On the Feasibility of the Link Abstraction in (Rural) Mesh Networks

Dattatraya Gokhale, Sayandeep Sen, Kameswari Chebrolu, Bhaskaran Raman
Indian Institute of Technology - Kanpur, INDIA 208016

Abstract—**Outdoor community mesh networks based on 802.11 have seen tremendous growth in the recent past. The current understanding is that wireless link performance in these settings in inherently unpredictable, due to multipath delay spread. Consequently, researchers have focused on developing intelligent routing techniques to achieve the best possible performance. In this paper, we are specifically interested in mesh networks in rural locations. We first present detailed measurements to show that the PHY layer in these settings is indeed stable and predictable. There is a strong correlation between the error rate and the received signal strength. We show that interference, and not multipath fading, is the primary cause of unpredictable performance. This is in sharp contrast with current widespread knowledge from prior studies. Furthermore, we corroborate our view with a fresh analysis of data presented in these prior studies. Based on our results, we argue that outdoor rural mesh networks can indeed be built with the link abstraction being valid. This has several design implications, and opens up a fresh perspective on a wide range of technical issues in this domain.**

I. INTRODUCTION

Community mesh networks based on IEEE 802.11 [1] (WiFi) technology are growing in popularity, both in terms of deployment [2], [3], [4], [5] as well as among researchers [6], [7], [8], [9]. This has been driven primarily by the fact that WiFi is a commodity technology.

Predictable performance of these networks is critical for real-time applications. In fact, for rural areas in developing regions, video-conferencing based applications have been reported to be the primary use of the network [3]. However, multi-hop mesh networks are not really known for predictable performance. In the current literature, the measurement work from the MIT Roofnet deployment [6], [10] is the primary extensive real-world study of performance in these networks. In this work, the authors suggest that unpredictability in performance comes right from the PHY layer [6]. That is, wireless links exhibit widely varying and unpredictable error rates due to large multipath-induced delay spreads in outdoor environments. The study finds little correlation between the received signal-to-noise ratio and the observed link error rate. The conclusions indicate that the very abstraction of a link breaks down: with packet error rate anywhere between 0% and 100%. Consequently, researchers have focused on optimizations at the routing layer [11], [12], [13].

In this paper, we focus on mesh networks deployed in rural regions. Such a consideration is important since such regions form a large fraction of the world today. Some examples of rural mesh networks include [3], [4], [5], [8]. We first motivate why the performance behaviour of links in such networks needs a revisit. We then present detailed measurements to show that

links can indeed have predictable performance. While such measurements have been done for WiFi-based Long Distance (WiLD) links earlier [7], [8], in this work we consider more traditional, short-distance links (e.g. deployed within a village).

We show that link error rate is strongly correlated with RSSI (received signal strength indicator) or the SNR (signal to noise ratio). There is a clearly identifiable RSSI threshold (close to the card sensitivity measured in controlled settings) beyond which the error rate is close to 0%. These observations are in contrast with those of Roofnet [6].

We provide strong evidence to indicate that external interference, and not multipath induced delay spread is the primary cause of unpredictable link behaviour and lack of dependence on the RSSI. In contrast, Roofnet concludes that it is unlikely that foreign 802.11 packets are responsible for the observed wireless packet errors [6]. We provide evidence for our view not only based on our own measurements, but also using a fresh analysis of the data from Roofnet itself [14].

Our analysis indicates that the conclusions in Roofnet are likely *incorrect*, not only for rural mesh networks, but also for urban mesh networks. This is significant, since the Roofnet study is widely cited, and is also used in follow-up work on routing protocols [7], [8], [9], [15].

Finally, based on short as well as long-running experiments, we show that the received signal strength indicator (RSSI) is stable, with only a small band of variation of 3-4 dB, for most links.

Our results have a wide variety of implications. The fact that external interference, or "RF-pollution", is the main factor which causes unpredictable performance, does not bode well for building mesh networks with predictable performance in urban settings, where such RF pollution is quite prevalent. However, in rural settings where there is little network connectivity to begin with, it is intuitive to expect that external interference will be minimal or non-existent. This intuition is also corroborated by practical experience in several deployments [7], [8], [3]. In such settings free of RF-pollution, our measurement results imply that outdoor links can be planned to have predictable performance by having RSSI above a threshold. Or in an already deployed network, links can be classified as existing or not existing based on the RSSI threshold. In other words, the link abstraction *does* hold, and forms a strong foundation on which to build applications with predictable performance. This forms the basis for a fresh perspective on technical issues in rural outdoor mesh networks, as articulated in our parallel work [16].

Furthermore, in our setting, we argue that routing metric

and routing protocol design as considered in existing mesh networking literature [11], [12], [13], [15], [17] are unlikely to be applicable in our setting, since these assume the *absence* of a link abstraction. Further, metrics such as ETX [12] & WCETT [17] try to distinguish between links which have intermediate loss rates (between 0% and 100%). In our setting, we encounter such links only while operating at or near the RSSI threshold. We show that trying to distinguish between such links can lead to unstable behaviour.

Paper outline: The next section (Sec. II) presents a brief background of the FRACTEL project and motivation for our work. Sec. III presents our detailed measurement study and analysis. We also present here, our interpretation of the data collected in [6], as well as evidence to support this interpretation. Subsequently, in Sec. IV, we describe our results on the stability of link behaviour over time. The results from our measurements have several implications, which we articulate in Sec. V. Finally, we present a few points of discussion in Sec. VI, and conclude in Sec. VII.

II. FRACTEL: BACKGROUND AND MOTIVATION

Community networks based on multi-hop 802.11 have been deployed around the world. There are two main categories of such outdoor mesh networks in the literature : (1) long-distance, with links up to several tens of kms, and (2) shorter distance, with most links below 500 m. Examples of the first kind of network are the Digital Gangetic Plains (DGP) project [18], and [8]. Examples in the second category are the MIT Roofnet project [19], and the various mesh deployments in [2].

In the FRACTEL (wiFi-based Regional/Rural data ACcess and TELephony) project, our goal is to build WiFi mesh networks to support data and multimedia applications, in rural settings. In a typical FRACTEL deployment, we wish to use the mesh to extend Internet connectivity from a single point to multiple nearby buildings. The single point may have connectivity from a wired connection, a satellite connection, or perhaps from a long-distance WiFi mesh like DGP itself. From this point, we may require to provide Internet connectivity to nearby buildings such as residential houses, community centres, schools, hospitals, or government offices. Examples of current deployments which fall under this category include [4], [5].

In many respects, FRACTEL resembles Roofnet more than long-distance mesh networks. Most links are expected to be short (up to 500 m). And importantly, we do not wish to use expensive, tall towers like in DGP [3]. However, the envisioned deployment environment for FRACTEL is quite unlike that of the Roofnet study: a dense urban setting is not our primary deployment target. In our setting, we expect a few buildings, two or three storeys tall in rural environments or sometimes even in semi-urban residential communities [20]. We hence expect the multipath behaviour in FRACTEL to be somewhere in-between that of Roofnet and DGP.

This then brings us to the question of whether the link characteristics in FRACTEL are going to be like that of Roofnet, or like that of DGP. The results from these two kinds of networks show starkly contrasting results. Table I provides a comparison of the main measurement results from Roofnet [6]

and DGP [7], along with the open questions in the context of FRACTEL. In the course of our study, we have not only found results contrasting those reported in [6], but also have re-analyzed the data from [6] itself, to arrive at alternative conclusions. The rest of this paper details our results and their implications.

III. ERROR RATE ANALYSIS

We now present our methodology, followed by our measurements, and then a fresh analysis of the data from Roofnet.

A. Experimental Methodology

Environment: For our measurements, we chose two kinds of environments: a rural village & a residential university campus. We chose one location within the village and five locations within the campus. For each *location*, we fixed one transmitter *position*, and varied the receiver *position* over six different choices. We thus have a total of 36 links.

The link lengths were in the range 150-400 m. The distances involved in our study are similar to those in the Roofnet study [6], except that the Roofnet study had a significant number of links over 500 m. Our choice of the environment as well as the transmitter/receiver positions have been driven by what we expect in a typical FRACTEL deployment in a rural scenario. The locations in the residential campus also had building heights and building density akin to the village location (densely packed houses, with most buildings at most two or three storeys tall, and scattered trees). The campus environment specifically helps us in studying interference versus multipath as being the cause for packet error rates.

It is worth noting that although we have only six locations in our measurement set, the *uniformity* of results across these six gives us confidence in our results and their implications. Fig. 1 shows the village location and the five on-campus locations. For each location, the transmitter positions are marked with a small circle in the figure. The receiver positions were all within the ellipse marked for each location. A brief description of each of the locations is as below.

(1) Village (Vill): we had a 400 m x 400 m area with densely packed houses (1-2 storeys tall), and a few scattered trees. *(2) Student Dormitory (Dorm):* a student residential dormitory, with four rows of three storey tall dorm rooms & a few scattered trees in the vicinity. *(3) Housing Apartment Area (Apt):* several rows of two-storey housing apartments on campus, with dense tree cover for a residential area. *(4) Housing Apartment Area-2 (Apt2Dorm):* at this location, we had the transmitter at another housing apartment area and receivers at the dorm area; there were a few scattered trees in-between. *(5) Academic Area Corridor (Acad):* within the university campus, with several buildings in the vicinity; for approximately a 100 m portion of the corridor, it was flanked by buildings about 3 m from the corridor, on either side; the rest of the portion of the corridor was open. *(6) Ground Area (Gnd):* near the *Apt* area above, we had a play ground, and an adjoining small forest like area with dense foliage.

Hardware: We used the Senao 2511CD plus ext2 PCMCIA cards, inserted into laptops (1.7GHz Intel Pentium, 512MB

	Typical link distances	Network architecture	Environment	Multipath effects	SNR or RSSI	External interference	Link abstraction
Long-distance mesh networks (e.g. DGP)	Up to few tens of kms	High gain diml./sector antennas on tall towers/masts	Rural setting studied in depth	Effect not apparent	Has strong correlation with link quality	Affects links performance	Valid
Rooftop mesh networks (e.g. Roofnet)	Mostly < 500 m	Mostly omni antennas on rooftops	Dense urban setting studied in-depth	Reported as a significant component	Not useful in predicting link quality	Reported as not significant	Not valid
FRACTEL	Mostly < 500 m	Would like to avoid tall towers	Rural, semi-urban	*To be determined*	*To be determined*	*To be determined*	*To be determined*

TABLE I
LONG-DIST. VS. ROOFTOP MESH NETWORKS, AND OPEN ISSUES IN FRACTEL

Fig. 1.   Maps showing experiment locations (courtesy Google Maps)

RAM) for our experiments. It is worth noting that both the Roofnet study [6] as well as the DGP study [7] used this same type of WiFi cards. Using the external connectors of the PCMCIA cards and RF cables, we connected 8 dBi omni antennas at the transmitter and the receiver (Roofnet too used similar 8 dBi omni-antennas). The antennas were mounted on a small tripod stand about 1.5 m tall.

Software: The laptops used Linux 2.6.11, and HostAP driver version 0.4.9. We instrumented the driver to pass various per-packet information to the user level at the receiver: RSSI, noise (silence) level, rate (1, 2, 5.5, or 11 Mbps), MAC header details, and whether or not the received packet passed the CRC check. The receiver was set in monitor mode for all of the experiments. The transmitter sent packets at a fixed transmit power and transmit rate, in adhoc, pseudo-ibss mode, with an inter-packet gap of 20 ms. In each experiment, the transmitter sent 6000 packets (overall duration of 2 min). Unlike [6], we specifically did not send packets back-to-back, which enables us to observe the effect of external interfering packets more directly[1]. Apart from the 2 min experiments, we had long-running experiments (24-48 hrs) at the *Apt* and *Apt2Dorm* locations.

Choice of transmitter/receiver positions: At each of the

[1] Inter-packet interval = 20 ms & per-packet transmit duration $\simeq$ 12 ms at most imply that we will most likely capture some foreign packets (if any) in the gap $\geq 8ms$ between two successive packets.

five locations, we first chose a convenient transmitter position. The transmitter was placed at an elevation, atop a building terrace in all cases except *Gnd*. At *Gnd*, the transmitter mast was placed at a clearing, on a 1.5 m tall tripod.

For the receiver position, we define a *good* position to be one where there was clear line-of-sight, and the mean RSSI was about $-70dBm$. The RSSI was calculated over an initial set of 1000 packets sent from the transmitter. We define a *medium* receiver position to be one which had some foliage or building in-between, and the RSSI thus calculated was about $-75dBm$. A *bad* receiver position is one where the RSSI was about $-80dBm$ or so, and there was no clear LOS. At each location, we chose a combination of good, medium, and bad positions.

B. FRACTEL Results

The primary characteristic for wireless links is the error rate, measured as a function of the RSSI [6], [7]. The correlation between the two to a large degree determines whether or not the link abstraction holds. An understanding of this is essential for any higher layer mechanism or application metrics like throughput. Hence we look at this aspect in-depth.

We observed that for many transmitter-receiver position pairs, there was some variation (up to 4 dB in most cases) in the RSSI, across the 6000 packets in the 2 min duration. Hence we decided to separate out the 6000 packets into 60 bins of 100 packets each. Note that these are bins of 100 *transmitted* packets, not 100 *received* packets. We also note that if a packet is received with an error in the CRC check, its reported RSSI reading is still reliable, since it is measured at the receiver's card.

For each of the 60 bins, we compute the average RSSI and the error rate. If in a bin we do not have any received packets (not even with CRC error), we take the RSSI to be $-100dBm$. We note that the noise, or silence level in most of our experiments was $-94dBm$ to $-95dBm$. So plotting the error rate against the RSSI is equivalent to plotting against the SNR.

When we originally plotted the error rate versus RSSI graph, we observed several points which showed a high error rate despite a high signal strength ($-75dBm$ or above). When we looked into our receiver side logs for these cases, we observed that there were lots of packets from external WiFi sources. This was especially so at the *Acad*, and *Gnd* locations. At the *Acad* location, there were many external WiFi sources in the vicinity. And at the *Gnd* location, the signals were likely coming from a long-range WiFi link setup by someone in the vicinity.
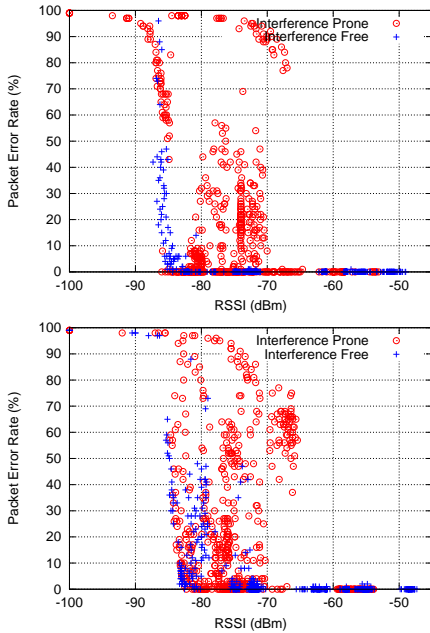
Fig. 2. Error rate vs. RSSI (1, 11 Mbps resply.)

We then separated out cases where our receiver logs showed foreign packets from those which did not. We term these as *interference-prone* and *interference-free* cases respectively. In our data at the interference-prone locations, we observed anywhere from about 500 to over 90,000 foreign packets over the 2 min duration.

We plot the error rate versus RSSI for the two categories of locations. Fig. 2 shows such plots for the experiments with transmit rates fixed at 1 Mbps and 11 Mbps respectively[2]. We plotted similar graphs for the 2 Mbps & 5.5 Mbps rates too, but do not show them for lack of space: the observations made below hold for these transmit rates too. Each point in each plot represents a 100 packet bin from a particular transmitter-receiver position pair.

We clearly see that at the interference-free positions, there is a very strong correlation between the signal strength and the observed error rate. There is a threshold signal strength above which the error rate is more or less negligible. For the 1 Mbps & 11 Mbps rates, this threshold is about $-86dBm$ & $-79dBm$ respectively. For 2 Mbps & 5.5 Mbps, this threshold was about $-85\ dBm$ & $-80\ dBm$ respectively.

To see how these threshold values compare with the scenario where there is no wireless channel, we used the same cards to perform a controlled experiment indoors when the transmitter and receiver cards were connected by an RF cable. In these controlled experiments too, the noise floor was between $-95dBm$ and $-94dBm$, much like in our outdoor experiments. The threshold RSSI values above which the error rate fell below 1% was $-88dBm$, $-86dBm$, $-82dBm$, and $-81dBm$ respectively for the four transmit rates. These are close to the card's sensitivity values. We observed similar values with other cards of the same make.

The RSSI threshold values observed in Fig. 2 are thus within

---

[2]Like in [6], [7], we turn-off the card's auto-rate fallback; the error rate measurements in fact serve as input for the design of any auto-rate mechanism.

1-2 dB of the threshold values observed in our RF cable experiment.

There are a few isolated points in Fig. 2 (around RSSI $\simeq -75dBm$, for 11 Mbps), where even at the interference-free locations, we have a significant error rate. A look at the receiver log for the experiment corresponding to this revealed that there was significant RSSI variation during the first 15-20 seconds of the experiment, which caused most of the packet losses. After this initial period, there were almost no RSSI variation or packet losses. We think this is likely due to experimental error.

In both the four plots, we see that at the interference-prone positions, correlation between the error rate and the RSSI breaks down, and there are several cases of high error rate even in the presence of high signal strength. In the figure, we can visually identify clusters of points for the interference-prone locations. The different clusters correspond to different locations, and the clustering is due to the different levels of interference at these locations.

For one of the positions in the *Apt* location, we were clearly able to see the dependence on external interference. We initially ran our 2 min experiments and observed interference at this position, and high error rates. We then identified the interference to be from a known WiFi source. We then had the interference source temporarily shut down. When we then re-ran the experiment, the error rates came down to close to zero.

We now compare the above results with a fresh analysis of the data from Roofnet.

C. Roofnet Data: A Fresh Analysis

Our experiments presented in the previous section show that external interference is the main culprit in causing high error rates at RSSI values above the receiver sensitivity. How does this compare with the conclusions from existing literature on measurements in this domain? The main measurement study available on outdoor mesh networks, which is widely cited, is the Roofnet study [6]. Hence it is imperative that we look at this in depth.

The authors in [6] make the following observations (among others). (1) Packet loss cannot always be attributed to low SNR; there are several cases where the loss rate is high even at high SNR (Fig. 14, 15 in [6]). (2) There is no correlation between the number of lost packets and the number of foreign (interference) packets (Fig. 18 in [6]). (3) Introducing delay spread causes high loss rates (Fig. 19 in [6]). (4) Prior measurements in urban environments have reported delay spreads of over $1\mu s$ [21], [22], which is well beyond the tolerance limits of the 802.11b receiver [23].

Based on the above observations, Roofnet suggests that, since external interference does not seem to be a factor, multipath induced delay spread in excess of $1\mu s$ is the cause of high loss rates at high SNR. This is contrary to our conclusion. On the surface it appears that this is because of the different experimental environments: dense urban versus rural/campus. That is, while the study in [6] has encountered significant multipath, our measurements have not. However, on a closer

examination of the Roofnet data, we have observed strong evidence to the contrary.

We have looked at the data reported in [6], available from [14]. A strange pattern we observed was that the noise floor measurements in this data were not only much higher than ours in many cases, but also showed significant variation across the various packets.

Now, the noise, or silence level as reported by the card, for a particular received packet, is the energy level as measured before the packet reception began. So irrespective of the presence or absence of multipath, the noise level should remain more or less constant for the various packets. This is what we observe in the data we collected in FRACTEL. Most packets show a noise level of $-94dBm$ or $-95dBm$.

To show how the Roofnet data shows very different behaviour, we choose the subset of the transmitter-receiver pairs where the average RSSI was over $-77dBm$ at the 11 Mbps data rate (well above the measured card sensitivity of $-81dBm$), and those which showed a loss rate between 20% and 80%. That is, where we have high signal strength but still high error rate. We have 26 such points. Fig. 3 plots the 5%-ile, 50%-ile (median), and 95%-ile noise levels for these points, in increasing order of the median noise level.


Fig. 3.   Noise levels in Roofnet data (11 Mbps)

Fig. 4.   Noise levels in Roofnet data (1 Mbps)

We wish to note that for these points, the average SNR too was very high (11 to 39 dB). We observe that the median noise levels are as high as $-86dBm$, with most median values at $-90dBm$ or above. We also observe that the 95%-ile minus 5%-ile (band where 90% of the values lie) can be as high as 16 dB! In contrast, the noise levels in our FRACTEL data was at most $-94dBm$, with a variation of only about 1 dB.

Fig. 4 plots a similar graph for 1 Mbps data points. Here we chose only those transmitter-receiver pairs where the average RSSI was over $-80dBm$ (again, well above the card sensitivity of $-88dBm$) and the error rate was between 20% and 80%. Once again, we see behaviour similar to the earlier plot.

Fig. 5 shows the variation of the RSSI and noise levels, as a function of the packet number, for the data point no: 13 (points are numbered 0 to 25) of Fig. 3. A CDF plot (not shown here) for the same data point showed that about 50% of the packets show a noise level of about $-90dBm$ or more!

The only explanation for the high and variable noise levels is the presence of several 2.4 GHz sources. More likely than


Fig. 5.   Per-pkt varn. of RSSI and noise levels

Col-1	Col-2	Col-3	Col-4	Col-5	Col-6	Col-7	Col-8	Col-9
Expt No:	Source	Mean RSSI (dBm)	Loss % from source	Mean Noise (dBm)	5%-ile Noise (dBm)	95%-ile Noise (dBm)	Noise Band (dB)	Max. Noise (dBm)
1	A	-89.74	100.0%	-93.26	-94	-90	4	-88
1	B	-75.59	0.5%	-92.1	-94	-88	6	-88
2	A	-85.23	99.2%	-92.53	-94	-85	9	-85
2	B	-74.68	18.3%	-89.34	-94	-85	9	-84
3	A	-80.69	63.2%	-90.85	-94	-80	14	-80
3	B	-75.73	37.2%	-85.16	-94	-80	14	-80
4	A	-75.25	39.8%	-93.06	-94	-92	2	-74
4	B	-75.11	61.3%	-90.18	-94	-75	19	-74

TABLE II
CONTROLLED INTERFERENCE EXPT. RESULTS

not, these are other WiFi sources: it is hard to imagine such a wide prevalence of non-WiFi 2.4 GHz sources.

We are now faced with a few follow-up questions. Why did Roofnet not observe correlation between the number of lost packets and the number of foreign packets? More importantly, how does the noise level reported by the card compare with the level of interference? Can such information be used for any interference-aware routing? We next present controlled experiments to understand these aspects.

D. Understanding Interference

The setup for our controlled experiment is depicted in Fig. 6. We have two transmitters A and B and one receiver R. R's card has two connectors for the external antenna (for diversity). We make use of these two connectors to connect A to R and B to R respectively. This is shown in the figure. R is in monitor mode. In this arrangement, A and B cannot hear one another, but R can hear both (i.e. a case of hidden nodes). Both transmitters were in pseudo-ibss mode, had auto-rate disabled, and were transmitting all packets at the 11 Mbps rate, with an inter-packet interval of 2 ms. Each experiment ran for about 2 min.


Fig. 6.   Setup for interference expt.

In this setup, A and B act as interference to one another. We fixed the attenuators such that the mean RSSI of B's packets at R was about $-75dBm$. We varied the attenuator at A across four experiments, such that we had mean RSSI values from A to be about $-90dBm$, $-85dBm$, $-80dBm$, and $-75dBm$.

From R's log, we calculate various statistics. These are summarized in Table II. The obvious aspect which stands out is that as A's RSSI increases, the loss rate of A's packets

decreases and that of B's increases (Col-3, Col-4 in the table). We make several subtle but significant observations below.

P1: First, the noise as well as the noise band are quite high in almost all of the cases in Table II (Col-5, Col-8). These are similar to the noise levels observed from Roofnet data in Sec. III-C.

So, *interference does cause the noise level to be high and variable.*

P2: Again focusing on experiment no: 2 of Table II, we see that so far as B's packets are concerned, there is a loss of 18.3% (Col-4). But then, the number of A's packets received are very few. With a loss rate of 99.2%, and a sending rate of 500 pkts/sec, A's packets are received at an average rate of only 4 pkts/sec. But this is sufficient to cause an 18.3% loss rate, which amounts to 91.5 lost pkts/sec! Even when we shut off B's transmissions and had only A transmitting, we observed that the number of received packets of A was low (it had about 99% loss rate). This was because the average RSSI from A is only around $-85dBm$, much below the sensitivity as measured earlier ($-81dBm$, see Sec. III-B).

So, *the packet loss rate can be high even when the number of observed interference packets is low.*

P3: Not related to our experiment above, it is easy to see why *the packet loss rate can be low even when the number of interference packets seen is high.* This can happen when the two transmitters can hear one another. So the interference simply backs off when transmission on the link of our interest starts.

Now, Roofnet used the following methodology to rule out external interference as a significant cause of packet error rates (Sec. 8 in [6]). On each link, they first measure the average rate at which foreign packets seen in a 90-sec duration. Then they measure the packet error rate seen on that link in the immediate next 90-sec period. A scatter plot of these two shows no correlation (Fig. 18 in [6]), and based on this they conclude that foreign WiFi sources are not a significant source of packet error rates.

P3 and **P2** taken together indicate how we may have no correlation between the foreign packet rate and the observed error rate, and still have all of the packet errors caused due to interference. This, taken along with the high and variable noise levels (Fig. 5) in Roofnet, leads us to conclude that external interference did play a major role in causing their error rates. This then also raises sufficient doubt on their conclusion of multipath delay spread being the main cause of packet errors in their environment.

E. Gauging the Level of Interference

P4: We now look at the question whether the card reported noise level be used to gauge the level of interference, using the results from the above controlled experiments. When we plotted the per-packet noise level of A's or B's packets, we observed a high degree of variation even in our highly controlled environment, similar to the variation in Fig. 5. Further, in Table II, we compare the card reported noise level (Col-5 to Col-8) with what we know to be the interference level (Col-3). Surprisingly, for B, the reported maximum noise level (Col-9) across the set of packets seems to correspond to the interference level (i.e. RSSI from A). However, there is no such relation for A's reported maximum noise levels. For example, the maximum noise level seen for A's packets in experiment 2 was only $-85dBm$, whereas A's interference B, we know was at $-75dBm$.

The variable noise floor can be explained as follows. The hardware for the Intersil Prism2 based cards maintains a noise floor based on an average of 256 samples [23]. The noise level reported by the card for a received packet, is the noise floor just before that packet's reception started. This value thus depends on the exact timing of the received packet, with respect to the interference traffic. This of course is unpredictable and variable.

What the variability implies is that, just reading the noise level to gauge the level of external interference, can be very error-prone, at least on this hardware.

P5: Can one estimate the link performance based on the average measured noise floor for packets? To explore this possibility, we plot the observed noise floor versus the error rate. We compute this within 100-packet bins for the same transmitter-receiver pair in Roofnet as for Fig. 3. Fig. 7 shows this plot. We see that for a given average noise floor, there is a wide range of error rates possible! This means that we cannot really estimate the expected interference, or the resultant link behaviour, based on the average noise floor either.


Fig. 7.   Avg. noise level versus error rate, 100 pkt bins

IV. STABILITY OVER TIME

We saw above that link behaviour is unpredictable in the presence of interference. This leaves us with the question of whether it is possible to build links with predictable performance in interference free environments.

Apart from the dependence of error rate on the RSSI (or SNR), the other element of predictability is the stability of the RSSI. That is, if the RSSI is (un)stable, the error rate can also be expected to be (un)stable. We are interested in knowing the stability at (a) small time scales, as well as at (b) large time scales. The former is important since it may affect routing decisions and the stability of any routing protocol dependent on link performance. The latter is of significance if we are trying to provision a link during a planned deployment. Or in an already deployed network, when we are trying to determine what the transmit power should be for two nodes to connect to one another.

To capture the short term variation in RSSI, we consider data from our 2 min experiments. For the various interference-free receiver positions in our experiments, we have calculated the variation of the per-packet RSSI. We express this variation in terms of the 5%-ile, the 50%-ile (median), and the 95%-ile. Fig. 8 shows the plot of these values, along with the mean RSSI,

for the various interference-free positions. The figure includes data from a total of fourteen different positions, and all four data rates. The points in the plot are sorted in increasing order of the median RSSI.
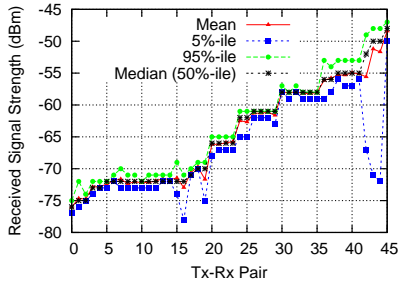

Fig. 8. Variation in the RSSI

The band between the 95%-ile and 5%-ile represents the band within which 90% of the packet RSSI values will lie. We see in the figure that for most of the experiments, this band is within 3-4 dB. For three of the cases, pair number 16, 19, and 36 in Fig. 8, the band is 6-7 dB. All of these were cases where we did not have clear LOS between the transmitter and the receiver. A look at the variation of RSSI with time revealed that in both cases there were periods of several seconds during which there was a marked drop in the RSSI. This likely indicates some person or obstacle in-between in that duration.

Fig. 8 has three cases (pair numbers 42, 43, 44) where the band was 18 dB, 23 dB, and 24 dB respectively. A look at the RSSI variation with time revealed that there were several data points which had about 20 dB lower RSSI. We have determined this to be a hardware quirk in the particular card make. Such sudden drops in RSSI can be seen even when the wireless channel is eliminated and the transmitting & receiving cards are connected via an RF cable. This hardware quirk has also reported in other studies [7]. But for these quirks, these data points too had a narrow RSSI band.

But for these exceptions, we can safely say that we can expect the RSSI variation, although dependent on the environments, to be within about 3-4 dB is most cases. In all of our LOS cases, we observed a band of at most 4 dB.

We also analyzed similar statistics for the Roofnet data, and our data at the interference-prone positions. We observed a similar pattern: the 95%-ile to 5%-ile band was within 5 dB for most links. There were a few links in the Roofnet data which showed larger bands (6 to 11 dB).

To understand RSSI variation over longer durations of time, we have run two separate experiments at the *Apt* and *Apt2Dorm* locations (see Fig. 1). In each case, we had one transmitter and three receivers at three different positions. This is marked in Fig. 1. At *Apt*, the experiment ran for a duration of 48 hours, while at *Apt2Dorm*, the experiment ran for 24 hours.

Table III shows a summary of the results from the six transmitter-receiver pairs. We see again that in the LOS cases, the 95%-ile minus 5%-ile band is within 4 dB. Even for one of the non-LOS link, the variation was small (2 dB).

In sum, over short time scales as well as larger time scales, we have a small variation band of about 3-4 dB in most cases: LOS as well as non-LOS. In LOS cases, the band never

Location	Rx-Posn.	LOS?	Durn. (hrs)	RSSI 95%-ile (dBm)	RSSI 5%-ile (dBm)	RSSI band (dB)
Apt	1	Yes	48	-66	-69	3
Apt	2	No (some foliage)	48	-69	-77	8
Apt	3	No (some foliage)	48	-76	-82	6
Apt2Dorm	1	Yes	24	-75	-77	2
Apt2Dorm	2	Yes	24	-70	-71	1
Apt2Dorm	3	No (some foliage)	24	-79	-81	2

TABLE III
LONG-TERM RSSI VARIATION

exceeded 4 dB. In a few non-LOS cases, a few positions showed bands larger than 4 dB.

The short term variation tells us when we *cannot* expect predictable behaviour. Note that the steep portion of the error rate versus RSSI plot (Fig. 2) is only about 4-6 dB wide. Given that the RSSI variation itself can be 4 dB or so, we cannot expect any stability (short term or long term) in the error rate if operating at or near this region.
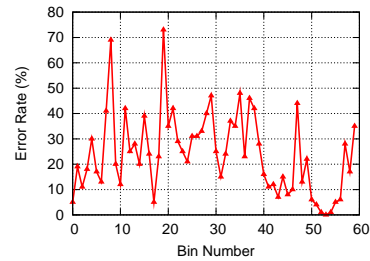

Fig. 9. Variation in the error rate

To illustrate the above aspect, we consider a receiver position in our data which showed an overall loss rate of 25%, at the 11 Mbps data rate. We picked this since the error rate is neither close to 0% nor to 100%. This is one of the positions in *Vill*, with an average RSSI of $-80.5 dBm$. The RSSI band for this position was $-82 dBm$ to $-79 dBm$. For this experiment data, we plot the observed error rate over 100 packet bins, as a function of the bin number (or equivalently time), in Fig. 9. This figure indicates why it is not safe to operate near the steep region of the error rate versus RSSI plot. There is significant variation in the error rate across just 100 packet bins, due to the variation in RSSI across the steep region of the curve in Fig. 2. We observed similar variation in other experiments too, where the error rate was between 0% and 100%.

V. DESIGN IMPLICATIONS

In this section, we describe the main implications of our measurement results. To summarize, the key novel points we make in the context of outdoor mesh networks are the following.

- Multipath does not show up in any significant manner, at least for the variety of deployment scenarios in which we have conducted experiments. All cases of packet error rates in our study can be attributed to external interference.
- RSSI is indeed a good predictor of link quality, with the threshold being within 1-2 dB of the threshold measured in a controlled environment using RF cables. Beyond the threshold, we can expect stable and low error rates.
- When operating at or near the threshold, loss rates can be unpredictable.

- RSSI variation is within a band of about 3-4 dB, over short as well as large time scales for most cases.
- External interference *is* a very significant factor and can cause high loss rates. In our experiments, it is the main factor which causes unpredictability in link performance.
- Gauging the level of external interference based on observed noise levels appears to be quite difficult, at least on the hardware we have used.

There are a wide variety of implications of the above points. We believe that they present a fresh perspective on a wide range of technical issues. We articulate this now.

A. The Link Abstraction

Much of the approach in building and managing outdoor community networks thus far has been based on the assumption that link abstraction is absent; that error rates are unpredictable due to multipath, which is not in our control. We have shown this to be untrue in our setting.

In the absence of external interference, our data on the long-term RSSI variation tells us how to achieve the link abstraction. Suppose we wish to build a mesh network. For a desired link between two nodes, we need to ensure that the RSSI threshold is above what is given in our error-rate versus RSSI plots. For e.g. from Fig. 2, this threshold would be $-79dBm$ for 11 Mbps links. Furthermore, we can expect an RSSI variation of 3-4 dB on larger time scales. To account for this, the RSSI must be set with a head-room of 3-4 dB higher than the above-mentioned threshold.

Such an approach can be taken also when determining what the transmit power should be at two nodes seeking to form a link, in an already deployed mesh network. For links formed as above, we can expect stable and predictable behaviour and the link abstraction will hold. The link would perform more or less like a wired link. This would simplify higher layer protocol design, and give a strong foundation on which to build applications which expect predictable performance.

B. Implications on Routing

Routing metrics: In the absence of the link abstraction, much work has focused on the design of appropriate routing metrics [11], [12], [13], [17]. These essentially seek to distinguish between links which have loss rates in-between 0% and 100%. This would happen in our setting if we were to operate at or near the threshold.

As shown in Fig. 9, such operation can lead to high variations in the error rate, which is unpredictable. This in turn would mean erratic behaviour at the routing layer if we use metrics such as ETX [12] or WCETT [17].

Opportunistic routing: In this technique, the approach to handle intermediate loss rates is to opportunistically use packet reception whenever it succeeds [15]. Such an approach can be used independent of whether the losses are caused due to multipath or due to external interference. But it appears quite daunting, if not impossible, to achieve any performance guarantees in such settings. This may be the best option if there is no way to control the external interference. Fortunately, the consideration of such adaptation is unnecessary in rural settings, since we do not expect external interference to be widespread.

Interference aware routing: There are several prior efforts which have focused on interference-aware routing (e.g. [24], [25]). Most of these seek to mitigate *internal* interference, i.e., interference among the links of the wireless mesh itself. The work in [26] seeks to gauge such interference and predict link performance. It also considers modeling external interference, based on the observed signal strength variation. However, in our experiments, we have not observed any significant *additional* RSSI variation due to external interference: there is already 3-4 dB variation even without interference.

Further, our measurements in Sec. III-E (**P4** & **P5**) indicate that gauging the level of external interference based on either the card reported noise level or even based on the average noise level (as suggested in [24]) can be error prone.

C. Implications on MAC

It is well known in the literature that the 802.11 CSMA/CA MAC is not suited for multi-hop mesh networks. It causes self-interference. That is, interference across multiple links in a path. Given the prevalence of external interference in our own measurements as well as in the Roofnet data, it appears all the more unlikely that CSMA/CA can achieve any predictable performance in such mesh networks. The use of RTS/CTS may not really help: as shown in Roofnet and as explained in Sec. III-D (**P2**), we can have several interference sources at interference range but not in reception range.

The feasibility of the link abstraction on the other hand, opens the door for TDMA-based MAC approaches. Prior work [27], [28], [29] has already shown prototypes of TDMA-based MAC implementations on WiFi hardware. However, multi-hop TDMA implementation and scheduling are still open issues. In parallel work, we have articulated how such issues pan out in FRACTEL, in the presence of the link abstraction [16].

VI. DISCUSSION

Multipath and delay spread: In the environments in which we have tested, multipath induced delay spread is clearly not a significant factor. And we have shown that in the Roofnet data too, external interference played a major role in causing error rates. But we stop short of ruling out multipath induced error rates in dense urban settings with several tall buildings, since the delay spread handling capabilities of 802.11b hardware are limited.

Roofnet cites [21], [22] which have measured delay spreads in such urban environments to be $1\mu s$. But then these studies have been done in the 910 MHz cellular band, not for the 2.4 GHz WiFi band. One has to be cautious while extrapolating such measurements across a wide range of frequencies. We would expect very different propagation behaviour for 2.4 GHz as compared to 910 MHz. Only a careful measurement can tell what the multipath delay spread values will be for 2.4 GHz in urban environments.

802.11g and 802.11a: Our measurements have been restricted to 802.11b. The results with respect to the lack of

the effect of multipath and the effect of interference will likely extrapolate to 802.11g too, since it operates in the same 2.4 GHz band. For 802.11a, which operates in the 5.2 GHz band, there are two aspects which will come into play.

First, we believe it is likely that multipath will have even less of an effect on 802.11a, since the attenuation levels at 5.2 GHz is higher than at 2.4 GHz. But only actual measurements can confirm this.

Second, more importantly, 802.11a has more frequencies of operation: 12 total, of which 8 are non-overlapping. This means that avoiding RF-pollution in this case is bound to be easier. In fact, if a community or the deploying entity so chooses, it may use a combination of 802.11a and 802.11b/g radios to further alleviate the issue.

For both 802.11a and 802.11g, the delay spread handling capabilities of the PHY layer are better than for 802.11b, at comparable data rates of operation (e.g. see [30]).

802.16: The upcoming IEEE WiMAX standard [31] offers much promise in the domain of outdoor long-distance networks. In this context, it is relevant to question whether 802.11 is indeed the right technology to use for outdoor mesh networks. It is hard to predict whether a technology will catch on, but 802.11 is the metaphorical "bird-in-hand". It has already gained widespread acceptance and several community mesh networks have already deployed it. 802.11 has also achieved economies of scale suitable for deployment in rural regions, where system cost is an important factor [18]. So there is sufficient motivation for exploring the limits of 802.11. In fact, our measurements show promise in the direction of using 802.11 to build mesh networks with predictable performance.

VII. CONCLUSION

The goal of the FRACTEL project is to build mesh networks for deployment in rural settings. The consideration of rural settings is significant; after all, a large fraction of the world's population is rural, especially in developing countries. We wish to achieve predictable link performance to enable real-time services. The PHY and link layer behaviour are critical to understand in this regard. We have undertaken a detailed measurement study for this. We find that the link abstraction is indeed possible to achieve, contrary to popular belief for outdoor mesh networks today. We not only analyze our own measurements, but also perform a fresh analysis of data from the popular Roofnet study. We find strong evidence to support our conclusion that external interference is the main cause of unpredictable link behaviour. Fortunately, such interference is not an issue in rural settings where network connectivity is sparse or non-existent to begin with.

Once we have the link abstraction in place, much of the currently advocated approaches to routing metrics and routing protocols are likely to be inapplicable in their current form. On the other hand, other issues such as multi-hop TDMA scheduling are likely to gain more significance.

REFERENCES

[1] "IEEE P802.11, The Working Group for Wireless LANs," http://grouper.ieee.org/groups/802/11/.
[2] "Wireless Community Network List," http://www.toaster.net/wireless/community.html.
[3] B. Raman and K. Chebrolu, "Experiences in using WiFi for Rural Internet in India," *IEEE Communications Magazine, Special Issue on New Directions In Networking Technologies In Emerging Economies*, Jan 2007.
[4] "Wray Community Communications," http://www.wrayvillage.co.uk/wraycomcomhome.htm.
[5] "DjurslandS.net: The story of a project to support the weak IT infrastructure in an low populated area of Denmark," http://djurslands.net/biblioteket/international/djurslands_net_english_presentation.ppt.
[6] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level Measurements from an 802.11b Mesh Network," in *SIGCOMM*, Aug 2004.
[7] K. Chebrolu, B. Raman, and S. Sen, "Long-Distance 802.11b Links: Performance Measurements and Experience," in *MOBICOM*, 2006.
[8] A. Sheth, S. Nedevschi, R. Patra, S. Surana, E. Brewer, and L. Subramanian, "Packet Loss Characterization in WiFi-based Long Distance Networks," in *INFOCOM*, May 2007.
[9] 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," in *USENIX NSDI*, Apr 2007.
[10] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," in *MOBICOM*, Aug/Sep 2005.
[11] D. S. J. De Couto, D. Aguayo, B. A. Chambers, and R. Morris, "Performance of Multihop Wireless Networks: Shortest Path is Not Enough," in *HotNets-II*, Oct 2002.
[12] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A High Throughput Path Metric for MultiHop Wireless Routing," in *MOBICOM*, Sep 2003.
[13] R. Draves, J. Padhye, and B. Zill, "Comparison of Routing Metrics for Static Multi-Hop Wireless Networks," in *SIGCOMM*, Sep 2004.
[14] "The experiment data used for the roofnet 2004 sigcomm paper," http://www.pdos.lcs.mit.edu/roofnet/roofnet-sigcomm04.tar.bz2.
[15] S. Biswas and R. Morris, "Opportunistic Routing in Multi-Hop Wireless Networks," in *SIGCOMM*, Aug 2005.
[16] K. Chebrolu and B. Raman, "FRACTEL: A Fresh Perspective on (Rural) Mesh Networks," in *NSDR*, Sep 2007, a Workshop in SIGCOMM 2007.
[17] R. Draves, J. Padhye, and B. Zill, "Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks," in *MOBICOM*, Sep 2004.
[18] P. Bhagwat, B. Raman, and D. Sanghi, "Turning 802.11 Inside-Out," in *HotNets-II*, Nov 2003.
[19] "MIT Roofnet," http://pdos.csail.mit.edu/roofnet/doku.php.
[20] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network," in *MOBISYS*, Jun 2006.
[21] D. C. Cox, "Delay Doppler characteristics of multipath propagation at 910 MHz in a suburban mobile radio environment," *IEEE Transactions on Antennas and Propagation*, vol. 20, no. 5, pp. 625–635, Sep 1972.
[22] E. S. Sousa, V. M. Jovanovic, and C. Daigneault, "Delay spread measurements for the digital cellular channel in Toronto," *IEEE Transactions on Vehicular Technology*, vol. 43, no. 4, pp. 1–11, Nov 1994.
[23] "ISL3873: Wireless LAN Integrated Medium Access Controller with Baseband Processor," Intersil Corporation. Application Note FN4868., 2000.
[24] A. P. Subramanian, M. M. Buddhikot, and S. Miller, "Interference Aware Routing in Multi-Radio Wireless Mesh Networks," in *WiMesh*, Sep 2006.
[25] T. ElBatt and T. Andersen, "Cross-layer Interference-aware Routing for Wireless Multi-hop Networks," in *IWCMC*, Jul 2006.
[26] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-Based Models of Delivery and Interference in Static Wireless Networks," in *SIGCOMM*, Sep 2006.
[27] M. Neufeld, J. Fifield, C. Doerr, A. Sheth, and D. Grunwald, "SoftMAC - Flexible Wireless Research Platform," in *HotNets-IV*, Nov 2005.
[28] P. Kumar, "Design, Implementation, and Evaluation of new MAC Protocols for Long Distance 802.11 Networks," Master's thesis, IIT-Kanpur, May 2006, *with Bhaskaran Raman*.
[29] B. Raman and K. Chebrolu, "Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks," in *11th Annual International Conference on Mobile Computing and Networking paper (MOBICOM)*, Aug/Sep 2005.
[30] "Cisco Aironet 802.11 a/b/g CardBus Wireless LAN Client Adaptor, Data Sheet."
[31] "IEEE 802.16 WirelessMAN," http://www.ieee802.org/16/.