

Behind-the-Scenes of IEEE 802.11a based Multi-Radio Mesh Networks: A Measurement driven Evaluation of Inter-Channel Interference



Sebastian Robitzsch, John Fitzpatrick, Seán Murphy and Liam Murphy

University College Dublin

Performance Engineering Laboratory

Dublin, Ireland

sebastian.robitzsch@ucdconnect.ie, (john.fitzpatrick, sean.murphy, liam.murphy)@ucd.ie

Abstract—To successfully develop IEEE 802.11a based wireless mesh network solutions that can achieve the reliability and capacities required to offer high quality triple play services the use of multiple radios in each mesh node is essential. Unfortunately, the co-location of multiple antennas in a single device leads to a number of interference problems. In this paper the impact of non-overlapping channel interference in IEEE 802.11a based multi-radio nodes is investigated. A detailed explanation of the performance decreases and their relation to radio settings is presented. The primary contribution of this paper is the discovery of a channel interference effect which is present over the entire 802.11a frequency space. This interference appears if two radios are located less than 50 cm from each other and both are attempting to operate as usual. The results were obtained by conducting experiments in a well planned testbed to produce reliable and reproducible results. The presented results incorporate multiple parameters including transmission power, modulation coding scheme, channel separation and physical layer effects such as adjacent channel interference, carrier sensing, retransmissions and packet distortion.

Keywords-Wireless LAN; Measurement; 802.11a; Multi-Radio Wireless Mesh Network

I. INTRODUCTION

In recent years Wireless Mesh Networks (WMNs) have become increasingly popular. This is primarily due to the high level of penetration achieved by Wireless Local Area Network (WLAN) as an access technology for end user devices and the widespread availability of low cost Wireless Fidelity (WiFi) hardware. Another important factor is that WiFi operates in the unlicensed Industrial, Scientific and Medical (ISM) radio spectrum, therefore, WMNs based on this technology can be deployed without requiring the purchase of expensive spectrum licenses. Moreover, the ability of WMNs to provide last mile communication infrastructure as a number of use cases such as campus, festival or conference deployments. This flexibility has driven a number of interest groups to investigating WMN deploying issues.

As mentioned in the previous published work, [1], interference between two interfaces on the same WMN can appear if both antennas are located relatively close to each other. Interference can still occur even if the well-known Adjacent Channel Interference (ACI) requirements,

as described by Angelakis et al. [2], Nachtigall et al. [3] and Cheng et al. [4], are fulfilled by setting the radios to a channel separation larger than one.

The focus in the initial work [1] was not to provide an exact set of data to derive necessary radio and interface parameters such as Transmission Power (TxPower), Modulation Coding Scheme (MCS) and channel separation, for a network with reliable and stable links. It was rather envisaged to go one step further than [2] and [4] to provide a first survey that the channel interference phenomenon is related to the three parameters TxPower, MCS and channel separation. Even in comparison to [3], the previous work [1] showed that focusing solely on antenna separation (distance) and channel separation does not give enough detail to fully understand the radio environment. The work presented in this paper expands on this previously published work with a more detailed analysis of the evaluation of obtained Received Signal Strength (RSS) and Noise Floor (NF) values, new measurements to address the question whether data frames or their Acknowledgements (ACKs) are affected by ACI or Inter-Channel Interference (ICI) and with a more detailed evaluation and discussion of the obtained measurement results starting to quantify the different investigated interferences.

As will be described in Section III-A, a testbed environment was established including scripts to conduct a number of experiments automatically. These experiments comprised every relation of TxPower, MCS and channel separation (as long as some interference could be investigated) for antenna separations up to 60 cm in 10 cm steps.

The remainder of this paper is structured as follows. The next Section II provides an overview of related work from other groups, this is followed by Section III which gives the detailed description of the experimental setup. Section IV describes the ACI effect and its use to evaluate the testbed hardware. Section V presents experimental results which investigates the relationship of RSS and Signal to Noise Ratio (SNR), the ICI as well as the relationship of Re-Transmission Rates (RTRs), Carrier Sense (CS) and application layer losses. The paper is then concluded in Section VI which summarises the paper.

II. RELATED WORK

In recent years the ACI effect has been well investigated by different independent researchers. However, no work appears to have been done for channel separations of more than one; this paper aims at addressing this shortcoming.

Angelakis et al. [2] qualified the effect of ACI in terms of throughput measurements within a single node equipped with multiple interfaces. In this paper the authors developed a mathematical model which showed that neighbouring 802.11a channels have a spectral overlap which produces a significant level of interference that can lead to lossy and unstable links in dual-radio equipped nodes. However, their experiments were conducted under laboratory conditions using attenuators and couplers to demonstrate the ACI effect. Due to this experimental environment the ACI effect on two transmitters with a channel separation of two could not be shown. The reason for this was based on the strength of the attenuators which they used. The level of these attenuators was too high and the transmit signal level was below the sensitivity threshold of a common WiFi card. Therefore, based on the obtained results it was not possible to conclude what level of channel separation is required to provide stable and reliable links. Even in their subsequent papers [5] and [6] this issue has not been investigated further.

Mishra et al. [7] assumed that the overlap between neighbouring channels in 802.11a is so low that it can be ignored for practical purposes. The authors conducted experiments for an 802.11b link and a channel separation of three which represents two orthogonal channels as 802.11a does. They did not observe any interference by measuring the throughput for a distance of 10m. Additionally, Mishra et al. defined an appropriate model for partially overlapping channels which calculates the level of interference caused by all non-orthogonal channels, as they appear in 802.11b/g. However, just measuring the throughput is not sufficient to conclude that there is no interference in cases where both radios are located closer than 10m to each other. The assumption of Mishra et al. that the only interference is due to small partially overlapping channels of 802.11a does not hold for distances below 30cm, as will be shown later in this paper.

This can easily be verified by previously published work, such as [2] [3] [4], which investigated that ACI issues must be taken into account when conducting multi-radio measurements with Institute of Electrical and Electronics Engineers (IEEE) 802.11a hardware. One of the most important pieces of work in relation to this issue are the results of Nachtigall et al. [3]. They demonstrated that the number of available non-interfering channels depends on both the antenna separation and the Physical Layer (PHY) modulation for a dual-radio scenario. They also stated that under their conditions only one channel can be used at the same time, which is not necessarily true.

In order to verify the testbed and provide reproducible results, Burchfield et al. [8] proposed three necessary steps based on an extended set of experiments in a real environment as well as under laboratory conditions using coax cables. In summary they recommend to check first for external networks using the same Medium Access Control (MAC) protocol. Secondly, the medium should be checked also for other interference which cannot be recognized by the chipset but still senses the medium, e.g., a microwave in case of deploying a 802.11b/g testbed. Finally, Burchfield et al. recommends the use a coaxial setup to verify the system's capabilities. However, a statistical evaluation after conducting the experiment seems to be another much more reliable step in terms of proving the mean values than confirm the experimental obtained graphs by replacing the wireless links with coax cable. To verify the claim by Burchfield et al. similar coax experiments were conducted. The results showed that even if both antenna outputs were connected via coax cable with each other it was possible to see regular beacons sent from another Access Point (AP) which was located approximately 30m away. Hence, the proposed verification by Burchfield et al. using coax cables does not seem to be 100% accurate.

The most promising work was done by Nachtigall et al. [3] who investigated the interference among the interfaces of a multi-radio node and found that the radios located close to each other interfere with each other significantly. They even stated that under their experimental conditions, an antenna separation of 15cm, only one channel within the entire 5.2GHz band can be used at the same time. They concluded their work with the statement that the number of available orthogonal channels depends on the antenna separation and MCS.

III. EXPERIMENTAL SETUP

This section provides a detailed configuration description of the experimental environment; it also provides an overview of how the logged data was processed. It also shows how the measurement process was designed in terms of duration and configuration using prior statistical evaluation, for example the method of independent replications. The use of the Fresnel formula to design the links properly is also described.

A. Testbed Configuration

The most important factor when planning the experimental testbed was to provide reliable and reproducible results with the least number of external dependencies as possible. To achieve this goal the testbed was deployed as shown in Figure 1.

All three machines (Node A, B and C) were x86 Intel based desktop machines running Ubuntu 32 bit server edition and kernel version 2.6.28. Node A was equipped with two wireless interfaces and both Node B and Node C each had




Figure 1. Experimental Setup

one; each wireless interface used RouterBoard R52 wireless 802.11a/b/g mini-Peripheral Component Interconnect (PCI) cards with the Atheros chipset AR5414. The R52 card is WLAN certified by the WiFi alliance which ensures the chipsets has been designed according to the standard and produces Orthogonal Frequency-Division Multiplexing (OFDM) signals with Frame Error Rates (FERs) defined by IEEE. This mainly gives the safety that any other WiFi certified WLAN card will behave like the R52 does and will give similar experimental results. To connect the the mini-PCI cards to the PCI slots in each machine, a RouterBoard 14 mini-PCI to PCI adapter was used. The connection between the cards and the 5.5 dBi omnidirectional antennas used 20cm length RG-178 U.fl to N pigtailed. Since the experiments being performed required that the antenna separation could be varied from 10cm up to 60cm, the connections from the interface cards to the antennas for interfaces A_1 and A_2 of Node A were extended using two 2m low-loss CLF-400 coaxial cables with N connectors on both sides.

In order to obtain application layer metrics, the traffic generation and bandwidth measurement tool Iperf version 2.0.4-3 [9] was used. Metrics related to the underlying wireless technology such as RSS and NF, were obtained from the radiotab header using TShark version 1.0.4 [10] to parse the received packets. For stability and usability reasons, the Multiband Atheros Driver for Wireless Fidelity (MADWiFi) version 0.9.4 revision 4023 [11] wireless driver was favoured over the ath5k driver.

The testbed was configured as shown in Table I. In all experiments the link between A_2C_1 operated on channel 36 to communicate while the link between A_1B_1 was changed to achieve the variation in channel separation. The IP routing tables in the nodes were configured such that A_1 communicates exclusively with B_1 and A_2 communicates exclusively with C_1 and vice versa. Additionally, for each established interface, e.g *ath0*, a virtual monitoring interface was created on the same physical interface. This was used to obtain the radiotab header from the data packets being received.

Each machine in the testbed was also connected to a switch over wired Ethernet and configured as part of the same subnet. This setup was required since a shell script

Table I
INTERFACE CONFIGURATION

Node	Interface	Type	IP
Node A	A_0	eth	192.168.0.1/24
	A_1	ath	10.0.2.1/24
	A_2	ath	10.0.3.1/24
Node B	B_0	eth	192.168.0.2/24
	B_1	ath	10.0.2.2/24
Node C	C_0	eth	192.168.0.3/24
	C_1	ath	10.0.3.2/24

operating on the sending side of each link was responsible for the configuration of the *ath* interfaces on both the Tx and Rx sides of the link. Hence, for each pair of wireless interfaces which make up a link a shell script is used to configure the interfaces to the required settings.

Furthermore, to synchronise both scripts one of the three machines acts as a synchronisation server. Every script simply creates a text file remotely in the /tmp directory of the synchronisation server which indicates the current state of the remote script. This communication was done using SSH sessions over the Ethernet connection.

Synchronisation between the scripts was necessary since after reconfiguring a wireless interface the time taken for layer two connectivity to be established is variable. Therefore, before starting the Iperf client during each experimental run, the script checks that the Iperf server is reachable using a single 64 byte ping. If the server is reachable, a check is performed to verify that the remote Iperf client is also ready. This check simply reads the synchronisation file from the synchronisation server. Only if both actions are completed successfully does the script begin the experiment and measurements. Due to some inaccuracies in the time taken for both interfaces to become active ($\approx 1-2$ s), the first and the last 10 IPerf samples are not considered when processing the results.

As it was required that TShark only capture data frames from the wireless interface in each node, the script automatically sets an appropriate Ethernet filter using the pcap library syntax. It is also worth noting that the shell script forces Iperf to bind on a particular IP address to ensure that the correct data is captured. This incoming packets are filtered based on the source/destination MAC addresses as well as the data type of the received packet; specifically only packets of type data are captured and both type ctl and type mgt are ignored. As mentioned earlier, the radiotab header is parsed to extract the required data; specifically the RSS (from the radiotab header field *dbm_antsignal*), NF (from the *dbm_antnoise* field) and the physical layer datarate (from the *datarate* field).

In all experiments performed the User Datagram Protocol (UDP) was used as the transport layer protocol. UDP was chosen over Transmission Control Protocol (TCP) and

Stream Control Transmission Protocol (SCTP) since it has no inherent congestion control mechanisms which would compensate for link degradation and hence distort the results. The UDP payload size for each experiment was set to a fixed value of 1400 octets at all times.

As recommended by Burchfield et al. [8], any research carried out in the 802.11 domain that is based on obtaining results from real deployments should be performed using their recommendations. To follow this approach, it can be confirmed that the room in which the experiments were conducted was free of any interference in the measured ISM band. This was verified in two ways; firstly by sniffing the medium for other WiFi radios. Secondly, in order to verify that there was little or no interference from non WiFi devices experiments were performed in which a single link between two 802.11 radios was set up. These experiments showed that the maximum possible UDP data rate without an RTRs could be achieved, thereby showing that there was no significant level of interference. However, the recommendation of Burchfield et al. [8] that verification of the testbed be carried out using coax cables between the radios was not conducted, as described in Section II.

B. Statistical Evaluation

In order to prove the reliability of the measured results the Confidence Interval (CI) was calculated for each sample mean using the method of independent replications, as described by Banks et al. [12]. This is shown in Equation 1 where $\bar{\mu}$ represents the sample mean, ν the Degree of Freedom (DF), α the chosen CI and $\sigma(\bar{\mu})$ the standard error or variance of the sample mean.

$$\bar{\mu} \pm t_{\frac{(1-\alpha)}{2}, \nu} \cdot \sigma(\bar{\mu}) \quad (1)$$

As proposed by Banks et al., the Student's t-distribution was used to define ν ; this is because every measured sample is independent and the common Gaussian distribution does not cover this case accurately. Since every measurement represents a set of non-normal data the number of samples used to calculate the sample mean and the corresponding CI should be at least 50 as recommended in Wang [13]. Hence, to calculate the throughput sample mean each measurement was performed over a period of 50s with an interval of 0.5s which gives a set of 100 values for the IPperf results. The corresponding DF of $\nu=1.99$ can be derived for a probability of 5% to exceed the critical value. Further mean values, e.g., RSS and NF, were taken directly from the radiotap header which appears per MAC frame and results in a DF of $\nu=1.96$ for total number of independent measured values above 100 per sample mean and the probability of exceeding the critical value.

C. Fresnel-Zone and Free Space Path Loss

In order for the results presented in this paper to be accurate and free from external influences, it was necessary




Figure 2. Loss Parameters within a System as shown in Equation 4

that the experimental environment had to fulfil some basic requirements. To achieve this there must be a direct Line of Sight (LOS) connection between each pair of transmitting and receiving antennas with no obstructions in the path which could cause inference due to reflections or shadowing effects. At a minimum the first Fresnel zone must be largely free from obstacles to avoid interference from reflected waves.

Equation 3 shows the simplified formula to calculate the n th Fresnel. Using this equation a radius F_n can be obtained which describes a zone that surrounds the direct LOS connection between both antennas that is completely free from obstacles, e.g., trees, hills or walls. Since in the experimental environment the only obstacle is the ground, Equation 3 just comprises the distance d between both antennas and is used to calculate the minimum antenna height such that the first Fresnel zone is clear. Therefore, the general simplified Fresnel formula is:

$$F_n = \sqrt{\frac{n \cdot \lambda \cdot d_1 \cdot d_2}{d_1 + d_2}} \quad (2)$$

$$F_n = \sqrt{\frac{n \cdot c \cdot d}{2f}} \quad (3)$$

To verify the accuracy of the results from the experimental environment the overall system loss a_s was computed analytically. A comparison between the actual attenuation experienced in the experimental set-up and the theoretical attenuation predicted was then made. To compute the theoretical attenuation that should be experienced, the following formula was used:

$$a_s = a_{L1} - g_1 + FSPL + a_z - g_2 + a_{L2} \quad (4)$$

Equation 4 is made up of the cable losses from Mesh Node (MN)₁ to antenna a_{L1} and from MN₂ to antenna a_{L2} , the gain of both antennas g_1 and g_2 , the free space path loss described as

$$FSPL = 92.4 + 20 \log(d) + 20 \log(f) \quad (5)$$

and some additional unpredictable losses a_z , as depicted in Figure 2. Examples for unpredictable losses are interferences due to multipath propagation especially for frequencies less than 10 GHz and losses due to atmospheric absorption.

IV. ADJACENT CHANNEL INTERFERENCE IN MULTI-RADIO NODES

There is currently a lot of interest in both the industrial and academic research communities for using IEEE 802.11a to provide backbone links for WMNs while simultaneously using IEEE 802.11b/g to provide user access. This means that each MN may have multiple 802.11a radios operating in close proximity and hence ACI issues in 802.11a are of particular importance.

Another important factor is the relatively short spacings that can be achieved between co-located antennas. It is simply not feasible to have antenna separations larger than 50 or 60 cm due to the dimensions of the MN casings and the goal of having non-intrusive and easily deployed MNs. For example it is not possible to have antenna separations of 2 m as proposed by Worldwide Interoperability for Microwave Access (WiMAX) deployment guidelines. The cumulative effect of both ACI and small antenna separations can have a significant affect on the performance of such systems and these issues must be addressed prior to deploying an 802.11a based WMN infrastructure. For this reason, the spectrum mask of the WiFi cards used in the work was examined to verify that it meets with the IEEE guidelines [14]. Figure 3 depicts the measured spectrum of a MikroTik R52 mini-PCI WiFi card. These measurements were performed using the IEEE recommended guidelines for spectrum measurements of 802.11a systems; specifically the spectrum analyser was set with a 100kHz Resolution Bandwidth (RBW) and a 30kHz Video Bandwidth (VBW) [14].

Although, Cheng et al. [4] had previously conducted this spectrum analysis, they considered a spectrum mask that should not exceed -20 dB at 11 MHz and -30 dB at 22 MHz. Furthermore, Cheng et al. stated that they used TxPower values of 30 dBm, 36 dBm and 99 dBm which were allegedly performed using the MADWiFi driver. Since the considered Power Spectral Density (PSD) limits and the chosen TxPower values do not fit to IEEE 802.11a, 802.11b or 802.11g specifications, the work and results obtained cannot be considered accurate. For operating in the 802.11a band the IEEE recommend a 20 MHz channel spacing, a maximal bandwidth of 18 MHz at 0 dB and offsets of at least -20 dB at 11 MHz, -28 dB at 20 MHz, and -40 dB at 30 MHz. Figure 3 depicts the measured PSD of the WiFi cards used for the experiments presented in this work. To accurately compare each of the different TxPower curves, the area around the centre frequency f_c of 5.2 GHz was normalised to $RSS_N = 0$. Therefore, for every TxPower curve the sample mean $\bar{\mu}_a$ between 5.191 GHz and 5.209 GHz was calculated and then added as a fixed value to the series of measurements. After applying the mean value to the whole curve the new mean value for the range between 5.191 GHz and 5.209 GHz was consistently above zero. To adjust the normalisation process to the correct mean value that will




Figure 3. 3D Spectrum Analysis of a MikroTik R52 IEEE 802.11a/b/g Card for 5.2 GHz Carrier with a 20 MHz Bandwidth

be added to the entire curve, half of the belonging standard deviation $\sigma(\bar{\mu}_a)$ was also subtracted, as shown in Equation 6.

$$RSS_N = x + \bar{\mu}_a - \frac{\sigma(\bar{\mu}_a)}{2} \quad (6)$$

Due to the small SNR values of the lower TxPower measurements, as shown in Figure 3, the PSDs for $f_c \pm [9 \text{ MHz}, 11 \text{ MHz}]$ could not be obtained as their values were already equal to or less than the noise floor. However, based on the obtained results shown in Figure 3 it is reasonable to say that the MikroTik R52 WiFi cards meets with the requirements of the IEEE guidelines.

V. RESULTS AND DISCUSSION

This section presents and describes results obtained from experiments performed on the experimental testbed described earlier.

A. Received Signal Strength and Noise Floor Level Measurements

In 802.11 the term Received Signal Strength Indicator (RSSI) is used as a generic unitless signal strength metric and is the only value used to describe the signal strength of a received packet. The RSSI is an arbitrary indication of the actual RSS with a range that is defined by each vendor individually depending on the level of granularity required. The lowest possible value is 0 and goes up to the highest arbitrary value (e.g. 70 for MADWiFi).

An RSSI of 0 always represents the NF inside the radio. To calculate the actual RSSI value MADWiFi takes the RSS value from the card and adds 96 dB (NF) to it. This corresponds to a maximal RSS of -26 dBm for MADWiFi since this value is equal to an RSSI of 70. In order to easily compare the RSS and CS threshold, and to be able to make accurate assumptions about the card and driver behaviour related to the signal strength, it is essential to use the RSS instead of the RSSI. However, caution must be taken when focusing solely on the RSS, as explained below.




Figure 4. Received Signal Strength and Noise Floor for a Radio With Constant TxPower and an Additional Interferer

Each WiFi chipset has an internal adaptation mechanism which adapts the noise floor level in response to the level of external interference. Specifically, the higher the level of interference, the lower the reported NF. It should be noted that this relationship is not linear and appears to use predefined thresholds which are set inside the chipset and are used to adapt the NF. This effect is depicted in Figure 4. It shows two radios placed at a distance of 10 cm from each other. The TxPower of the first radio is fixed while the TxPower of the second is incremented from 0 to 15 dBm. In Figure 4 the x axis shows the TxPower of A_1 as it increases, whereas the y axis depicts the RSS (upper surface with constant lines) measured by interface A_2 with a fixed TxPower as well as its NF (lower surface with dotted lines). From this Figure it is clear that both surfaces are equal. In particular, when the neighbouring interface A_1 increased its power from 3 dBm to 4 dBm the measured RSS and NF at radio A_2 in C_1 decreased by 12 dB.

Since the TxPower of A_2 was set at a fixed value of 15 dBm for the entire duration of the experiment, the drop in RSS is quite unexpected. Indeed it was assumed that it would report the same RSS for every received packet with only a small and relatively insignificant difference in Standard Deviation (StdDev). To investigate this issue in more detail the SNR, as given in Equation 7, for both transmissions is calculated as depicted in Figure 5.

$$SNR = RSS_m - NFL_m \quad (7)$$

As in the previous Figures the constant surface describes the radio interface sending with a fixed TxPower and the dotted line represents the radio which incrementally increased its TxPower. It is clear from Figure 5 that the SNR of the connection $A_2 C_1$ is constant over the entire measurement space. As expected the SNR surface of the connection $A_1 B_1$ increases incrementally as the TxPower of the sending interface A_1 is increased.

The relationship between the measured RSS and NF to the actual RSS can be described as:

$$RSS = RSS_m - |NFL_{def} - NFL_m| \quad (8)$$




Figure 5. Derived Signal to Noise Ratio Values from Figure 4

where RSS_m is the RSS reported by the card, NFL_m is the NF defined by the vendor (i.e. -96 dBm) and NFL_m is the actual reported NF.

B. Conducted Measurements

As already shown in the previously published work [1], there are interference issues between neighbouring antennas beyond a channel separation of one. This was one of the main contributions of Robitzsch et al. and due to the importance of the problem this has been investigated further in this work. The presented results within this paper however just incorporates TxPower settings. Since the distance between every transmitter and receiver were fixed, both values, TxPower and RSS, are interchangeable. Additionally, throughput measurements are known to give a good indication of how well a link operates. However, they do not tell the whole story if they are evaluated without any further measured parameter such as RTR or application layer loss.

To discover the reason the maximum theoretical throughput cannot be reached in some cases other metrics must also be considered to investigate possible sources of interference. In particular, the parameters RTR, MAC layer loss and application layer loss give a more detailed view of why the maximum throughput was not achieved. The importance of the additional metrics will become clearer after this section when a detailed description of the obtained graphs is provided. For instance, if the throughput of a connection is always 60% lower than expected the question is whether this is due to RTRs caused by interference or whether the card backed-off a large number of times. Additionally, the relation between the MCS to the level of interference gives a strong indication as well whether this MCS is suitable under these conditions and provides enough robustness.

Figure 6 to Figure 13 show a complete set of throughput measurements for an antenna separation of 20 cm including the corresponding RTRs which gives an indication of the weak link performance. As can be seen in Figure 7, radio A_2 was transmitting with a constant TxPower of 15 dBm on channel 36, whereas the TxPower of interface A_1 was

increased incrementally from 0 dBm up to 15 dBm on channel 40. To make a comparison between the different results easier, the surface of interface A_2 (fixed power surface and channel 36) was mapped to the corresponding TxPower settings of A_1 . For instance, if A_1 (channel 40) was operating with TxPower 5 the corresponding measured throughput value for A_2 was mapped to the same TxPower on the x-axis, however, A_2 was still sending with 15 dBm.

From Figure 6 it can be seen that once A_1 has reached a TxPower of 2 dBm both radios achieve the same throughput. This could indicate that both radios are fairly sharing the medium as if only a single channel is being used; this occurs because the side-band from each causes the carrier sensing mechanism in the neighbouring interface to detect the channel as busy due to ACI. It could also indicate that the PSD of the side-band from each is not high enough to cause CS in the other radio but rather that there is large number of RTRs due to corrupted received frames in either B_1 , C_1 or in both. To find the source of this problem the aggregated throughput is depicted in Figure 14. This shows that a throughput value of 42 Mbps for a channel separation of one and a PHY rate of 54 Mbps is achieved. This value is close to the maximum possible throughput for a single channel when using UDP and the highest MAC Service Data Unit (MSDU) size.

The RTR in Figure 7 is always around 10% which indicates that there are quite high levels of interference. This interference causes a reduction in the maximum possible throughput of approximately 3 Mbps. Additionally, it can be confirmed that in these experiments the application layer loss was always close to zero which indicates that there was a small number of expired retransmissions. Combined these results show that it is not only the card's internal CS mechanism that causes ACI in terms of less successfully transmitted packets, but also distortion of the transmitted packets from the neighbouring radio sending a packet after the CS mechanism detected the medium to be idle. This claim can be verified by the previous work of Angelakis et al. who show the impact of the sideband of a 802.11a 20 MHz wide OFDM carrier, which can cause ACI if the antennas are located quite close, e.g., 20 cm. However, whether it was the data packet or the ACK frame that was affected by the interference cannot be obtained from this set of results. This will be examined in the next Section V-C.

It was expected that by increasing the channel separation by one the ACI effect would no longer be present and that the maximum expected throughput should be achieved on each channel. However, as shown in Figure 8, this is surprisingly not the case. In comparison to the previous Figure 6 both surfaces have less variation as if a low pass filter had been applied. However, as the TxPower of A_1 is increased the throughput of A_2 decreases significantly and when a TxPower of 8 dBm is reached both surfaces are very similar.




Figure 6. Application Layer UDP Throughput for Channel Separation of 1 and Antenna Separation of 20 cm




Figure 7. Retransmission Rate for Channel Separation of 1 and Antenna Separation of 20 cm

The aggregated throughput in Figure 14 confirms that T_{agg} is still significantly lower than double the maximum throughput of a single channel that would be expected when using two independent non-interfering channels. However, the throughput value for the highest TxPower and MCS reaches 47 Mbps which is just slightly above the theoretical maximum throughput of a single connection. The corresponding RTRs for this experiment, as depicted in Figure 9, indicates that after the TxPower value of 8 dBm is reached the RTR drops to almost 0%; however, only for interface A_1 which increased its TxPower. Furthermore, interface A_2 always sent packets without requiring any retransmissions, except when the 64-Quadrature Amplitude Modulation (64-QAM) MCSs was used (i.e. 48 Mbps and 54 Mbps where the RTR goes up to 1%). Hence, this shows that after reaching a specific ratio between a radio's TxPower and its neighbouring TxPower there is no longer interference from the neighbouring radio; there is however interference from an as yet uninvestigated source.

The channel separation was then increased to three, the corresponding throughput and RTR values are presented in Figures 10 and 11, respectively. As can be clearly derived from Figure 10, the performance of the link $A_2 C_1$ is no longer significantly affected by the neighbouring connection $A_1 B_1$. However, as shown in Figure 14 the expected total




Figure 8. Application Layer UDP Throughput for Channel Separation of 2 and Antenna Separation of 20 cm




Figure 11. Retransmission Rate for Channel Separation of 3 and Antenna Separation of 20 cm




Figure 9. Retransmission Rate for Channel Separation of 2 and Antenna Separation of 20 cm




Figure 12. Application Layer UDP Throughput for Channel Separation of 4 and Antenna Separation of 20 cm




Figure 10. Application Layer UDP Throughput for Channel Separation of 3 and Antenna Separation of 20 cm




Figure 13. Retransmission Rate for Channel Separation of 4 and Antenna Separation of 20 cm

aggregated throughput is still not achieved. Rather an aggregated throughput of 62 Mbps with a TxPower of 15 dBm and a PHY rate of 54 Mbps is obtained. Taking the corresponding RTRs into account from Figure 11 it can be seen that there is still interference from A_2 in A_1 which causes packet distortion reflected by the high RTR of the link $A_1 B_1$. Once A_1 was configured with a TxPower of 13 dBm both links share the medium equally but still with a low aggregated throughput.

Since both sending interfaces still interfere with each other and achieve throughput levels lower than the theoretical

maximum, the channel separation was increased to four with A_2 operating on channel 36 and A_1 on channel 52. From the corresponding throughput and RTR results shown in Figures 12 and 13 respectively, it can be seen that both links no longer significantly affect each other and that the aggregated throughput is now almost twice the throughput of a single connection, as depicted in Figure 14.

To complete the set of experiments for an antenna separation of 20 cm and proving that both links $A_2 C_1$ and $A_1 B_1$ no longer interfere the missing aggregated throughputs for channel separation of five, six and seven, i.e., channel 36-56,




Figure 14. Aggregated UDP Throughput for Antenna Separation of 20 cm and TxPower 15 dBm




Figure 15. Aggregated UDP Throughput for Antenna Separation of 10 cm and Tx Power 15 dBm

36-60 and 36-64, respectively, are illustrated in Figure 14. As can be seen from these results, the aggregated throughput is always equal to twice the throughput of a single 802.11a UDP connection.

As mentioned earlier, experiments were conducted for an antenna separation of 20 cm; however a full set of experiments were also performed for the 10 cm case. Unfortunately, the stability and reliability of the links were extremely poor meaning that the obtained results were extremely difficult to interpret. As can be seen in Figure 15 the obtained results are however good enough to show the same general trends as seen in the 20 cm case. It can be argued that under these conditions the higher the channel separation, the better the aggregated throughput of the system. However, only when a channel separation of seven is reached - which means using channel 36 and 64 - does the aggregate throughput T_{agg} reach twice the throughput of a single UDP transmission. This shows that there are other significant factors to be considered other than just the side-band of a 20 MHz channel, as the side-band will only impact the adjacent channels. In order to explain these results it is assumed that each 802.11a chipset produces some interference over the entire 802.11 frequency spectrum that smoothly flattens out at higher channel separations to the frequency the adjacent radio is operating on.

Results for an antenna separation of 30 cm showing T_{agg} for a TxPower of 15 dBm are presented in Figure 16. As




Figure 16. Aggregated UDP Throughput for Antenna Separation of 30 cm and Tx Power 15 dBm




Figure 17. Aggregated UDP Throughput for Antenna Separation of 40 cm and Tx Power 15 dBm




Figure 18. Aggregated UDP Throughput for Antenna Separation of 50 cm and Tx Power 15 dBm

depicted, when a channel separation of two is used the aggregate throughput achieved is equal to twice the possible throughput of a single 802.11a UDP link; this indicates that there is no more interference from the neighbouring radio as was the case for channel separations from four to seven in the 20 cm case (Figure 14).

C. Retransmissions, Carrier Sensing and Application Layer Loss

As shown in the previous Section V-B, under certain conditions two neighbouring transmitting radios can cause CS, RTRs and packet distortion. This was shown to be due to the effect of a high RTR which leads to a throughput decrease; however it is still not clear whether it is because the transmitted packet or the ACK was destroyed. To investigate




Figure 19. Throughput Measurement for TxRx Case, Channel Separation of 3 and with NoAck Policy Enabled




Figure 20. Application Layer Loss for TxRx Case, Channel Separation of 3 and with NoAck Policy Enabled

this issue further another experiment was conducted.

The same experimental environment as described in Section III-A was used, except in this case the channel configuration was changed to $A_2 C_1$ on channel 52 and $A_2 C_1$ on channel 64. The direction of link $A_2 C_1$ was reversed to $C_1 A_2$. Later on the term Transmission-Receiving (TxRx) will refer to the reversed testbed setting and Transmission-Transmission (TxTx) to the settings where both neighbouring antennas are transmitting data. Note, that for the sake of simplicity only TxPower values of 1, 8 and 15 dBm were considered for interface A_1 . Furthermore, the No Acknowledgement (NoAck) policy of the MADWiFi driver was used to force the card to send and not wait for ACKs. Taking these settings and recording a log of the application layer loss provided by IPerf, it was possible to investigate whether the actual sent packet was destroyed or the returning ACK (in the cases where the RTR was between 5 to 10%). To answer this question, Figure 19 and Figure 20 were produced which depict the logged data for the throughput of a TxRx use case where B_1 and A_2 were the transmitters and the corresponding application loss chart, respectively. With regard to the sent ACKs in the experiments described in the previous section, it can be clearly seen from both Figures that these management packets from B_1 to A_1 are not destroyed by the neighbouring antenna A_2 for any of the TxPower settings in A_2 . This can be seen in Figure 19 where the throughput is equal to that of a single connection with no interference for the lowest MCS. Furthermore, the application loss presented in Figure 20 shows no losses at all for channel 52. Additionally, it can be seen that if both connections use the same TxPower of 15 dBm that by using 64-QAM with both Coding Rates (CRs) 2/3 and 3/4 A_1 destroys close to 80% of the sent packets of C_1 to A_2 . What has still not been investigated so far is the question whether both radios affect the data packets sent by the neighbouring interface for higher channel separations than one. In order to do so further experiments have been conducted. Figure 21 shows the chart for a TxTx set-up where A_2 again stays constant at a TxPower of 15 dBm




Figure 21. Throughput Measurement for TxTx Case, Channel Separation of 3 and with NoAck Policy Enabled

and A_1 changes its TxPower to 1, 8 and 15 dBm. Taking the corresponding application loss Figure 22 into account, it is worth noting that as A_1 increases its TxPower it got more often access to the medium which leads to a higher throughput. At the same time the throughput of A_2 drops down to the same level as A_1 when both operating with the same TxPower. Therefore, they consequentially share the medium equally over all TxPowers, as depicted in Figure 23 which shows the aggregated throughput. That no packet was destroyed on either A_1 's or A_2 side can be confirmed by taking the application loss, as depicted in Figure 22. It is clearly observable from this figure that almost no packet was affected except for the case where A_1 sent at 1 dBm and with MCS 16-Quadrature Amplitude Modulation (16-QAM). However, it can be confirmed that by switching to 64-QAM there was no throughput value logged at all by IPerf. Therefore, both MCSs 16-QAM and 64-QAM are not robust enough to not get affected by the neighbouring interface if the neighbour's Equivalent Isotropically Radiated Power (EIRP) is much high than the outgoing power from the own antenna. Since the previous provided aggregated throughput Figures 15 to 18 show solely TxPowers of 15 dBm this application loss for 16-QAM and 64-QAM MCSs can be ignored.




Figure 22. Application Layer Loss for TxTx Case, Channel Separation of 3 and with NoAck Policy Enabled




Figure 23. Aggregated Throughput for TxTx Case, Channel Separation of 3 and with NoAck Policy Enabled

VI. CONCLUSION AND FUTURE WORK

The use of multiple radio equipped nodes within a WMN is the most promising approach for significantly increasing network performance.

A key problem with this approach however, is that antenna separations of less than 50 cm have a significant impact on the performance which can be achieved. The presented work here shows that it is not only ACI which has an impact but also ICI (channel separations of more than one). This problem does not appear to have been investigated previously and is the primary contribution of this paper.

The results were obtained based on experiments performed in a real testbed environment which was evaluated to produce reliable and reproducible results. This evaluation used CIs calculations and prior offline planning by using the Fresnel formula and statistical methods to design the testbed. The results presented show that by increasing the channel separation between co-located radios that the level of ICI decreases. All presented results take the radio parameters TxPower, MCS, channel separation and physical layer effects into account to explain the performance degradation due to carrier sensing, packet distortion and backing-off.

The results obtained will be used to develop an algorithm that takes the radio parameter settings, external dependencies and some prior knowledge as an input and provides the

optimal global configuration of nodes in a WMN so that ICI will be minimised.

ACKNOWLEDGEMENTS

The authors of this paper would like to thank Mathias Kretschmer and Christian Niephaus for their initial work on the awarded paper.

The support of the Irish Research Council for Science, Engineering and Technology (IRCSET) and the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged.

This work was partially funded by the European Commission within the 7th Framework Program in the context of the ICT project Carrier-Grade Mesh Networks (CARMEN) [15] (Grant Agreement No. 214994). The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the CARMEN project or the European Commission.

REFERENCES

- [1] Sebastian Robitzsch, Christian Niephaus, John Fitzpatrick, and Mathias Kretschmer. Measurements and Evaluations for an IEEE 802.11a Based Carrier-Grade Multi-radio Wireless Mesh Network Deployment. In *2009 Fifth International Conference on Wireless and Mobile Communications*, pages 272–278. IEEE, 2009. ISBN 978-1-4244-4679-7. doi: 10.1109/ICWMC.2009.52.
- [2] Vangelis Angelakis, Stefanos Papadakis, Vasilios Siris, and Apostolos Traganitis. Adjacent channel interference in 802.11a: Modeling and testbed validation. In *2008 IEEE Radio and Wireless Symposium*, pages 591–594. IEEE, 2008. ISBN 978-1-4244-1462-8. doi: 10.1109/RWS.2008.4463561.
- [3] Jens Nachtigall, Anatolij Zubow, and Jens-Peter Redlich. The Impact of Adjacent Channel Interference in Multi-Radio Systems using IEEE 802.11. In *2008 International Wireless Communications and Mobile Computing Conference*, pages 874–881. IEEE, August 2008. ISBN 978-1-4244-2201-2. doi: 10.1109/IWCMC.2008.151.
- [4] Chen-Mou Cheng, Pai-Hsiang Hsiao, H. T. Kung, and Dario Vlah. Adjacent Channel Interference in Dual-radio 802.11a Nodes and Its Impact on Multi-hop Networking. In *IEEE Globecom 2006*, pages 1–6. IEEE, 2006. ISBN 1-4244-0357-X. doi: 10.1109/GLOCOM.2006.500.
- [5] Vangelis Angelakis, Nikos Kossifidis, Stefanos Papadakis, Vasilios Siris, and Apostolos Traganitis. The effect of using directional antennas on adjacent channel interference in 802.11a: Modeling and experience with an outdoors testbed. In *6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and*

- Wireless Networks and Workshops*, pages 24–29. IEEE, 2008. doi: 10.1109/WIOPT.2008.4586029.
- [6] Vangelis Angelakis, Konstantinos Mathioudakis, Emmanouil Delakis, and Apostolos Traganitis. Investigation of Timescales for Channel, Rate, and Power Control in a Metropolitan Wireless Mesh Testbed. In Paul Cunningham and Cunningham Miriam, editors, *ICT-MobileSummit 2009 Conference Proceedings*, 2009. ISBN 978-1-905824-12-0.
- [7] Arunesh Mishra, Vivek Shrivastava, Suman Banerjee, and William Arbaugh. Partially overlapped channels not considered harmful. In *SIGMETRICS '06/Performance '06: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 63–74, New York, NY, USA, 2006. ACM. ISBN 1-59593-319-0. doi: <http://doi.acm.org/10.1145/1140277.1140286>.
- [8] R. Burchfield, E. Nourbakhsh, J. Dix, K. Sahu, S. Venkatesan, and R. Prakash. RF in the Jungle: Effect of Environment Assumptions on Wireless Experiment Repeatability. In *2009 IEEE International Conference on Communications*, pages 1–6. IEEE, 2009. doi: 10.1109/ICC.2009.5199421. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5199421>.
- [9] NLANR and DAST. Iperf. URL <http://iperf.sourceforge.net/>.
- [10] Gerald Combs. TShark - The Wireshark Network Analyser. URL <http://www.wireshark.org>.
- [11] MADWiFi. Multiband Atheros Driver for WiFi. URL <http://madwifi-project.org/>.
- [12] Jerry Banks, John Carson, Barry L Nelson, and David Nicol. *Discrete-Event System Simulation, Fourth Edition*. Prentice Hall, December 2004. ISBN 0131446797.
- [13] F K Wang. Confidence interval for the mean of non-normal data. *Quality and Reliability Engineering International*, Volume 17:257–267, 2001. doi: 10.1002/qre400.
- [14] IEEE. IEEE 802.11-2007 Wireless LAN Medium Access Control and Physical Layers Specifications, June 2007.
- [15] CARMEN. Carrier Grade Mesh Networks. URL www.ict-carmen.eu.