The Myth of Non-Overlapping Channels: Interference Measurements in IEEE 802.11

Paul Fuxjäger, Danilo Valerio, Fabio Ricciato
Telecommunications Research
Center Vienna (ftw.)
A-1220 Vienna, Austria
Email: {fuxjaeger,valerio,ricciato}@ftw.at

***Abstract*— It has become a widely accepted assumption that multiple IEEE 802.11b/g transmissions in physical proximity can coexist without interfering each other. This is claimed to be the case when using separate channels with a minimum distance of 25 MHz, e.g. channel 1 and 6, which are often referred to as *non-overlapping*. In contrast we show that in practice cross-channel interference can be present also between non-overlapping channels if the interfering transmitter is in the proximity of the receiver. This phenomenon is known as the "near-far effect" in wireless communications. On IEEE 802.11 this has two main effects: frame corruption due to increased interference noise and channel blocking due to spurious carrier detection. The problem can be particularly serious when using IEEE 802.11 technology to build multi-hop mesh networks.**

Through an extensive set of experiments with off-the-shelf certified WiFi chipsets we demonstrate the presence and the detrimental effects of cross-channel interference between non-overlapping channels. We adopt an incremental approach: we first consider the case of unacknowledged broadcast packets, then we extend to regular UDP streams, finally we provide preliminary results for multi-hop TCP flows.

I. Introduction

Over the last 10 years the widespread success of the IEEE 802.11 standard family has led to an ever increasing number of wireless network deployments worldwide. End-user equipment prices are dropping, WiFi chipsets are being integrated into home entertainment components and in various types of mobile devices like notebooks, PDAs, mobile phones, etc. Driven by the demand for higher throughput and connectivity, a large number of research publications are being dedicated to the enhancement of IEEE 802.11 based wireless LANs. Current hot research topics include routing protocols in support of self-organizing networks, MAC layer refinements, automatic channel selection and power adaption schemes aimed at increasing the spatial reuse of the unlicensed ISM spectrum. The vast majority of the contributions in this field are based on the assumption that multiple wireless transceivers can operate in close physical proximity, without interfering as long as the frequency separation between operating carriers is larger than 25 MHz.

The multi-channel architecture proposed in [1] uses a distributed channel assignment algorithm and considers non-overlapping channels as part of independent interference domains. The same assumption is done by Alicherry *et al.* [2] in the mathematical formulation of a joint channel assignment and routing problem. In [3] a study of the capacity of multi-channel wireless networks is presented with the assumption that multiple interfaces on a single node are capable to transmit and receive data simultaneously on non-overlapping channels. Raman [4] focuses on the problem of channel allocation in IEEE 802.11-based mesh networks. He uses multiple channels to divide the network into channel subgraphs that are considered fully decoupled concerning the medium contention and the transmission scheduling. Burton [11] mathematically evaluates the channel overlapping in infrastructured IEEE 802.11 networks and claims that a better efficiency can be achieved by reducing the minimal channel separation so as four channels can be used instead of three. A similar line of argument is presented in [5]. However both contributions only consider the interference between access points without taking into account the interference caused by clients, thus neglecting the near-far problem. Several further studies (e.g [6]–[10]) have proposed dynamic channel assignment schemes in the context of self-organizing networks. Most of them simply neglect any inter-channel effect.

In this contribution we will show that the assumption of perfect independence between non-overlapping channels does not always hold in practice. By means of simple experiments with commercially available hardware, we demonstrate the presence of detrimental interference effects also between channels that are nominally "non-overlapping". We found that the level of interference varies with physical distance, concurrent link-load, modulation rate, framesize, transmission power, receiver sensitivity and design, antenna patterns, etc. Such effects should be taken into account in the future research on enhanced IEEE 802.11.

II. Related Work

Very few works so far have reported the problem of interference between non-overlapping channels in IEEE 802.11. In [12] the authors performed some measurement on a 802.11b long distance point-to-point link. They noticed mutual interference between channel 1 and 11. They accounted this effect to the leakage in the near field of the antenna from the pigtail RF connector. Draves *et al.* [13] propose a new metric for routing in multichannel networks based on the loss rate and the link-bandwidth. By experimenting with TCP they found that two flows interact with each other, resulting in

1-4244-0860-1/07/25.00 © 2007 IEEE

considerable reduction of the total throughput. The authors of [14] performed experiments of multiple channel usage in a 802.11b wireless mesh backbone. They noticed that antenna proximity does have a strong impact on overall performance. In [15] the authors analyze the impact of antenna separation on channel orthogonality and show that if the antennae are in close proximity there is no interference-free channel pair. They find also that the internal electronic circuitry itself represents an additional source of crosstalk between the interfaces on the same motherboard.

All the above papers have incidentally noticed the effect of inter-channel interference during their experiments but none of them provided any detailed investigation. In this paper we directly address the problem and provide a comprehensive analysis of causes and effects. Based on an extensive set of ad-hoc experiments we investigate the impact of several factors like physical distance, channel separation, modulation scheme, carrier sensing and type of traffic.

III. IEEE 802.11 Specifications

The IEEE 802.11 specifications include a detailed description of PHY/MAC operational requirements. While there is a single MAC specification, several amendments for new physical layers have been standardized during the last years (e.g.. 802.11b, 802.11g, etc) with different modulation schemes for higher data rates.

The PHY layer embeds two components: The Physical Medium Dependent (PMD) system, that defines the transmitting and receiving schemes through the wireless medium, and the Physical Layer Convergence Protocol (PLCP) that represents an interface between MAC protocol and different physical media.

A. Physical Medium Dependent system (PMD)

In this paper we consider the two transmission techniques known as High Rate/Direct Sequence Spread Spectrum (HR/DSSS) and Orthogonal Frequency Division Multiplexing (OFDM), defined respectively in the 802.11b and 802.11g amendments. According to the standard specifications the 2.4 GHz ISM band is divided into 11 channels (for the FCC domain) or 13 channels (for the ETSI domain). Each channel has a nominal bandwidth of 22 MHz and the spacing between two neighboring central frequencies is only 5 MHz, as sketched in Figure 1. Two channels separated by at least 25 MHz are said to be "non-overlapping". We can identify several pairs of non-overlapping channels. A single triple exist in the FCC domain (1,6,11) and few more in the ETSI domain.

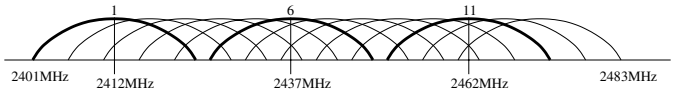

Fig. 1. IEEE 802.11b/g channelization scheme.

The coupling (interference) between different channels plays an important role in the interface design. In order to reduce the inter-channel interference the standard specifies a spectral mask for each modulation scheme and sets limits on the maximum out-of-band power relative to the peak power level, as sketched in Figure 2

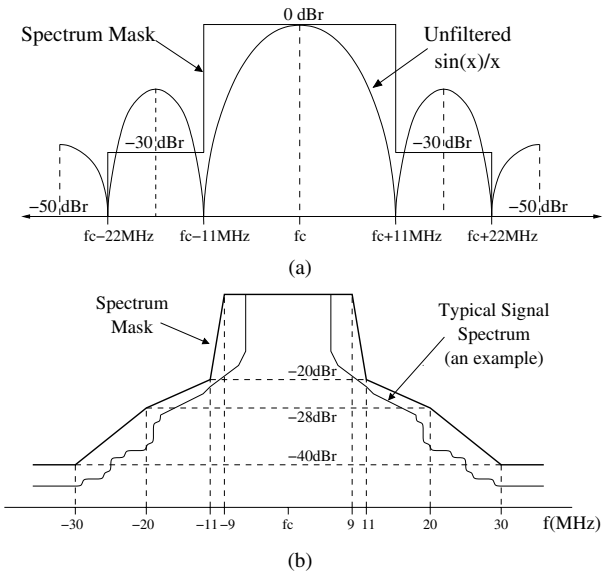

Fig. 2. IEEE802.11 spectrum mask for 802.11b (a) and 802.11g (b).

B. Physical Layer Convergence Protocol (PLCP)

The PLCP maps the IEEE 802.11 MAC frames into a format suitable to the specific medium. It adds to each MAC-PDU a preamble and a header. The preamble is used to synchronize the transmitter and the receiver, while the header contains some physical parameters used by the PLCP. The IEEE 802.11b standard mandates that the PLCP header is transmitted at 1 Mbps or 2 Mbps for long and short PLCP frame formats respectively. This scheme is also supported by IEEE 802.11g for backward compatibility. Additionally two 802.11g nodes can exchange PLCP header via OFDM at rate 6 Mbps.

The PLCP includes a threshold-based carrier sensing scheme. The received energy in the channel is measured during the header transmission and is mapped by the microcode into a nonnegative integer value called RSSI (Receive Signal Strength Indicator). The latter has a maximum value of RSSI_MAX (≤ 255). Note that the RSSI granularity, the value of RSSI_MAX and the mapping relation between energy and RSSI are not specified and therefore vendor specific. As an example Table I reports the values for a specific chipset.

IV. Testbed Setup

A. Transceiver type and software environment

The testbed consists of four laptops equipped with off-the-shelf INTEL PRO 2200BG mini-pci cards. This chipset (also known as Calexico/2) has been integrated in Pentium-M based portable computers since its first introduction in 2004. We chose a popular brand since we expect that the results are very

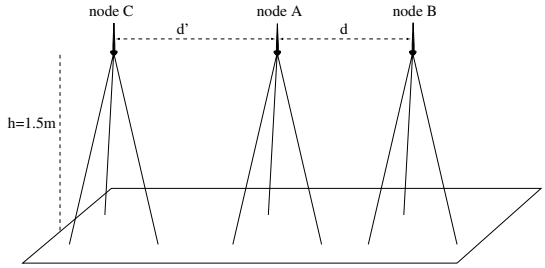

Fig. 3. Experiment topology.

dependent on vendor-specific design issues [17], therefore the adoption of a popular hardware type implies that the results are representative of a scenario found in practice.

The operating system environment was GNU/Linux (kernel-version 2.6.15-26) and the driver (version 1.1.4) for the wireless cards was taken from the IPW2200 project website [18] which is officially supported by the chipset manufacturer. The setting of the operational parameters (ad-hoc mode, transmission-power, modulation-rate, channel number, etc.) was accomplished via the `iwconfig` command. We remark that although this command provides also an option to change the receiver sensitivity (`sens`), this had no effect on any of our measurement results, rising doubt about the effectiveness of this option.

Similarly to other vendors the firmware source code is not publicly available, therefore other performance parameters like Energy Detection Threshold (EDT) or Clear Channel Assessment (CCA) mode cannot be manipulated. This prevented us from investigating the impact of these values on our measurements. Generally speaking, very little information is publicly available about the internal structure of commercial chipsets.

B. Antenna characteristics and propagation environment

We placed all nodes in a large open-space office environment. During the simulation campaign this location remained unpopulated. During the experiments we shut down any ISM band transmitter in the vicinity. External 2 dBi antennae with omnidirectional radiation patterns in the horizontal plane were mounted on tripods 1.5 m above the floor. A linear line-of-sight topology was used: antennae are placed along a straight line with no intermediate obstructions, as in Figure 3.

C. Measurement methodology

The interference effects shown in this paper are plotted over the RSSI difference between frames of interest and frames originating from other interfering transmissions on a non-overlapping channel. In subsection V-B we will motivate this choice showing that RSSI is a more adequate metric to describe the quality of wireless links than the physical distance.

The following method is used to retrieve the RSSI values. We used the tool packETH [19] to generate UDP packets at the transmitting node with 1472 bytes of random payload. We

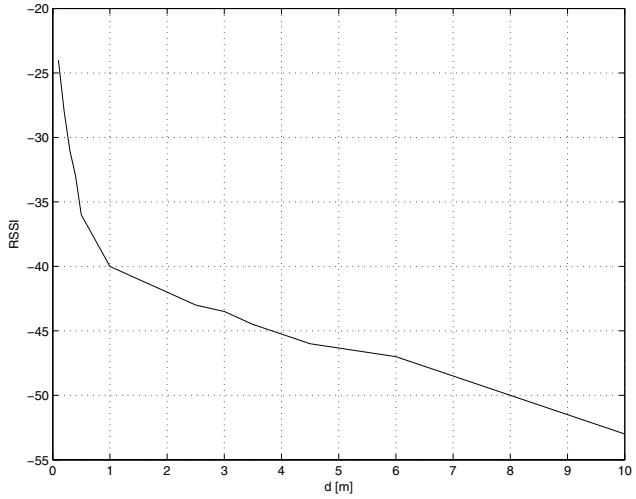

Fig. 4. Averaged RSSI vs. physical distance between node A and B.

enforced the transmission of CCK-modulated frames, for a maximum data rate of 11 Mbps. UDP packets were sent at the maximum rate to the broadcast MAC address, in order to avoid MAC acknowledgments. At the measuring station we used `ipwstats`, a script which is included in the IPW2200 driver package [18], to output the RSSI value of every successfully received data frame. In order to eliminate short term fluctuations RSSI values are averaged over at least 10^4 consecutive frames. We remark that the `ipwstats` provides dBm values of the RSSI, converting the RSSI scale defined in the standard by means of the RSSI_to_dBm mapping criterion chosen by Intel (see Table I). In Figure 4 we plot the averaged RSSI versus the physical distance between the transmitter and the receiver.

TABLE I
RSSI SETTING FOR CHIPSET INTEL 2200BG.

RSSI_MAX	92
Energy Range	$-85dBm - 20dBm$
Noise Level	$-85dBm$
RSSI_to_dBm	$RSSI - 112$

V. EXPERIMENTS

The claim of perfect separation between non-overlapping channels (e.g. X, Y) implies that *none* of the following two detrimental effects are observed in IEEE 802.11:

1) **Spurious Carrier Sensing** A station operating in channel X with packets in its transmission queue defers channel access because of activity on channel Y.
2) **Increased Interference Noise** A station receiving on X fails in successfully decoding the frame because of excessive interference originating from transmissions on channel Y.

In the next subsections we show experimental results where both effects are clearly visible, thus proving that the physical

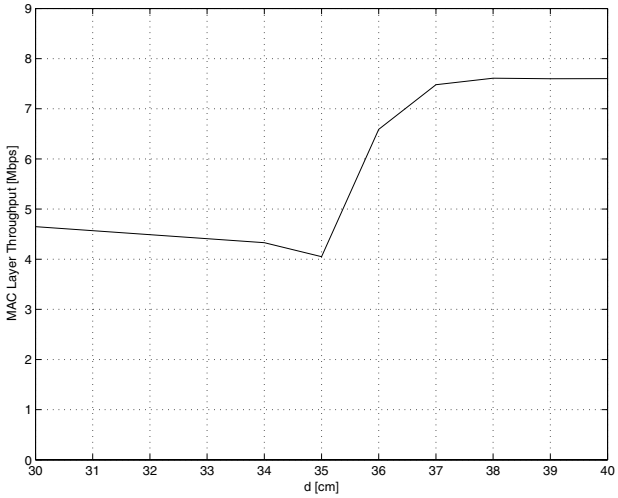

Fig. 5. MAC throughput vs. node distance for nodes broadcasting in channel 3 and 8.

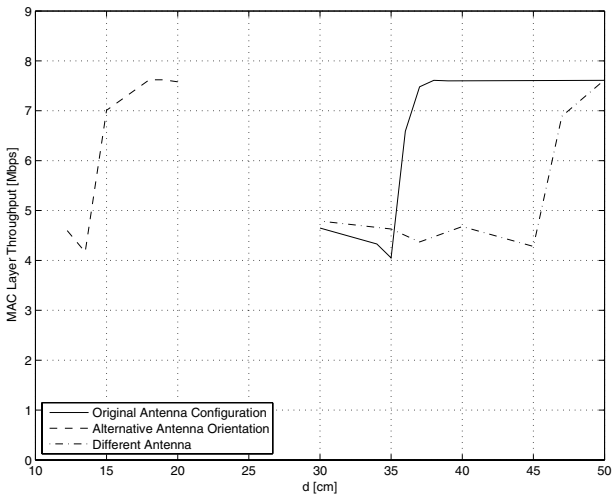

Fig. 6. MAC throughput vs. node distance for three different antennae configurations.

separation between so called "non-overlapping" channel does not hold in general.

A. Experiment 1: Spurious carrier sensing

The status of the carrier sensing mechanism which triggers the deferral of a transmission cannot be tracked directly, since there is no information available on how to read this data from the 2200BG firmware registers. Consequently we had to resort to an indirect test.

We let two nodes (A and B) transmit a continuous flow of MAC broadcast frames at maximum rate, with the setting described in Section IV-C, and we measure the actual outbound throughput at each station. The two nodes operate on two non-overlapping channels (we used 3 and 8). We repeat the experiment by changing the distance d between the two nodes. We change the distance d in step of 1 cm. Note that the use of non-acknowledged broadcast transmissions ensures that a drop in the throughput can be directly attributed to the spurious carrier sensing effect. In other words the use of non-acknowledged transmission rules out additional detrimental effects caused by interference noise on DATA and ACK frames. The effect of noise will be later considered in the next subsections.

In Figure 5 we plot the throughput from node A as a function of d. At distance above 38 cm the throughput is maximum (note the 11 Mbps data rate at the PHY correspond to $\approx$7.6 Mbps at the MAC layer). Below 38 cm the CCA (Clear Channel Assessment) mechanism starts to defer transmissions and the MAC layer throughput quickly decreases, reaching a local minimum of 4 Mbps at 35 cm. Essentially, the behavior at $d < 35$ cm is equal to the case when both nodes use the same channel. While this effect is negligible when considering single radio nodes which are usually separated by more than 35 cm, it identifies a serious problem for the design of wireless *multi-channel* nodes. Integrating two IEEE 802.11b/g transmitters tuned on non-overlapping channels in one single box with few centimeters of antenna separation might cost poor performances. Note that the throughput slightly increases again for very small distances: we conjecture that in this case the stronger cross-talk is beneficial because the carrier sensing can better synchronize the two sources in a deterministic manner.

B. Experiment 2: RSSI as discriminant parameter

We repeated the above experiment introducing a slight variation in the antennae orientation. Additionally, in a third set of experiment we changed both antennae. The new plots for the throughput are given in Figure 6. Remarkably the critical distance shifted dramatically (from 35 cm to 45 cm), despite the testbed configuration has changed only minimally.

For the same set of experiments, we plot in Figure 7 the throughput versus the RSSI values measured on the interfering channel (i.e the operating channel of the competing node). Note that the curves are flipped because of the inverse relationship between the RSSI and the physical distance. All the different curves now overlap very well with a common critical RSSI value of -29 dBm. This result suggest that RSSI is a better discriminant parameter than the physical distance as far as the impact of inter-channel interference is concerned.

This behavior provides an alternative description of the spurious carrier sensing phenomenon. Using the 2200BG chipsets with default transmission power of 20 dBm, spurious carrier sensing between non-overlapping frequencies starts to happen regularly when the total attenuation in the communication path (including cable-losses, path-loss, antenna gains, etc.) is less than 49 dB.

C. Experiment 3: Frame decoding from non-overlapping channel

In this experiment we tuned two nodes (A and B) to non-overlapping channels (3 and 8) and let one of them broadcasting a continuous stream of frames on his channel,

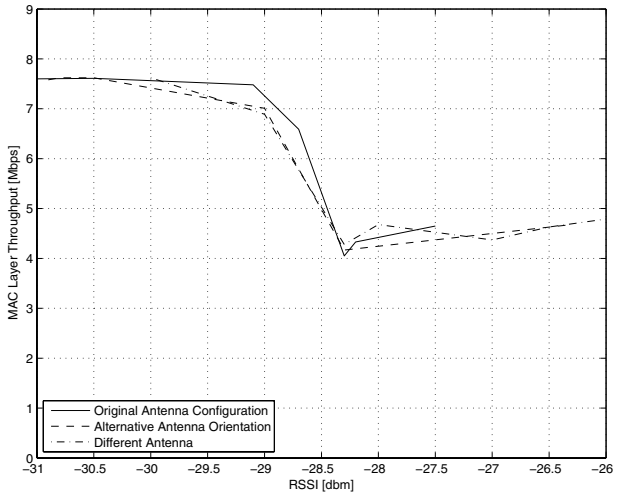

Fig. 7. MAC throughput vs. RSSI.

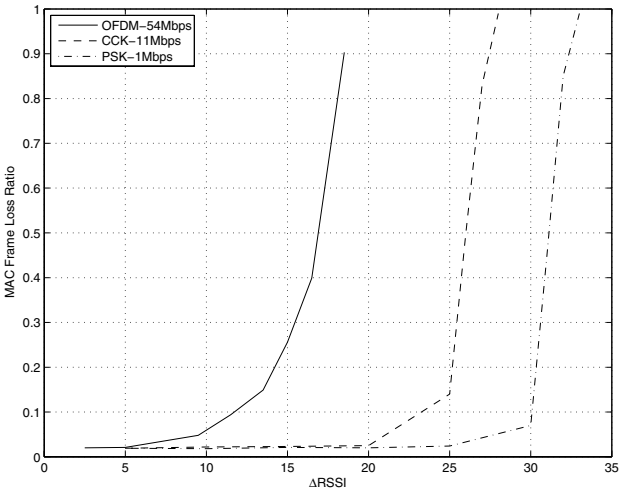

Fig. 9. Frame loss ratio vs. ΔRSSI for different modulation schemes.

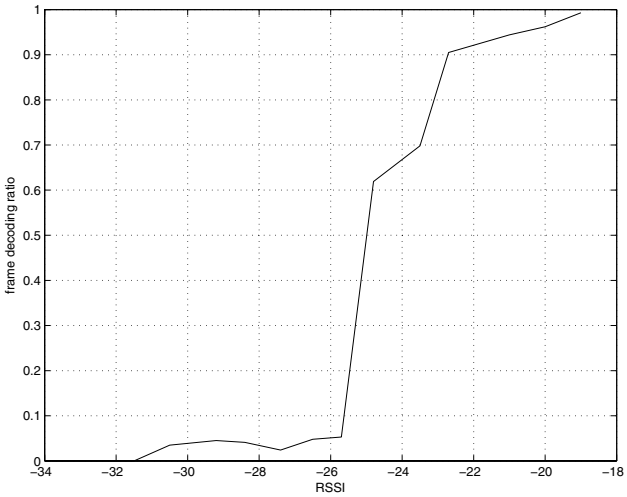

Fig. 8. Ratio of successfully decoded frames vs. RSSI.

with DBPSK at 1 Mbps. We observed that the other node was able to successfully decode frames on the other channel. In other words, the residual power received on channel 8 from a signal transmitted on channel 3 was sufficient to successfully decode the frames. We verified that no background scanning was active on the receiver node[1]. The experiment was repeated by changing the distance between the nodes, so as to vary the RSSI, and each time measured the average fraction of successfully decoded frames. The results are plotted in Figure 8. The curve yield a critical value around -25 dBm, above which the decoding ratio is above 70%. Again this behavior is in contrast with the assumption of full decoupling between non-overlapping channels.

[1]Most commercial cards perform automatic background scanning of other channels in support of the roaming functionality and auto-association.

D. Experiment 4: Frame reception errors

In the scenarios considered so far we have observed that transmissions in one channel can be sensed and even decoded on a different channel, despite non-overlapping. Expectedly such cross-channel interference disturbs the frame reception in case of simultaneous transmissions on non-overlapping channels, thus increasing the decoding failure rate, i.e MAC frame loss. In the next experiment we aim at quantifying this phenomenon.

With reference to Figure 3 we set node C to broadcast maximum-size frames on channel 3 at maximum rate. Node A acts as the receiver, and is tuned to channel 3 as well. Node B acts as an interference source, transmitting the same type of broadcast frames on channel 8. The use of maximum rate broadcast frames (unacknowledged) maximizes the channel usage. The only transmission gaps are due to the Distributed Interframe Frame Space (DIFS) and the random backoff interval. The level of interference can be quantified by the difference of the RSSI values measured on the two channels at the receiver node A, i.e ΔRSSI= RSSI$_{int}$-RSSI$_{rec}$, wherein the value of RSSI$_{int}$ can be measured offline simply by tuning the node A on the interfering channel. The performance metric in this case is the ratio of reception failures (frame losses) at node B. In Figure 9 we plot the measured values for different 802.11b/g modes.

Obviously the effect of interference increases with the absolute data rate: 1 Mbps DBPSK 802.11b is the most robust and 54 Mbps 802.11g OFDM-64QAM mode the most susceptible to inter-channel noise. In the latter case at ΔRSSI=17 dB half of the incoming packets cannot be successfully decoded due to the competing transmission on a non-overlapping channel. The progression of loss ratio versus ΔRSSI does not significantly change when using more distant channel pairs. In Figure 10 we compare three different channel pairs with spacing of 25, 30 and 35 MHz for the 54 Mbps 802.11g mode. Only a small shift of 1-2 dB can be observed.

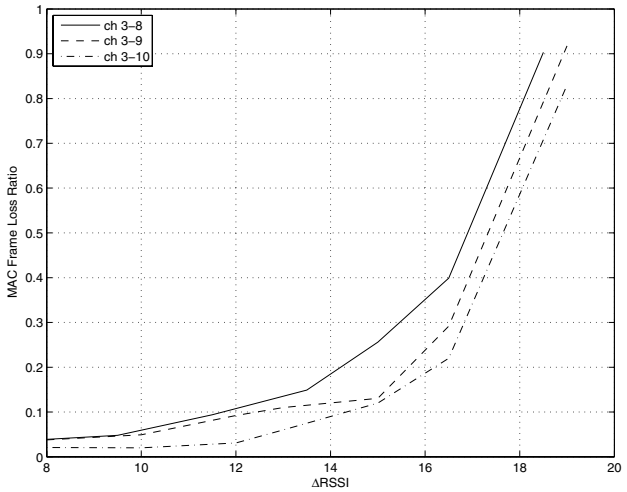

Fig. 10. Frame loss ratio vs. ΔRSSI varying the channel separation using OFDM 54Mbps.

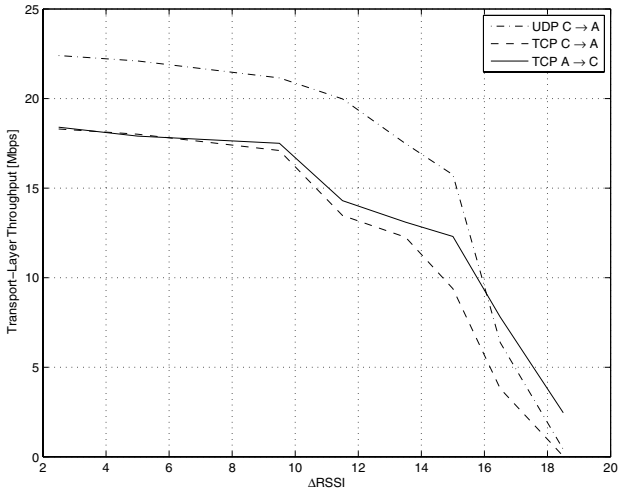

Fig. 11. Transport-Layer throughput vs. ΔRSSI using different transport protocols for OFDM 54Mbps.

E. Experiment 5: Frame loss for unicast flows

The previous experiment has shown that the frame loss at the MAC layer increases in the presence of simultaneous broadcast transmissions on a non-overlapping channel. In the following test, we investigate the impact of frame loss on the throughput of regular UDP and TCP flows. To generate UDP/TCP flows we used iperf [20], a popular tool for bandwidth monitoring applications.

The iperf client running on node C is generating UDP unicast frames (1472 bytes payload) at maximum rate towards the iperf server at node A on channel 3. Note that in this case, the unicast transmission implies acknowledged mode. At the same time node B is continuously broadcasting frames on channel 8. Similarly to the previous experiment the interference between the two channels disturbs the reception at node A. But in this case every lost frame triggers a retransmission from node C as the ACK timeout expires. This means a lower throughput for both UDP and TCP. However the impact on TCP is more severe as indicated in Figure 11. A significant throughput degradation starts to occur when ΔRSSI exceeds 10 dB. In Figure 11 we also include the curve for a reverse TCP flow, i.e TCP DATA sent from node A to node C. Interestingly the degradation effect is still observable since the interfering transmission by node B is affecting the reception of TCP ACK.

F. Experiment 6: Influence of interference load

In all previous experiments the interfering node B was continuously transmitting broadcast frames. This setting maximizes the channel usage and hence the amount of interference noise. In the next experiment we investigate the impact of interfering traffic load onto the TCP throughput. We repeated the previous experiment setting a ΔRSSI=17 dB and gradually reducing the packet rate transmitted by the interfering node B. In Figure 12 we plot the resulting throughput versus the

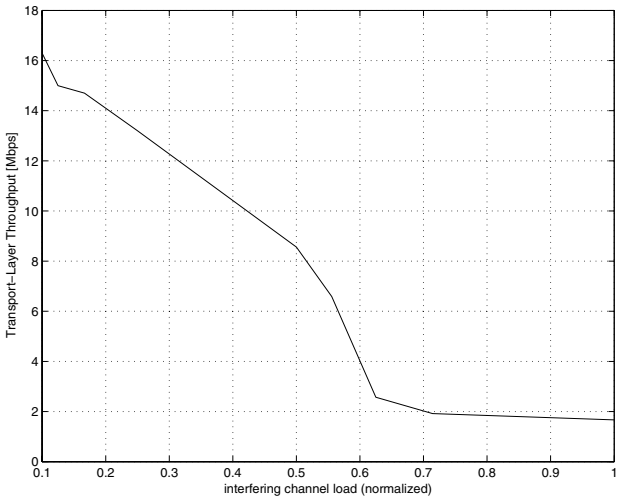

Fig. 12. Transport-Layer throughput vs. normalized channel-load using OFDM 54Mbps.

normalized load on the interfering channel. It can be seen that the relative loss in throughput is severe when the interfering load is above 0.6. For lower load values the relationship between the two variables is approximately linear.

G. Experiment 7: Preliminary experiment with TCP

In all the experiments considered so far the interfering source consisted of broadcast packets. In order to investigate the impact of interference in more realistic scenarios we perform a few preliminary experiments with point-to-point TCP flows. In order to eliminate any near-field radiation leakage effect by the circuitry of the two interface cards we split the central node A into two physically separated PCs, namely A_1 and A_2, each with its own antenna and connected via wired LAN. The distance between the antennae of A_1 and A_2 was set to 50 cm and the transmission mode was set

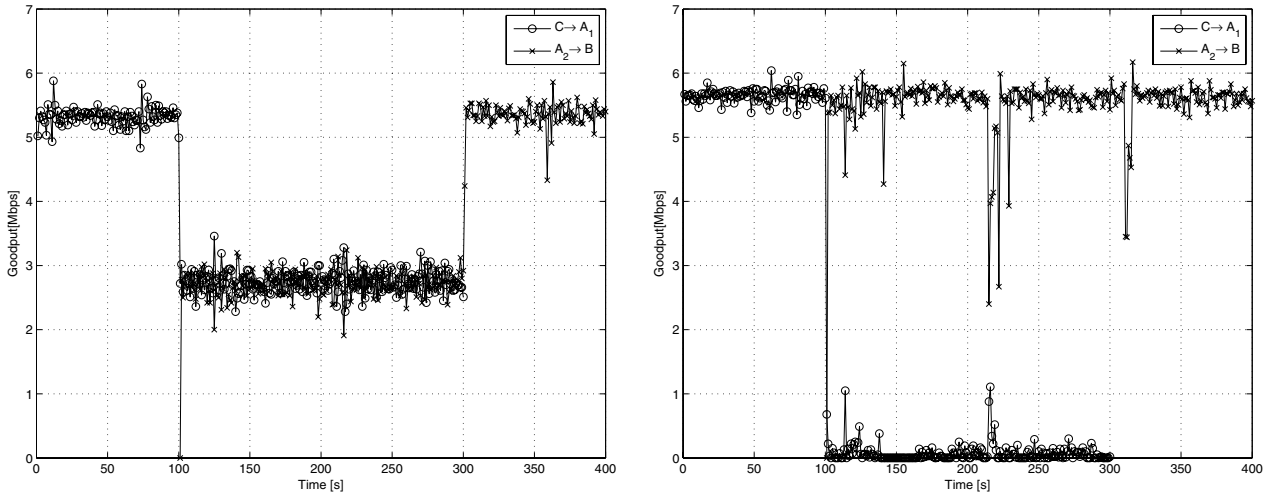

(a) Channels 3-3

(b) Channels 3-8

Fig. 13. Independent TCP flows.

to CCK-11Mbps. In a first experiment we launch an infinite TCP transfer from node C to A_1. The Maximum Segment Size (MSS) was set to 1480 Bytes. On the receiving node A_1 we measured the goodput in timebins of 1 sec. After 100 sec we launch a separate TCP download from A_2 to B and measure the goodput on the receiving node B. Finally after 300 sec we terminate the first TCP flow. In Figure 13 we report the performance in terms of goodput for both flows using two different channel configurations. In 13(a) the two flows coexist on the same channel: it can be seen that the two flows share the channel capacity equally, achieving an average goodput of around 2.7 Mbps which is roughly half of the goodput achieved by each of the flow in isolation (5.3 Mbps). In this scenario the Carrier Sensing mechanism in the contending stations is able to coordinate the access to the medium. Surprisingly, when we use different (non-overlapping) channels the overall performance gets *worse*. When the second flow is started the transmission of the long data packets by A_2 (towards B) interferes with the reception of the packets at A_1 (from node C) causing the corruption of the incoming packets. Note that in this case the received cross-channel power is below the Energy Detection Threshold, therefore the Carrier Sensing remains idle. The global effect is a severe starvation of the $B - A_1$ flow. This case is highly illustrative of the type of counter-intuitive effects that might take place in a real network due to interference across "non-overlapping channels".

In a second experiment we launch a single infinite TCP transfer from node B to C with node A acting as an IP relay: The IP packets received at A1 are passed via Ethernet to A_2 and forwarded via radio to C. In figure 14 we report the measured goodput for two different channel configurations. Similarly to above, the use of non-overlapping channels leads to inferior performance, with roughly 40% loss in goodput due to cross-channel interference that cannot be mitigated by the Carrier Sensing scheme.

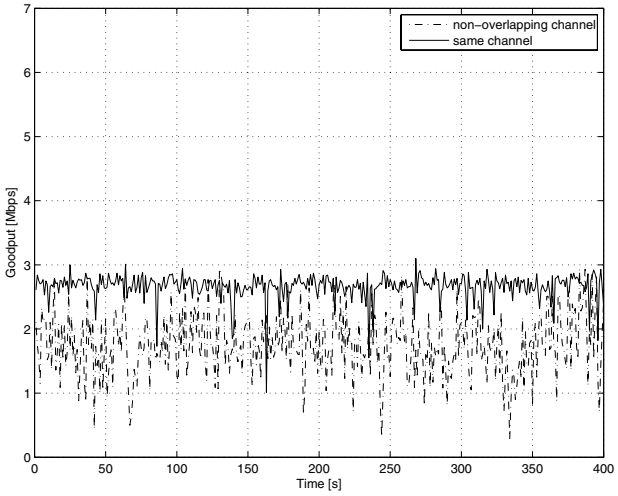

Fig. 14. Single TCP flow with IP relay.

VI. Conclusions

In this work we reported empirical evidence that so called "non-overlapping" channels in IEEE 802.11 are not completely decoupled. In all the experiments we have considered a physical distance between the nodes ranging from tens of centimeters to few meters. While the interference between non-overlapping channels might be negligible for larger distances, it identifies a serious problem for the design of wireless multichannel mesh networks. In fact such applications typically involve simultaneous transmission and reception by the same node on different channels. Our results suggest that current off-the-shelf IEEE 802.11 chipsets might not be ready to be integrated in a single box with few centimeters of antenna separation. These findings point to a serious mismatching between some routing and channel assignment schemes proposed in previous research works, that assume full separation

between non-overlapping channels. The near-far problem can be mitigated to some extent by refinements to the RF design, e.g. better filters. While such approach would certainly raise the cost of the equipment, it is unlikely to completely solve the problem.

With reference to multichannel mesh networks, and particularly to distributed self-organized systems, an interesting point for future research would be to devise new schemes for joint routing and channel assignment based on actual measurements of RSSI *in all channels*, including non-overlapping pairs.

In our future work we plan to investigate further the effect of channel interference on multiple concurrent TCP flows.

REFERENCES

[1] A Raniwala, T. Chiueh, "Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Network", in proc. IEEE INFOCOM, Miami, Florida, March 2005.

[2] M. Alicherry, R. Bhatia, L.E. Li, "Joint Channel Assignment and Routing for Throughput Optimization in Multiradio Wireless Mesh Networks", in proc. ACM MOBICOM, Cologne, Germany, August 2005.

[3] P.Kyasanur, N. Vaidya, "Capacity of Multichannel Wireless Networks: Impact of Number of Channels and Interface", in proc. ACM MOBICOM, Cologne, Germany, August 2005.

[4] B. Raman, "Channel Allocation in 802.11-based Mesh Networks", in proc. IEEE INFOCOM, Barcelona, Spain, April 2006.

[5] A. Mishra, V. Shrivastava, S. Banerjee, W. Arbaugh, "Partially Overlapped Channels Not Considered Harmful", ACM SIGMETRICS Performance, Saint-Malo, France, 2006

[6] J. So, N. Vaidya, "Multi-channel MAC for Ad-Hoc Networks: handling Multi-Channel Hidden Terminal using a Single Transceiver", in proc. MOBIHOC, 2004.

[7] P. Bahl, R. Chandra, J. Dunagan, "SSCH: Slotted Seeded Channel Hopping for Capacity: Improvement in IEEE 802.11 Ad-Hoc Wireless Networks", in proc ACM MOBICOM, 2004.

[8] R. Maheshwari, H. Gupta, S. Das, "Multichannel MAC Protocols for Wireless Networks". To appear in SECON 2006

[9] P. Kyasanur, N. Vaidya, "Routing and Interface Assignment in Multi-Channel Multi-Interface Wireless Networks", in proc. WCNC, 2005.

[10] S. Wu, C. Lin, Y. Tseng, J. Sheu "A new Multi-channel MAC Protocol with On-Demand Channel Assignment for Multihop Mobile Ad-hoc Networks", in proc. International Symposium on Parallel Architectures, Algorithms and Networks (ISPAN), 2000.

[11] M. Burton, "Channel Overlap Calculations for 802.11b Networks", http://www.cirond.com/white_papers/fourpoint.pdf

[12] K. Chebrolu, B. Raman, S. Sen,"Long-Distance 802.11b Links: Performance Measurements and Experience", in proc. ACM MOBICOM, Sep 2006, Los Angeles, USA.

[13] R. Draves, J. Padhye, B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks", in proc. ACM MOBICOM, 2004.

[14] S. Liese, D. Wu, P. Mohapatra, "Experimental characterization of an 802.11b wireless mesh network", in proc. ACM 2006 international Conference on Communications and Mobile Computing, Vancouver, British Columbia, Canada, July 2006

[15] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, L. Krishnamurthy, "Experimenting with a Multi-Radio Mesh Networking Tetsbed", in proc. WINMEE Apr 2005, Riva del Garda, Italy

[16] IEEE std 802.11, Wireless LAN Media Access Control (MAC) and Physical layer (PHY) Specifications," 1999 http://standards.ieee.org/getieee802/.

[17] A. Di Stefano, A. Scaglione, G. Terrazzino, I. Tinnirello, V. Ammirata, L. Scalia, G. Bianchi, C. Giaconia, "On the Fidelity of IEEE 802.11 Commercial Cards", in proc. First International Conference on Wireless Internet (WICON'05), Hungary, July 2005.

[18] Intel®PRO 2200BG Driver for Linux, http://ipw2200.sourceforge.net/

[19] packETH: a Linux GUI packet generator tool for ethernet, http://packeth.sourceforge.net/

[20] Iperf 1.7.0 - The TCP/UDP Bandwidth Measurement Tool, http://dast.nlanr.net/Projects/Iperf/