802.11b/g Link Level Measurements for an Outdoor Wireless Campus Network

Giuseppe Bianchi, Fabrizio Formisano, Domenico Giustiniano
Università di Roma "Tor Vergata", Italy*
{giuseppe.bianchi,fabrizio.formisano,domenico.giustiniano}@uniroma2.it

Abstract

Outdoor WLAN communication is envisioning an increasing interest, due to the massive emergence and deployment of outdoor wireless Mesh Networks. This paper provides a report on an extensive measurement campaign carried out in an outdoor Wireless LAN Campus scenario. A public domain driver, namely MADWiFi, has been modified in order to achieve a high (per-node and per-frame) measurement granularity, and in order to distinguish the different causes that contribute to link-layer errors. The main result of our experimental investigation is that, unlike 802.11b, which appears a robust technology in most of the operational conditions, 802.11g may lead to severe inefficiencies when employed in an outdoor scenario.

1 Introduction

Outdoor WLAN communication and networking has been recently boosted by the emergence of WLAN Mesh Networks [1, 2]. An IEEE 802.11 WLAN Mesh is a network where Access Points (AP) are interconnected through wireless links based on the 802.11 technology themselves. WLAN Mesh networks can provide Internet access over a wide area, typically but not necessarily outdoor, with minimal infrastructure expenditure. They have been recently deployed in both the commercial world by specific vendors (e.g. Tropos Network, Firetide, Nortel, BelAir, etc), community networks ([3]) and for academic research purposes (e.g. the MIT RoofNet [4], etc). 802.11 WLAN 802.11 Mesh networks are also being standardized in the frame of the recently activated 802.11s Task Group [5].

Literature work reports several outdoor measurement studies for the 802.11b technology. Among those, perhaps most known is the well documented experimental measurement campaign carried out for the MIT RoofNet Outdoor Mesh deployment [6]. This work shows that the great majority of outdoor Mesh Links seem to be characterized by an"intermediate" delivery probability ratio (quote from [6]). In other words, in most cases, an outdoor link quality does not result to be neither clearly bad nor clearly good and shows a marginal dependence on the SNR measured by the hardware WLAN interface. These results were explained by considering multi-path as the main reason of frame loss in outdoor channels.

Unlike 802.11b, to the best of our knowledge, little outdoor measurement work has been carried out for the widely diffused IEEE 802.11g standard. A goal of this paper is to fill this gap by providing an extensive experimental measurement campaign in an outdoor scenario employing 802.11g links, and by comparing the results with that achieved with 802.11b. The measurement results provided in this paper have been carried out in the terrestrial area covered by the University of Rome "Tor Vergata" Campus.

One might expect 802.11g to be more robust than 802.11b in the presence of a multi-path rich environment, thanks to its adoption of OFDM as modulation technology. However, our experimental results show that 802.11b links outperform 802.11g ones. This can be explained by the fact that 802.11g, despite its support for OFDM, is natively designed for indoor scenarios. Most likely the usage of a short PLCP preamble and the limited tolerance of the cyclic prefix employed by the 802.11g OFDM symbols, leads, in an outdoor scenario (where packet detection, timing synchronization and channel state estimation become critical factors), to an increased number of physical errors which dramatically impair performance.

On a methodological side, we remark that providing a convincing experimental measurement campaign is quite tricky. Indeed, there are multiple factors which may influence the results gathered from the components and devices employed in the experimental deployment, among which the native filtering and smoothing mechanisms provided by the network card drivers when collection of measurement information is aimed at. To this purpose, rather than bounding our investigation to driver-dependent performance figures, we have modified the software code of a well known public-domain driver (the open-source Multi-

*This work has been supported by the Italian Ministry for University and Research (MIUR) under the PRIN project TWELVE (http://twelve.unitn.it)

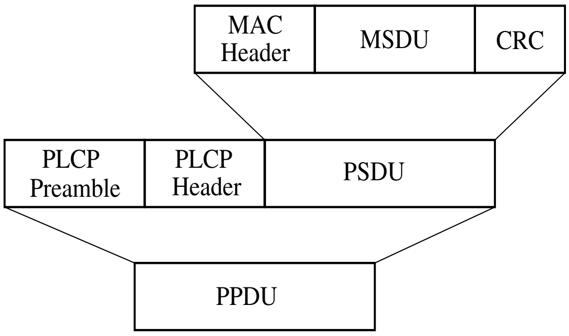

Figure 1. PPDU IEEE 802.11 frame format

band Atheros Driver for WiFi - MADWiFi [7] - for WLAN Atheros [8] chipsets) to collect high-granularity measurements (on a per-frame basis and at both transmitting and receiving sides), and to derive low-level performance figures such as per-frame SNR and per-frame error typology.

The rest of the paper is outlined as follows. Section 2 briefly reviews the IEEE 802.11 frame formats, and discusses the various kinds of frame errors. Section 3 details the driver modifications and the resulting statistic gathering/parsing methodology employed. The trial scenario is described in Section 4, and the main results of our extensive experimental campaign are discussed in Section 5. Finally, conclusions are drawn in Section 6.

2 IEEE 802.11 PHY

Goal of this section is to shortly review the basic concepts that will be used throughout this paper to critically discuss the experimental findings and assess the causes that lead to errored frames. We assume the reader to be familiar with both 802.11b and 802.11g PHY basics.

Figure 1 reviews the format of the transmitted PHY Protocol Data Unit (PPDU), which is common to each 802.11a/b/g PHY standard. The PPDU frame consists of a PLCP (Physical Layer Convergence Procedure) preamble, a PLCP header and a Physical Service Data Unit (PSDU). Each PSDU consists of the MAC header, the frame body (MSDU) and of a 32 bit Cyclic Redundancy Check (CRC). Extra bits (Tail/Pad bits), not reported in the Figure, are appended after the CRC when OFDM is employed as modulation scheme (802.11a/g).

The PLCP preamble is carefully designed to enable synchronization. IEEE 802.11g typically uses the ERP-OFDM mode for the PLCP format. With the ERP-OFDM preamble, it takes just 16 μs to train the receiver after first detecting a signal on the RF medium with respect to the 144 μs for IEEE 802.11b. Failure in frame detection and/or synchronization results in a physical layer (PHY) error.

The PLCP header carries the essential information needed by the receiver to properly decode the rest of the frame. This includes the frame size as well as the rate (modulation/coding scheme) at which the PSDU is transmitted (1, 2, 5.5 and 11 Mbps for the Barker/CCK 802.11b PHY; 6, 9, 12, 18, 24, 36, 48, and 54 Mbps for the OFDM 802.11a/g PHY). Note that the PLCP header is in any case transmitted according to a given (fixed) modulation/coding scheme (basic rate). Inability to properly decode the PLCP header (CRC16 failure in 802.11b, parity bit failure in 802.11a/g) results, again, in a PHY error.

MAC CRC check is performed only if the frame has been properly synchronized and the PLCP header is correctly received. Note that the presence of a CRC error notification on a received frame indirectly says that no PHY errors occurred in the PLCP. It is important to stress once again that the employed rate impacts the CRC error ratio (the higher the rate for a given SNR, the higher the CRC error probability), while it is irrelevant for what concerns PHY errors. This consideration is essential to properly understand the following discussion in section 5.

A closer look at the two 802.11b/g PHY technologies reveals that they significantly differ in terms of robustness to Inter-Symbol Interference (ISI). The IEEE 802.11b standard receiver is composed of a RAKE receiver and an equalizer with a certain (but implementation dependent) multipath robustness [9]. Anyway there exist a maximum delay T_m that may be tolerated by the RAKE receiver. With standard receivers, this value coincide with the symbol length,which is 1 μs for Barker codes (corresponding to a 1 Mbps and 2 Mbps physical rate) and 0.73 μs for CCK codes (i.e. 5.5 and 11 Mbps). On the other hand, in 802.11g OFDM, the PLCP header plus the PSDU is divided into blocks of size N=80, referred to as an OFDM symbol. The OFDM resistance to multi-path interference results from the increased symbol duration for each individual carrier (as compared to other modulation schemes with the same data throughput) and from the use of a cyclic prefix (i.e. a guard interval) preceding each OFDM symbol. The cyclic prefix helps in coping with Inter-Symbol Interference between pair of OFDM symbols since up to $\mu = 16$ over N =80 samples of each received OFDM symbol affected by multi-path interference can be discarded without any loss relative to the original information sequence. With μ = 16 and a sampling interval of $T_s = 1/(20 * 10^6)$, the corresponding maximum temporal delay T_m for which ISI is removed is $T_m \leq \mu T_s = 16 * 20MHz = 0.8\mu$s, in the assumption of standard feed-forward equalization techniques (as implemented in off-the-shelf basebands). However, the 802.11g multi-path tolerance is generally much smaller due to OFDM symbol level offset in the synchronization performed during the PLCP preamble. The impact of this issue is that the "effective" cyclic prefix is reduced from 16 down to 10 samples, which corresponds to $T_m = 0.5\mu sec$.

Figure 2. "Tor Vergata" map and outdoor links assessed

3 Statistics gathering approach

The test-bed implementation is based on 802.11 b/g wireless interfaces that use the fourth generation Atheros chipset based on AR5213 MAC/baseband [8], which is driven by the open-source Multiband Atheros Driver for WiFi (MADWiFi) [7]. This driver natively filters and smooths the internally collected statistics, and exposes to the upper layers only running averages. For example, in each AP, Signal-To-Noise Ratio (SNR)[1] measurements are not distinguished on the basis of the transmitting nodes, this being particularly critical when a node receives packets from multiple independent transmitting stations. In a Mesh network context, this would be especially critical. Since multiple links are active on a single network node, the driver would not distinguish the quality perceived on these different links, but would expose only an aggregate - hence meaningless - SNR. Moreover, statistics are averaged on subsequently received frames. In fact, the MADWiFi driver embeds an Exponentially Weighted Moving Average (EWMA) filter, whose default weight[2] is set to $\alpha = 0.1$ (i.e. its effect is somewhat analogous of taking a running average over the latest 10 samples). Since the Atheros chipset is indeed capable of providing per-frame measurements, we have modified the MADWiFi driver in the kernel space to by-pass the native filtering and smoothing mechanisms. Parsing and processing of the per-frame measurement samples have been thus delegated to our own software scripts developed in user space.

[1]More precisely, cards manufactures usually provide signal quality indicator (RSSI, Receiver Signal Strength Indicator), which is not linearly related to actual SNR. Indeed, the 802.11 standard itself does not define it in absolute way, but let vendors to implement it freely. For an extensive discussion refer to the enlightening white paper: Joshua Bardwell, *You Believe You Understand What You Think I Said...*, available online at: http://www.connect802.com/download/techpubs/2004/you_believe_D100201.pdf

[2]An EWMA filter (i.e. a single pole Infinite Impulse Response filter) is defined as $y_n = (1 - \alpha)y_{n-1} + \alpha x_n$, where y_n is the running average, x_n is the current measurement, and α is the filter weight.

Physical rate	1, 6, 11, 12, 24 Mbps
MSDU	1500 bytes
Physical preamble	144 μsec in 802.11b 16 μsec in 802.11g
ACK Timeout	48 μsec
Maximum number of retry	11
CWmin	31
CWmax	1023
DIFS	50 μsec

Table 1. 802.11 configuration values

Another main issue was the analysis of erroneous frame with CRC failure in the PSDU. If the field `rs_status` in the driver is set to zero, then the frame was correctly received; otherwise the error information is indicated. Particularly, if it denotes a CRC failure, the MADWiFi driver updates the number of CRC errors before discarding the frame. A driver modification has also allowed for passing on to the driver and elaborating field informations of frames with CRC errors, that would be normally immediately dropped.

4 Measurement scenario and link/traffic settings

Measurements have been carried out at the University Campus of Rome "Tor Vergata" (Figure 2). 8 point-to-point outdoor links (see Figure 2) have been set-up in distinct time frames. Measurements, with only one exception, have been independently carried out for both directions of the deployed links, thus providing a total of 15 link measurements (results not shown in this paper for reason of space indeed confirm that link performance may significantly differ in the two directions). These links have been set-up in correspondence with the various buildings we are planning to connect in our future Mesh deployment. Measurements have been held over the roof of the university buildings. The links considered in our study differ in terms of distance (ranging between 100 and 250 meters) and obstruction (from partially obstructed by surrounding obstacles to almost free-space). The height of the buildings varies from 8 to 15 meters. The APs deployed over the roofs are laptops running the Linux operating system with kernel version 2.6. The laptops are equipped with 802.11 b/g compliant cardbus driven by the AR5213 MAC/baseband chipset from Atheros via the MADWiFi driver and transmitting with a 5 dBi external antenna. The deployed antennas are omni-directionals. The

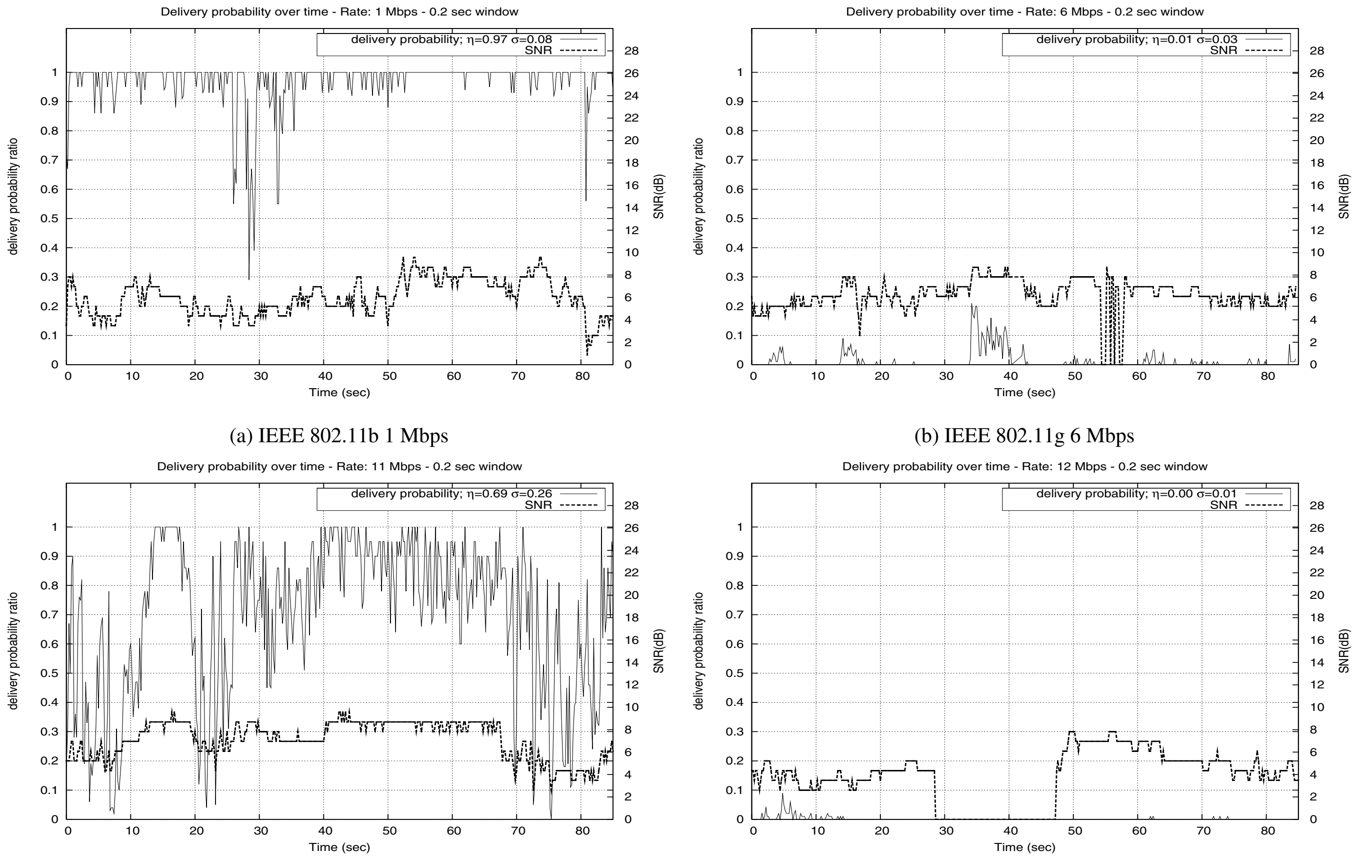

(a) IEEE 802.11b 1 Mbps

(b) IEEE 802.11g 6 Mbps

(c) IEEE 802.11b 11 Mbps

(d) IEEE 802.11g 12 Mbps

Figure 3. Delivery probability ratio of a typical Mesh Link

transmitted EIRP power is set to 20 dBm, to comply with the European regulation.

Traffic has been generated through a series of unicast[3] ICMP Echo requests of size 1500 bytes, disabling the corresponding ICMP Echo reply to avoid data traffic traveling in the opposite direction. Each measurement was performed over a 90 seconds period of time. The generation rate of ICMP packets was set to 100/second (totalizing a 1.2 Mbps expected network layer goodput).

In all experiments, the automatic rate selection mechanisms implemented into the driver has been disabled, and the link rate has been set to a fixed static value. The RTS/CTS mechanism was disabled. The MAC retry limit has been set to the value 11 (i.e. a frame is dropped after a first unsuccessful transmission plus further 11 unsuccessful consecutive retransmissions). For 802.11b baseband, we further fixed the physical PLCP preamble to 144 μsec (long preamble, sent at 1Mbps), while we used the widely accepted ERP-OFDM mode for 802.11g baseband (unfortunately DSSS-OFDM mode is not yet supported by Atheros WLAN devices). Table 1 summarizes the 802.11 settings.

[3] We believe that unicast traffic gives a more realistic and practical characterization of a wireless link, since it provides information on both forward data frame and backward ACK frame transmission, as experienced by a normal data transfer.

5 Experimental Results

5.1 Delivery Probability Ratio

We define the Delivery Probability Ratio (DPR) as the probability that a single asynchronous two ways handshake DATA/ACK is successfully concluded (i.e. regardless of the fact that a transmission or a retransmission is considered for a given frame). DPR and SNR plots versus the measurement time are reported in Figure 3. These figures have been obtained for a single outdoor link among the ones considered, and report the DPR in the four cases of a) 802.11b at 1 Mbps, b) 802.11g at 6 Mbps, c) 802.11b at 11 Mbps, d) 802.11g at 12 Mbps. The average DPR (indicated with η) and its standard deviation (indicated with σ), measured over the whole 90s measurement time are also reported in

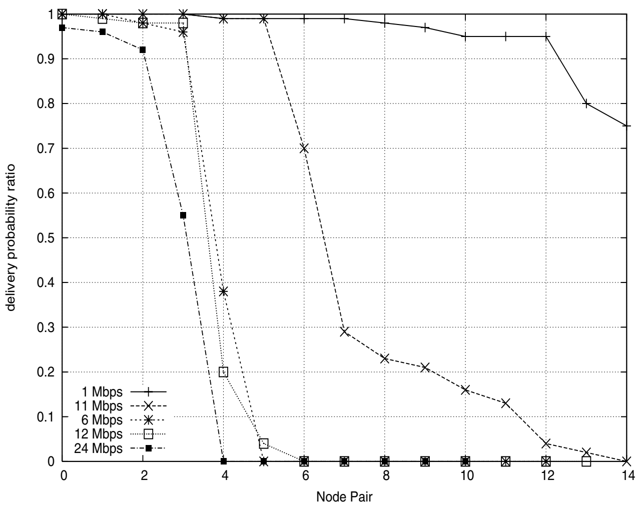

Figure 4. Delivery probability ratio as function of the node pair

the figures. Plot lines represent the running DPR (solid line) and SNR (dashed line) values, averaged over a 200 milliseconds interval. In comparable SNR conditions, results emphasize the ineffectiveness of 802.11g. We remark that 802.11g at 6 Mbps (case b) achieves dramatically worse performance (actually an average DPR close to 0!) with respect to 802.11b at its full 11 Mbps rate.

Since this poor performance might be incidental for a very specific link, we have repeated this experiment on all available 15 one-way links. Results are schematically reported in Figure 4, where, for each rate case, links are ordered from left (best performing) to right (worst performing). Indeed, the Figure confirms that, as expected, 802.11b 1 Mbps links provide in all cases a DPR close to 100% (again, we stress that this is on a per-transmission basis: this performance will be further improved by the MAC layer retransmission of errored frames). Perhaps surprisingly, this Figure confirms that for all considered links, 802.11b at 11 Mbps always outperforms 802.11g at 6 Mbps. Another very interesting result which emerges from the plot is the scarce dependence of the DPR versus to the 802.11g rate (only in the case of high modulation constellations, e.g. 24 Mbps / 16QAM, performance impairments start to emerge). This seems to imply that physical errors, at least for intermediate and bad quality links, largely predominate over CRC errors. Goal of the next section is to quantify this statement.

5.2 Received frames and causes of errors

Figure 5 aims at classify the possible frame error causes. Similarly to Figure 3, results have been obtained for a single outdoor link among the ones considered, and report performance in the four cases of a) 802.11b at 1 Mbps, b) 802.11g at 6 Mbps, c) 802.11b at 11 Mbps and d) 802.11g at 12 Mbps. We have selected a link where 802.11g results to be uneffective (i.e. a ”typical” link - see comment to Figure 4).

In the x-axis, the Figure reports the SNR at which a frame has been received[4]. The y-axis reports the percentage of successfully received frames over the total transmitted, as well as the percentage of errored frames classified into three possible error types: CRC error, PHY error and lack of received ACK.

As expected, the Figure shows that the percentage of successfully transmitted frames increases with the measured SNR. This is evident only in the 802.11b case, as in the 802.11g mode very few frames were correctly received and, as the Figure shows, this occurs only for the case of 6 Mbps and for an SNR close to 10 dB.

For what concerns the error distribution, let’s first analyze the 802.11b case, with specific reference to the 11 Mpbs case (in the 1 Mbps case a successful frame transmission occurs in most of the cases). We note that at small SNR, PHY and CRC errors appear dominant over missing ACKs (note that the SNR evaluation for an ACK error is based on the corresponding data frame value at the receiver side). As the SNR grows we see that both types of errors tend to reduce while, as expected, the amount of missing ACKs remains basically independent on the measured SNR at the receiver (and thus, in percentage, grows). An interesting remark is that, unlike what expected, the amount of PHY errors remains always greater than the CRC errors.

Let’s now analyze the 802.11g case. Here the error distribution is very different, and there is a clear predominance of PHY errors, with CRC errors being a fairly small percentage. Comparison between the 802.11g 6 Mbps case and the 802.11g 12 Mbps case shows why performance are loosely dependent on the 802.11g rate (which is common to all considered links, as shown in Figure 4). In fact a rate decrease “transforms” CRC errors into successfully delivered packets, but clearly rate variations do not affect the probability that a PHY error occurs. The obvious conclusion is that the 802.11g PHY layer, and especially the shorter physical PLCP preamble of IEEE 802.11g and the small multipath tolerance of the cyclic prefix implemented in the IEEE 802.11g OFDM symbol, are the main limiting factors for the exploitation of 802.11g in outdoor scenarios. The (non trivial) quantification of the different impact of preamble

[4]This was not trivial to obtain from our measurement results. In fact we first note that a PHY error do not result in a received frame and thus in a corresponding SNR measurement at the receiver. In order to properly compute the number of frames which have been affected by PHY errors we thus needed to correlate the statistic logs taken at the transmitter with the ones taken at the receiver. From this comparison we can determine how many transmitted frames have not been received and thus have been affected by PHY errors. Second, since we did not have an explicit SNR measurement for PHY-errored frames, we computed this value by averaging the SNR measured at the received over 200 ms time windows and assigning this average SNR to all the frames with PHY errors in the same window.

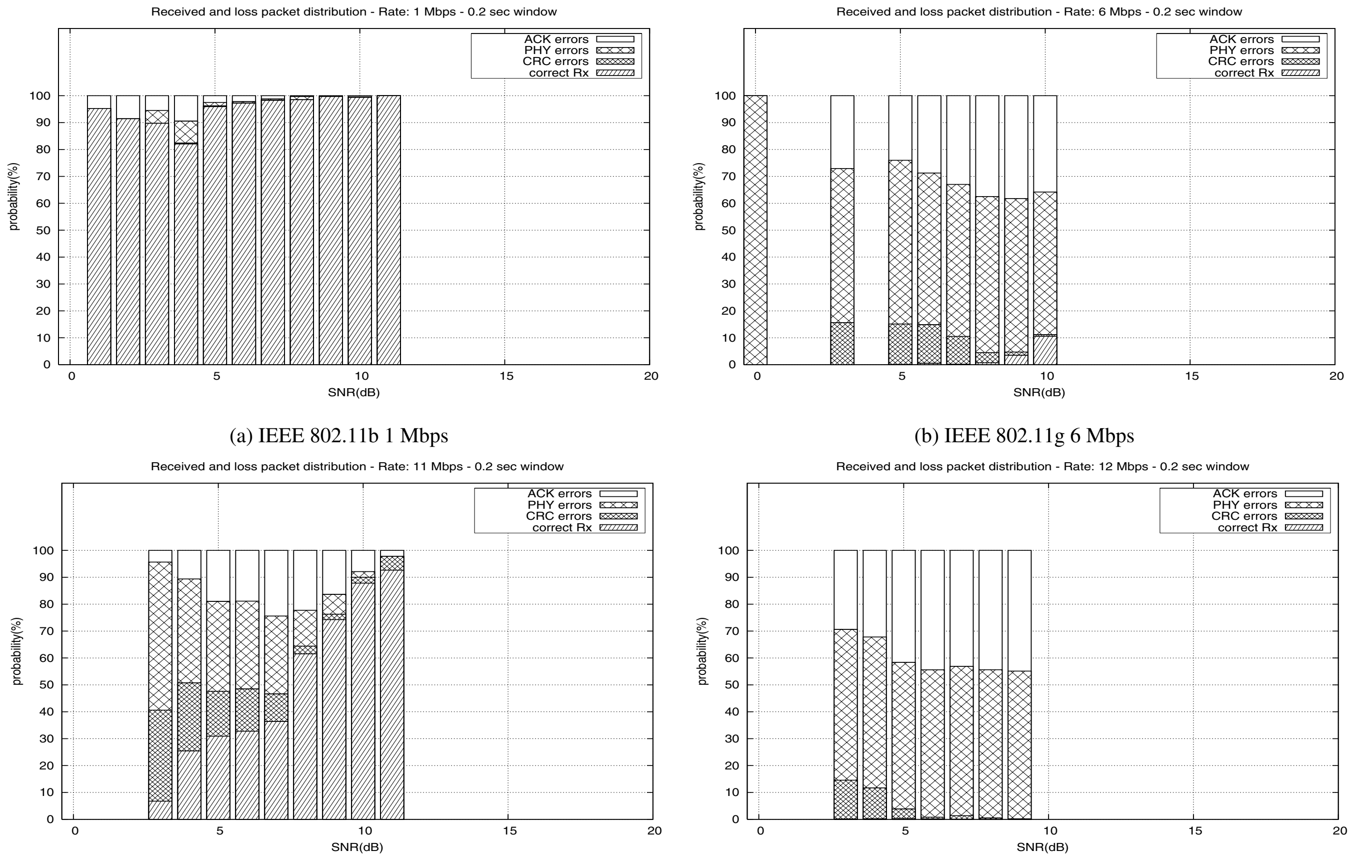

(a) IEEE 802.11b 1 Mbps

(b) IEEE 802.11g 6 Mbps

(c) IEEE 802.11b 11 Mbps

(d) IEEE 802.11g 12 Mbps

Figure 5. Typical results of the correctly received and loss packet distribution

and OFDM cyclic prefix on the performance degradation is object of current research work.

6 Conclusion

In this paper, we have documented an extensive measurement campaign carried out in a WLAN outdoor campus scenario. Both 802.11b and 802.11g links have been considered. Per-frame measurements have been collected and analyzed to quantify the link performance and the detailed distribution of the frame errors for both IEEE 802.11b and IEEE 802.11g mode. Our results show that 802.11g results poorly performing (compared to 802.11b) in an outdoor scenario, for reasons mainly imputable to PHY errors, namely short PLCP preamble and limited tolerance of the cyclic prefix.

References

[1] Ian F. Akyildiz, Xudong Wang, Weilin Wang, *Wireless Mesh Networks: a survey*, Computer Networks, 47 (2005), pp. 445-487.

[2] Raffaele Bruno, Marco Conti, Enrico Gregori, *Mesh networks: commodity multihop ad hoc networks*, Communications Magazine, IEEE, March 2005.

[3] Rudi van Drunen, Jasper Koolhaas, Huub Schuurmans, and Marten Vijn, *Building a wireless community network in the Netherlands*, USENIX/Freenix Conference, June 2003.

[4] J. Bicket, D. Aguayo, S. Biswas, R. Morris, *Architecture and evaluation of an unplanned 802.11b Mesh network*, MobiCom 2005, Cologne, Germany.

[5] Wi-Mesh Alliance, http://www.wi-mesh.org

[6] D. Aguayo, J. Bicket, S. Biswas, G. Judd, R. Morris *Link-level measurements form an 802.11b Mesh Network*, SIGCOMM 2004, Portland, Oregon USA.

[7] http://madwifi.org

[8] http://www.atheros.com/pt/AR5004XBulletin.htm

[9] J. G. Proakis: *Digital Communications*. McGraw-Hill, Inc., Third edition, 1995.