Packet Loss Characterization in WiFi-based Long Distance Networks

Anmol Sheth
University of Colorado, Boulder
anmol.sheth@colorado.edu

Sergiu Nedevschi, Rabin Patra, Sonesh Surana, Eric Brewer
University of California, Berkeley
{sergiu,rkpatra,sonesh,brewer}@cs.berkeley.edu

Lakshminarayanan Subramanian
Intel Research, Berkeley and New York University
lakshmi@cs.nyu.edu

***Abstract*— Despite the increasing number of WiFi-based Long Distance (WiLD) network deployments, there is a lack of understanding of how WiLD networks perform in practice. In this paper, we perform a systematic study to investigate the commonly cited sources of packet loss induced by the wireless channel and by the 802.11 MAC protocol. The channel induced losses that we study are external WiFi, non-WiFi and multipath interference. The protocol induced losses that we study are protocol timeouts and the breakdown of CSMA over WiLD links.**

Our results are based on measurements performed on two real-world WiLD deployments and a wireless channel emulator. The two deployments allow us to compare measurements across rural and urban settings. The channel emulator allows us to study each source of packet loss in isolation in a controlled environment. Based on our experiments we observe that the presence of external WiFi interference leads to significant amount of packet loss in WiLD links. In addition to identifying the sources of packet loss, we analyze the loss variability across time. We also explore the solution space and propose a range of MAC and network layer adaptation algorithms to mitigate the channel and protocol induced losses. The key lessons from this study were also used in the design of a TDMA based MAC protocol for high performance long distance multihop wireless networks [12].

I. Introduction

Many developing regions around the world are in dire need for low-cost connectivity solutions to provide network coverage. These regions have low telephone penetration rates (roughly 2% in Africa) [9], and rural areas with their low user density cannot support the cost of cellular base-stations or fiber (unlike urban areas). Satellites provide excellent coverage, but bandwidth is extremely expensive, typically more than US$2000 per Mbps per month. Additionally, although WiMAX [17] has been suggested as another potential solution, and may prove useful down the road, it suffers from two problems: (a) It is currently very expensive; (b) WiMAX, so far, has been intended for carriers (like cellular) and is thus hard to deploy in the "grass roots" style typical for developing regions.

WiFi-based Long Distance (WiLD) networks [8] are emerging as a low-cost connectivity solution and are increasingly being deployed in developing regions in both urban [1] and rural settings. The primary cost gains arise from the use of very high-volume off-the-shelf 802.11 wireless cards, of which over 140 million were made in 2005. These links exploit unlicensed spectrum, and are low power and lightweight, leading to additional cost savings [4].

Despite the promise of low-cost connectivity, the performance of WiLD networks in the real world has been abysmal. This poor performance is primarily triggered by the high loss variability observed on WiLD links. Figure 1 shows the loss rate measured over two of our links ("K-P" and "B-R") over a period of 3 hours on different days. The loss rate was averaged over 30-second intervals for a 1 Mbps unidirectional UDP CBR traffic flow with the MAC-layer ACKs turned off and retries set to zero.

The two main characteristics that we observe are: 1) WiLD links demonstrate high variability of loss rate; and 2) the loss rate can be highly asymmetric across a link. Bursts vary in magnitude as well as duration. For example, on the K to P link, loss bursts ranged in magnitude from 15–80% and the duration of bursts also varied from a transient high burst to a long burst lasting over 25–30 minutes. In contrast, the reverse path (P to K) had almost 0% loss for the entire duration. In addition to the high variability of the loss rate, there is also a residual loss that is always present and remains constant over long time periods. This residual loss ranges between 0–10% and varies with each link. Although Figure 1 shows only two links in our testbed, the above behavior is characteristic of all our urban links. In contrast, our rural links consistently show loss rates close to zero with a maximum of less than 2%. We explore these differences further and point out that many WiLD links have one end in an urban area. In addition, the losses shown

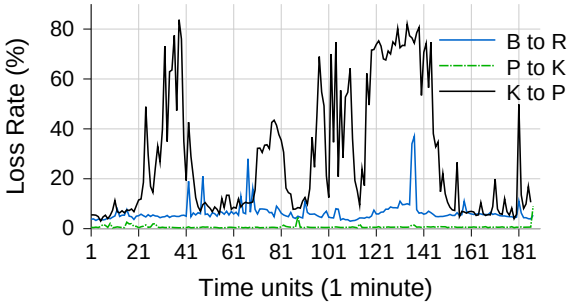

Fig. 1. Packet loss variation over a period of about 3 hours. The loss rate was averaged over 30-second intervals for a 1440 byte UDP CBR flow of 1 Mbps at 802.11b datarate of 11Mbps

This work is supported in part by the National Science Foundation under Grant No. 0326582.

[1]In urban regions in Africa, satellite-based Internet providers use WiLD networks as a distribution network to reach out to the end-users within the region.

here are only those due to the channel; the 802.11 protocol itself also induces losses.

A. Our Contributions

In this paper, we perform a detailed measurement study to analyze the packet loss characteristics and the sources of packet loss in WiLD network settings. We categorize the sources of packet loss into two broad categories: (a) *channel losses* induced by the long distance wireless channel; (b) *protocol-induced losses* due to the 802.11 MAC protocol. Our study is based on a real-world WiLD network deployment consisting of 6 links with lengths varying from 2–20 km. Unlike existing WiLD deployments [5], our testbed includes both rural and urban links. In addition to the real deployment, we also perform detailed experiments using a wireless channel emulator, which enables repeatable controlled experiments. These key contributions were also used to drive the design and implementation of the TDMA based MAC protocol for WiLDNet [12].

This paper makes four important contributions:

Channel loss characterization: We analyze three well known causes for channel losses in wireless environments, namely, *external WiFi interference*, *non-WiFi interference* and *multipath interference* (Sections III–V). Among these, we show that external WiFi interference is the most significant source of packet losses in WiLD environments and the effect of multipath and non-WiFi interference is not significant. This is in contrast to the results of the Roofnet mesh network [1] where the authors observed multipath to be the most significant source of packet loss.

Protocol-induced losses: The stock 802.11 MAC protocol is ill-suited for WiLD links due to the breakdown of CSMA over long distances and propagation delays (Section VI). Here, we pinpoint the fundamental shortcomings of the 802.11 MAC protocol.

Loss variability analysis: We classify the loss patterns over time into two basic categories: *bursts* and *residual* loss. We further classify bursts into short and long bursts. We make three important observations (Section VII): (a) Although the burst arrival patterns can be approximately modeled based on a Poisson process, the duration and magnitude of a burst are harder to predict; (b) The residual loss characteristics over certain links are stationary, while some others exhibit non-stationary behavior even over daily timescales; (c) The loss variability observed in our urban links significantly differs from that under rural settings as observed in prior work [5].

Loss remedies: Having identified external WiFi interference as the primary source of losses in WiLD links, we discuss three potential remedies to mitigate these losses (Section VIII): (a) 802.11 frequency channel adaptation; (b) 802.11 PHY datarate adaptation, and (c) adaptive FEC. We evaluate the effectiveness of each of these remedies.

The focus of our packet loss characterization study is significantly different from other wireless-based loss measurement studies [1], [14]. The work done by Raman et al. [5] is the only other measurement-based study of WiLD deployments of which we are aware. However, the two studies are orthogonal: we focus on loss variability characterization, determining the impact of different sources of losses and remedies for loss alleviation, their work focused more on performance analysis of 802.11 network at various layers in the network stack and the effect of other parameters (weather, SNR, payload, datarate) on loss. Our work also differs from mesh networks like Roofnet [1] in that WiLD networks, as we show, have very different loss characteristics, with loss much more due to external interference than multipath effects.

Link	Distance (km)	Environ.	Antenna height(m)
K-P	20	Urban	50
B-R	8	Urban	30
M-P	2	Urban	40
T-A	11	Rural	20
T-S	13	Rural	25
W-N	15	Rural	20

TABLE I
LIST OF OUR URBAN AND RURAL WILD TESTBED LINKS.

II. EXPERIMENTAL METHODOLOGY

We perform our packet loss characterization measurements on a WiLD network testbed comprising of links in both rural and urban environments. Table I summarizes some of the urban and rural links in our deployments. The links range from 2–20 km in length.

The two main characteristic of WiLD links that differentiate them from links in a multi-hop urban mesh deployment [1] are the longer distances and the use of high-gain directional antenna (24 dBi, 8 degree beamwidth). The two endpoints of each link have direct line-of-sight (LOS). In multihop settings, nodes have one radio per fixed point-to-point link to each neighbor, which can independently operate on different channels.

In addition to the testbed, we also use a wireless channel emulator (Spirent 5500 [16]) to study each source of packet loss in isolation. The emulator allows us to place the two ends of the link in separate RF-isolated boxes (80dB isolation) and then emulate in real time the RF channel between them. The Spirent 5500 accurately emulates radio channel characteristics with channel loss, fast and slow fading and delay spreads. This enables us to emulate links of any length or loss profile with repeatable results. We perform tests by connecting the channel emulator to the same radios used in our WiLD deployments.

We use Atheros 802.11 a/b/g radios for all our experiments. The wireless nodes are 266 MHz x86 Geode single board computers running Linux 2.4.26. We use *iperf* to measure throughput. All our results are based on CBR UDP traffic streams. Unless otherwise stated, for all our experiments we turn off MAC-layer ACKs and set the maximum retries limit to zero. This allows us to measure the real channel loss rate in absence of any MAC-layer acknowledgments and retries.

We instrument the stock Atheros *madwifi* driver to log fine-grained information for each frame received and transmitted. In addition to capturing all the frames on the link, we also capture and log frames being transmitted by external WiFi sources. This is achieved by creating a virtual network interface set in "monitor mode" on the same channel as the primary interface. This technique is equivalent to using two physical network interfaces, one being the primary and the other a passive monitor. To summarize, we collect the following information for every frame: complete 802.11 MAC header and IP payload, received signal strength, PHY layer transmit datarate, timestamp, PHY and CRC errors, and the noise floor immediately after the frame is received. We also modify the Atheros driver to pass up frames with CRC and PHY errors.

Using the WiLD testbed and the channel emulator, we explore two categories of loss: *channel losses* induced by the wireless channel and *protocol-induced losses* by the 802.11 MAC protocol. Specifically, for channel-induced losses we investigate: a) External WiFi interference, b) External non-WiFi interference and c) Multipath interference. For 802.11 protocol induced losses, we investigate: a) Timeouts due to propagation delay, and b) Breakdown of CSMA over long distances.

III. External WiFi Interference

In this section, we investigate external WiFi interference as a potential source of packet loss in WiLD links. Any WiFi traffic that is not a part of the primary WiLD link is categorized as external WiFi interference. Based on the measurements performed on our WiLD testbed and the wireless channel emulator, we show three key results:

- In the presence of external WiFi interference, the loss rate is strongly correlated with the amount of external traffic received on the same and adjacent channels. In contrast, due to the omni-directional antennas used in the Roofnet deployment [3], no such strong correlation was observed.
- Packet loss due to external WiFi interference is far more significant in WiLD deployments than local mesh networks.
- The loss due to external WiFi interference depends on the relative power level between the primary and external traffic, their channel separation, and the rate of external interference.

A. Correlation of loss rate and external WiFi traffic

Figure 2 shows the loss rate across all (rural and urban) our WiLD links. We observe that the loss rate of the urban links vary across a wide range (4–70%). In contrast, all the rural WiLD links have a very small loss rate. The maximum loss rate observed in all our rural WiLD links was 2%.

To study this contrast between the rural and urban links, we collected detailed packet level MAC traces. By parsing the MAC header source and destination fields, we are able to count the number of frames received from external WiFi sources. In the traces collected over all our rural links we see negligible external WiFi traffic. However, significant

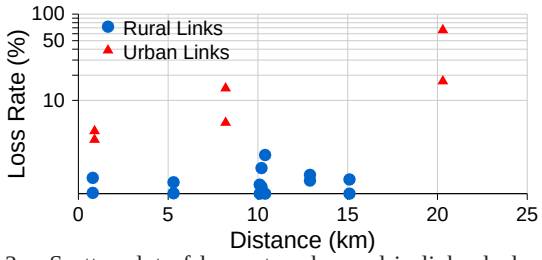

Fig. 2. Scatter plot of loss rates observed in links deployed in urban and rural areas (note: loss rate is plotted in logscale)

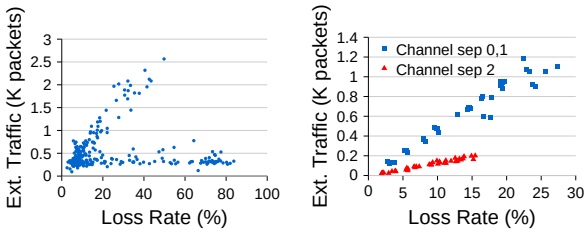

(a) In WiLD link (P-K) (b) In Wireless emulator

Fig. 3. Correlation between loss rate and external interference. Traffic is 5Mbps UDP CBR packets of 1440 bytes each at 802.11b PHY datarate of 11Mbps.

amount of external WiFi traffic was captured from the traces collected in the urban WiLD deployment.

Figure 3(a) shows a scatter plot between the loss rate and the absolute number of external WiFi traffic frames received on an urban link (K → P) for a period of 6 hours. The figure shows that a subset of the loss rate samples are strongly correlated with the external traffic. For the other subset of the samples, the loss rate increases even when there is no significant increase in WiFi traffic on the same channel.

To investigate this further, we perform a controlled experiment using the wireless channel emulator. To model interference from an external traffic source, along with the primary link traffic we introduce a controlled interference source at the receiver. The traffic rate of the interference source was varied from 0.1 to 1 Mbps and the traffic rate on the primary link was kept fixed at 5 Mbps. Figure 3(b) shows a scatter plot of the loss rate and the total number of frames received from the external interference source. From the graph, we observe that for a given loss rate, the amount of external traffic captured by the monitor device depends on the channel separation of the primary and interference source.

The above observed trend is the same as that in Figure 3(a). At a channel separation of 0 and 1, the receiver can receive both the primary link traffic as well as the frames from the interference source. Hence, the loss rate is directly correlated with the amount of external WiFi traffic captured by the monitor interface. At a channel separation of 2, the receiver is not able to receive the frames from the external interference source. However, the signal spillage of the interference source in the primary channel is sufficient to cause frame corruption. From the traces we observed that almost 100% of the lost frames contained CRC errors.

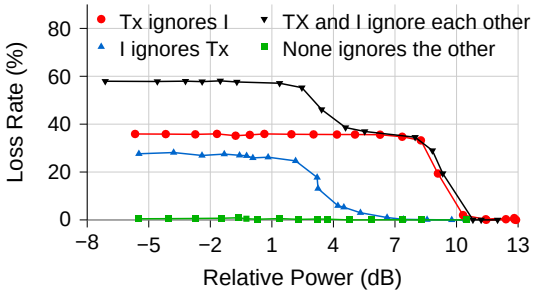

Fig. 4. Losses due to different hidden terminal effects. Both main and interfering traffic is 1440 byte UDP CBR packets at 11Mbps PHY datarate of 802.11b.

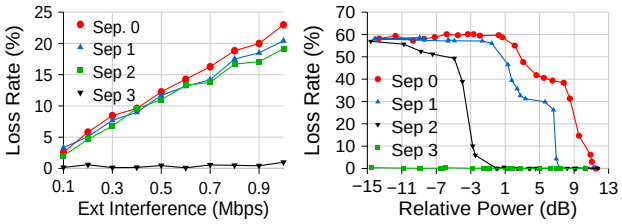

(a) Varying interference rate (b) Varying interference power

Fig. 5. Loss rate at different channel separations. Both main and interfering traffic is 1440 byte UDP CBR packets at 11Mbps PHY datarate of 802.11b.

B. *Effect of hidden terminals in WiLD networks*

Unlike WiLD deployments, where we have observed significant correlation between loss rate and external interference, it has been observed that there is no significant correlation in outdoor mesh-network deployments (Roofnet [3]). In a mesh-network deployment, an external interference source (I) that is within range of the omni-directional transmitter (Tx) would be able to sense the medium to be free and backoff its transmission. However in WiLD links, the transmissions are highly directional and the propagation delays are higher. These factors in combination exacerbate the *hidden terminal* problem in WiLD networks. The transmitter and the interference source can erroneously sense the medium to be free leading to collisions whenever they are out of range of each other (because of the directional nature of transmission) or when they cannot sense the medium to be busy in time to backoff (because of the longer propagation delays).

Collisions at the receiver can manifest in two different situations: a) When I doesn't hear Tx, and initiates a transmission when the medium is busy with an ongoing packet transmission from Tx, and b) When Tx doesn't hear I, and causes a collision by interrupting an ongoing packet transmission from I.

To isolate the above two cases and measure the performance degradation due to each case, we perform controlled experiments using two WiFi links. We simultaneously send packets from both Tx (512 Kbps traffic) and I (3Mbps), and measure the packet loss rate on the primary link ($Tx \rightarrow Rx$) with MAC-layer ACKs disabled.

To create the situation where Tx cannot hear I, we disable the Clear Channel Assessment (CCA) at Tx, which simply causes Tx to ignore I. We also eliminate propagation delay between Tx and I so that I's CCA works perfectly. We reverse the operations to create the situation in which I cannot hear Tx, but Tx hears I perfectly.

We then run four experiments, reflecting the losses in four situations: when Tx can't hear I, when I can't hear Tx, when neither can hear each other (representative of cases in WiLD networks), and when both Tx and I hear each other (representative of most cases in urban mesh networks).

Figure 4 shows the loss rate for each of the above four cases. In the case where I ignores Tx, to overcome the interferer completely (achieve 0% loss), packet transmissions from the Tx have to be 7dB stronger than the interfering transmissions. This threshold, at which the primary link is loss free, is much higher (12dB) in the case where Tx ignores I. When neither of Tx and I can hear each other, both the above two types of collisions are possible. Hence the loss rate is the sum of the losses generated by the above two types of collisions. However, when both Tx and I are in range of each other, resembling a mesh-network, losses due to collisions are close to zero. In this case, CSMA ensures that the two transmitters, Tx and I, share the medium properly.

From the above experiment we conclude that the effect of hidden terminals, causing collisions at the receiver, are greatly exacerbated in WiLD networks compared to urban mesh networks.

C. *Effect of relative power and rate of external interference*

To study the effect of relative power and rate of the external WiFi traffic on the loss of the primary link, we perform two experiments using the wireless channel emulator.

In the first experiment, we fix the relative power between the interference source and primary WiLD link, and vary the rate of the external interference source. The received signal strength of the interfering source was approximately 6dB higher than the primary link traffic. From Figure 5(a) we observe that for channel separations of 0, 1 and 2, the loss rate increases as the rate of the external interference increases. However, beyond a channel separation of 2, there is no significant interference from the external WiFi traffic source and the loss rate is almost zero.

Figure 5(b) shows the variation in loss rate for different relative power levels of the interference source and WiLD link. In this experiment, we fix the power level of the primary WiLD link traffic and vary the relative power of the primary link to the power of the interferer from -15dBm to +13dBm. The primary link CBR traffic rate is fixed at 512 Kbps, while the interferer transmits at a rate of 3 Mbps.

We observe that when the interference source is on the same channel, even a 12dB lower signal could lead to packet loss on the primary WiLD link. When the interference source is significantly higher than the WiLD link (6dB and beyond), the loss rate is very high ($\geq$50%) for channel separations 0, 1 and 2. This corresponds to the

situation where any collision results in the capture of the packet on the primary link. Beyond a channel separation of 2, we do not observe any loss on the primary link.

D. Implications

- We conclude that external WiFi interference is a significant source of packet loss in WiLD networks. Any deployment of WiLD networks in dense urban deployments has to take into account external WiFi interference.
- When calculating the link budget for urban links, it is beneficial to over-provision the received power. A high signal strength could potentially immunize the WiLD link from external WiFi traffic.
- MAC layer adaptation algorithms like adaptive channel switching, rate adaptation, and adaptive FEC could significantly reduce the loss due to external WiFi interference. In section VIII we evaluate each one of these as potential remedies to mitigate external WiFi interference.

IV. Non-WiFi Interference

The 802.11b communication protocol operates in the 2.4 GHz shared ISM band. This frequency band is shared with a host of other non-802.11 devices, such as microwave ovens, cordless phones, baby monitors, etc. Most of these non-802.11 devices do not follow a channel-access protocol which could lead to a significant amount of interference caused by these devices.

In Sheth et al. [15], the authors were able to detect and measure non-WiFi interference by sampling the noise floor of the Atheros chipset. The authors observed that in presence of external non-WiFi noise, the noise floor linearly increases with increasing noise. We performed the same experiment on our WiLD testbed, where we sample the noise floor for every packet received. In presence of external noise causing high loss, we would expect the noise floor to be correlated with the loss rate. However, based on extensive measurements carried out on the urban links we do not see any correlation between noise floor and loss rate. In fact, the noise floor remains mostly constant with minor 1–2 dB variations.

In addition to the above test, we also check for wide-band non-WiFi noise. A wide-band noise source would cause interference across the entire 802.11 spectrum. Ideally, this can be measured using a spectrum analyzer and detecting a rise in power across the entire spectrum. However, using a spectrum analyzer is infeasible on the outdoor WiLD links. Thus, to detect wide band noise in our WiLD deployment we synchronize the two ends of a link to rotate across channel 1, 6 and 11 periodically. The sender generates 1 Mbps UDP CBR traffic on each channel and the receiver measures the loss rate on each channel. In presence of any wide-band noise, we would expect to observe a correlation among loss rates across all three channels. However, based on long-term experiments performed on three urban links, we determined that there was no statistically significant correlation, and thus no significant broadband noise.

V. Multipath Interference

Multipath interference is a well known source of packet loss in WiFi networks [1], [6]. It occurs when a RF signal takes different paths from a source to a destination node. Hence, along with the primary line-of-sight signal, the receiver also receives multiple secondary reflections that causes inter-symbol interference (ISI) which might lead to packet corruption. Based on the experiments performed on our WiLD deployments, we conclude that unlike urban mesh deployments, the order-of-magnitude lower delay spreads in WiLD deployments significantly reduces the interference due to multipath.

The two factors contributing to lower delay spreads in WiLD networks are the long distance between the two end hosts and the line-of-sight deployment of the nodes. The strong line-of-sight component in WiLD deployments ensures that the attenuation of the primary signal is only due to path loss, and most of the secondary paths are due to reflections from the ground. In comparison to our WiLD deployment, an urban mesh-network deployment (like Roofnet) has shorter and many non-line-of-sight links.

Table 6 shows the delay between the primary path and secondary path assuming the antenna is mounted at a height of 30 meters and reflection is only from the ground. The two delays are computed for a secondary path reflecting at the midway point and at the quarter point respectively between the transmitter and the receiver. As the length of the link increases, the primary and the secondary path travel almost the same distance, and hence the delay between the primary and secondary reflection reduces. This reduces the probability of inter-symbol interference. As seen from the table, there is an order-of-magnitude difference between the delay in WiLD links and medium range mesh-network style links. Aguayo et al. [1] also observed that the RAKE receiver is able to tolerate delay spreads upto 0.3–0.4 μsec.

Dist. (km)	Delay spread (μsec)
0.5	(4.75, 3.59)
1.0	(2.4, 1.80)
8.0	(0.3, 0.22)
16.0	(0.15, 0.11)
100.0	(0.02, 0.01)

Fig. 6. Delays between a primary and secondary reflection at midway and quarter-way point.

Our hypothesis was that most of the loss in our urban deployment was due to external WiFi interference. Hence, in absence of external interference the WiLD links deployed in the rural areas should not have any loss. Figure 2 validates our hypothesis, where rural links have a very low loss as compared to urban links.

VI. 802.11 Protocol-Induced Losses

In this section we study the limitations of the standard 802.11 MAC protocol over point-to-point WiLD links. The 2P protocol [13] proposes modifications to the stock 802.11 MAC protocol to enable synchronous send and receive in point-to-multipoint WiLD links. However, in this section

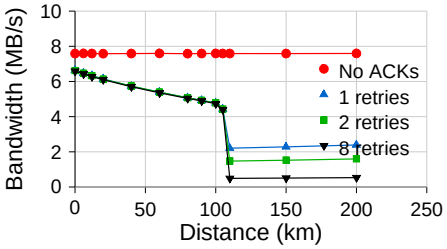

Fig. 7. Under-utilization due to the 802.11 link recovery mechanism. Traffic is 1440 byte UDP CBR packets at 11Mbps PHY datarate of 802.11b

we argue that the 802.11 protocol suffers from fundamental limitations that make it unsuitable even for just point-to-point long distance links. The two main limitations of the protocol are the link-layer recovery mechanism and the breakdown of CSMA over long distances.

A. Link layer recovery mechanism

The 802.11 MAC uses a simple stop-and-wait protocol, with each packet independently acknowledged. The receivers are required to send an ACK within a tight time bound (ACKTimeout), or the sender has to retransmit. Because of this, with increasing link distance, the sender has to wait for a longer time for the ACKs to return (proportional to propagation delay). This decreases channel utilization. Also, if the time it takes for the ACK to return exceeds the ACKTimeout parameter, the sender will retransmit unnecessarily.

These problems can be understood by performing a simple experiment using the wireless channel emulator. We configure the emulator to introduce a propagation delay and vary it to emulate links ranging from 0-200 km. We set the ACK timeout value to the maximum possible (746μs) in Atheros chipsets, corresponding to a distance of 110 km (other cards like Prism 2.5 have lower limits). Figure 7 shows the performance of the 802.11 stop-and-wait link recovery mechanism over increasing link distance. With the MAC-layer ACKs turned off (No ACKs), we achieve a throughput of 7.6 Mbps for 1440-byte CBR traffic source using the 11 Mbps datarate. When MAC ACKs are enabled, the sender has to wait for an ACK after each transmission, and this leads to decreasing channel utilization with increasing link distance. After 110 km, the propagation delay exceeds the maximum ACK timeout and the sender always times out before the ACKs can arrive. We notice a sharp decrease in received bandwidth, as the sender retries to send the packet over and over again (even though the packets were most likely received), until the maximum number of retries is reached.

B. The Breakdown of CSMA

The 802.11 MAC protocol was originally designed for short-range indoor broadcast environments where both the transmitter and receiver nodes are in range of each other. The CSMA/CA channel-access mechanism of 802.11 requires each node to sense the medium before transmitting and initiate the transmission only if the channel is idle. However, on longer distance links, it is possible that the two nodes will begin transmission within the window defined by the propagation delay. Thus, in presence of bidirectional traffic there are frequent collisions between the two end points of a WiLD link and the throughput of the link severely degrades as the distance is increased.

C. Implications

•*TDMA based WiLD MAC protocol*: As discussed above, an un-synchronized channel-access mechanism like CSMA causes severe collisions even for plain long distance point-to-point links. This necessitates a MAC protocol that synchronizes the transmissions from both the end points of the link. Although Raman et al. [13] already motivate the need for a TDMA based MAC protocol for point-to-multipoint topologies, we observe that such a synchronized MAC protocol is required even for point-to-point WiLD links.

•*Adaptive link recovery*: An alternate approach that mitigates the under-utilization of the medium due to the large timeouts and propagation delay is to relax the constraint of having only a single un-acknowledged frame. We propose a sliding-window based flow-control approach, in which the receiver acknowledges a set of frames at once (bulk ACKs). Thus recovery mechanism also allows uni-directional traffic flow within any TDMA time slot, and hence avoiding collisions of the data frames and acknowledgments. The MAC protocol for WiLDNet [12] combines TDMA based slot allocation with adaptive link loss recovery using bulk ACKs.

VII. Loss Variability

In this section, we analyze the variability of packet loss over time on the WiLD links. We first propose a simple mechanism we use to classify loss periods as either bursts or residual losses and then individually describe the loss characteristics for bursts and residual losses.

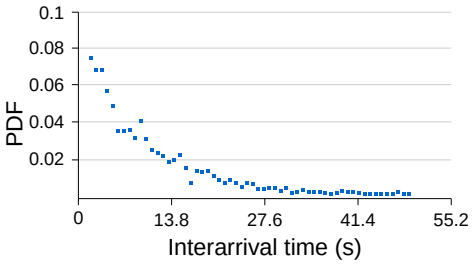

Fig. 8. Prob. distribution of inter-arrival time of bursts (R-B link)

A. Burst-Residual Separation

We observe that all the links in our testbed exhibit a bi-modal loss variation over time where the loss-rate at any given time can be classified into into two categories: *bursts* and *residual* losses. While bursts refer to time-periods with sharp spikes in the loss rate, residual losses refer to the losses that constantly occur in the underlying channel over time. Unlike previous studies on WiLD links in rural

environments [5], we observe a non-zero residual loss-rate in most of our links in urban environments.

To classify each time-period into either a bursty or residual loss-period, we determine a *demarcation region* for the loss distribution on a given link. We estimate parameters p_1 and p_2 ($> p_1$) such that a significant majority ($> 99\%$) of the loss samples fall in the regions $[0, p_1)$ and $(p_2, 1]$. All loss periods with the loss-rate in the range $[0, p_1)$ are classified residual and those in the range $(p_2, 1]$ are classified bursty. The remaining samples are considered transition phases. If adjacent loss periods of a transition period are bursty, then the transition phase is also classified as bursty.

B. Burst characteristics

To analyze burst characteristics, we need to measure the variability of three parameters associated with bursts: duration, arrival pattern and magnitude.

Burst duration and arrival: Based on the duration of bursts, one can classify a burst as either as a *short burst* or a *long burst*. Across our links, we observe a majority of the bursts to be short bursts that last for less than 0.3s – the median loss rate is less than 1s across most links. However, in certain links, especially those in urban environments, we observe a continuous burst period that can last up to 70s. The characteristic arrival pattern that we observed for long bursts is that a single long burst is followed by a string of other long bursts separated by short time-periods (in the order of a few seconds). Overall, the entire string of long bursts that occur together in time lasts for several minutes representing time periods where the underlying channel experiences very high loss rates. Based on the results in Section III, we conclude that these elongated bursts occur due to interference from external WiFi traffic sources.

We next focus on the arrival pattern of short bursts. Figure 8 shows the distribution of inter-arrival times between bursts with short durations for the R-B link in our testbed. For this link, we observe that the underlying distribution of inter arrival time resembles an exponential distribution with a mean inter-arrival time of 15s. In addition, we observe that the inter-arrival time distribution is stationary across various time-periods. These observations suggest that the underlying arrival process can potentially be modeled based on a Poisson arrival process. We observe a similar behavior across all the links in our testbed.

Burst-loss magnitude: We found burst magnitudes to be very hard to predict. For both short spikes and long-duration bursts, the loss-rate varied across the entire spectrum between $10 - 60\%$. Even within a single burst, we observed the loss-rate across episodes to fluctuate rapidly. Given that our links operate in static environments, such wild fluctuations in very short periods appear to be triggered due to external WiFi interference as opposed to multi-path fading channel conditions.

C. Residual loss characteristics

Every link in an urban environment in our testbed exhibits a non-trivial residual loss rate where packet losses occur at regular intervals as opposed to bursts. The residual loss rate varies between $1 - 10\%$ in our urban links in the testbed. However, residual loss-rates are negligible in our rural links. Based on analyzing the loss distributions over different timescales for different links, we make two observations. First, except for one specific link (K-P), we observed that the residual loss distribution is stationary over hourly time scales while on the K-P links, the distribution is time-varying. Second, we observe that the residual loss rate on any link remains roughly constant over a few minutes even in the presence in short bursts during such periods.

D. Implications

In summary, we make three observations. First, we can classify the loss sample at any time period into three categories: short burst, long burst or residual. Second, while the arrival of short bursts can be approximately modeled based on a Poisson arrival process, the arrival of long bursts are highly correlated in time and not memory-less. Finally, unlike rural links which exhibit negligible residual losses, we observe a non-negligible residual loss-rate in urban environments.

VIII. Remedies

Having identified external WiFi interfernce as the main source of packet loss in WiLD networks, in this section we outline the potential remedies to mitigate external WiFi interference. We evaluate adaptive frequency selection, rate adaptation and adaptive forward error correction (FEC) algorithms as the potential remedies. For each, we simulate the adaptation algorithms and measure the improvements gained for real loss traces from our testbed and experiments performed on the wireless channel emulator.

A. Frequency Channel Adaptation

A simple solution to mitigate external WiFi interference could be to select an alternate less congested channel and switch to that channel. To motivate this simple technique we perform a channel switching experiment on our WiLD deployment on the K-P link. The source and destination switch between channel 1 and 11 synchronously every 30 seconds. Figure 10 shows the variability of loss rate across the two channels for a period of about 2 hours. We can observe that both channel 1 and 11 show bursts that stretch upto a few minutes. It is important to note that by averaging the loss rate over 30 seconds we are not capturing the transient changes in the channel conditions.

Given the above loss trace across the two channels, table 9 compares different channel switching algorithms by the achieved loss rate and the no of channel switches required. In the base case (No adapt), where the channel is fixed at either channel 1 or 11, the average loss rate across the entire trace is either 9.2 or 8.3%. If the receiver has complete knowledge of the loss rate on both channels 1 and 11 at the beginning of a time interval (Oracle), then switching to the least lossy channel at any given time achieves the lowest loss rate (at 6.8%); but this comes at a cost of frequent switches of the channel. Adding a

	Loss	No
No adapt	(9.2, 8.3)	0
Lowest rate	6.8	40
Oracle (5%)	7.01	26
Change $\geq$ 10%	7.76	8

Fig. 9. Channel switching algorithms (loss rate and no. of switches)

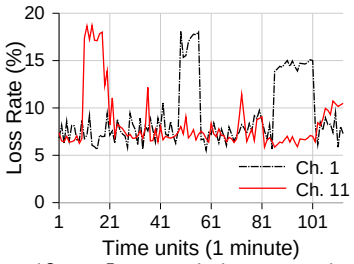

Fig. 10. Loss variation over time across channels 1 and 11 on link K-P; loss rate averaged every 1 minute.

small hysteresis of 5% (Oracle 5%) for channel switching reduces the number of switches from 40 to 26 without increasing the average loss rate significantly. In absence of knowledge of loss rates on other channels, we can use the simple approach of jumping to the alternate channel when the loss rate on the current channel exceeds a threshold (e.g. 10% in Change $\geq$ 10%).

Although the reduction in loss rate shown in Table 9 by the different algorithms is only of the order of 1-2%, the advantages of channel switching could be significant in presence of long or high-loss bursts.

Implications of channel switching: Even though adaptive channel switching seems to be a viable solution, large scale WiLD mesh deployments require careful channel assignment to avoid interference between multiple radios mounted on the same tower [12], [13]. Switching the frequency channel on one link could lead to a cascading effect requiring other links to also change their operating channel. Hence, although it could mitigate interference, it is not always possible to switch a frequency channel in a large scale deployment.

B. Rate Adaptation

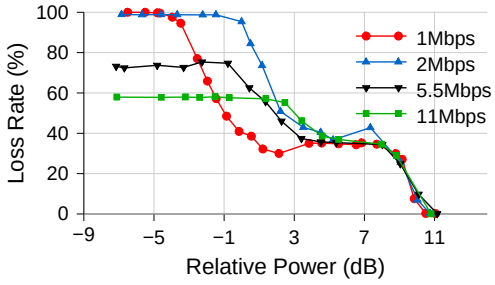

Fig. 11. Loss rate for 802.11b encoding rates at varying relative power of transmitter compared to interferer. Traffic is 1440 byte UDP CBR packets at 11Mbps PHY datarate of 802.11b.

Figure 11 shows the variation of loss rates as the relative power of the primary transmitter is increased with respect to that of the interference source for different 802.11 datarates.

We observed that in presence of external WiFi interference, data rate adaptation could either degrade the performance further or cause no effect on the loss rate. From figure 11 we see that when the received signal strength of the primary transmitter is higher than that of the interference source (from 0 to 12 dB), there is no difference in the loss rate for different 802.11b datarates. Whereas, when the interferer is stronger than the transmitter, reducing the data rate actually exacerbates the performance. This is because the increased transmission time of the frame increases the probability of a collision with the external traffic.

Implications for datarate selection: Most of the 802.11 radios have built in rate-adaptation algorithms which selects a lower rate with resilient encoding on experiencing high loss. However, the above analysis shows that in the presence of loss due to external WiFi interference, it is not worthwhile to adapt the data rate. Rather, we propose using other techniques such as adaptive FEC and link-layer retransmissions to mitigate the loss.

C. Adaptive Forward Error Correction

As discussed in the previous two sections, both channel and rate adaptation may not be feasible in large-scale WiLD networks. Furthermore, they only provide coarse-grain adaptations, which may not be suitable for QoS specific applications like video streaming. In this section we propose adaptive FEC as a solution to achieve fine-grained control. With an estimate of the channel loss variability, adaptive FEC allows addition of the "right" amount of redundancy to cope with the channel losses.

We evaluate a simple Reed-Solomon based adaptive FEC mechanism. Time is divided into slots (25 ms) and at the end of each slot the receiver informs the transmitter of the loss observed in the previous slot. Based on this link information, the transmitter adjusts the redundancy for the next round. To deal with transient spikes in loss rate, the sender maintains a moving window average of the loss rate (WinSize = 10). The application traffic is assumed to be a CBR traffic source (1.8 Mbps) ; there is sufficient bandwidth per slot to introduce 100% redundancy.

	Loss
No FEC	19.98
Oracle FEC	0
Adapt FEC (Win = 10)	4.78

Fig. 12. Effect of adaptive FEC over the trace shown alongside

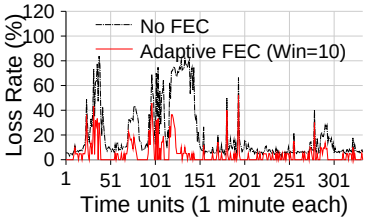

Fig. 13. Comparison of loss rate observed with and without adaptive FEC. Adaptive FEC can significantly reduce the loss rate during periods of long bursts

Figure 13 shows a loss trace on the M-P link. The traffic source was a 1.5 Mbps UDP CBR traffic generator and the loss rate was averaged over 1 minute for a duration of approximately 6 hours. Here again, the MAC-layer ACKs were turned off and retries set to zero. From the above figure we observe that the link was extremely bursty with bursts as high as 70–80% lasting for 20–30 mins. Table 12

shows the performance comparison of the adaptive FEC algorithms. We measure the average loss rate at the end of every slot.

The baseline case is when there is no FEC being applied (No FEC). In this case, the average loss rate across the entire 6 hour period is 19.98%. If the exact loss rate could be predicted for each slot (Oracle), then the loss rate is 0. However, in practice the channel loss rate cannot be predicted accurately, especially since the loss rate is determined by the external WiFi interference. The simple approach where a moving average window of the loss rate is maintained significantly reduces the loss rate to 4.78% (Table 12). Figure 13 shows the loss rate along with the original loss rate. From the figure we observe that the above simple approach can tolerate long bursts of high loss rate. However, FEC cannot adapt to transient high bursts.

IX. Related Work

While there have been several research works on packet loss characterization and methodologies, here, we only focus on those works which are closely related to our work.

Other WiLD deployments: Raman et al. [2] were among the first to deploy a WiLD network consisting of approximately 10 links and lengths ranging from 1–16 km. In [5] they present a detailed performance study of WiLD links. This is the only other performance study of which we are aware. They study the behavior of WiLD links for varying packet sizes, data rates, link lengths, SNRs and weather conditions. Based on their study the authors also experienced high loss due to external interference. In this paper, we present a comprehensive study of the most common sources of packet loss by the wireless channel and the stock 802.11 protocol. In [13] Raman el al. also present modifications to the stock 802.11 MAC protocol to enable point-to-multipoint synchronous transmission and reception in WiLD networks. WiLDNet [12] uses lessons from these measurement studies to design a TDMA based MAC protocol with a synchronization mechanism that is more robust in lossy conditions and with adaptive link loss recovery using bulk ACKs and FEC.

Other measurement based studies: Aguayo et al. [1] present a detailed link layer measurement for a outdoor 802.11 mesh deployment, in which they identify the sources of packet loss. Our study indicates that WiLD deployments are faced with a different set of problems as compared to an outdoor 802.11 mesh deployment.

A large number of measurement based studies have also been carried out to study the source of packet loss in indoor large scale 802.11 deployments [7], [10], [11], [14]. The authors in [11], [14] study the performance of 802.11 in a conference setting, where a large number of clients are using the wireless network. The authors observed both short- and long-term variability in link quality and performance degradation under heavy usage of the wireless network. The authors also point out that rate fallback exacerbates the link quality, leading to a higher number of retransmissions and dropped frames.

X. Conclusions

We perform a detailed study of channel induced (WiFi, non-Wifi, and multipath interference) and protocol induced (timeouts, breakdown of CSMA) losses in WiLD settings. Our main result is that most of the losses arise due to external WiFi interference on same and adjacent channels. This result is in contrast to loss studies of urban mesh networks, where multipath is reported to be the most significant source of loss. We also show that 802.11b protocol limitations make it unsuitable not just for point-to-multipoint links, as claimed in prior work, but also unsuitable for simple point-to-point links. In addition, we analyze the loss variability in both urban and rural links and show that urban links suffer from a higher degree of residual loss. Finally, we propose and analyze the effectiveness of three remedial strategies to mitigate the losses caused by external WiFi interference.

References

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *ACM SIGCOMM*, 2004.

[2] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 Inside-out. *ACM SIGCOMM CCR*, 34:33–38, January 2004.

[3] S. Biswas and R. Morris. Opportunistic Routing in Multi-Hop Wireless Networks. *Hotnets-II*, November 2003.

[4] E. Brewer. Technology Insights for Rural Connectivity. *Wireless Communication and Development: A Global Perspective*, October 2005.

[5] K. Chebrolu, B. Raman, and S. Sen. Long-Distance 802.11b Links: Performance Measurements and Experience. In *ACM MOBICOM*, 2006.

[6] M. V. Clark, K. Leung, B. McNair, and Z. Kostic. Outdoor IEEE 802.11 Cellular Networks: Radio Link Performance. *IEEE ICC*, 2002.

[7] T. Henderson, D. Kotz, and I. Abyzov. The Changing Usage of a Mature Campus-wide Wireless Network. In *MOBICOM*, 2004.

[8] IIT Kanpur. Digital Gangetic Plains. http://www.iitk.ac.in/mladgp/.

[9] International Telecommunications Union. World Telecommunications/ICT Development Report 2006. 2006. http://www.itu.int/ITU-D/ict/publications/wtdr_06/.

[10] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the Real-world Performance of Carrier Sense. In *ACM SIGCOMM E-WIND Workshop*, 2005.

[11] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding Link-layer Behavior in Highly Congested IEEE 802.11b Wireless Networks. In *ACM SIGCOMM E-WIND Workshop*, 2005.

[12] R. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. *NSDI*, 2007.

[13] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *ACM MOBICOM*, August 2005.

[14] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan. Measurement-based Characterization of 802.11 in a Hotspot Setting. In *ACM SIGCOMM E-WIND Workshop*, 2005.

[15] A. Sheth, C. Doerr, D. Grunwald, R. Han, and D. Sicker. MOJO: A Distributed Physical Layer Anomaly Detection System for 802.11 WLANs. In *MOBISYS*, New York, NY, USA, 2006. ACM Press.

[16] Spirent Communications. http://www.spirentcom.com.

[17] WiMAX forum. http://www.wimaxforum.org.