

Realizing High Performance Multi-radio 802.11n Wireless Networks

Sriram Lakshmanan*, Jeongkeun Lee[†], Raul Etkin[†], Sung-Ju Lee[†] and Raghupathy Sivakumar*

*School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, Georgia, USA

[†]Hewlett-Packard Laboratories, Palo Alto, CA, USA

Abstract—We explore the design of a high capacity multi-radio wireless network using commercial 802.11n hardware. We first use extensive real-life experiments to evaluate the performance of closely located 802.11n radios. We discover that even when tuned to orthogonal channels, co-located 802.11n radios interfere with each other and achieve significantly less throughput than expected. Our analysis reveals that the throughput degradation is caused by three link-layer effects: (i) triggering of carrier sensing, (ii) out of band collisions and (iii) unintended frequency adaptation. Using physical layer statistics, we observe that these effects are caused by fundamental limitations of co-located radios in achieving signal isolation. We then consider the use of beamforming antennas, shielding and antenna separation distance to achieve better signal isolation and to mitigate these problems. Our work profiles the gains of different physical isolation approaches and provides insights to network designers to realize high-performance wireless networks without requiring synchronization or protocol modifications.

I. INTRODUCTION

Wireless network deployments based on the IEEE 802.11 Wi-Fi standard, continue to increase rapidly. Recently, cellular data-offload techniques that use Wireless LANs to convey traffic generated in cellular networks are contributing to explosive growth in traffic volumes in Wireless LANs [1]. As a consequence, achieving high capacity in wireless LANs is becoming increasingly important. In parallel to these developments, the maturity of WiFi technologies in local area networks, the ease of deployability and use of unlicensed spectrum have encouraged their use in wireless backhaul networks, where the capacity requirements are even more challenging to achieve. For instance, a backhaul network for oil exploration [2] must transport several Gigabits per second of traffic over distances of several kilometers! Thus, realizing very high throughput WiFi networks that can carry several Gigabits per second of data, is becoming critical to address the needs of a variety of application scenarios.

To realize high throughput in WLANs, the IEEE 802.11 (a,b,g and n) standards allow multiple channels for use by different links in the same vicinity. These channels are non-overlapping in the spectrum and can be used simultaneously in space and time without interference. They are called *orthogonal* channels. There are 3 orthogonal channels in the 2.4 GHz band and 12 orthogonal channels in the 5 GHz band. Channel orthogonality improves the aggregate capacity by allowing more than one link to operate in a given region. Traditionally, Access Points (APs) in a WLAN are equipped with a single radio each, and operate on orthogonal channels. Since the throughput of such networks is limited by the capacity of a

single radio, the use of multiple radios on a single AP has been explored [3], [4]. However, it is observed that channel orthogonality is not achieved in practice in a multi-radio link [3], [4], [5] due to out-of-band interference between closely spaced 802.11a/b/g radios. Recently, the 802.11n standard [6], which incorporates several physical and link layer enhancements, was ratified for use toward high-throughput wireless LANs. With this development, the key question is whether co-located 802.11n radios behave similar to legacy 802.11 radios.

Given the importance of understanding radio co-location among 802.11n radios, we study whether co-located 802.11n radios operate effectively when tuned to orthogonal channels. We use a testbed of 802.11n radios and perform extensive experiments in several settings to evaluate the performance of co-located 802.11n radios. We first discover that co-located radios on orthogonal channels do not operate concurrently; even when two co-located 802.11n radios are tuned to channel frequencies separated by more than 500 MHz, they still do not provide the sum of the throughputs of the individual radios. More importantly, these findings are not restricted to a specific hardware or setting but occur across a variety of hardware and test conditions. On analyzing the link layer statistics provided by our hardware, we identify that the performance degradation stems from three main problems at the link layer: (i) out-of-band carrier sense triggering, (ii) out-of-band collisions, and (iii) unintended radar frequency adaptation.

We analyze the behavior of the radios microscopically, using spectrum analysis to understand the underlying causes for performance problems. We observe that the signal emission outside the bandwidth of operation is negligible at large distances (greater than few meters) from the radio but is significant at short distances. These signals cause an increased received signal strength on radios tuned to other channels with three effects: (i) triggering carrier sense to inhibit radio transmission, (ii) causing collisions at a receiver and (iii) triggering frequency adaptation algorithms designed to avoid interference to radars. In particular, the impact of radio co-location on *dynamic frequency adaptation to radars* has not been previously identified, to the best of our knowledge. Identifying the above link layer effects forms our first contribution.

At a deeper level, we identify three key reasons for the above effects: (i) transmit and receive filter imperfections, (ii) generation of image frequencies at the transmitter and (iii) saturation-induced distortion at the receiver. We show that the combination of the above factors causes out-of-band interference to be more pronounced at some frequencies than others.

As a consequence, we make a counter-intuitive observation that *throughput is not a monotonic function of channel spacing at close distances*; using channels spaced farther apart can lead to reduced throughput compared to using closely spaced channels. Our insights call for a re-consideration of channel assignment algorithms in multi-radio systems and illustrate that naïve use of frequency spacing for channel assignment can be harmful. This forms our second contribution.

Finally, we consider three approaches to achieve better isolation and reduce co-location problems: metal shielding, antenna separation and directional antennas. We specifically focus on approaches that work without requiring synchronization among the radios or protocol modifications that current multi-radio aggregation solutions [5] require. Our experiments reveal that each of these approaches improves the aggregate throughput but incurs different trade-offs. Further, using a combination of these approaches, the channel orthogonality can be improved close to expected values. The insights from our experiments are useful for the design of practical multi-radio 802.11n based wireless networks. This forms our third contribution.

The rest of this paper is organized as follows: Section II presents the motivation and the experimental characterization of co-located 802.11n radios. Section III quantifies the impact of three approaches to mitigate the problem. Section IV presents the related work and Section V concludes the paper.

II. CHARACTERIZING THE PERFORMANCE OF CO-LOCATED 802.11N RADIOS

The goal of this work is to realize effective multi-radio operation in 802.11n networks. To study the essential design considerations of co-located 802.11n radios, we perform a variety of experiments with commercial 802.11n equipments. We use different spacings between the radios, different channels and different number of radios. Whenever appropriate, we also compare the performance across different hardware manufacturers to verify whether the observations hold true across different 802.11n equipments.

A. Baseline Performance

1) *Experimental Setup*: We first analyze the baseline performance of two 802.11n links where both the APs and the clients are placed at close distances. The APs are enterprise grade dual-radio 802.11n APs (E-MSM 422) from HP [7]. We use Iperf as the traffic generating application, UDP as the transport protocol and study the downlink throughput from the APs to the clients for runs of 100 seconds. For the clients, we use HP laptops equipped with internal Intel 4965agn 802.11n cards. The laptops run Windows Vista. The experiments are performed in enterprise indoor scenarios where the 5 GHz band is not used by any other wireless node. For increased confidence, before each experiment, multiple scans of the spectrum are performed to ensure that there are no other transmitters on the channels. The APs use frame aggregation with a limit of 25 frames.




Fig. 1. Setup for baseline experimentation

2) *Results*: In this setup, the two APs are separated by 50cm (measured from the center of the APs) and the two clients are separated by 50cm as illustrated in Figure 1. The AP-client distance is 300cm. We vary the channels used on the two links from the following: (i) 36,40 (adjacent channels) (ii) 36,48 (moderately spaced channels) (iii) 36,100 (far spaced channels), (iv) 36,149 (very far spaced channels) and (v) 36,165 (maximally spaced channels). The unicast throughput of Iperf with UDP is measured first with only one link active at a time. The sum of the throughputs when each link works in isolation is noted as the expected throughput. Both links are then activated simultaneously and the throughput is measured at the receivers. The sum throughput of the two links is plotted in Figure 2 along with the ideal expected throughput. We observe that even the state of the art 802.11n radios do not operate effectively when placed together. The throughput degradation can be as high as 61% and occurs even for far-separated channels.




Fig. 2. Baseline results: Channel pairs

We now explore the scaling of the throughput with the number of radios. We set the separation to 25cm and consider the APs in the following arrangements. i.e. ends of a line of 25cm length and the end-points of an equilateral triangle of side 25cm, ends of a square with side 25cm. We set the channels to be sufficiently orthogonal, i.e. channels (36,100) for the two link case and channels (36,100,165) for the three link case. We disable the transmission of 802.11 ACKs to study the effect of the transmitters in isolation. The resulting sum throughput is plotted in Figure 3 for the single radio, two radios and three radios cases. We observe that the throughput degrades by 43% and 71% of the expected for two and three radios, thereby highlighting the poor scalability with increasing number of radios.

We have performed experiments with equipments from other vendors (Apple Airport Extreme-N AP, Linksys WRT320N AP, Linksys USB adapter WUSB600N) and have observed similar performance degradation. Thus, our experiments reveal




Fig. 3. Baseline results: Number of Radios

that *even the state of the art 802.11n radios exhibit throughput degradation when co-located. The degradation increases with number of radios and occurs across equipments from different vendors.*

B. Analysis of Co-located 802.11n Radios

Given the poor performance of co-located 802.11n radios, we design experiments to analyze the underlying reasons for the poor performance. We answer three key questions: (i) Does the signal leakage problem in 802.11 a/g radios [4], [5] also exist in 802.11n radios? If so, is the impact greater or less? (2) What radio-specific factors does this leakage depend on? (3) Are there other reasons besides signal leakage that contribute to performance degradation? We organize the discussion into two parts: the analysis of the link layer statistics such as the number of packet transmitted, retransmissions, etc. and the analysis of the signal spectrum.

1) *Experimental Setup:* We use the topology shown in Figure 1 throughout this section with Iperf over UDP as the traffic source. In addition to the statistics provided by Iperf, we use the web-user interface of the E-MSM 422 APs to obtain statistics such as packet errors, number of retransmissions, etc. Since WiFi chipsets do not provide fine-grained signal power estimates on different parts of the spectrum, we use the WiSpy spectrum analyzer [8] to study the spectral details. WiSpy uses a frequency tunable chipcon radio (cc2500) to identify the spectral power in different frequencies and bandwidths. WiSpy is shipped with a software called Chanalyzer, which provides a visual display of the spectrum. The Chanalyzer allows the resolution (the spectral width over which the power is computed) and the frequency sweep range to be configured. We use 100KHz steps to obtain a good frequency resolution. We perform experiments on different Wi-Fi channels available in the 5 GHz band, also called the Unlicensed National Information Infrastructure (UNII) band.

2) *Link-layer Analysis:* We analyze the link layer statistics to identify the reasons for the throughput degradation. We organize the discussion into two parts (i) operation in the UNII-1 and UNII-3 bands (channels 36 to 48 and channels 149 to 165), (ii) operation in the UNII-2 band (channels 52 to 140).

i. Unintended carrier sensing and collisions in the UNII-1 and UNII-3 bands

We focus on the two parallel 802.11n links scenario in Figure 1. The results are presented in Table I for an experimental run of 15 seconds with the two APs on channels 36 and 40

(The results are similar for other pairs of adjacent channels). As shown in Section II, the throughput on each link is reduced to almost 50% of the throughput of each link in isolation.

TABLE I
PACKET-LEVEL ANALYSIS

Performance Parameters	Individual		Simultaneous	
	Link 1	Link 2	Link 1	Link 2
Throughput (Mbps)	72.2	67.2	36.5	35.2
# of Unicast frames transmitted	87747	87564	48279	48346
Packet delivery ratio at MCS 15	100	99.82	99.51	99.34
# of Retries	10	52	7	29

We observe that the throughput trend is strongly correlated with the statistics of the *transmitted* frames which show a corresponding 50% value of the number of frames transmitted in isolation. We next observe the number of retransmissions in each of the cases. Recall that retransmissions occur due to DATA or ACK packet losses and can stem from either channel impairments (low signal to noise ratio) or collisions. From Table I, we see that the number of retransmitted packets is very small compared to the number of transmitted packets. This indicates that the effect of channel impairments and collisions is very small. This is also verified by observing that the packet delivery ratio at the maximum rate is close to 100% . If there were collisions, the packet delivery ratio at the maximum rate would have been low because packet losses trigger rate adaptation in IEEE 802.11. Hence, carrier sensing across orthogonal channel transmissions causes severe performance degradation among co-located radios. Since carrier sense prevents simultaneous transmission and collisions occur only when more than one simultaneous transmissions exist, the first and perceptible effect is carrier sensing.

In addition, when the direction of one of the flows in our basic setup is reversed, collisions occur between the transmissions on orthogonal channels. The senders of the two flows can not carrier sense each other as complete as they were in very close proximity, which leads to hidden interference. Thus, *both out-of-band carrier sensing and out-of-band collisions contribute to performance degradation, among co-located 802.11n radios.*

ii. Mis-triggered DFS in the UNII-2 band

The UNII-2 band differs from the UNII-1 and UNII-3 bands since military and weather radars operate in the UNII-2 band (channel 52 to 140). WiFi radios are allowed to operate only in radar-free channels in this band. Since additional algorithms need to be implemented, many commercial APs do not support operation in this band. Hence, we use a modified firmware to operate the E-MSM 422 AP on channels 52 to 140. We use the same setup as before but set link 1 on channel 116 and link 2 on channel 165 (i.e., a separation of 240 MHz). We analyze the spectrum using the WiSpy probe placed at a distance of 10cm.¹ We first activate link 1 alone and the spectral plot at the beginning of a 60 second run is shown in the upper part of Figure 4. The plot shows the spectrum centered around the

¹A similar effect is observed up to a distance of 25cm

expected frequency range of channel 116. Next, we activate link 2 and show the spectral plot in the lower part of Figure 4. To our surprise, the spectrum shifts completely from channel 116 to channel 40! This is a non-intuitive result since the configuration setting on link 1 was fixed to channel 116 and not to channel 40.

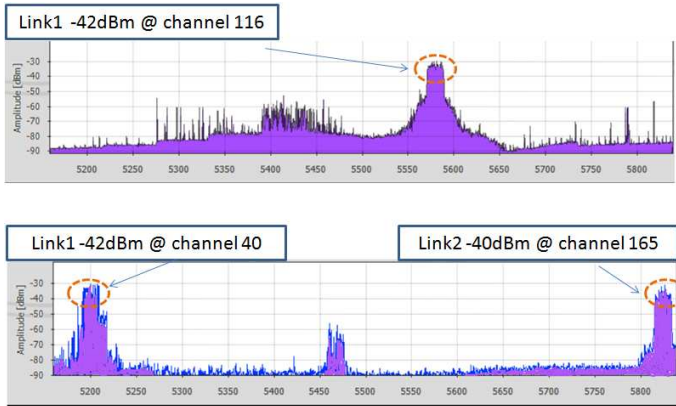


Fig. 4. Spectrum shifts automatically due to DFS operation after link 2 is activated. Spectral power on y-axis vs. WiFi channel on x-axis

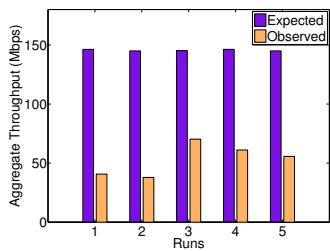


Fig. 5. Aggregate Throughput on DFS channels

This shifting of the operating channel occurs whenever there is significant signal leakage into channels from 52 to 140. When the same experiment is repeated in non-DFS channels out of the UNII-2 band, such a phenomenon does not occur. The unique feature of channels 52 to 140 is that they support Dynamic Frequency Selection (DFS) algorithms to react to radar signals. We verified that there are no radars in the vicinity before every experiment by using a channel scan on a single radio tuned to each of these channels. Hence the underlying reason is that *the radios mistake the signal leakage from co-located radios to be radar signals*. This leads to *unintended triggering of DFS algorithms*, causing the operating channel to change and the performance to be reduced drastically.

The expected and the observed throughputs for 60 second runs is plotted in Figure 5. Our results reveal that the performance degradation can be very severe leading to loss of connectivity and significantly reduced throughput compared to the expected values. The underlying reason is as follows: when a Wi-Fi device detects a radar signal, it must stop its operation in the current channel and migrate to a different channel within 10 seconds [9]. This directly contributes to the communication disruption time. Further, if the new channel is also in the UNII-2 band, the device must wait for at least 60 seconds [9] before resuming its communication to ensure

there is no radar signal in the channel; these wait times add up to the communication disruption time.

3) *Physical Layer Analysis*: Our experiments reveal that there are three dominant effects that contribute to the lack of concurrent transmissions among co-located radios: (i) the transmit signal leakage into adjacent bands, (ii) the generation of spurious signal components in the receiver and (iii) insufficient image frequency reduction at the transmitter and/or receiver. The transmit signal leakage has been identified in [4], [5]. However, to the best of our knowledge, the causes of spurious signals or the image frequency have not been identified in 802.11n radios. We present the gist of our analysis here and refer the interested reader to [10] for the details.

In practice, the transmission power from each radio must be restricted to a given bandwidth and center frequency. The power emitted in other frequencies of the band is called “Out-of-Band (OOB) emission [4]” and must be minimized. The FCC provides regulations on the power permissible in different parts of the spectrum when operating on a Wi-Fi channel in the United States. As per the specifications, for a class A digital device (which includes WiFi radios), the spectrum measurement is to be conducted at a distance of 3m and the OOB power must be restricted to different power limits depending on the frequency separation from the center frequency of operation. However, it is not possible to perfectly restrict the signals to a given bandwidth and achieve zero OOB. Commercial WiFi devices are designed to meet the regulations at a distance of 3m from the transmitter. However, practical multi-radio systems would have antennas separated by a distance less than 3m, where the regulations do not apply. Thus, commercial WiFi radios that follow the emission guidelines for single radio operation cause emission related problems in a multi-radio setting when placed close to one another.

We first present in Figure 6, a plot of the spectrum from the 802.11n radio of the E-MSM 422 AP observed using the Wi-Spy spectrum analyzer. We see that compared to the ideal case where the spectrum is clearly defined and limited to the bandwidth of interest, there are three distinguishing characteristics:

i. *Signal leakage into adjacent channels*

The signal leaks into channels adjacent to the channel that the radio is tuned to. This phenomenon has been observed in 802.11a/b radios [5], [4]. The reason for this leakage is known to be the imperfection of the transmit filter [4], since it is impossible to completely eliminate power in adjacent bands and obtain a sharp filter response. Thus, we observe the leakage phenomena even in ‘the state of the art’ 802.11n radios. Chipset vendors define ‘adjacent channel rejection’ as a ratio of the emitted power in the tuned channel to that in the adjacent channel and claim that their chipsets implement the rejection filter in compliance with the OOB regulation. From our experiments, we observe that the adjacent channel rejection does not vary with the channel tuned to but is more sensitive to the distance from the transmitting antenna (particularly at close distances).




Fig. 6. Leakage, image frequency and non-linear distortion observed at a distance of 25 cm

TABLE II
LOCATION OF SIGNIFICANT IMAGE COMPONENTS (I.E. IMAGE POWER GREATER THAN 10 dB ABOVE NOISE FLOOR AT 25CM)

Channel	Center Frequency (MHz)	Image Frequency (MHz)	Frequency Separation (MHz)
36	5180	5480	300
40	5200	5470	270
44	5220	5450	230
48	5240	5630	390
60	5300	5560	260
64	5320	5550	230
100	5500	5570	70
136	5660	5200	460
140	5700	5180	520
153	5760	5710	50
157	5785	5690	95
165	5825	5210	615

ii. Image frequency in non-adjacent channels

We set one radio to transmit on channel 36 and observe the spectrum at the WiSpy receiver placed at a distance of 25cm. Figure 6 shows the resulting spectrum. We see that the interference keeps reducing as the frequency separation increases within the same band but surprisingly, there is an increased interference perceived around channel 100. Interestingly, the spectral width of this interfering signal is the same as the main transmission but its power levels are lower. Similarly, when the transmitted signal is at channel 153, this component occurs around channel 48. When the operating channel is shifted to 165, the interference occurs around channels 44 and 52. Thus, for each transmitting channel, there is a unique and non-adjacent channel around which a strong component is observed. We deem this component as *image frequency*.

We tabulate the location of the significant image component for different channels on the transmitting radio in Table II. We define a component as significant only if the image power measured at 25cm is greater than 10 dB above noise floor. The absence of a channel in the Table means that the image component for that channel is not significant. The frequency separation between the center and the image frequency, although relatively similar for the frequencies in the same band (i.e. channels 36 to 48), is not constant for all channels. We believe that this is due to differences in the method used to generate intermediate frequencies and carrier signals depending on the channel. We have observed that the image component scales with distance in the same manner as the carrier frequency. Further, with bonded channels, the bandwidth of the image component is also doubled as expected.

iii. Non-linear distortion at the receiver

When the WiSpy antenna is placed at very close distances up to 5cm from the transmitting antenna, significant interference is produced at multiple locations in the spectrum. This effect is illustrated in Figure 6. We reason that these occur due to the non-linear distortion from the saturation of the amplifiers at the receiver. The amplifiers operate in a linear manner only when the input signal power is restricted between the minimum power (P_{min}) and the maximum power (P_{max}). At close distances, these restrictions are not met, leading to very high power signals fed into the radio. This shifts the operating points of the amplifiers and makes the operation non-linear [4]. Thus, we observe non-linear distortion effects when 802.11n radios are placed at very close distances. However, the effects are eliminated very fast as the separation between radios is increased.

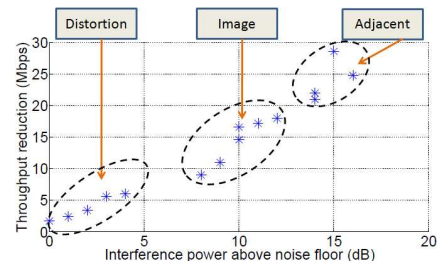


Fig. 7. Correlation between spurious power and throughput

Correlation to throughput: We set the two links to operate on different pairs of channels that include adjacent channels, channels with image components and channels with non-linear distortion components. We compute the throughput degradation compared to the ideal throughput in each case. We compute the correlation coefficient between the interference power obtained from the spectral analysis and the throughput reduction compared to the ideal. This plot in Figure 7 clearly indicates a good correlation between the interference sources we identified and the actual throughput achieved. In particular, the values close to the left part of the figure occur due to non-linear distortion, whereas the values on the right are due to adjacent channel leakage. The values in the middle part are contributed by image frequency components.

C. 802.11n vs. 802.11n technologies:

The MSM 422 APs allow setting the radio in the 802.11a mode or the 802.11n mode. We perform experiments under

TABLE III
POWER REJECTION RATIO OF 802.11A AND 802.11N AT 5 GHz

Power Rejection	Minimum (dB)	Typical (dB)
OFDM, 6 Mbps	16	36
OFDM, 54Mbps	-1	21
HT20, MCS 0	16	33
HT20, MCS 7	-2	7
HT20, MCS 8	16	34
HT20, MCS 15	-2	5
HT40, MCS 0	16	42
HT40, MCS 7	-2	26
HT40, MCS 8	16	42
HT40, MCS 15	-2	24




Fig. 8. Aggregate Throughput: 802.11a vs. 802.11n

identical settings in an indoor environment as in Figure 2. We observe the resulting throughput normalized to the maximum PHY rate for both 802.11a and 802.11n in five consecutive runs and plot the results in Figure 8. Since the use of multiple antennas at the transmitter is the main difference between 802.11a and 802.11n as far as signal power is concerned, we attribute this phenomenon to the use of multiple antennas in 802.11n and the combination of the signals from multiple antennas at the receiver. This also corroborates with the datasheet that specifies the adjacent channel power rejection to be higher for the 802.11n Modulation and Coding set compared to the 802.11a Modulation and Coding set. The datasheet [11] specifies the power rejection values that we have reproduced in Table III. Notably, for the same modulation and coding, the power rejection for 802.11n two-stream mode (HT20, MCS 15) is 5 dB which is significantly lower than the power rejection for the 802.11a case of 21 dB (OFDM, 54 Mbps).


III. SOLUTION STRATEGIES FOR EFFECTIVE OPERATION OF CO-LOCATED 802.11N RADIOS

As mentioned in Section II, the out of band leakage power emission triggers carrier sense on radios tuned to orthogonal channels and prevents the radios from transmitting concurrently. This problem can be solved using different approaches such as (a) better filter design to reduce out of band emission or (b) protocol modifications such as adjusting carrier sense thresholds. However, redesigning the filters within the chipset to achieve larger reduction at short distances involves re-architecting the radio which is difficult. While modifying the carrier sense threshold might help in some cases [5], it affects the protocol correctness. For instance, when a higher carrier sense threshold is used to overcome false inhibition by adjacent radios, the power from a legitimate user (who is far-

away) may be masked by the new threshold. Consequently, the fairness and operational correctness of CSMA/CA will likely be affected. Hence we focus on methods to achieve better isolation among the antennas without requiring protocol or radio chipset modifications. We consider two main approaches: (i) antenna physical isolation and (ii) antenna directionality.

A. Physical Isolation

In this approach, we explore how better physical isolation between antennas of radios can be achieved. We consider two dimensions namely, increasing the separation distance between the antennas and the use of shielding materials that can attenuate the signal power leaked to adjacent radios.



(a) At 25cm

(b) At 100cm

Fig. 9. Signal leakage vs. distance

1) *Separation Distance*: We first consider the effect of different separation distances between the antennas of two radios using the same setup in Figure 1. We set one radio to channel 48 transmitting MAC layer DATA packets (without ACK) at the maximum rate of 130 Mbps (using the Multicast option in IEEE 802.11n). We set the AP radios to operate using their internal omnidirectional antennas. We measure the spectrum at both 25cm and 100cm separations using the WiSpy spectrum analyzer and the plots are shown in Figure 9. The figure clearly highlights that the power in the adjacent band changes from -76 dBm to -94 dBm i.e., by about 18 dB (a factor of 64) when the separation distance is increased from 25cm to 100cm, indicating an approximately inverse cube dependence with distance (at the short distances). Interestingly, the power in the desired channel changes from -38 dBm to -52 dBm i.e., by 14 dB. The effect of the distance is larger for the adjacent channel power than the tuned channel indicating clearly that there are some non-linear effects that occur at close distances. Thus, increasing antenna separation reduces leakage power. Additionally, we note that the tested radio design complies with FCC channel orthogonality requirement enforced for distance at and beyond 3 meters. The 802.11n transmit spectral mask requires a power reduction of 28 dBm at 20 Mhz spacing, which is always observed in our measurements.

Next, we fix one radio to operate on channel 36 and set the second radio on channels 40, 44, 48 and 149 respectively in subsequent experiments. For each channel pair, we test 25cm, 50 cm and 100cm separations between the antennas of the two radios. We generate multicast traffic from the two radios and measure the aggregate throughput as the number of bits per second successfully transmitted by the two radios to observe

the level of carrier sense between the radios. Figure 10 shows the measured aggregate throughput against the channel separation for different distances. While the expected throughput of a single radio without carrier sense is around 50 Mbps, the two links together should provide twice this rate if they are fully concurrent. We observe that adjacent channels are not fully orthogonal even at 100cm. But channels separated by 60 MHz (channel separation of three) achieve good orthogonality at 50cm and beyond. At very close distances, only the channels in the extreme ends of the 5 GHz Wi-Fi spectrum (separated by 565 MHz) are orthogonal. Thus, *distance separation provides complete concurrency only after 100cm*. However, achieving such separation in practice might increase the overall system volume and the deployment complexity.

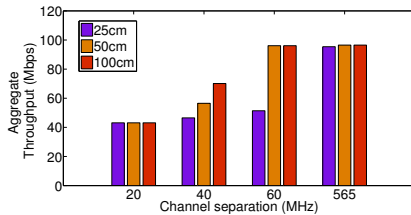


Fig. 10. Throughput for different separation distances

2) *Metal Shielding*: Shielding is a process of reducing the electromagnetic fields using barriers made of conductive and/or magnetic materials [12]. In general, the amount of signal isolation depends on the material, the thickness, the size of the shielded volume, the frequency of the radiation and the presence of apertures. High-frequency radiations are reduced to 37% of their incident power at a depth of the barrier called the skin depth. The skin depth at 5 GHz for all metal shields is in the order of 10 μm . The absorption loss is known to be 99% at five skin depths or a depth of 50 μm [12]. In practice, solid sheet metal shields of a depth greater than several μm are used to prevent signal leakage from radio equipment.

For our experiments, we test solid sheet metal shields between the antennas of co-located radios. The setup is the same as in the previous sections with two 802.11n radios separated by 25cm. We ensure that the thickness of the barrier is always sufficient to ensure complete absorption. However, the radiative losses depend on the dimensions. We use three different sheet metal shields of the following dimensions: 18.5cm \times 18.5cm \times 0.4cm, 18.5cm \times 18.5cm \times 0.8cm, and 18.5cm \times 37cm \times 0.4cm. Essentially, we have a basic shield, one with twice the height and another with twice the thickness (We note that all dimensions are much greater than the wavelength and the skin depth). We study the throughput under two channel settings in each of these cases. Radio 1 is always set to channel 36. Radio 2 is set to channel 40 or channel 48 to study the case of adjacent and well separated channels.

The throughput is plotted in Figure 11. We observe that for channels 36 and 48, the throughput improves with shielding and the height of the shield is a more important factor than the thickness of the shield. We conjecture that the smaller thickness (0.4cm) is sufficient to reflect signal away from

one radio but signals can bend around the shield if the height and width are not sufficient. However, for the adjacent channel case of channels 36 and 40, the improvements are minimal. This is because the interference power is very strong and the shield is not sufficient. This effect is also clearly illustrated by the WiSpy spectrum plot shown in Figure 12. Thus, *the use of metal shields helps reduce the interference power across co-located radios*. While more sophisticated shielding structures can be constructed, we believe that our metal sheet experiments provide a first order characterization of the effectiveness of shielding.

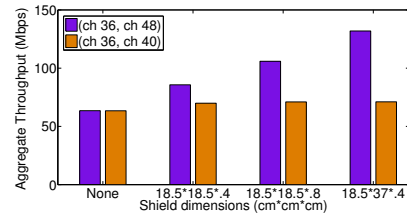



Fig. 11. Throughput vs. shield dimensions



(a) Without shield

(b) With 18.5*18.5*0.8 cm^3 shield

Fig. 12. Effect of metal shielding

B. Antenna Sophistication

Next, we consider the use of directional antennas with controllable orientation such that the directional gains help reduce leakage power between radios. We use two antenna prototypes to understand the effects of directional antennas. Our first test equipment is the Phocus Array from Fidelity-Comtech [13], which is an eight-element electronically steerable antenna array fitted with an 802.11g radio. The Phocus array is shipped with a set of directional beams that can be electronically controlled using commands. Our second equipment is the integrated directional MIMO panel antenna from LairdTech (S245112PT), consisting of three directional antennas together in a single antenna structure [14]. This antenna is specially designed for 802.11n radios and has a 13.5 dBi gain and 20 degree beamwidth at 5.5 GHz. Two of the three antennas use the same vertical polarization and the third antenna uses a horizontal polarization.

Our experiments are intended to answer the following questions: (1) Does directionality occur even at very short distances? (2) Do directional antennas reduce signal power leakage among closely located radios? (3) Do directional antennas used in conjunction with 802.11n radios provide improved concurrency?

1) *Spectrum Results*: Conventional knowledge on directional antennas suggests that directionality holds only in the far-field region i.e. at distances greater than 2λ where λ is the wavelength of the radiation. The far-field occurs at a distance of 12cm from the transmitting antenna for 5 GHz. However, antennas of co-located radios are typically separated by closer distances, where directionality is not guaranteed. Hence, we first observe the spectrum at a fixed distance using the Phocus array antenna. We first place the WiSpy probe at a fixed distance of 25cm from the Phocus Array. We then vary the orientation of the Phocus Array consecutively across 360 degrees. We present the spectrum at two specific orientations of zero degrees (towards the WiSpy receiver) and 180 degrees (away from the WiSpy receiver) in Figure 13. It is interesting that the reduction in the power at the center frequency is only around 15 dB when the antenna orientation is changed by 180 degrees. However, the adjacent channel rejection changes by about 47 dB for the same scenario. The gain due to directionality on adjacent channel power is greater than the main channel gain due to the non-linearities of the receiver. Thus, *directional antennas are very effective in mitigating out-of-band signal leakage*. We believe non-linear effects at the receiver are not limited to any specific radio. It was reported that receiver blocking and inter-modulation products affect a WiFi receiver when placed close to a WiMAX transmitter [4].

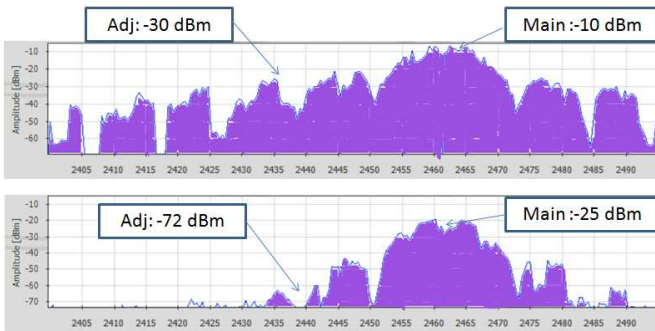


Fig. 13. Spectrum vs. antenna orientation (0 degree and 180 degree)

2) *Throughput Results*: We perform experiments with the MIMO directional antenna. As seen from the radiation patterns, there is an isolation of more than 20 dB between the direction of the main-lobe (zero degrees) and the backlobe (180 degrees). We perform experiments by fitting the antennas to two 802.11n APs placed in an outdoor environment at a height of 150cm from the ground. The antennas are separated by a distance of 50cm from each other and the two radios are tuned to channels 36 and 40 respectively. We run Iperf UDP flows and vary the orientation of the antenna attached to the second AP. The expected and the observed aggregate throughputs are plotted in Figure 14. The figure shows that even with adjacent channels, directional antenna gains can be used to improve the concurrency. The improvement in throughput depends on the relative orientation.




Fig. 14. 802.11n Directional antennas

C. Combination of Approaches

We study how a combination of the above approaches of using directional antennas, spacing and metal shield performs. We consider nine orthogonal channels 36, 40, 44, 48, 149, 153, 161, 165. We use the same shield dimensions used before and place the shield between the antennas. Similarly, we place the directional antennas at around 180 degrees to each other. We test the throughput of two channels at a time and identify the subset of channels that do not mutually interfere. This number is presented in Table IV for different strategies. We observe that the baseline case yields only two times the throughput of the single channel case at an antenna spacing of 25cm. But when shielding and directional antenna are used in isolation, the throughput jumps to four times and five times respectively. This corresponds to the ability to use non-adjacent channels across the spectrum. Similarly, when both shielding and directional antennas are used, the throughput does not increase further since the gains are not sufficient to allow adjacent channel operation. However, when the spacing is increased to 50cm and both directional antenna and shielding are used, the throughput reaches close to the ideal throughput of nine channels.

TABLE IV
COMBINATION OF APPROACHES

Spacing	Antenna	Shielding	Ratio of maximum throughput to single channel throughput
25cm	Omni	No	2
		Yes	4
	Dir	No	5
		Yes	5
50cm	Omni	No	4
		Yes	5
	Dir	No	8
		Yes	9

D. Implications

Our experiments reveal the following insights about achieving signal isolation in multi-radio systems.

- Concurrency of co-located radios with omni-directional antennas improves when the spacing between antennas is greater than 100cm. For smaller distances, the leakage power is significant and causes interference.
- The use of metal shields provides improvement in concurrency when the channel separation is greater than 60MHz and provides minimal benefits for adjacent channels. The benefits depend on the height and the width of the shield and less on the thickness of the shield.
- Directional antennas provide directional gains even at short distances and reduce the effect of leakage powers

significantly. The magnitude of the reduction depends on the directional gain and relative beam orientations.

- In practice, the dimensions of the multi-radio system must be minimized to produce compact APs. Given volume or weight constraints, using directional antennas is more preferable than using shielding, which in turn is more preferable than just increasing antenna separation. Similarly when the weight of the system is to be minimized, directional antennas and spacing must be used in preference to shielding, since solid metal shields increase the overall system weight.

IV. RELATED WORK

The out of band emission of signal or signal leakage problem has been studied in different contexts. In [15], the authors study a three node, two-hop testbed, with the common node having two 802.11 radios. They study only the two-hop behavior of the network and conclude that if a single node contains two wireless cards alone, these cards will not be able to receive or transmit traffic at the same time. In [3], the authors identify the effect of interference across two wireless interfaces on the same node, each using a different channel. Similarly, in [16], [17], [18], the authors argue that it is not possible to simultaneously use two radios on the same node. In [4], the authors highlight the problems of using co-located WiFi, Bluetooth and WiMAX radios. They show that beamforming can potentially help reduce the problem but do not present any analysis or solution. Similarly [19] studies co-located 802.11a radios with large antenna separation distance greater than 1m. However, the focus of our work is on shorter separation distances in practical multi-radio systems.

Perhaps, the most relevant work is Glia [5], where the authors present a software solution for aggregated use of 802.11 radios between two nodes. Glia does not study 802.11n specific issues among co-located radios. Glia is developed for a point-to-point link, requires synchronization across radios, removal of 802.11 ACKs and protocol modifications. These requirements are difficult to achieve in multi-point to point and multi-hop network environments. Our work provides a different approach to enabling concurrent WiFi radios that does not require synchronization across radios or 802.11 protocol modifications.

Along the same vein, an integrated multi-radio product has been developed in the industry. Xirrus [20] uses 16 radios together, with a directional antenna attached to each radio to communicate with independent clients. Being a commercial product the interaction among the radios and the underlying insights are not yet known to public.

In summary, none of the above works study the problem of aggregation in co-located 802.11n radios. To the best of our knowledge ours is the first work to perform a detailed experimental characterization of the problems, the impact of directional antennas and the effect of shielding with co-located 802.11n radios.

V. CONCLUSIONS

We present a detailed characterization of closely spaced 802.11n multi-radio links. We identify that: (1) Out of band signal leakage triggers unintended carrier sensing, collisions and (radar) frequency adaptation in 802.11n radios, causing throughput degradation. (2) At the radio level, filter imperfections, image frequencies and non-linear distortion cause varying interference across the spectrum. Consequently, throughput is not a monotonic function of the frequency separation between co-located radios (3) Metal shielding, antenna separation and antenna directionality can be used to mitigate these effects. There is a trade-off between the degree of effectiveness and the practicality of the system, which is different for each of these approaches.

REFERENCES

- [1] A. Balasubramanian, R. Mahajan, and A. Venkataramani, "Augmenting mobile 3G using WiFi: Measurement, system design, and implementation," in *ACM MOBISYS*, 2010.
- [2] S. Savazzi and U. Spagnolin, "Synchronous ultra-wide band wireless sensors networks for oil and gas exploration," in *International Symposium on Computer and Communications*, July 2009.
- [3] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for IEEE 802.11 wireless networks," in *IEEE BROADNETS*, Sep. 2004.
- [4] J. Zhu, A. Waltho, X. Yang, and X. Guo, "Multi-Radio Coexistence: Challenges and opportunities," in *IEEE ICCCN*, 2007.
- [5] S. Kakumanu and R. Sivakumar, "Glia: A practical solution for effective high datarate WiFi-arrays," in *ACM MOBICOM*, Sep. 2009.
- [6] IEEE standard for Wireless Local Area Networks: 802.11n. [Online]. Available: <http://www.ieee802.org/11>
- [7] HP dual-band dual radio access points. [Online]. Available: http://h17007.www1.hp.com/us/en/products/wireless/HP_E-802_11n_Dual_Radio_Access_Point_Series/index.aspx
- [8] Wi-Spy USB spectrum analyzer. [Online]. Available: <http://www.metageek.net/products/wispy>
- [9] Cisco report on dynamic frequency selection. [Online]. Available: <http://www.cisco.com/en/US/docs/routers/access/3200/software/wireless/RadioChannelDFS.pdf>
- [10] S. Lakshmanan, J. Lee, R. Etkin, S.-J. Lee, and R. Sivakumar. (2010) Enabling concurrent operation of co-located 802.11n radios:gnan technical report. [Online]. Available: <http://www.ece.gatech.edu/research/GNAN/archive/tr-mrd.pdf>
- [11] Atheros AR 5008 chipset datasheet. [Online]. Available: <http://www.atherosxspan.com/modules/articles/files/AR9001AP-3NX2Bulletin.pdf>
- [12] C. R. Paul, *Introduction to Electromagnetic Compatibility (Second Edition)*. Hoboken, NJ, USA: John Wiley and Sons, 2006.
- [13] Phocus WiFi antenna array. [Online]. Available: <http://www.fidelity-comtech.com>
- [14] Directional 802.11n MIMO antenna - S245112PT. [Online]. Available: <http://www.lairdtech.com/WorkArea/DownloadAsset.aspx?id=3999>
- [15] Joshua Robinson et al., "Experimenting with a multi-radio mesh networking testbed," in *WINMEE*, 2005.
- [16] S. Liese, D. Wu, and P. Mohapatra, "Experimental characterization of an 802.11b wireless mesh network," in *IWCMC*, 2006.
- [17] Yunxin Liu et al., "An experimental study on multi-channel multi-radio multi-hop wireless networks," in *IEEE GlobeCom*, 2005.
- [18] C.-M. Cheng et al., "Adjacent channel interference in dual-radio 802.11a nodes and its impact on multi-hop networking," in *IEEE GlobeCom*, 2006.
- [19] Vangelis Angelakis et al., "The effect of using directional antennas on adjacent channel interference in 802.11a: Modeling and experience with an outdoor testbed," in *WINMEE*, 2008.
- [20] Xirrus 16 radio WiFi-array. [Online]. Available: <http://www.xirrus.com/products/arrays-80211abg+n.php>