This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings.

An explanation for unexpected 802.11 Outdoor Link-level Measurement Results

Domenico Giustiniano*, Giuseppe Bianchi*, Luca Scalia** and Ilenia Tinnirello**
*Università degli Studi di Roma - Tor Vergata, Dipartimento di Ingegneria Elettronica, Italy
**Università di Palermo, Dipartimento di Ingegneria Elettrica, Elettronica e delle Telecomunicazioni, Italy

***Abstract*—This paper provides experimental evidence that "weird"/poor outdoor link-level performance measurements may be caused by driver/card-specific antenna diversity algorithms unexpectedly supported/activated at the WLAN transmitter side. We focus our analysis on the Atheros/MADWiFi card/driver case, and we observe that the transmit antenna diversity mechanisms remain by default enabled when the available antennas are not homogeneous in terms of gain or, even worse, when only a single antenna is connected. This may cause considerable performance impairments (large frame loss ratio), in conditions frequently encountered in outdoor link deployments. The negative impact of transmit antenna diversity is not limited to the transmission of broadcast frames (where a cyclic shift between the "two" assumed antennas is performed), but under certain circumstances it can severely affect the delivery of unicast frames as well, and despite the fact that in this case the ACK receptions may provide a feedback about the best receiving antenna. While, as obvious, driver developers are expectedly fully aware of the existence of such mechanisms, we believe that the scientific research community has very limited awareness of the implications these mechanisms have on the measured link-level performance. Indeed, to the best of our knowledge, ours is the first research paper which explicitly raises this issue.**

I. Introduction

With the boost of 802.11-based wireless Mesh networks [1], and with the further adoption of 802.11 as technology for long-distance links, the experimental performance assessment of outdoor Wireless LAN deployments [2]–[6] has become increasingly important. Indeed, 802.11 outdoor links may exhibit critical performance in terms of achievable link quality. For instance, [2] shows that most of the links in an outdoor 802.11b Mesh deployment are characterized by an intermediate delivery probability ratio, i.e. 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 cause of frame loss in outdoor channels. For longer-distance links (up to 37 km in length and with highly directive antennas), the experimental assessment of 802.11b links was carried out in [3]. Here, the error rate was instead shown to be a sharp function of the SNR, as expected from theoretical results. More recently, outdoor measurement work was carried out also for the 2.4 GHz 802.11g technology [4], which was found to underperform 802.11b, due to a higher than expected error rate at the physical layer.

Because of the availability of open-software driver implementations and of their high configuration/customization possibilities, two WLAN card brands are being mostly employed by the research community: i) 802.11b Prism NICs equipped with the HostAP driver (e.g., used in [2], [3]), and ii) 802.11a/b/g Atheros NICs with the MADWiFi driver [7] (e.g., used in [4], [6], [8]–[13]). Specifically, this latter card/driver pair is undoubtedly used in the majority of the most recent works and nowadays can be somehow considered as the "de-facto" standard for 802.11 experimental activities. With such an amount of researchers relying on such equipments, it is of paramount importance to understand whether these card/driver pairs do have operation modes which might eventually (and unexpectedly) impact the experimental insights derived.

The key finding of this paper is that, for the Atheros/MADWiFi driver/card pair, the implemented transmit antenna selection (diversity) algorithms appears to be a primary cause of the poor frame delivery probability experienced in some outdoor link conditions. To this purpose, we recall that the MADWiFi driver allows to support two antenna ports and to dynamically choose the operating one on the basis of a simple (if compared with literature proposals such as [14]–[17]) transmit antenna selection algorithm. The algorithm, which is enabled by default, aims to improve the link-level performance by appropriately select the transmit antenna which correspond to the best signal path experienced at the receiver. Now, when only a single antenna is connected (a frequent configuration choice in experimental trials), or if one of the two antennas is not appropriate (as in our experiments, where the second antenna was for 5 GHz 802.11a transmissions), the transmit diversity algorithm remains enabled. Hence, the transmitter works with two highly heterogeneous antennas: a good one (the proper antenna connected) and a very poor one (the low-gain - or even missing - one).

As shown in the rest of the paper, whenever one antenna works deterministically worse than the other one, the dynamic antenna selection schemes may have dramatic consequences. These are most evident in the case of broadcast transmission, as the MADWiFi transmit diversity algorithm appears to cyclically (periodically) switch between the two antennas, thus resulting in half of the frames being likely lost. A more subtle situation occurs for unicast transmissions. For such frames, the algorithm's operation (actually, as discussed in [18], a distinct algorithm residing in the Hardware Abstraction Layer provided by the card manufacturer) is apparently smarter, as it appears to exploit the feedback provided by the reception of ACKs. Nevertheless, we show that under certain channel conditions, a substantial switching between antenna ports can also occur with unicast frames, thus leading again to a significant performance degradation.

For reasons of complexity and space availability, this paper focuses on “just” the specific case of Atheros/MADWiFi. However, preliminary results presented in a companion work [18] imply that a similar problem may also emerge also in the case of the Intel/ipw2200 card/driver. Hence, we believe that raising awareness on the existence of such possibly unexpected driver operation can be extremely useful for the WLAN networking community involved in experimental activities. In fact, unlike the developers' community[1], we believe that most of the scientific research community involved in experimental activities is still largely unaware of the possible strong dependency of the measured WLAN performance on some quite specific algorithms implemented in the driver (such as the transmit diversity one here dissected). We argue that lack of appropriate knowledge of the performance effects induced by an unexpected driver/card operation can easily mislead and affect the conclusions that can be drawn from an experimental campaign.

II. Measurement Scenario

The reference scenario of our experimental study is the outdoor wireless network of the University Campus of Rome Tor Vergata. The network is composed of 9 point-to-point outdoor links, differing in terms of distance (ranging between 50 and 205 meters) and obstruction (from partially obstructed by surrounding obstacles to almost free-space). Owing to the well known link asymmetry (see e.g., [19], and indeed verified also from our results), measurements have been independently carried out for both directions of each deployed link, thus providing a total of 18 link measurements. Each link has been tested in a separate time frame, with all the other links inactive to avoid RF (Radio-Frequency) interference.

The wireless nodes deployed over the campus roofs were net4826 Soekris boards [20], with a Pyramid Linux distribution [21] running a 2.6.18 kernel. Such boards have been equipped with AR5212 Atheros 802.11 a/b/g compliant mini-pci cards presenting two antenna ports, to which we connected two rubber duck external omni-directional (on the horizontal plane) antennas, devised respectively for 802.11b/g and 802.11a transmissions. The first antenna had a gain of 5 dBi at 2.4 GHz, and the second one had a gain of 3 dBi at 5 GHz. The card driver was a customized version of the MADWiFi one, extended to allow statistic collection at both transmitter and receiver sides, and their subsequent off-line cross-correlation to reveal specific per-frame loss events not natively provided by the MADWiFi driver (such as PHY errors - details about the measurement methodology can be found in [4]). For the measurement results presented in this paper, we are mostly interested in the Delivery Probability Ratio (DPR) and per-frame measured RSSI (Receiver Signal Strength *Indicator*) values. The DPR is the probability that a transmitted frame is successfully received. In the case of unicast frames, the DPR is measured irrespective of retransmissions, i.e. a retransmitted frame is counted as an independent transmission (in other words, in the unicast case, the DPR is defined as the probability that a single asynchronous two ways handshake DATA/ACK is successfully concluded. In the case of broadcast frames, no ACK is transmitted and here, unlike the unicast case, the DPR is measured at the receiver as the probability that the DATA frame is correctly decoded. If ambiguity occurs, to distinguish the DPR measured for unicast frames from that measured for broadcast frames we will use for this latter the notation DPR-RX (DPR at the receiver). Regarding RSSI, we recall that it is an estimate of the signal power at the receiver and is provided by each manufacturer on a proprietary scale. Atheros NICs measure RSSI in terms of SNR referred to the noise floor power. Thus, in what follows, we will simply refer to SNR. To obtain per-frame SNR measurements, we disabled the smoothing filter natively provided by the driver. For convenience of plotting (and for further elaborations as shown when discussing the broadcast measurement cases), we provided a custom smoothing on the collected measures. Unless otherwise stated, each plotted sample is obtained as the average taken over consecutive non-overlapping time window set to the default value of 200 msec[2].

Links were tested through the generation of ICMP *echo requests*, with the corresponding ICMP *echo reply* disabled 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 frames was set to 100 frames per second (i.e. approximatively up to 1.3 Mbps goodput). The ICMP datagram size was set to the unusual value of 1601 bytes, to easily detect, during post-processing, possibly interfering frames (indeed a very rare occurrence - in which case we have discarded the experiment). In addition to this somehow naive interference control, durig the trial set-up we have assessed, through a spectral analyzer, the interference level by evaluating the overall adjacent/co-channel interference in absence of our link transmissions. Interfering signals have been found on some link just around the 2.47 GHz frequency: based on this we selected a transmission channel (namely, channel five) far away from this frequency. As such we can safely exclude RF interference from being a cause of frame losses in our measurements. In all experiments, the automatic rate selection and the RTS/CTS mechanism have

[1] Indeed, after having spent a considerable amount of time/effort to unveil and understand, on our own, the causes underlying the “weird” measurement results presented in this paper, we found out a posteriori that a few notes and/or trouble tickets related to the problems emerging in the broadcast case - see e.g., http://madwifi.org/changeset/1430 - had been actually issued on the MADWiFi developers' site. Most likely, as it happened in our own case, this, as well as other warnings, it has remained unnoticed by other researchers actively involved in WLAN experimental activities. In any case we are not yet aware of warnings related to the unicast case, even in the developer's community (probably because the unicast algorithm resides in the Hardware Abstraction Layer - HAL - which is separately provided by Atheros, and not part of the MADWiFi specification).

[2] For unicast frames, we have verified that the smoothing time scale does not affect the measurement results. Furthermore, the selected window size guarantees both a sufficient high granularity and number of data per window.

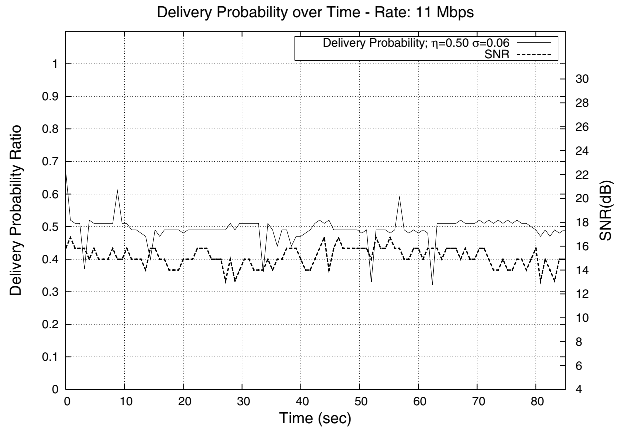

Fig. 1. DPR-RX and link quality for a selected link - 0.8 sec windowing

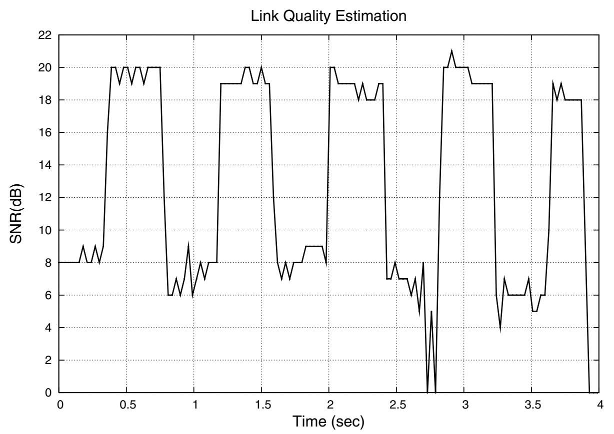

Fig. 2. Link quality for the same selected link - 40.96 msecs windowing

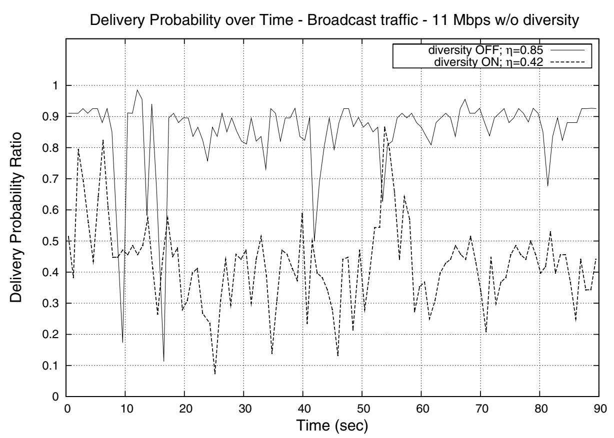

Fig. 3. Impact of transmit diversity on broadcast traffic

been disabled, and the MAC retry limit has been set to the fixed value 7.

III. Experimental results

The large amount of tested links (18) gave us a quite large base of different channel conditions (in terms of resulting DPR and SNR and link asymmetry, etc). In what follows, for reasons of space, we present and discuss results regarding a subset of links where the anomalies induced by the driver/card transmit diversity algorithms are most evident.

A. Transmit diversity on broadcast data frames

The following results are presented for 802.11b at 11 Mbps rate, but the same results have been found for other rates [22][3]. Fig. 1 reports two performance metrics, gathered in a 90-seconds experiment, for a given outdoor link. The first metric is the time-varying DPR-RX. The label in the figure also indicates the DPR-RX mean value (η=0.50) and the standard deviation (σ=0.06) taken along the whole measurement time. The second performance figure is the SNR. In the specific case of Fig. 1, the DPR-RX as well as the SNR were averaged over 800 ms windows. The figure suggests that the considered link exhibits an intermediate performance, with 50% of the frames being corrupted despite of the stable SNR samples (mostly in the range from 13 to 16 dB).

Fig. 2 replots the SNR values obtained by the *same* experiment, but in this case averaged with a time window set to 40,96 ms (40 times the IEEE 802.11 1.024 ms Time Unit - TU). This much shorter time window reveals a periodic fluctuation of the SNR. In particular, it shows that the measured SNR switches every $\approx$ 400 ms (more precisely, 400 TU, i.e., 4 beacon intervals) from a high value to a much lower value (about 10-15 dB less).

The almost perfect 50% DPR-RX highlighted in figure 1 is thus readily explained as the average between the almost 100% DPR-RX experienced during the ”good” periods (thanks to the SNR in the order of 20 dBs, above the receiver threshold), and the close to 0% DPR-RX experienced in the ”bad” periods (owing to a very low SNR in the order of 6-8 dBs). We remark that, by fixing the link rate, a large amount of outdoor links will happen to be in such intermediate conditions, whenever the SNR fluctuates above and below the receiver sensitivity.

By changing the link under test, we expect that the resulting DPR-RX may hence change, although remaining in a sort of intermediate performance state, based on the actual SNR values experienced in both periods. Even if the difference in the SNR between the ”good” and ”bad” periods remains constant in the order of 10-15 dBs – as duly verified by different experiments – the SNR experienced in the ”good” state may not be sufficient to guarantee a 100% DPR-RX. This is experimentally confirmed in figure 3 (label “diversity ON”) which shows results for a link experiencing an about 42% average DPR.

Being aware of the transmission diversity algorithm implemented in the MADWiFi driver, it is straightforward to justify the measured data as induced by the abrupt change in the transmission power resulting from the periodic switching between the antenna ports. As a confirmation of the fact that this ”weird” measurement plot is actually caused by the transmit diversity algorithm, we looked inside the MADWiFi documentation for a way to disable it. In particular, we found the

[3]We remark that in such a preliminary work we had not yet discovered the existence of the transmit diversity mechanism here addressed, and thus, with no other convincing explanation available, we have erroneously attributed the outcomes of our findings to the presence of some proprietary power control mode implemented in the card at NIC level to save energy consumption and increase battery life.

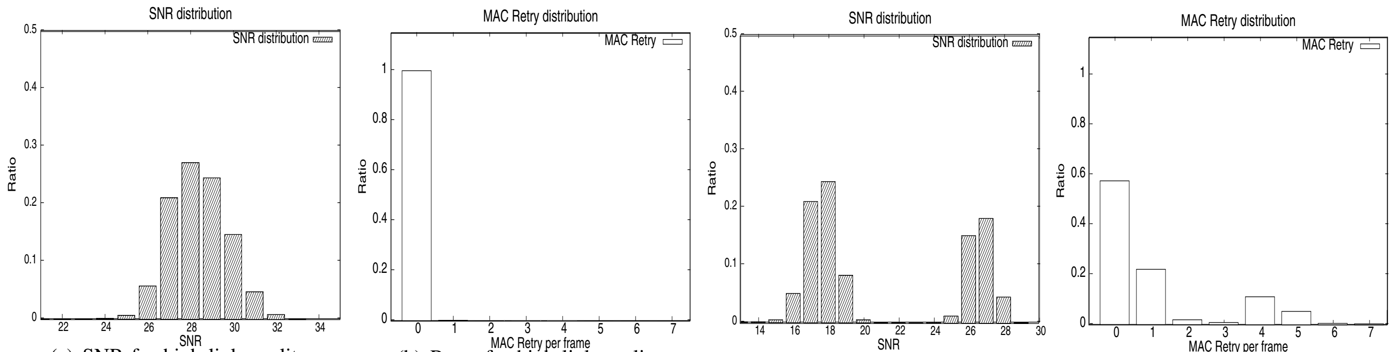

(a) SNR for high link quality (b) Retry for high link quality (c) SNR for low link quality (d) Retry for low link quality

Fig. 4. Retry and SNR distribution at the receiver

sysctl dev.wifi0.txantenna=1 or *sysctl dev.wifi0.txantenna=2* settings to force, respectively, the use of the first or second antenna connected to the card (being *0* the default setting which enables transmit diversity). Figure 3 shows the resulting DPR performance in the case of transmit diversity disabled, and experimentally confirms that the resulting DPR (85%) is about the double of that experienced with diversity enabled (42%).

These results allow to dissect the transmit diversity algorithm's operation in the presence of broadcast frames, as well as its rationale: since no feedback (in terms of received ACK) is available for broadcast frames, and since different end-users may experience different channel conditions, the transmitter has no way to assess which is the best available radio channel among the two available. Thus the most obvious strategy is to periodically switch between the two antennas to achieve a sort of average channel conditions. This results in a poor strategy when one of the two antennas has a persistently lower gain (such as in our case when the 5 GHz antenna was used for 2.4 GHz transmissions).

B. Transmit diversity on unicast data frames

A completely different behavior was detected for unicast frames, but as discussed below also in this case we realized that transmit antenna diversity was playing a significant role.

First of all, on several links where the broadcast frames were showing intermediate performance levels, the link quality measured with unicast traffic was good (DPR close to 100%), thus excluding cyclic attenuation phenomena like the ones revealed for broadcast frames. However, an anomaly was shown to emerge on lower quality links, namely links where the frame loss ratio was not negligible.

Figure 4 compares the SNR distribution measured at the receiver (figure 4(a)) and the corresponding retry distribution (figure 4(b)) for a high quality link, versus the SNR (figure 4(c)) and the retry (figure 4(d)) distributions for a low quality link. The SNR distribution is computed by counting the occurrences of received frames with a given SNR value. The retry distribution is computed as the probability that a frame retransmitted for the i-th time (with i ranging from 0 - first transmission attempt - to 7 - last transmission attempt after which the frame is dropped) is successful.

For the *high* quality link, we see, from figure 4(b) that all the frames are successfully received at the first transmission attempt. We also see, from figure 4(a), that the SNR distribution is, as expected, Gaussian shaped and centered at about 28 dBs. Surprises emerge in the case of the *low* quality link. Here, the SNR distribution plotted in figure 4(c) appears to follow a bi-modal shape, apparently suggesting that frames are transmitted according to two different transmission power levels separated of about 10 dB. Even more interesting is the retry distribution reported in figure 4(d), which shows a non monotonic behavior, and specifically suggests that the probability to receive a frame transmitted for the first or second time, as well as fifth or sixth time, is greater than the probability to receive it during the third or fourth transmission (or seventh/last transmission).

Assuming that, again, transmission diversity is the cause for such an operation, it is straightforward to conclude that the specific algorithm run by the card/driver consists in switching from an antenna to the other when two consecutively transmitted frames are lost (i.e., no return ACK is received). Note that this algorithm is smarter than the one employed for the broadcast frames, as it takes advantage of the feedback provided by the ACK frames. Moreover, this algorithm justifies why, in good channel conditions, no antenna switching occurs (as no frame losses emerge, the algorithm remains stuck to the antenna that provides good channel conditions). However, this algorithm shows significant weaknesses with low link quality: since in such conditions two consecutive losses can occur even when the "good" antenna is chosen, the algorithm frequently switches to the "bad" antenna, thus further reducing the delivery performance.

The experimental confirmation that this operation is induced by the enabled transmit diversity is provided in figure 5, which compares, for a same link, the retry distribution of successful frames with transmit diversity activated (figure 5(a)) and disabled (figure 5(b)). In this latter case the retry distribution is regular and monotonic, as intuitively expected. The DPR performance comparison between the case of diversity enabled and disabled is reported in table I, for two selected links and for two link rates, showing that disabling transmit diversity leads to a significant performance improvement in all the

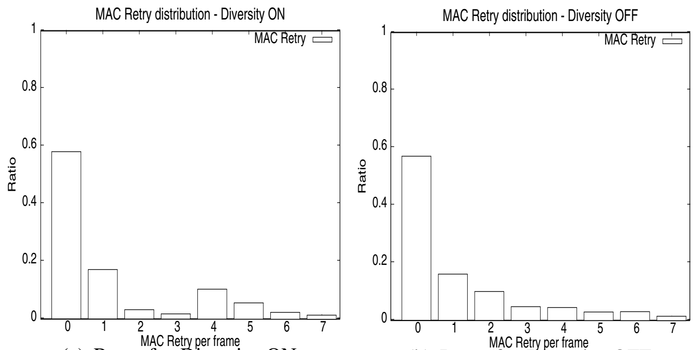

(a) Retry for Diversity ON (b) Retry for Diversity OFF

Fig. 5. Retry distribution

	Link 1		Link 2	
	1 Mbps	11 Mbps	1 Mbps	11 Mbps
Diversity ON	0.36	0.10	0.48	0.68
Diversity OFF	0.51	0.21	0.92	0.80

TABLE I
IMPACT OF DIVERSITY ON TWO SELECTED LINKS (UNICAST TRAFFIC)

considered cases.

IV. CONCLUSIONS

In this paper we raise awareness on the fact that a WLAN driver/card pair widely used by the research community, namely MADWiFi/Atheros, employs a transmit antenna diversity scheme which is shown to significantly affect link level performance under specific circumstances. In fact, this scheme is enabled by default even when the 802.11 station is equipped with either a single antenna or with two non homogeneous antennas (e.g. a 2.4 GHz antenna and a 5 GHz one as in our trial).

We have presented experimental results which show that the delivery of broadcast frames can be significantly affected by such diversity mechanisms, leading to a situation where all links experience a sort of intermediate (neither good nor bad) state. Such diversity mechanisms appear to affect also the unicast frame delivery, although in this case the resulting performance impairment is more complex to predict: it depends on the native quality of the deployed link, and becomes critical only when low quality links are considered.

To a more general extent, we believe that the importance of this paper stays in its attempt to raise awareness on these (and possibly other, still to be disclosed) issues regarding unexpected driver/card operation modes. We deem possible that other researchers in our field may be mis-interpreting their experimental findings simply because of lack of knowledge of the actual (versus the theoretical) operation of the equipments used in the trials. This is especially critical as, to the best of our knowledge, ours is the first research paper that raises such an issue (which, as preliminarily shown in [18], seem to emerge also in other driver/card cases), and as such it is likely that a significant fraction of the research community might not yet be duly aware of the related criticalities in terms of reliability of the measurement results.

REFERENCES

[1] Ian F. Akyildiz, X. Wang, W. Wang, *Wireless Mesh Networks: A survey*, Computer Networks Journal (Elsevier), vol. 47, no. 4, pp. 445-487, March 2005.
[2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, R. Morris, *Link-level measurements from an 802.11b Mesh Network*, in Proc. of ACM Sigcomm 2004, pp. 121-132, Portland, OR, USA.
[3] K. Chebrolu, B. Raman, and S. Sen, *Long-Distance 802.11b Links; Performance Measurements and Experience*, in Proc. of ACM Mobicom 2006, pp. 74-85, Los Angeles, CA, USA.
[4] G. Bianchi, F. Formisano, D. Giustiniano, *802.11b/g Link Level Measurements for an Outdoor Wireless Campus Network*, in Proc. of IEEE WoWMoM 2006 EXPONWIRELESS Workshop, pp. 525-530, Niagara Falls, ON, Canada.
[5] J. Camp, J. Robinson, C. Steger, and E. Knightly, *Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network*, in Proc. of ACM MobiSys 2006, pp. 96-109, Uppsala, Sweden.
[6] A. Sheth, S. Nedevschi, R. Patra, S. Surana, L. Subramanian, E. Brewer, *Packet Loss Characterization in WiFi-based Long Distance Networks*, in Proc. of IEEE Infocom 2007, pp. 312-320, Anchorage, AL, USA.
[7] The MadWiFi driver, http://madwifi.org
[8] The Berlin RoofNet project http://www.berlinroofnet.de
[9] A. Miu, H. Balakrishnan, C. E. Koksal, *Improving Loss Resilience with Multi-Radio Diversity in Wireless Networks*, in Proc. of ACM Mobicom 2005, pp. 16-30, Cologne, Germany.
[10] G. Judd, and P. Steenkiste *A Simple Mechanism for Capturing and Replaying Wireless Channels*, in Proc. of ACM Sigcomm 2005 E-WIND Workshop, pp. 58-63, Philadelphia, PA, USA.
[11] S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, and J. Deng, *Methods for Restoring MAC Layer Fairness in IEEE 802.11 Networks with Physical Layer Capture*, in Proc. of ACM REALMAN 2006, pp. 7-14, Florence, Italy.
[12] T. Ireland, A. Nyzio, M. Zink and J. Kurose *802.11g Long-distance Measurement: Antenna Placement and Orientation* in Proc. of ICST WiOpt 2007 WinMee Workshop, Limassol, Cyprus.
[13] C. M. Cheng, P. H. Hsiao, T. H. Kung, D. Vlah, *Adjacent Channel Interference in Dual-radio 802.11 Nodes and Its Impact on Multi-hop Networking*, in Proc. of IEEE Globecom 2006, pp. 1-6, San Francisco, CA, USA.
[14] S. Sanayei and A. Nosratinia, *Antenna Selection in MIMO Systems*, IEEE Communications Magazine, vol. 42, no. 10, pp. 68-73, October 2004.
[15] A. K. Miu, H. Balakrishnan, C. E. Koksal, *Improving loss resilience with multi-radio diversity in wireless networks*, in Proc. of ACM Mobicom 2005, pp. 16-30, Cologne, Germany.
[16] E. Vergetis, E. Pierce, M. Blanco, R. Guerin, *Packet-Level Diversity - From Theory to Practice: An 802.11-based Experimental Investigation*, in Proc. of ACM MOBICOM 2006, pp. 62-73, Los Angeles, CA, USA.
[17] Chen-Mou Cheng et Al., *Transmit Antenna Selection Based on Link-layer Channel Probing*, in Proc. of WoWMoM 2007 EXPONWIRELESS Workshop, pp. 1-6, Helsinki, Finland.
[18] D. Giustiniano, I. Tinnirello, L. Scalia, A. Levanti *Revealing Transmit Diversity Mechanisms and their Side-Effects in Commercial IEEE 802.11 Cards*, to appear in Proc. of IEEE QoS-IP 2008, Venice, Italy.
[19] D. S. J. De Couto, D. Aguayo, J. Bicket, R. Morris, *A High-Throughput Path Metric for Multi-Hop Wireless Routing*, in Proc. of ACM Mobicom 2003, pp. 134-146, San Diego, CA, USA.
[20] Soekris engineering, http://www.soekris.com/net4826.htm
[21] Pyramid Linux OS, http://pyramid.metrix.net/
[22] D. Giustiniano, G. Bianchi, *Broadcast Link Quality Measurements in 802.11 Networks*, in Proc. of IEEE WoWMoM 2007 EXPONWIRELESS Workshop, pp. 1-6, Helsinki, Finland.