Field Measurements of 802.11 Collision, Noise and Hidden-Node Loss Rates

Douglas J. Leith, David Malone
Hamilton Institute, NUI Maynooth, Ireland.

***Abstract*—In this paper we present the first field measurements taken using a new approach proposed in [1] for measuring link impairments in 802.11 WLANs. This uses a sender-side MAC/PHY cross-layer technique that can be implemented on standard hardware and is able to explicitly classify lost transmission opportunities into noise-related losses, collision induced losses, hidden-node losses and to distinguish among these different types of impairments on a per-link basis. We show that potential benefits arising from the availability of accurate and reliable data are considerable.**

I. Introduction

In 802.11 there are a variety of reasons why a packet can fail to be transmitted successfully including collisions, hidden nodes, weak signal strength, etc. However, feedback to the sender of a frame is limited to the presence or absence of an acknowledgement. This limited feedback makes schemes for understanding the of the 802.11 channel particularly interesting; information on the cause of failed transmission is required to effectively adapt to the environment.

In this paper we present field measurements taken using a new approach proposed in [1] for measuring link impairments in 802.11 WLANs. This uses a sender-side MAC/PHY cross-layer technique that can be implemented on standard hardware. The scheme is able to classify lost transmission opportunities as noise-related losses, collision induced losses, hidden-node losses and to distinguish among these different types of impairments on a per-link basis. As the technique is based at the sender, where most choices about packet transmission are made, it can be used to adapt parameters such as PHY rate and contention window. We show that potential benefits arising from the availability of accurate and reliable data are considerable.

We note that it is the current lack of such measurements that underlies the poor performance of many approaches currently implemented in commodity hardware. For example, at present rate adaptation is commonly based on the number of transmission retries (e.g. a typical approach might involve lowering the rate after n retries and increasing the rate after m successful transmissions). However, since the number of retries is affected not just by channel noise but is also closely linked to the number of contending stations (with associated collision related losses), this can easily lead to poor performance [2]. Analogous problems occur in the presence of hidden nodes, e.g. see [3]. The availability of a measure of the loss rate specifically induced by channel noise would potentially allow much more effective rate adaptation algorithms to be employed. Similarly, channel selection algorithms are fundamentally related to channel impairments and typically depend upon the availability of an appropriate channel quality metric, which can then be optimised by a suitable search over available channels.

The paper is organised as follows. In Section II we introduce aspects of the 802.11 MAC that are important for our measurement technique, and then describe the link impairments that we aim to identify. In Section III we detail our measurement methodology and then in Sections IV–VI we present baseline and field measurements that show how these measurements give insight into the state of an 802.11 channel.

This material is based upon works supported by the Science Foundation Ireland under Grant Nos. 07/IN.1/I901 and 07/SK/I1216a.

II. Link impairments in 802.11

A. CSMA/CA protocol

In 802.11 WLANs, the basic mechanism controlling medium access is the *Distributed Coordination Function* (DCF). This is a random access scheme, based on Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). In the DCF Basic Access mode, a station with a new packet to transmit selects a random backoff counter. Time is slotted and if the channel is sensed idle the station first waits for a Distributed Inter-Frame Space (DIFS), then decrements the backoff counter each PHY time slot. If the channel is detected busy, the countdown is halted and only resumed after the channel is detected idle again for a DIFS. Channel idle/busy status is sensed via:

- CCA (Clear Channel Assessment) at the physical level which is based on a carrier sense threshold for energy detection, e.g. -80dBm. CCA is expected to be updated every physical slot time. It aims to detect transmissions within the interference range.
- NAV (Network Allocation Vector) timer at MAC level which is encapsulated in the MAC header of each 802.11 frame and is used to predict the end of a received frame on air. It is naturally updated once per packet and can only gather information from stations within the decoding range. This method is also called virtual carrier sense.

The channel is detected as idle if the CCA detects the channel as idle and the NAV is zero. Otherwise, the channel is detected as busy. A station transmits when the backoff counter reaches zero. The 802.11 protocol uses a half-duplex process where an acknowledgement (ACK) is always sent by the receiver upon the successful receipt of a unicast frame. The ACK is

sent after a period of time called the Short Inter-Frame Space (SIFS). As the SIFS is shorter than a DIFS, no other station is able to detect the channel idle for a DIFS until the end of the ACK transmission. The DCF allows the fragmentation of packets into smaller units. Each fragment is sent as an ordinary 802.11 frame, which the sender expects to be ACKed. However, the fragments may be sent as a burst. That is, the first fragment contends for medium access as usual. When the first fragment is successfully sent, subsequent fragments are sent after a SIFS, so no collisions are possible. In addition, the medium is reserved using virtual carrier sense for the next fragment both at the sender (by setting the 802.11 NAV field in the fragment) and at the receiver (by updating the NAV in the ACK). Burst transmission is halted after the last fragment has been sent or when loss is detected. A similar bursting technique is used in 802.11e's TXOP feature, though greater flexibility in setting the NAV is permitted.

B. Link impairments

The manner in which link impairments are manifested is closely linked to the interaction between MAC and PHY operation. We distinguish three main types of link impairment when using the 802.11 DCF.

1) Collisions: Collisions are part of the correct operation of CSMA/CA. Collisions occur when two or more stations have simultaneously decremented their backoff counter to 0 and then transmit. Note that collisions can only occur on data packet transmissions. The level of collision induced packet losses is dependent on both load and the number of stations active in the network. For example, 802.11b with four saturated nodes has a collision probability of around 14%, while 20 saturated have a collision probability of around 40% [4]. We denote by p_c the probability that a transmitted data frame is lost due to a collision.

2) Hidden nodes: Frame corruption due to concurrent transmissions other than collisions are referred to as hidden node interference. We denote by $p_{h,data}$ the probability that a data transmission fails to be received correctly due to hidden node interference. Similarly, we denote by $p_{h,ack}$ the probability that an ACK transmission is lost due to hidden node interference. A lost data packet or a lost ACK both lead to a failed transmission and so we combine data and ACK losses into an overall hidden node error probability p_h.

3) Noise errors: Frame corruption due to sources other than transmissions by other 802.11 stations are referred to as noise losses. We denote by $p_{n,data}$ (respectively, $p_{n,ack}$) the probability that a data (respectively, ACK) frame is lost due to noise related errors. Since data and ACK losses both lead to a failed transmission we lump these together into a combined noise loss probability p_n.

III. Measurement methodology

Similarly to the approach proposed in [1], we make use of the following properties of the 802.11 MAC:

1) Time is slotted, with well-defined boundaries at which frame transmissions by a station are permitted.
2) The standard data-ACK handshake is affected by all types of link impairment considered and a sender-side analysis can reveal any loss.
3) When fragmentation is enabled, second and subsequent fragment transmissions are protected from collisions and hidden nodes by the NAV values in the fragments and ACKs. We treat hidden nodes that are unable to decode either NAV value as channel noise.

We may also use TXOP packet bursting with alternative NAV settings.

A. Estimating Noise Errors

Consider a station sending fragmented packets to a given receiver. Each fragment is immediately acked by the receiver when it arrives, allowing detection of loss. Fragments are sent back to back with a SIFS interval between them. Hence, second and subsequent packets are protected from collisions. Importantly, fragment ACK frames update the NAV and so the fragment-ACK handshake is akin to an RTS-CTS exchange from the point of view of hidden nodes Hence, second and subsequent fragments are also protected from hidden node collisions. That is, while the first fragment will be subject to collisions, noise and hidden node errors, subsequent fragments are only subject to noise errors and we have that

$$\mathbb{P}[\text{fragment 2 success}] = A_S/T_S = (1 - p_n), \qquad (1)$$

where the station transmits T_S second and subsequent data frames and of these A_S are successful because an ACK is received. We can therefore directly estimate the probability of noise errors p_n from the fraction of second and subsequent fragments with no ACK,

$$p_n = 1 - A_S/T_S \qquad (2)$$

Since the impact of noise losses may depend on the frame length (longer frames typically having higher probability of experiencing bit errors), we equalize the length of the fragments we send and transmit fragments of length equal to the packet size used for regular data transmissions. The frame loss rate estimated from fragment measurements can then be reliably applied to estimate the loss rate for other transmissions.

B. Estimating Hidden Node Interference

We now distinguish frame losses due to hidden node interference. To achieve this we exploit the fact the NAV value in the header of ACK packets echoes the value in the header of the corresponding data packet. Hence, by changing the NAV value in the first fragment packet header we can change the NAV in the ACK so that it covers just the ACK transmission i.e. does not protect the second fragment. Note that this is essentially equivalent to the TXOP functionality in 802.11e. Such second fragments are then subject to noise or hidden node interference, but not to collision losses. That is,

$$\mathbb{P}[\text{TXOP 2 success}] = A_1/T_1 = (1 - p_h)(1 - p_n), \qquad (3)$$

where the station transmits T_1 second fragment frames that are not protected by the ACK NAV and of these A_1 are successful

Name	Details
AP1	Samsung NC10 netbook, Debian Lenny, 2.6.26 kernel, Atheros chipset (radio 10.3, MAC 6.1, PHY 10.2)
AP2	Apple Airport 802.11a/b/g, firmware v. 7.4.2
AP3	Cisco 1200 series
AP4	Linksys/Cisco WRT160N

TABLE I
ACCESS POINT DETAILS

because an ACK is received. We can now use our estimate of p_n (based on fragment loss measurements, see equation (2)), to allow estimation of the probability p_h of hidden node losses as:

$$p_h = 1 - (A_1 \cdot T_S)/(A_S \cdot T_1). \quad (4)$$

C. Estimating Collision Rate

The first packet in a fragment or TXOP burst is subject to collisions, hidden nodes and noise errors. Suppose that over some time period the station contends and transmits such first in a burst data frames T_0 times and of these A_0 are successful because an ACK is received. Assuming that these sources of frame loss are independent, if the station transmits the probability of success over the link is:

$$\mathbb{P}[\text{success}] = A_0/T_0 = (1 - p_c)(1 - p_h)(1 - p_n). \quad (5)$$

Finally p_c can be estimated from Eq. (5) and (3):

$$p_c = 1 - (T_1 \cdot A_0)/(T_0 \cdot A_1). \quad (6)$$

D. Experimental apparatus

We implemented the above link quality measurement methodology via a modified Linux Madwifi driver (based on 10.5.6 HAL, 0.9.4 driver). We performed sender-side measurements using an Asus 700 laptop running Debian Lenny with 2.6.26 Linux kernel. The laptop is equipped with an Atheros 802.11 a/b/g chipset (radio 14.2, MAC 8.0, PHY 10.2). We disabled Atheros's Ambient Noise Immunity feature which has been reported to cause unwanted side effects [5]. Transmission power was fixed and antenna diversity disabled. We took measurements using a number of different APs, see Table I. AP2 operated in 802.11a mode, all other APs in 802.11b/g mode. AP1 and AP4 are located in Hamilton Institute, AP2 in a domestic residence in Dublin, AP3 in the university library at NUI Maynooth. We use multiple AP types, because these are the APs available in the field, but we have also used it to eliminate the possibility that a measurement is due to a peculiarity of a particular AP.

IV. BASELINE MEASUREMENTS

A. OFDM noise floor

We first present measurements with the client laptop on which measurements are being taken located in close proximity to the access point. For AP2 and AP4 the radio environment is clean with, as determined using a spectrum analyser, no other transmitters in the channel. For AP1 and AP3 where measurements are taken in, respectively, an office and a a library environment, other network users are active. Figure 1 plots typical measurements of packet loss rate as the PHY rate/modulation is gradually increased from 1Mbps to 54Mbps. Packet losses are categorised into first packet losses, losses of second fragments in a burst (i.e. with NAV protection) and losses of second packets in a TXOP burst (i.e. without NAV protection). Measurements are shown for four different APs as detailed in Table I. Since the channel is clean for AP2 and AP4 , we expect the loss rates of first and second packets to be similar and indeed this is the case. For AP1 and AP2 the rate of first packet losses is higher than the rate of second packet losses, which is attributed to colliding transmissions from other network users. The rates of the two types of second packet loss are similar, consistent with a low-noise, hidden-node free link.

Since the laptop and AP are co-located, we might expect to see a loss rate for second packets very close to zero for all PHY rates (and similarly low loss rates for first packets, except for AP3 measurements). However, while we measure loss rates close to zero at the 802.11b rates (1, 2, 5.5, 11Mpbs), it can be seen from Figure 1 that there a significant second packet loss rate of 1-5% is consistently observed for the 802.11a/g rates. While the measurements presented in Figure 1 cover a range of different AP hardware and software, the client laptop is common to the measurements and so might be the source of the observed behaviour. However, we have also taken measurements with a number of different client hardware/software configurations and observed a similar loss-rate floor and so this seems unlikely. One clue is that we find that this behaviour is only observed with unicast traffic and not broadcast, and so it seems to be associated with lossy reception of MAC ACKs. Recalling that the 802.11a/g rates use OFDM modulation while the 802.11b rates use CCK modulation, with associated differences in transmitter and receiver processing, and we speculate that the observed behaviour may be associated with this difference, e.g. related to calibration of the gain thresholds for triggering decoding that are used in the OFDM but not the CCK receiver.

B. Path attenuation

Figure 2 illustrates the impact of increasing the distance between the client laptop and the AP. It can be seen that the loss rate remains low as the PHY rate is increased until a "cliff" is reached at which point the loss rate rises to a high value. This reflects the quantisation present in the set of PHY rates supported by 802.11g, and the exponential dependence of loss-rate on SNR. Similar behaviour has, of course, been widely observed in other measurement studies and measurements are included here firstly as additional validation of our measurement methodology and secondly to provide a comparison against which to compare the interference-related loss measurements presented in the next section.

C. External interference: microwave oven

Figure 3 illustrates the impact of external (non-802.11) interference generated using a microwave oven. These measurements were taken using AP2 in a clean environment (as confirmed using a spectrum analyser) with the client

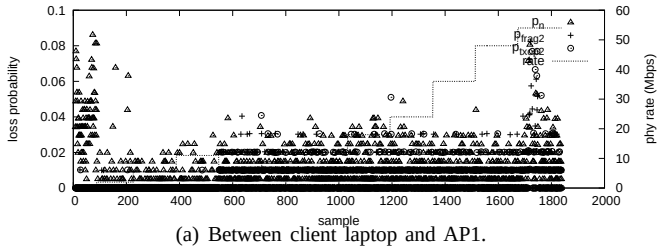

(a) Between client laptop and AP1.

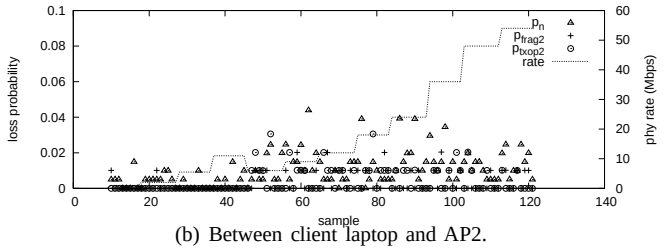

(b) Between client laptop and AP2.

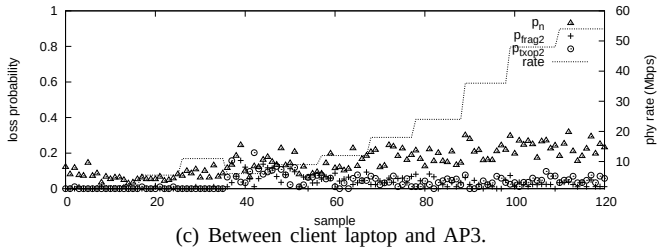

(c) Between client laptop and AP3.

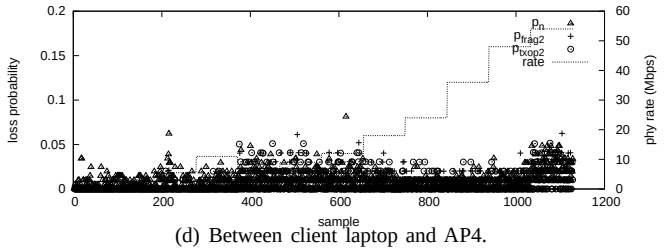

(d) Between client laptop and AP4.

Fig. 1. Measured packet loss rates vs PHY rate. Client station co-located with AP, clean channel. Data points plotted are averages over 100 packets. Legend: p_n losses on first packet in a burst, p_{frag2} on second fragment in a burst, p_{txop2} on second packet in a TXOP burst.

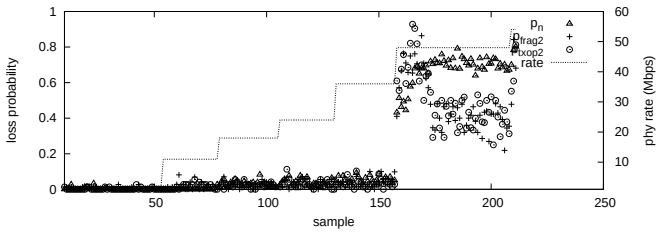

Fig. 2. Measured loss rates vs PHY rate. Client station located ≈15m from AP2.

laptop located ≈1m from the AP. The oven is operating for approximately the middle 120 measurement samples for each PHY rate, and not operating for the first and last 20 samples. The latter samples, without interference, provide a baseline against which to compare the impact of the interference. It is interesting to compare the measurements in Figure 3 to those in Figure 2. One immediate observation is that the loss-rate of second packets in Figure 3 is not monotonic in PHY rate, decreasing significantly at rates above 2Mbps. A second is that

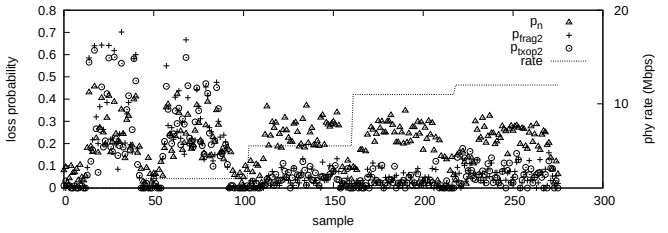

Fig. 3. Measured loss rates vs PHY rate with client laptop located 1m from AP5. Interference is generated using a microwave oven.

the loss-rate experienced by first packets is roughly constant with PHY rate, in contrast to the situation with path attenuation induced losses.

Although we lack sufficient details of the interference generated by the specific oven used here to obtain a definitive explanation of these features, we note that the microwave oven radio power (i) has an on-off duty cycle and (ii) during an on cycle the power is expected to be sufficient to corrupt packets at all available transmit rates. First packets in a fragment or TXOP burst are transmitted at random times (due to the stochastic nature of CSMA/CA channel access) and so can be expected to randomly sample the on and off interference cycles, yielding a high loss rate for first packets that is insensitive to PHY rate. However, the transmission of second packets is conditioned on the successful transmission of a first packet, and second packets are transmitted back to back (with a SIFS space) to the first packet. The first packet therefore acts as a "probe" packet for second packets. At higher PHY rates, where packet transmission durations are sufficiently short, success of the first packet indicates an interference off cycle and a subsequent second packet may often be transmitted quickly enough to complete before the next on cycle starts, yielding a lower loss rate for second packets than for first packets. At the lower 1 and 2Mbps PHY rates, we infer that packet transmissions take a sufficiently long time that the transmission of two back to back packets takes longer than the length of an off cycle and so the loss rate of second packets is increased above that for first packets.

D. Discussion

In summary, our baseline measurements indicate at least two qualitatively different patterns of loss rate versus PHY rate. Namely,

1) Loss rate is small for low PHY rates, with a "cliff" or sharp increase in loss rate above a threshold PHY rate. This is expected behaviour in conventional channel models such as AWGN channels.
2) Loss rate is insensitive to PHY rate. This appears to be associated with on-off interference/noise where during an on cycle packets are corrupted regardless of the PHY rate used, while during an off cycle packets are transmitted successfully at all PHY rates.

In addition, there is observed to be a loss-floor of around 1-2% when OFDM modulation is used. As we will see later, these qualitative patterns can prove helpful in diagnosing sources of packet loss in field measurements where the environment is less well controlled than in these baseline tests.

V. Field measurements: NUIM Library

In this section we present field measurements taken at the National University of Ireland Maynooth library, corresponding to AP3 in Table I. The client laptop is located outside the library building, approximately 30m from the AP, and we note that this is a production network with other active users. Figure 4(a) shows typical measurements of loss rate vs PHY rate. It can be seen that the measured loss rates exhibit quite complex behaviour that appears to consist of a mix of the behaviours observed in our baseline measurements in the previous section. At PHY rates of 36Mbps and above, the loss rate for both first and second packets increases to be close to one (data points for these extremely high loss rates are not shown in Figure 4(a) due to difficulties associating with the AP) – this "cliff" appears to be associated with path attenuation. At PHY rates below 36Mbps, the loss rate for first packets appears insensitive to the PHY rate and is around 20% – this may indicate on-off interference. The loss rate for second packets is similar to that for first packets for PHY rates of 5.5Mbps and above, but significantly different at PHY rates of 1 and 2Mbps. At these lower rates the loss rate for second packets in TXOP bursts is around 60% while that for second fragments in a burst is less than 5%. This difference in loss rates for the different types of second packets indicates the presence of hidden node interference; we are also able to observe beacons from at least one other AP. To our knowledge, this is the first time that measurements of this type have been presented.

In view of the complex nature of the radio environment, we can expect that PHY rate control algorithms will experience difficulties in selecting a good rate at which to operate. Figure 4(b) plots measurements of completion time vs connection size for uploads to the Hamilton Institute web server from the client laptop via the same AP (we expect similar behaviour in the download direction, but can only directly control the PHY rate on the client laptop since we do not have administrative access to the AP). Measurements are shown using (the default in the Madwifi driver) SampleRate algorithm and for fixed PHY rates of 11Mbps and 18Mbps, and include both the measured completion times for individual connections and the mean completion time over multiple runs. It can be seen that by using either the 11Mbps or 18Mbps rates the mean completion times can be reduced by a factor of 2-3 compared to SampleRate. We expect that the performance of the rate control algorithm might be improved by making use of the additional information provided by our link quality measurement methodology, but leave investigation of this as future work.

VI. Field measurements: Hamilton Institute

We took a second set of field measurements within the Hamilton Institute, corresponding to AP4. These measurements were taken within an office environment. Figure 5 shows typical measurements of loss rate vs PHY rate. Measurements are shown for AP4 operating on channels 1, 6 and 11. It is known that the building security system generates substantial radio interference in channel 11 and this is reflected in the

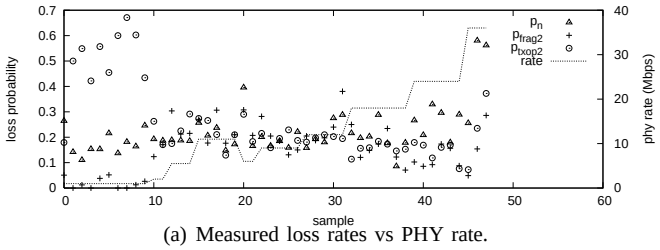

(a) Measured loss rates vs PHY rate.

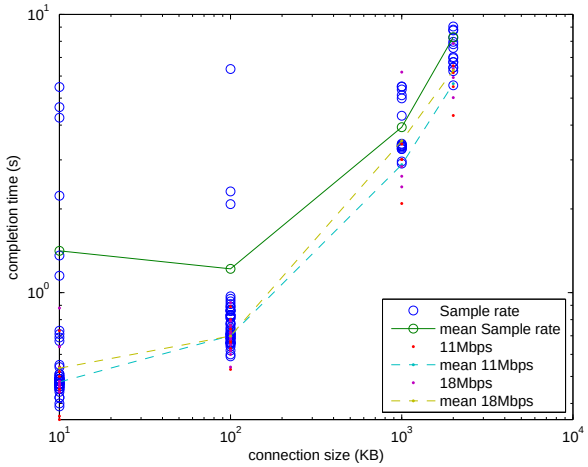

(b) Measured completion times vs connection size.

Fig. 4. Measurements taken at NUIM library.

high loss rates experience by both first and second packets at all PHY rates, see Figure 5(a). A wireless testbed operates on channel 1 in an adjacent room and it can be seen from Figure 5(c) that the resulting colliding transmissions lead to an increased loss rate for first packet transmissions, although both types of second packet experience similar loss rates which indicates that there are no significant hidden node effects.

Figure 6 shows the corresponding RSSI measurements. It can be seen that these provide a poor basis for selecting the best channel on which to operate AP4 — channel 11 has the lowest RSSI yet from Figure 5 we know that this channel experiences a far higher loss rate than channel 6. This discrepancy arises because the RSSI measurements are derived from the signal strength of MAC ACKs received by the client laptop and so are conditioned on successful transmission of a data packet while the sender-side measurements in Figure 5 take account of both failed and successful data packet transmissions. To explore this issue further, we took measurements of completion time vs connection size for uploads to the Hamilton Institute web server from the client laptop. These measurements are shown in Figure 7 for channels 1 and 6 from which it can be seen that use of channel 6 yields approximately a factor of 3 reduction in flow completion time compared to channel 1. Measurements are not shown for channel 11 as the high level of packet loss led to TCP timeouts and grossly long completion times. This suggests that the additional information provided by the measurements in Figure 5 might be used to assist in making better channel allocation choices at APs, or AP association choices at client stations.

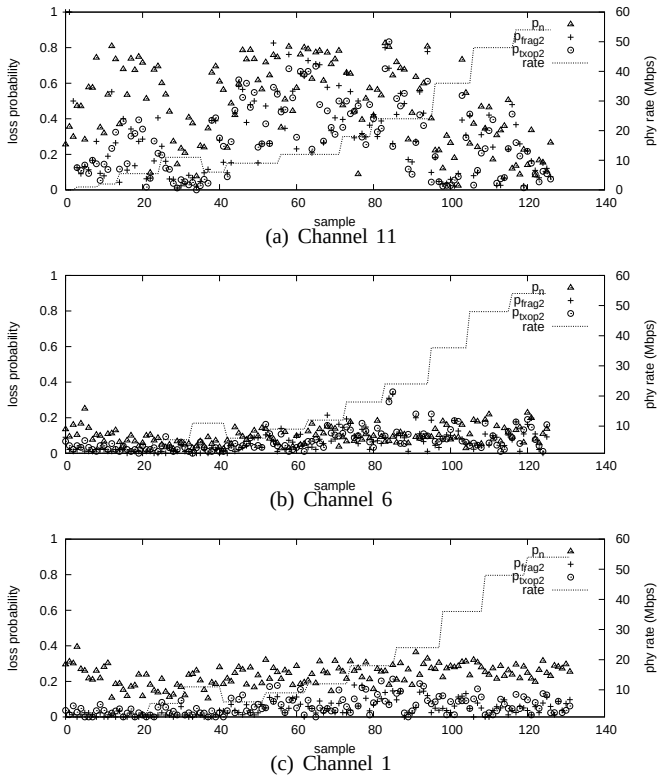

(a) Channel 11

(b) Channel 6

(c) Channel 1

Fig. 5. Measured loss rates vs PHY rate, Hamilton Institute.

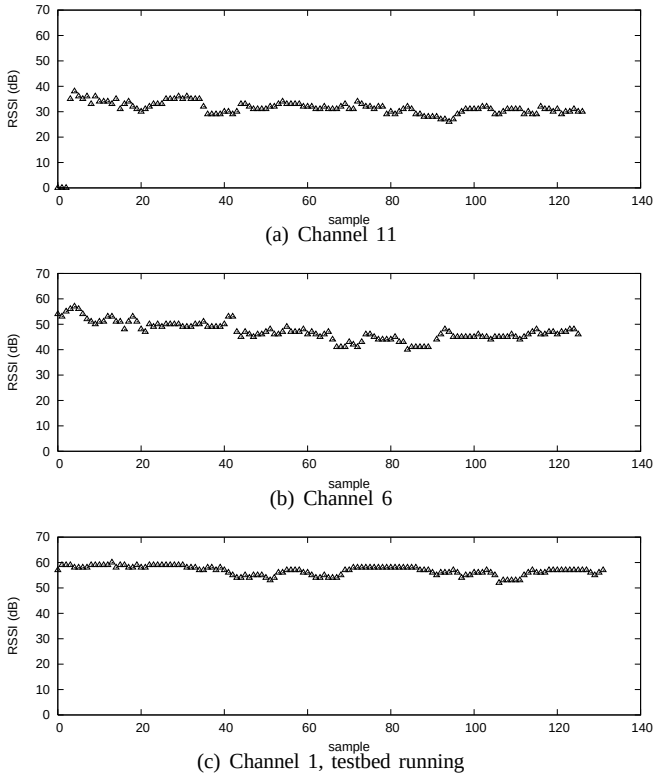

(a) Channel 11

(b) Channel 6

(c) Channel 1, testbed running

Fig. 6. Measured RSSI vs PHY rate, Hamilton Institute.

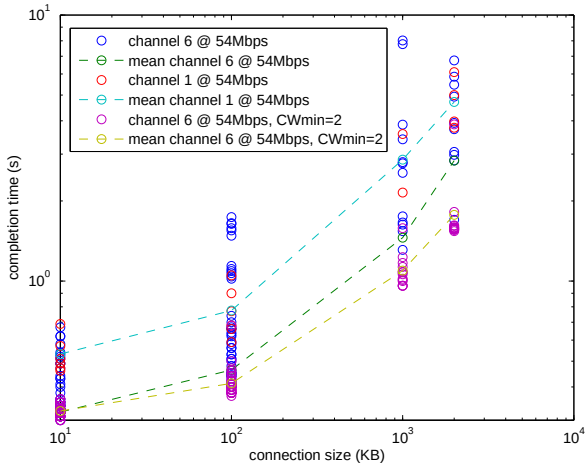

Fig. 7. Measured completion times vs connection size and choice of channel.

VII. Conclusion

In this paper we have presented field measurements of 802.11 collision, noise and hidden-node loss rates. We see that the measurement technique from [1] is able to offer insight into the causes of lost packets. While we identify collision, noise and hidden-node related losses, we are able to identify effects that seem to be related to OFDM reception and on-off noise. We show that these measurements could be used to drive rate control or channel selection decisions that offer considerable performance benefits.

We would like to use this technique to estimate how common various impairments are in production 802.11 environments. We would also like to implement this measurement technique on a other 802.11 chipsets to determine if issues such as the OFDM loss-floor are chipset specific.

References

[1] D. Giustiniano, D. Malone, D.J. Leith and K. Papagiannaki, "Measuring transmission opportunities in 802.11 links", *IEEE Trans on Networking*, to appear.

[2] K. Ramachandran et al., "Scalability analysis of Rate Adaptation Techniques in Congested IEEE 802.11 Networks: An ORBIT Testbed Comparative Study", *Proc. IEEE WoWMoM*, 2007.

[3] S. Wong, et al., "Robust Rate Adaptation for 802.11 Wireless Networks", *Proc. ACM MobiCom*, 2006.

[4] G. Bianchi, "Performance analysis of IEEE 802.11 distributed coordination function", *IEEE J. Sel Area Comm*, 18(3):535–547, Mar. 2000.

[5] I. Tinnirello, D. Giustiniano, L. Scalia, and G. Bianchi, "On the side-effects of proprietary solutions for fading and interference mitigation in IEEE 802.11b/g outdoor links. *Elsevier Computer Network Journal*, 2009