the I-factor denoted by $I_{\text{theory}}(i, j)$ as the amount of overlap between channels i and j . In this idealized model we assume that the transmitted signal’s power distribution has the exact form of the transmit spectrum mask (Figure 4). Since it is advantageous for a wireless card designer to use the same filter for transmitting a signal and band-limiting the reception, we use the receiver filter to be the same as the transmit spectrum mask.

Then, I-factor in this idealized model can be computed by instantiating $\tau = 5|i - j|$ in MHz in Equation 1 (separation in MHz) and $S_T(f)$ and $B_R(f)$ in the same equation as follows:

$$B_R(f) = S_T(f) = \begin{cases} -50dB & \text{if } |f - F_c| > 22\text{MHz} \\ -30dB & \text{if } 11\text{MHz} < |f - F_c| \leq 22\text{MHz} \\ 0dB & \text{otherwise} \end{cases}$$

where F_c denotes the channel center frequency.

The discrete I-factor can also be obtained empirically through measurements as follows. If P_i denotes the power received at a given location of a particular signal on channel i and P_j denotes the received power for the same signal and at the same location on channel j then $I_{\text{measured}}(i, j)$ can be calculated as $\frac{P_i}{P_j}$. This essentially gives the fraction of a signal’s power on channel j that will be received on channel i and can be empirically obtained through simple measurements for any given wireless technology.

The plot in Figure 6 shows that the theoretical and measured I-




Figure 6: Theoretical vs measured I-factor for two 2.4 GHz on a 2.4 GHz 802.11 channels.

factor values for two interfering 802.11b wireless channels match fairly well (measurement performed using commodity hardware). We also show the theoretical I-factor at a 2.4 GHz 802.11 channel due to interference from a 802.16 transmission in Figure 7. This was calculated using the 802.11 transmit spectrum mask (Figure 4) and the 802.16 transmit spectrum mask (Table 1). Note, measured I-factor for this case was not available due to lack of 802.16 hardware.

3.1 Inferring interference effects

A wireless signal attenuates in strength as it travels from the transmitter to the receiver. If the transmitter and the receiver are tuned to the same channel, this is the only form of signal attenuation visible at the receiver. However, if the receiver and the transmitter can be tuned to different channels (with different center frequencies) an additional signal attenuation is perceived at the receiver, as given by the I-factor. Thus, both physical distance as well as spectral distance between the transmitter and the receiver is responsible for signal attenuation. Hence both of these factors should be considered when the interference effects of a signal is defined.

Let the transmitter T and the receiver R be separated by a distance, d and operating on channels i and j respectively. If P_t is the transmitted power of the signal, then by enhancing the two-ray ground propagation model [19, 20] with the notion of I-factor, the received power is given by:

$$P_r = \frac{P_t C_t I(i, j)}{d^k} \quad (2)$$

where C_t denotes a set of constants capturing various radio properties such as the antenna gains and height, etc. and the path loss parameter k , is typically between 2 and 4. If N is the ambient noise at the receiver, then the signal is correctly received at the receiver if the Signal to Noise Ratio (SNR) exceeds a Carrier Sense Threshold, Th . In this case, R is considered to be within the transmission range of T . On the other hand, if the SNR falls below this threshold, the signal is not correctly decoded at the receiver. Instead the received power adds to the noise at the receiver. In this case we consider R to be in the interference range of T but not within its transmission range.

Now consider the case where two transmitters, T and T' , are attempting to communicate with two receivers R and R' respectively. Assume that both receivers are within the transmission ranges of both transmitters. In such a case, both transmissions cannot happen simultaneously (due to interference effects) thereby reducing parallelism. However, since both transmitters are in range of the receivers, the contention to transmit can be resolved using standard standard MAC level mechanisms, such as RTS-CTS.

When a receiver R' is in the interference range of one transmitter T , but not in its transmission range, such contention resolution is not possible. This is because MAC-level mechanisms, such as




Figure 7: Theoretical I-factor for two 2.4 GHz on a 2.4 GHz 802.11 channel due to an interfering 802.16 channel.




Figure 8: Transmission range (normalized) and channel separation for 802.11 networks.

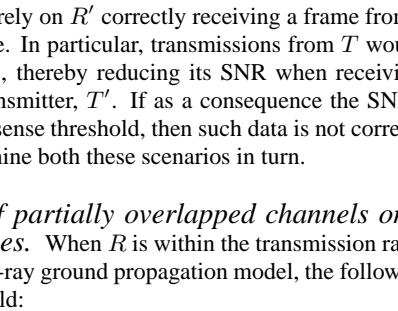


Figure 9: Effect of partial overlap on the bit error rate.

RTS-CTS, rely on R' correctly receiving a frame from T , which is not the case. In particular, transmissions from T would add to the noise at R' , thereby reducing its SNR when receiving data from another transmitter, T' . If as a consequence the SNR falls below the carrier sense threshold, then such data is not correctly received.

We examine both these scenarios in turn.

Impact of partially overlapped channels on transmission ranges. When R is within the transmission range of T , under the two-ray ground propagation model, the following condition needs to hold:

$$Th < \frac{P_r}{N} = \frac{P_t C_t I(i, j)}{N \cdot d^k} \text{ i.e.,} \quad (3)$$

$$d < \sqrt[k]{\frac{P_t C_t I(i, j)}{N \cdot Th}} \quad (4)$$

Note that this equation implies that a transmission made on any channel can potentially be correctly received by tuning the receiver to a neighboring, partially-overlapped channel.

If $i = j$, i.e., both the transmitter and the receiver are operating on the same channel, then $I(i, j) = 1$, and the transmission range is given by $d_{ii} = \sqrt[k]{\frac{P_t C_t}{N Th}}$. Otherwise, if the transmitter is on channel i and the receiver on channel j , then the transmission range is given by $d_{ij} = \sqrt[k]{\frac{P_t C_t I(i, j)}{N_a}}$. It follows that

$$d_{ij} = I(i, j)^{1/k} d_{ii} \text{ where, } d_{ii} = \sqrt[k]{\frac{P_t C_t}{N Th}} \quad (5)$$

Similar models can be derived for other radio propagation models as well. In Figure 8 we illustrate how the transmission range varies with channel separation for the two-ray ground propagation model (normalized against the interference range for channel separation of zero). Note, that y-axis is in a logarithmic scale. We can observe that the interference range decay is quite fast. For example, channels 1 and 4 (with separation of 3) can be used without direct interference when separated by a mere distance of about 3

meters (assuming the transmission range on the same channel is 300 meters). This implies that good spatial re-use is possible by employing a set of partially overlapping channels. We present more detailed arguments for capacity improvements using partially overlapped channels in Section 4.

Impact of partially overlapped channels on noise and bit errors. If a signal is received at R with power below the carrier sense threshold, it cannot be correctly decoded at the receiver and instead adds to the noise. If R is attempting to receive a signal from another transmitter, such a transmission from T lowers its SNR and may contribute to losses.

Let $P_{ii} = \frac{P_t C_t}{d^k}$ be the received power on channel i with the transmission made on the same channel. Let P_{ij} denote the power received on channel j . Clearly, $P_{ij} = I(i, j) \frac{P_t C_t}{d^k}$. Considering this transmission as noise, we note that the partial overlap has reduced the signal strength by $I(i, j)$. That is $1/I(i, j)$ concurrent such transmissions on partially overlapped channels would bring up the noise to the same level as a single transmission on the same channel.

For example, consider the case when the receiver is on channel 4 of the 2.4 GHz band and transmitters are on channel 6. To match the interference effect of a single transmitter on channel 4, we would need $1/I(6, 4) = 10^{-1.147} = 14$ transmitters in channel 6. Even in densely deployed wireless environments, such occurrences are rare. This is because 14 transmitters on the same channel (channel 6, for example) can transmit simultaneously if and only if none of them are in range of each other. (If two of these transmitters are in range of each other, then normal channel contention mechanisms, such as RTS-CTS, will only allow one of the transmissions to proceed at any time.) Given that the receiver R is in interference range of all 14 transmitters, it is difficult to find a configuration in which none of these transmitters are in mutual range.

We now study the effect of noise from such interfering transmitters on partially overlapped channels on a given receiver, by examining its bit error rate.

Bit error rates: In order to model the impact of interfering transmitters on bit error rate, we need to use some model for the modulation scheme in use. Consider an 802.11-based environment where the modulation scheme used is DSSS based binary phase shift keying (BPSK). For such a modulation scheme, the bit error rate of the channel is given by $p_b = \text{erfc}(2 * E_b/N_o)$ where p_b gives the probability of a bit being received in error, E_b is the energy per bit of a transmission and N_o is the background noise level [19].

Let T be a transmitter in interference range of R , while T' be a transmitter in transmission range of R . The received power of the interfering transmission from T as received by R is calculated using the two-ray ground propagation model as discussed in the previous section. We assume that both T and T' use the same transmission power P_t and have similar radio characteristics (such as antenna gain, etc.) denoted by the constant C_t . We are interested in calculating the bit error rate at R , tuned to channel j , for the following scenario: T' is transmitting some data to R on channel j , while a transmission from T on channel i is causing interference. Based on the two-way ground propagation model, the error model for BPSK, and the definition of $I(i, j)$, this is given by:

$$p_b = \text{erfc}\left(\frac{2 * P_t C_t}{d_{T'R}^k (N_o + I(i, j) * \frac{P_t C_t}{d_{TR}^k})}\right)$$

where, $d_{T'R}$ and d_{TR} are the distances between T' and T from R respectively.

We show the effect on p_b by using realistic values for the var-

ious parameters. We use a transmit power of 100 mW, a carrier sense threshold of -102 dBm and an indoor ambient noise level of -110 dBm (based on measurements) mimicking typical 802.11 radio properties [21]. Using a receiver sensitivity of about -87 dBm, this gives a transmission range of about 30m. We use a maximum value of 30 m for $d_{T'R}$, the distance between the transmitter T' and receiver R . This value reduces the energy per bit E_b to the minimum possible as per the receiver sensitivity thresholds, thus allowing us to observe the worst possible impact on the bit error rate. The physical distance between the interferer T and the receiver R is varied between 10 m to 100 m while maintaining the maximum possible separation of 30 m between T' and R . Figure 9 plots the bit error probability p_b for various configurations, when $k = 2$. Note that each plotted line is cut off at the distance for which T' 's transmission is correctly received at R (i.e., R is now in transmission range of T). The 'Same Channel' case shows the effect on the bit error rate when T is operating on the same channel as R . The 'Ambient Noise Only' case shows the bit error rate due to the ambient noise without the presence of any interference on any channel. It is easy to observe that the bit error rate falls rapidly with increase in channel separation as well as with increase in distance. (Note that y-axis is in the log scale.)

4. CAPACITY IMPROVEMENTS

We now discuss the impact of partial overlapped channels on total capacity of wireless environments. In particular, we compare the achieved capacity to the scenario when only non-overlapping channels are used.

Let us consider that in a given spectral band there are a total of M channels of which N are non-overlapping. For 802.11b, $M = 11$ and $N = 3$. Note, that we are considering the same spectral band in both scenarios. Hence, the total bandwidth being compared are the same in both cases.

We first examine non-overlapping channels. Now, consider a wireless environment where we have a set of nodes V within a certain specified region which share the set of N non-overlapping channels. We define a link between any two nodes $u, v \in V$ which are interested in communicating with each other through basic wireless transmissions (without any higher level routing). We represent this link by a directed edge $e = (u, v)$ as shown in Figure 10.

For simplicity, let's assume that every node has a single radio with similar radio characteristics such as transmit power, receiver sensitivity thresholds (minimum power required to receive a packet), etc. Thus, each node has a fixed transmission range R . Let d_{uv} denote the distance between nodes u and v . (u, v) is a wireless link if and only if $d_{uv} \leq R$, this implies that the node u can transmit to node v .

Consider a node u that makes a transmission to node v on a specific channel i as shown in Figure 10. If there are, say, n nodes contending for the medium along with node u , we would expect u node to get roughly $1/n$ share of the capacity of the wireless medium using standard methods to resolve contention at the link layer. For example, techniques such as the distributed coordination function (DCF) of 802.11 can be used. The number of nodes which contend with u is a crucial factor in determining the long term throughput of node u .

One way of reducing this contention experienced by node u is to partition the set of contending nodes among the N non-overlapping channels. For ease of exposition, let's assume that the nodes are distributed randomly within the specified region of interest. Let ϕ denote the node density, that is the number of nodes within a unit area. Let $\lambda(N)$ denote the expected number of nodes that contend




Figure 10: The interference region for node u transmitting to node v .

with node u , which is given by calculating the number of nodes within the region of contention:

$$\lambda(N) = R^2 \frac{\phi}{N} = \lambda(N)$$

We now examine the case where we use the M partially-overlapped channels. When using partially overlapped channels, a node on channel i would contend with nodes within a region given by $I(i, j)^{1/2}R$ on a partially overlapped channel j . Here, $I(i, j)$ is the I-factor function discussed in Section 3. Using these M channels and the notion of I-factor we can compute expected number of nodes that contend with node u as $\lambda(M)$, given by

$$\lambda(M) = \sum_{j \in \{1 \dots M\}} (R(I(i, j))^{1/2})^2 \frac{\phi}{M}$$

We evaluate the values of $\lambda(N)$ and $\lambda(M)$ for the channel structure present in the 2.4 and 5GHz bands. Here, channels are spaced 5 MHz apart, while each channel has a width of about 20 MHz. Thus, N non-overlapping channels give us $M = 5N - 4$ partially overlapped channels. We evaluate the ratio $\frac{\lambda(M)}{\lambda(N)}$ as

$$\frac{\lambda(M)}{\lambda(N)} = \frac{N}{M} \sum_{j \in \{1 \dots M\}} I(i, j) = \frac{1.2N}{5N - 4}$$

Now for a reduction in the number of contending nodes using partially overlapped channels we would expect $\lambda(M) < \lambda(N)$, that is, $\frac{1.2N}{5N-4} < 1$ or $N > \frac{4}{3.8}$, which is trivially true for partially overlapped channels to exist (as $N \geq 2$). For example, in the 2.4GHz band there are 3 non-overlapping channels and thus $N = 3$. This gives an expected reduction in the link layer contention by a factor of 3.05.

In Section 5.1, we discuss modifications to an existing channel assignment scheme in the context of wireless LANs. The main difference between a wireless LAN and the ad-hoc wireless environment here is the centralized nature of the links. Clients send wireless traffic to designated devices called access points (APs). This makes the links AP-centric in nature. However, such a centralized link structure is still a special case for the analytic reasoning presented here. For environments which exhibit high interference due to higher node densities, our enhancements to the channel assignment scheme discussed in Section 5.1 bring improvements in application level TCP/UDP throughputs by a factor of about 3.0 matching our theoretical result here ($\frac{5N-4}{1.2N} \approx 3.05$). This shows that using the partially overlapped channels as modeled by I-factor, wireless nodes can experience less contention at the link level which translates to better throughput for the higher layers of the networking stack.

5. APPLICATIONS

In this section we discuss how the proposed partially-overlapped channel model can be employed in improving spectrum utilization in two different scenarios: (i) WLANs and (ii) multi-hop wireless

Algorithm 1 Randomized_Compaction (X, T, k)

X = set of access points,

T = set of range sets for each client,

k = number of channels

$\theta : X \rightarrow \{1 \dots k\}$ is the resultant channel assignment

- 1: X' be a random permutation of X .
 - 2: Let $X' = \{x_1, x_2, \dots, x_i\}$.
 - 3: Set $\forall x \in X, \theta(x) = -1$ /* indicates an unassigned AP */
 - 4: **while true do**
 - 5: $val \leftarrow \text{OBJ}(T, \theta)$
 - 6: **for** $i = 1 \dots |X|$ **do**
 - 7: $\theta(x_i) \leftarrow \text{Compaction_Step}(x_i, \theta, T, k)$
 - 8: **end for**
 - 9: **if** $\text{OBJ}(T, \theta) = val$ **then**
 - 10: **stop**
 - 11: **end if**
 - 12: **end while**
-

mesh networks. In both cases, we start with an existing channel assignment algorithm proposed in prior literature — a centralized greedy-style approach for WLANs [22] and an LP-based formulation for mesh networks. Both of these algorithms use a *boolean indicator variable* to model presence or absence of interference in the wireless environment. The goal of these algorithms was to increase overall utilization of the wireless spectrum using non-overlapping channels only. Our proposed modifications in both these cases require change in just the definition of this indicator variable to capture the notion of interference due to partially overlapped channels. Subsequently, by allowing assignment of partially overlapped channels, we show in this section that significant performance gains are possible.

We make two additional comments prior to discussing our proposed modifications. First, when we compare performance of non-overlapping channels to that of partially overlapped channels, we use the same spectral band. In particular the width of the spectral band is identical in both cases. Second, the choice of a boolean indicator variable to indicate interference is not ours, but were made by these prior algorithms. We believe a more efficient design of these algorithms is possible by using a more continuous representation of interference, e.g., using a real number in the $[0,1]$ range. However, we intentionally use the same construct as in these prior pieces of work in the non-overlapping assignment as well as in the partially-overlapped assignment cases. Further results (not shown in this section) illustrate that the relative performance of the two approaches remain the same irrespective of use of boolean or continuous indicator variables.

5.1 Channel Assignment In Wireless LANs

Wireless LANs have seen significant deployment as the last hop connectivity solution in various indoor environments. They operate in the so called ‘infrastructure’ mode (as opposed to ad-hoc) where a centralized entity called an access point (AP) acts as a link level gateway for a client’s traffic. A wireless client *associates* to an AP within communication range in order to obtain network service from it. An AP together with its associated clients form a basic service set (BSS). All entities belonging to a BSS operate on a single channel.

The problem of channel assignment in WLANs deals with allocating channels to APs so as to maximize performance by eliminating interference among neighboring BSSs’. Typically, a wireless LAN either uses a static channel assignment or the APs use simple heuristics such as searching for the least congested channel




Figure 11: A WLAN example.

[23]. While a careful assignment of channels to APs can improve performance over heuristics, load balancing of clients among available APs is important for achieving significant gains. Given the increase in the density of APs in an average neighborhood, careful assignment of channels and balancing of client load has become an important problem.

In [22], authors describe a client-driven approach for channel assignment, that uses set-theoretic constructs to model interference. We refer to this algorithm in this paper as *Randomized Compaction*. We first give a brief discussion of the concepts presented therein. Subsequently, we show the advantages of using partial overlap among channels by making corresponding enhancements to the techniques. **Background:** An AP along with its *associated* clients that form a BSS operate on a single channel. Two interfering APs (and associated clients) operating on the same channel can lead to severe performance penalties. For example, consider the topology shown in Figure 11 with two clients and three APs. Client C_i is associated to AP AP_i , $i = 1, 2$. Here, clients C_1 and C_2 could interfere with each other if operating on the same channel. Thus, a good channel assignment strategy would assign AP_1 and AP_2 to different channels.


In [22], authors model this problem as a *conflict* set coloring formulation. The term *conflict* refers to the interference suffered by a client. For ease of exposition, we discuss a simplified version of this model here, however, our enhancements using partial overlap were performed over the original formulation. Let (X, C) denote a wireless LAN, with X as the set of APs, and C as the set of clients. Each client c is associated with a range set, say r_c . This is the set of APs within communication range of the client. For example, for client C_1 in Figure 11, $\{AP_1\}$ comprises its range set.

Channel assignment is performed in a centralized manner using the *Randomized Compaction* algorithm presented as Algorithm 1 which optimizes a min-max objective function denoted by OBJ . This objective function captures the total interference experienced by each client as a min-max value to address issue of fairness among clients. The total interference or conflict experienced by a client can be computed as $cf_c = \sum(\eta(x) + 1)$ where $\eta(x)$ is the total number of clients associated to AP x . This sum is taken over all APs $x \in r_c$ which operate on the same channel as c . Let $\vec{CF} = \{cf_1, \dots, cf_{|C|}\}$ denote a *conflict vector*, which is the total interference or conflict experienced by each client, arranged in non-increasing order of value. A given channel assignment is said to be a min-max assignment if its corresponding conflict vector \vec{CF} has the same or lower lexicographical value than any other feasible channel assignment. That is,

Optimize: $OBJ(X, C) = \vec{CF}$ as the best min-max lexicographical value.

The randomized compaction algorithm (Algorithm 1) starts with a random permutation of APs (Step 1). In Step 5 and 9, the OBJ is computed to evaluate a specific channel assignment. During the ‘compaction’ step (Step 7), the algorithm progressively chooses the ‘best’ channel among all available channels for an AP according to this objective function. The algorithm invokes the compaction step for each AP in succession according to the random permutation performed in Step 1.


The quantity cf_c which captures the total interference suffered by a client, is the primary place which distinguishes a wired net-



(a) UDP throughput

(b) MAC Collisions

Figure 12: Results for high interference topologies.



(a) UDP throughput

(b) MAC Collisions

Figure 13: Results for low interference topologies.

work from a wireless one. This function captures the interference suffered by clients due to the broadcast nature of wireless transmissions. However, in its current form, it assumes independence among channels. Next, we discuss how our model of I-factor for partial overlap enables this algorithm to make full use of them.

Partial Overlap Based Enhancements: We briefly discuss our enhancements which incorporate the notion of partially overlapped channels. All our enhancements are confined to the cf_c function, discussed earlier. This function captures the total interference experienced by client c . The randomized compaction algorithm which uses this objective function requires no modifications.

The basic concept behind the enhancements stems from the observation that the interference function cf_c assumes independence among the available k channels. This is modified to consider interference from clients/APs on partially overlapped channels according to the I-factor model. Let $POV(x, x_{ch}, y, y_{ch})$ be an indicator boolean function which denotes if the nodes x and y , operating on channels x_{ch} and y_{ch} respectively, interfere with each other. This function can be computed using the I-factor based model discussed in Section 3. Based on the POV function, the client interference cf_c can be computed by summing up $\eta(x) + 1$ (the interference from the BSS x) over all APs x that interfere with client c on any partially overlapped channel as given by values of the POV function. With these changes, the randomized compaction algorithm now performs channel assignment and balancing of client load using all available partially overlapped channels.

The POV function can also be computed empirically by performing a *scan* operation (defined in the 802.11 standard) of all channels to determine which nodes interfere (on which channels). Such a scan operation is already performed periodically by each client participating in the randomized compaction algorithm [22] to maintain the freshness of their range and interference sets. Thus, modifications to a practical implementation of this technique can be easily performed.

Simulation Results: Here, we perform a quantitative study of using partially overlapped channels with the randomized compaction

algorithm through packet-level simulations. The simulations were performed over a modified version of the NS-2 simulator. This study compares the performance of the randomized compaction algorithm alongside the enhancements discussed earlier.

We performed modifications to the NS-2 simulator to improve the wireless link behaviors and incorporate realistic channel models. Our modifications include fixing of the binary interference model present in NS-2, an implementation of a bit error rate model and an implementation of partially overlapped channels according to the I-factor model.

The simulation parameters used closely match the configurations discussed in [22]. Various topologies were generated by varying the number of clients over a fixed region of coverage while maintaining a fixed client:AP ratio. For each topology generated, two types of simulations were performed; one with TCP flows and the other with UDP. FTP download applications over large file sizes were used to create the TCP flows. The UDP flows were generated using a constant bit rate (CBR) traffic generator with rates high enough to saturate the medium. The packet size for all traffic was set at 1024 bytes and the bit rate for the medium was set at 11Mbps.

We observed that using 11 partially overlapped channels brought about improvements in TCP/UDP throughputs by a factor between 1.6 to 3.0 depending on the client densities. If all these eleven channels had been non-overlapping we would have expected a best-case improvement of $11/3 = 3.67$. Thus, we find that 11 partially overlapped channels improve the throughputs significantly over using 3 non-overlapped ones in the 2.4 GHz band. We also observe that with partially overlapped channels, the client throughputs reach to about 70% of their maximum attainable values as compared to using 11 non-overlapping ones. This gap of about 30% is attributed to the partial overlap present among the channels.

Throughput Improvements: We first discuss the throughput results over 28 randomly generated topologies with 200 clients and 50 APs. Half of these had 'high' interference (average of 8 APs in range for each client), while the others had 'low' interference (4 APs in range). The interference was controlled by setting the transmit power of the clients/APs and the receiver sensitivity thresholds (which is the minimum signal power required for packet reception).

Figures 12(a) and 13(a) compare the per-client UDP throughput in a *normalized* manner for high and low interference topologies respectively. Both plots were normalized by the same maximum data point for comparison. The X-axis shows the topology number. Both plots show that partially overlapped channels perform consistently better. For example, in Figure 12(a) for topology number 0, randomized compaction using non-overlapping channels gave a normalized per-client UDP throughput of 0.3, while using partial overlap a throughput of 0.82 was achieved which gives an improvement factor of 2.73.

Overall, for the high interference topologies (Figure 12(a)), we observe an average improvement factor of 2.57. Note that the availability of 11 non-overlapping channels could have *maximally* improved the throughputs by a factor of $11/3 = 3.67$. The sparse topologies observe an improvement by a factor of 1.58. This shows that the actual factor of improvement depends on the local density of interference. Such improvements at the application level metrics are brought about by reduced interference or contention at the link layer. We measure such contention by observing the MAC level collisions encountered by the TCP/UDP flows. Figures 12(b) and 13(b) plot the MAC level collisions observed during the same experiments (normalized by the same constant). These collisions directly reflect the amount of contention observed by the clients. As can be observed from Figures 12(b) and 13(b), the number of collisions reduce in the partial overlap case by an average factor of

2. This reduction shows how partially overlapped channels can reduce such contention thereby improving the application perceived metrics.

Effect of Client Density: Increased client densities mimic the so called *chaotic* network deployments [24] that are beginning to emerge in the form of hotspots and unmanaged networks. Figures 15 and 14 plot the UDP/TCP throughputs respectively over various topologies with increasing density of clients/APs. The relative improvement brought about by using partially overlapped channels increases with density. For example, the factor of improvement using partial overlap grows from 1.25 for 120 clients (Figure 14) to about 1.9 for 360 clients. Similar improvements can be seen for the UDP throughputs (Figure 15). This is because the effect of reduced contention on application level metrics is magnified as the topologies become dense. This observation is also evident from Figure 16 which plots the MAC level collisions observed during the UDP simulations. This figure shows that the number of collisions increase much faster with density when using only non-overlapped channels. For example, with 120 clients the number of collisions using partial overlap were about 30% less than the non-overlapped case. With 360 clients, the number of collisions using partially overlapped channels were lessened by a factor of 1.8.

The results of our enhancements to an existing channel management approach show the wide applicability of this I-factor model which captures partial overlap among channels in a quantitative fashion. These results demonstrate the significant gains that can be achieved by considering the possibilities that arise from performing spatial and spectral re-use in a joint manner.

5.1.1 Increasing Capacity of WLANs

The *Randomized Compaction* algorithm of the previous section addressed the channel assignment problem given a specific deployment of APs in the environment. In this section we briefly examine the joint question of deployment strategy and channel assignment for APs from the viewpoint of increasing total capacity of a WLAN.

We assume that all the APs in a WLAN network (ESS) are connected to a high-speed backbone. Hence capacity achievable for clients are only limited by the capacity of their wireless environments. We also assume that clients are located randomly in a square region. Additionally their location can change over time. Under such assumptions, the APs need to be laid out in a regular pattern. In this section we study the capacity of two possible regular patterns for AP distribution in a large region, and comparing the impact of using non-overlapping and partially overlapped channels.

In particular we consider the 2.4 GHz band with 3 non-overlapping channels (1, 6, and 11) and compare its performance to using 4 partially overlapped channels (1, 4, 7, and 11). Our deployment strategy for this scenario had the following two goals: (i) guarantee that any point on the physical space is covered by at least one AP, and (ii) to reduce interference on the topology, the minimum separation of APs using channels i and j be governed by their transmission ranges as given by d_{ij} and d_{ij} in Equation 5.

We use 3-clique (based on an equilateral triangle) and 4-clique (a square with additional diagonal edges) as two regular patterns to cover the physical space (with an AP on each vertex of these shapes). These two patterns are optimally colorable with 3 and 4 channels.

It turns out that when 4 partially overlapping channels are used with the 4-clique pattern, the possible separation between neighboring APs can be made smaller than the 3-clique pattern. This allows us to deploy more APs in the same physical space using 4 partially-overlapping channels than the 3 non-overlapping channel case. Note that in both cases we do not violate either of the two




Figure 14: TCP Throughput (normalized) versus density.




Figure 15: UDP Throughput (normalized) versus density.





Figure 16: MAC level collisions (normalized) versus density for UDP flows.



(a) UDP throughput

(b) TCP throughput

Figure 17: Greater spatial re-use with overlapped channels in WLANs.

goals outlined earlier. This is because partially-overlapped channels allow better spatial re-use as discussed in Section 3. Consequently it follows that the aggregate throughput achieved by clients using the 4 partially-overlapped deployment strategy is much higher.

We show this for TCP and UDP based experiments (conducted using similar *ns* simulations) in Figure 17. In these experiments, the clients were placed at locations chosen uniformly at random within a square region with the APs laid out as discussed above. The number of clients were varied for different experiments from 500 to 1000. Using just four partially overlapped channels the improvements in throughputs was about 50 to 80%, as can be observed in the plots. The improvements increase as the number of partially-overlapped channels were increased from 4 to 11 within the same spectrum band.

If we used the denser 4-clique pattern and optimally colored them using 3 non-overlapping channels (so that the number of APs deployed was the same as used in the 4 partially overlapped channel scenario), then the aggregate throughput actually falls marginally below the 3-clique, 3 non-overlapping channel scenario. This is due to increase in interference between APs and clients in such a deployment. We also illustrate this performance degradation in Figure 17.

5.2 Channel Assignment in Mesh Networks

Wireless mesh networks are increasingly becoming popular as an Internet access technology for enhancing or extending connectivity to residents and local businesses where wires are a costly solution. These networks are formed by a collection of access points or mesh routers which are devices with potentially more than one radio interface. Such multi-hop wireless networks have a few routers designated as Internet gateways and the goal is to route traffic from clients (associated to other routers) to the ‘best’ gateway. Thus, a careful assignment of channels to these potentially interfering set of mesh routers with the goal of creating low interference paths to the set of gateways is important to realize the full potential of this technology. Thus, there has been much recent interest in addressing

the joint problem of channel assignment and routing in the context of mesh networks [3, 4, 7, 6, 5].

In [6], authors present a mathematical formulation which includes a set of linear programs (LP) to perform joint channel assignment and routing for a mesh network with the objective of improving client throughput. This problem is divided into three subproblems, namely that of channel assignment, flow routing and link flow scheduling. The problem of link flow scheduling deals with coordinating transmissions among an interfering set of edges on each channel. Flow routing addresses the problem of maximizing the per-client throughput carried by the network to the gateways. This also addresses client fairness issues. Channel assignment deals with partitioning the flow assignments obtained by solving the flow routing problem into specific channels so as to minimize the overall interference.

Here, we discuss our enhancements which lets their mathematical formulation take advantage of the partially overlapped channels. Due to space constraints we briefly discuss our enhancements to the linear program which addresses the problem of scheduling interfering links. The basic idea behind these changes is the key invariant that has led to similar enhancements to the other subproblems as well. First, we briefly discuss the concepts behind the modifications to the link flow scheduling problem. Later we compare the results from using our full set of enhancements over an experimental mesh topology which show an average performance improvement factor of 1.6.

Link Flow Scheduling: A mesh network is modeled as a directed graph G over the set of mesh routers as the vertices V . Directed edges comprising the set E indicate wireless links in the mesh backbone. The problem of link scheduling refers to how a set of interfering links can be scheduled to transmit such that the frequency of the transmissions that interfere or collide is reduced. This is achieved by dividing the time into specific slots. The algorithm that specifies which edge transmits on which slot is called a *schedule*. The period of a schedule T is the number of slots after which the slot assignments repeat.

A schedule S specifies if an edge $e \in E$ is allowed to transmit on a channel i in a specific slot $\tau \in \{1, 2, \dots, T\}$ where T is the period of the schedule. Let $X_{e,i,\tau}, e \in E, i \in F(e), \tau \geq 1$ be the indicator variable where $X_{e,i,\tau}$ is 1 if and only if link e is active in slot τ on channel i according to the given schedule S . Here $F(e)$ is the set of common channels among the end-nodes of the edge e . For each edge e , let $ISSET(e)$ denote the set of edges that interfere with it. This set can be computed by considering the radio characteristics such as transmit power and a propagation model much like the set of contending nodes that was computed in Section 4.

A schedule S is said to be interference free if for any time slot no two interfering edges (as given by $ISSET$) are allowed to transmit. Thus, on a given slot τ , if the edge e transmits, none of the other edges $e' \in ISSET(e)$ can be allowed to transmit. This constraint is

necessary for any schedule to be interference free and is stated as follows:

Necessary Condition for Link Flow Scheduling: For any slot τ ($1 \leq \tau \leq T$) any valid interference free edge communication schedule S on a channel i must satisfy

$$X_{e,i,\tau} + \sum_{e' \in I(e)} X_{e',i,\tau} \leq q \quad (6)$$

The average flow carried by an edge according to this schedule S can be computed by summing up this inequality over all $\tau \in \{1 \dots T\}$, that is $f(e(i)) = \sum_{\tau \in \{1 \dots T\}} X_{e,i,\tau}$. Here, $f(e(i))$ is the total amount of traffic carried by the edge e on channel i .

In Equation 6, q denotes an upper bound on the number of concurrent transmissions among edges that belong to $ISET(e)$ given that edge e is not transmitting in a slot τ , that is, $X_{e,i,\tau} = 0$. The value of q depends on the geometric properties of the interference region for an edge e . This is because the concurrent transmissions are among edges which are ‘close’ to e as given by $ISET(e)$. Note that this constraint is replicated for each channel as an artifact of the assumption of independence among channels.

We modify this constraint to contain all indicator variables $X_{e',j,\tau}$ such that edge e' on channel j would interfere with edge e on channel i .

$$X_{e,i,\tau} + \sum_{e' \in I(e)} \sum_{j \in \{1 \dots M\}} POV(e, e', i, j) * X_{e',j,\tau} \leq 1$$

In the above equation, $POV(e, e', i, j)$ is an indicator constant which evaluates to 1 if the edges e on channel i interferes with the edge e' on channel j . This can be computed using physical distance between the nodes belonging to edges e and e' and taking into account the reductions due to the I-factor function $I(i, j)$. This binary notion of interference is fundamental to linear program based formulations. Although the I-factor based model provides a more continuous notion of interference (in terms of the ratio of the received signal strengths), nevertheless we modify it for the purposes of the LP.

Replacing $q = 1$ in Equations 5.2 and 6, we obtain the sufficient condition for link scheduling. This is supported by an approximation algorithm presented in [6] which gives an interference free schedule. By operating on Equation 5.2, this algorithm will be able to take advantage of partially overlapped channels to obtain an interference free schedule.

It might appear that our enhancements to the link flow scheduling constraint could reduce the flow of an edge e by inducing further terms on the left hand side of the Equation 5.2. This misconception is cleared by observing that the number of partially overlapped channels M is typically a factor of about 5 times greater than the number of non-overlapped ones. This gives the linear programs a greater number of channels which can significantly reduce the amount of contention suffered by a link e . For example, assuming the channel structure in the 2.4 and 5 GHz bands, the number of partially overlapped channels can be computed as $M = 5N - 4$ where N is the number of non-overlapping ones. An analytic reasoning presented in Section 4 shows that on average the number of nodes/links contending with a transmitter is reduced by factor of about 3.

We next present an evaluation of our overall enhancements to the mathematical modeling presented in [6]. For each vertex $v \in V$, let $flow(v)$ denote the amount of flow that was generated on that node attributing to clients. For each such vertex, the flow routing problem attempts to route an amount of flow given by $\gamma \cdot flow(v)$. The objective function is to maximize γ . The authors in [6] do this to provide fairness proportional to their traffic demands from the

network. Table 2 shows that normalized throughput achieved by using $N=3,4$ and 5 non-overlapping channels along with the performance of using corresponding $M = 5N - 4$ partially overlapped channels, which can be compared to the scenario of $M = 5N - 4$ non-overlapping ones. This acts as an ‘upper bound’ to the improvements given by using M partially overlapped channels. From the table we can observe that partially overlapped channels improve the per-client throughput by an average factor of 1.6. We further note that as N increases, the gap between the ‘upper bound’ values and those achieved by using the partially overlapped channels reduces. It is also important to realize that, both sets of channels (partially overlapped and the non-overlapped ones) utilized the exact same frequency space. Thus, the performance gains that we achieve do not come at the cost of any additional spectrum.

6. RELATED WORK

A wide range of research literature has addressed the problem of channel assignment in wireless networks. However, the very notion of channels has implied independence among them. Thus, not surprisingly much of the prior work on channel assignment for various wireless technologies including cellular networks, wireless LANs and mesh networks has inherently utilized a set of non-overlapping channels. Below, we summarize such research on channel assignment as applied to various different technologies.

Mesh Networks: Channel assignment in wireless mesh networks is addressed jointly along with the problem of routing and balancing of client load. Proper assignment of channels is crucial to building low interference topologies which are then exploited by the routing algorithms. Work in [6, 7] addresses this joint problem by formulating it as a set of linear programming constraints. In [3, 4], authors build centralized routing and channel assignment algorithms and evaluate their performance over various topologies.

Wireless LANs: With the significant deployment that wireless LANs have seen during the recent years, careful utilization of available channels has become an important problem. In [8], authors model a wireless LAN as a weighted graph with the set of vertices denoting the access points. They present centralized and distributed algorithms for performing channel assignment via graph coloring. Work in [10] also presents graph theoretic algorithms and implementation issues are addressed in [11].

Cellular Networks: Cellular networks are still the most widely used and commercialized form of the wireless communication technologies. Due to the commercial value of spectrum in addition to scarcity, there is a large amount of literature [12, 15, 16, 17, 14, 13] which has addressed the problem of channel assignment for various cellular technology standards such as GSM, CDMA, etc.

Another topic of related interest is capacity of wireless networks which has been primarily analyzed for single-radio [25] and multi-radio wireless environments [5] specifically using non-overlapping channels only. Our work in this paper briefly sketches necessary extensions for partially-overlapped channels and uses the same analysis approach as done in [26].

7. CONCLUSIONS

Channel assignment and management is an important problem that arises in almost any wireless environments, e.g., WLANs, multi-hop networks, and cellular networks. Most prior approaches to channel assignment have restricted solutions to use only non-overlapping channels. We illustrate in this paper that such a restriction is inefficient since it leads to poor utilization of the wireless spectrum. Proper utilization of the spectral resources require use of partially-overlapped channels. To meet this objective, it is necessary to prop-

Number of Channels (N, M)	N Channels No-overlap	M Channels Partial Overlap	M Channels No-overlap
3, 11	0.27	0.42	0.88
4, 16	0.36	0.63	1
5, 21	0.45	0.81	1

Table 2: Aggregate throughput achieved (normalized to best-case scenario within these set of experiments).

erly model interference effects of partially overlapped channels and then apply such models in systematic design, or in some cases, enhancement of channel assignment algorithms.

In this paper we define appropriate models for partially overlapped channels that are applicable for a number of different wireless technologies. Subsequently we use this model to enhance two prior published channel assignment algorithms (one in the context of WLANs and the other in the context of wireless mesh networks) and quantify their benefits.

We believe that our work is the first attempt to systematically exploit partially overlapped channels in wireless environments. Apart from future directions in defining more sophisticated models, there seems to be a large scope for re-visiting a large body of prior literature in channel assignment and management techniques and examining how each such can be enhanced under the partial overlap model.

Finally, in this paper we only explored a restricted set of partial overlap, primarily in the context of *discrete* channels as defined by common technologies such as 802.11 and 802.16. An interesting line of future research can be in exploring how partial overlap can be exploited in designing *continuous channels*. Such a notion can possibly be of significant interest in developing algorithmic channel access mechanisms for software-defined radios.

8. ACKNOWLEDGEMENT

We wish to acknowledge help from Yuan Yuan during the experiments, the constructive feedback on an early version of this work from Thyagarajan Nandagopal and the guidance from our Shepherd, Dina Papagiannaki that brought this paper together into its final form.

The first author wishes to dedicate this paper to his late father, Dr. K.L.P. Mishra.

9. REFERENCES

- [1] Pradeep Kyasanur and Nitin Vaidya, "Routing and interface assignment in multi-channel multi-interface wireless networks," in *IEEE WCNC*, 2004.
- [2] Jungmin So and Nitin Vaidya, "Routing and channel assignment in multi-channel multi-hop wireless networks with single network interface," *Technical Report, University of Illinois at Urbana Champaign*, 2005.
- [3] Ashish Raniwala, Kartik Gopalan, and Tzi cker Chiueh, "Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks," *MC2R*, April 2004.
- [4] Ashish Raniwala and Tzi cker Chiueh, "Architecture and algorithms for an ieeee 802.11-based multi-channel wireless mesh network," in *IEEE Infocom*, 2005.
- [5] Pradeep Kyasanur and Nitin Vaidya, "Capacity of multi-channel wireless networks: Impact of number of channels and interfaces," in *ACM MobiCom*, 2005.
- [6] M. Alicherry, R. Bhatia, and L. Li, "Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks," in *ACM MobiCom*, 2005.
- [7] M. Kodialam and T. Nandagopal, "Characterizing the capacity region in multi-radio, multi-channel wireless mesh networks," in *ACM MobiCom*, 2005.
- [8] A. Mishra, S. Banerjee, and W. Arbaugh, "Weighted coloring based channel assignment for wlangs," *MC2R*, 2005.
- [9] Youngseok Lee, Kyoungae Kim, and Yanghee Choi, "Optimization of ap placement and channel assignment in wireless lans," in *IEEE LCN*, 2002.
- [10] Eduard Garcia Villegas, Rafael Vidal Ferr, and Josep Paradells Aspas, "New algorithm for distributed frequency assignments in ieeee 802.11," in *European Wireless*, 2005.
- [11] Eduard Garcia Villegas, Rafael Vidal Ferr, and Josep Paradells Aspas, "Implementation of a distributed dynamic channel assignment mechanism for ieeee 802.11 networks," in *Proceedings of IEEE PIMRC*, 2005.
- [12] I. Katzela and M. Naghsineh, "Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey," *IEEE Personal Communications*, June 1996.
- [13] K. Naik, D.S.L. Wei, and S. Olariu, "Channel assignment in cellular networks with synchronous base stations," in *Proceedings of the 2nd ACM PE-WASUN*, 2005.
- [14] A. Baiocchi, N. P. Magnani, V. Palestini, and F. Sestini, "Application of dynamic channel allocation strategies to the gsm cellular network," *IEEE JSAC*, vol. 15, Oct. 1997.
- [15] B. Krishnamachari, S. Wicker, R. Bejar, and C. Fernandez, "On the complexity of distributed self-configuration in wireless networks," *Journal of Telecommunication Systems*, 2003.
- [16] F. Mazzini, G. Mateus, and J.M. Smith, "Lagrangian based methods for solving large-scale cellular network design problems," *Journal of Wireless Networks*, 2003.
- [17] W. K. Lai and G. G. Coghill, "Channel assignment through evolutionary optimization," *IEEE Trans. on Vehicular Technology*, vol. 45, no. 1, Feb. 1996.
- [18] A. Mishra, E. Rozner, S. Banerjee, and W. Arbaugh, "Exploiting partially overlapping channels in wireless networks: Turning a peril into an advantage," in *ACM/USENIX Internet Measurement Conference*, 2005.
- [19] T.S. Rappaport, *Wireless Communications: Principle and Practice, Second Edition*, Prentice Hall, 2002.
- [20] Kaixin Xu, Mario Gerla, and Sang Bae, "How effective is ieeee 802.11 rts/cts handshake in ad hoc networks?," in *IEEE GLOBECOM*, November 2002.
- [21] "Enterasys roamabout wireless high-rate pc card," <http://www.enterasys.com>.
- [22] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh, "A client-driven approach for channel management in wireless lans," in *IEEE Infocom*, 2006.
- [23] Jim Geier, "Assigning 802.11b access point channels," *Wi-Fi Planet*, 2004.
- [24] Aditya Akella, Glenn Judd, Srinivasan Seshan, and Peter Steenkiste, "Self management in chaotic wireless deployments," in *ACM MobiCom*, 2005.
- [25] P. Gupta and P.R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, 2000.
- [26] V.S. Anil Kumar, Madhav V. Marathe, S. Parthasarathy, and A. Srinivasan, "Algorithmic aspects of capacity in wireless networks," in *ACM Sigmetrics*, 2005.