On the Feasibility of the Link Abstraction in Wireless Mesh Networks

Bhaskaran Raman, Kameswari Chebrolu, Dattatraya Gokhale, and Sayandeep Sen

***Abstract*—Outdoor community mesh networks based on IEEE 802.11 have seen tremendous growth in the recent past. The current understanding is that wireless link performance in these settings is 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. While our initial measurements focus on 802.11b, we then use two different PHY technologies as well, operating in the 2.4-GHz ISM band: 802.11g and 802.15.4. These show similar results too. 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, including at the MAC and routing layers, and opens up a fresh perspective on a wide range of technical issues in this domain.**

***Index Terms*—IEEE 802.11, IEEE 802.15.4, interference, link abstraction, link-level measurements, rural networks, WiFi, wireless mesh networks.**

I. Introduction

COMMUNITY mesh networks based on IEEE 802.11 [1] (WiFi) technology are growing in popularity, both in terms of deployment [2]–[5] as well as among researchers [6]–[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, literature on multihop mesh networks suggests that their performance is unpredictable. 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 originates from the PHY layer itself [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]–[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]–[5], and [8]. We first motivate why the performance behavior of links in such networks needs a detailed study. We then present extensive 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 received signal strength indicator (RSSI) or the signal-to-noise ratio (SNR). While prior work [14] has shown correlation between RSSI and link performance (in terms of UDP throughput) in an urban setting, we have studied the link error rate at various transmit rates of operation. In our study, 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, not multipath-induced delay spread, is the primary cause of unpredictable link behavior and lack of dependence on the RSSI. In contrast, Roofnet concludes 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 [15].

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]–[9], [16].

Next, based on short as well as long-running experiments, we show that the RSSI is stable, with only a small band of variation of 3–4 dB, for most links.

Our initial measurements are focused on IEEE 802.11b. We then also undertake in-depth studies using two other technologies which operate in the 2.4-GHz ISM band: 802.11g and

Manuscript received June 01, 2008; revised November 09, 2008; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor Shivkumar Kalyanaraman. Current version published April 15, 2009. This work was supported in part by Airtight Networks and the Tata Telecom Centre at IIT–Bombay, India.

B. Raman and K. Chebrolu are with the Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai 400 076, India (e-mail: br@cse.iitb.ac.in).

D. Gokhale is with the Indian Navy, India.

S. Sen is with the Department of Computer Science, University of Wisconsin-Madison, Madison, WI 53706 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2009.2013706

802.15.4 [17]. These indicate similar results, too: that RSSI can be used as a predictive tool for link performance, in the absence of external interference.

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 nonexistent. This intuition is also corroborated by practical experience in several deployments [3], [7], [8]. Likewise, external interference is unlikely in remote deployments of 802.15.4-based mesh networks (e.g., BriMon [18], Volcano monitoring [19]). 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 [20].

Furthermore, in our setting, we argue that routing metric and routing protocol design as considered in existing mesh networking literature [11]–[13], [16], [21] are unlikely to be applicable in our setting since these address issues that arise in the *absence* of a link abstraction. That is, metrics such as ETX [12] and WCETT [21] try to distinguish between links that 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 behavior.

Apart from predictability at the PHY layer, this paper also discusses another element of multihop mesh networks that is important for predictable performance—the MAC protocol. It is well known that the CSMA/CA MAC is not really suited for such scenarios and that time-division multiple-access (TDMA)-based approaches give better predictability in performance. The challenge here is to implement such approaches on the already popularly accepted WiFi-based platforms. We have been successful in achieving fine-grained synchronization (few tens of microseconds (μs) error) across multiple hops, using off-the-shelf 802.11 hardware.

Paper Outline

Section II presents a brief background of the wiFi-based Regional/Rural data ACcess and TELephony (FRACTEL) project and motivation for our work. Section III presents our detailed measurement study and analysis. Also presented is our interpretation of the data collected in [6] as well as evidence to support this interpretation. Subsequently, in Section IV, we describe our results on the stability of link behavior over time. We then present our measurements on 802.11g and 802.15.4 in Section V. The results from our measurements have several implications, which we articulate in Section VI. Finally, we present a few points of discussion in Section VII and conclude in Section VIII.

II. FRACTEL: Background and Motivation

Community networks based on multihop 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 kilometers, and (2) shorter distance, with most links below 500 m. Examples of the first kind of network are the Digital Gangetic Plains (DGP) project [22], and [8]. Examples in the second category are the MIT Roofnet project [23], and the various mesh deployments in [2].

In the FRACTEL 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 in a village. 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 centers, schools, hospitals, or government offices. Examples of current deployments that fall under this category include [4] and [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 as 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 stories tall in rural environments. We hence expect the multipath behavior 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 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 reanalyzed 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

1) Environment: For our measurements, we chose two kinds of environments, a rural village and a residential university campus (Indian Institute of Technology—Kanpur). 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.

TABLE I
LONG-DISTANCE VERSUS ROOFTOP MESH NETWORKS, AND OPEN ISSUES IN FRACTEL

	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 kilometres	High gain dirnl./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*

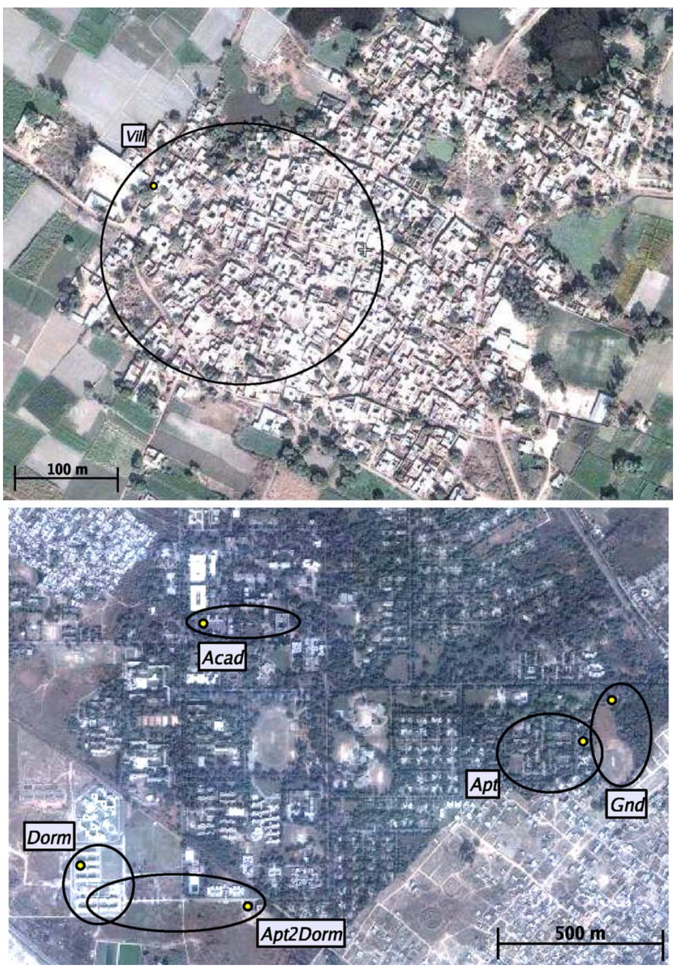

Fig. 1. Maps showing experiment locations (courtesy of Google Maps).

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 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 similar to the village location (densely packed houses, with most buildings at most two or three stories 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 location follows.

(1) Village (Vill): We had a 400 m $\times$ 400 m area with densely packed houses (1–2 stories tall) and a few scattered trees. *(2) Student Dormitory (Dorm):* This was a student residential dormitory, with four rows of three-story-tall dorm rooms and a few scattered trees in the vicinity. *(3) Housing Apartment Area (Apt):* This comprised several rows of two-story 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):* This was 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 playground and an adjoining small forest like area with dense foliage.

2) Hardware: We used the Senao 2511CD plus ext2 PCMCIA 802.11b cards, inserted into laptops (1.7 GHz Intel Pentium, 512 MB RAM) for our experiments. It is worth noting that both the Roofnet study [6] and the DGP study [7] used these same types 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 also used similar 8-dBi omni-antennas). The antennas were mounted on a small tripod stand about 1.5 m tall.

3) 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 *ad hoc*, 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 enabled us to observe the effect of external interfering packets more directly.[1]

[1]The per-packet transmit duration $\simeq 12$ ms at 1 Mbps, and $\simeq 1.5$ ms at 11 Mbps; hence, $\mathrm{inter\text{-}packet\ interval} = 20$ ms implies that we will most likely capture some foreign packets (if any) in the $\mathrm{gap} \geq 20 - 12 = 8$ ms between two successive packets.

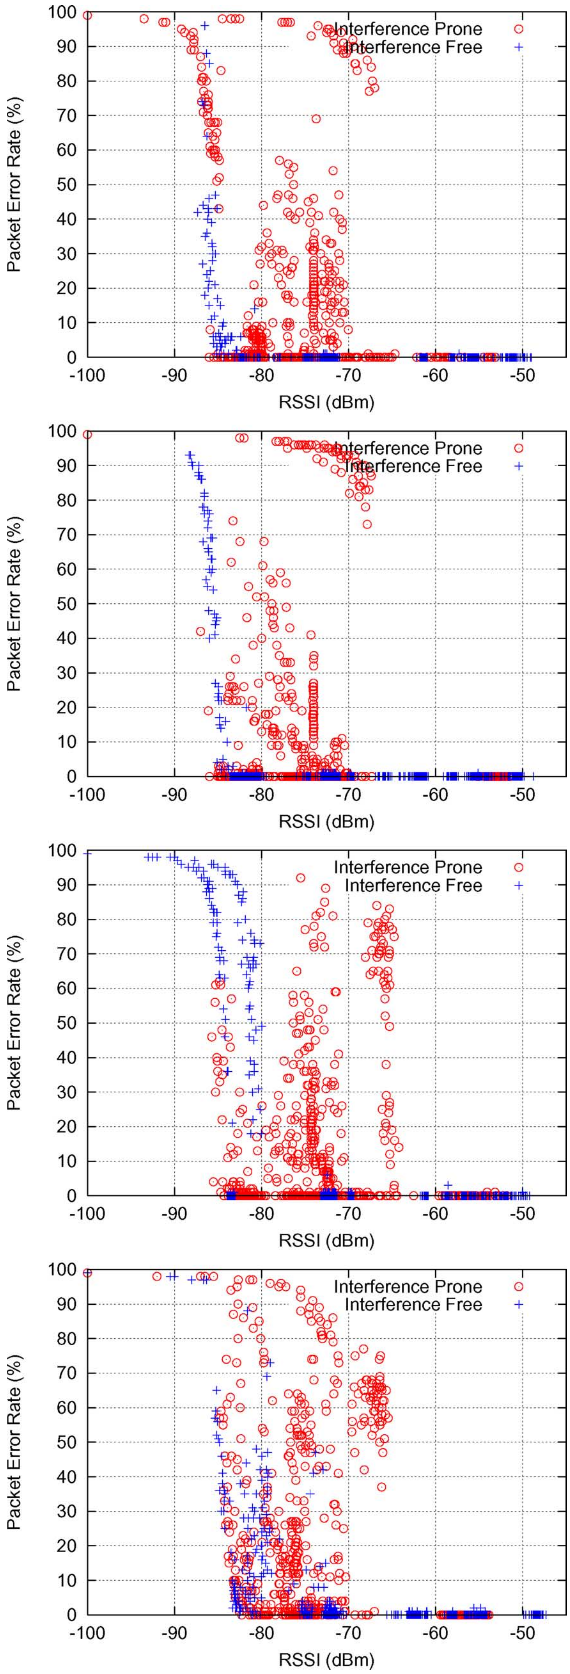

Fig. 2. Error rate versus RSSI (1, 2, 5.5, and 11 Mbps, respectively).

Apart from the 2-min experiments, we had long-running experiments (24–48 hrs) at the *Apt* and *Apt2Dorm* locations.

4) Choice of Transmitter/Receiver Positions: At each of the 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 (LOS) and the mean RSSI was about −70 dBm. The RSSI was calculated over an initial set of1000 packets sent from the transmitter. We define a *medium* receiver position to be one that had some foliage or building in between and the RSSI thus calculated was about −75 dBm. A *bad* receiver position is one where the RSSI was about −80 dBm 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 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. Therefore, we decided to separate out the 6000 packets into 60 bins of 100 packets each. Note that these are bins of 100 consecutively *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 −100 dBm. We note that the noise or silence level in most of our experiments was −94 to −95 dBm. Therefore, 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 that showed a high error rate despite a high signal strength (−75 dBm or above). When we looked into our receiver side logs for these cases, we observed that there were several 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. At the *Gnd* location, the signals were likely coming from a long-range WiFi link setup by someone in the vicinity.

We then separated out cases where our receiver logs showed foreign packets from those that 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, 2, 5.5, and 11 Mbps, respectively.[2] Each

[2]As in [6] and [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.

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-, 2-, 5.5-, and 11-Mbps rates, this threshold is about -86, -85, -80, and -79 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 -95 and -94 dBm, much like in our outdoor experiments. The threshold RSSI values above which the error rate fell below 1% was -88, -86, -82, and -81 dBm, 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 1–2 dB of the threshold values observed in our RF cable experiment.

There are a few isolated points in Fig. 2 (around $\mathrm{RSSI} \simeq -75$ dBm, 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 s 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.

A feature we can observe in the 5.5-Mbps plot in Fig. 2 is that there are two steep lines among the interference-free points. These two lines are separated by about 3–4 dB, especially at high error rates. There are two possible reasons for this. First, there is short-term variation in the RSSI of about 3–4 dB (characterized further in Section IV). Furthermore, at high error rates, the computation of the average RSSI (among 100 packets) has an inherent error: We have RSSI samples only for packets that were received. It is likely that the RSSI for the packets that were *not* received was lower due to RSSI variation.

In all 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 identified the interference to be from a known WiFi source. We then had the interference source temporarily shut down. When we reran the experiment, we observed negligible error rates ($< 1\%$). Readings from both these experiments are included in Fig. 2.

For each of the other positions, prior to experiment setup, we sought to find a channel of operation free from external interference. However, for the interference-prone data points marked in Fig. 2, we were unable to find a free channel or identify the WiFi source(s) causing external interference (with the exception of the *Apt* location mentioned above).

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 [6, Figs. 14 and 15]. 2) There is no correlation between the number of lost packets and the number of foreign (interference) packets [6, Fig. 18]. 3) Introducing delay spread causes high loss rates [6, Fig. 19]. 4) Prior measurements in urban environments have reported delay spreads of over 1 μs [24], [25], which is well beyond the tolerance limits of the 802.11b receiver [26].

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 μs 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 [15]. 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 -94 or -95 dBm.

To show how the Roofnet data shows very different behavior, we chose the subset of the transmitter–receiver pairs where the average RSSI was over -77 dBm at the 11-Mbps data rate (well above the measured card sensitivity of -81 dBm) and those that 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 5th percentile, 50th percentile (median), and 95th percentile noise levels for these points, in increasing order of the median noise level.

We wish to note that for these points, the average SNR also was very high (11–39 dB). We observe that the median noise levels are as high as -86 dBm, with most median values at -90 dBm or above. We also observe that the 95th percentile

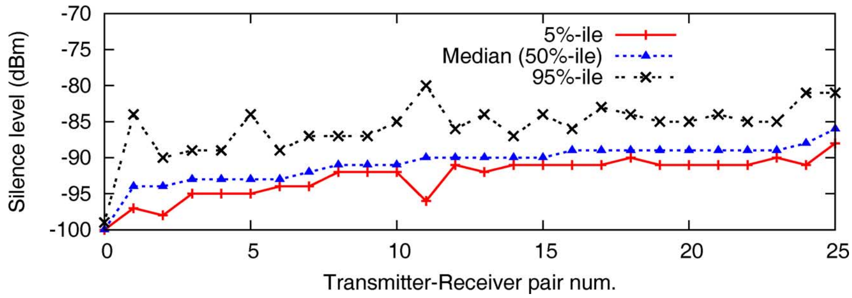

Fig. 3. Noise levels in Roofnet data (11 Mbps).

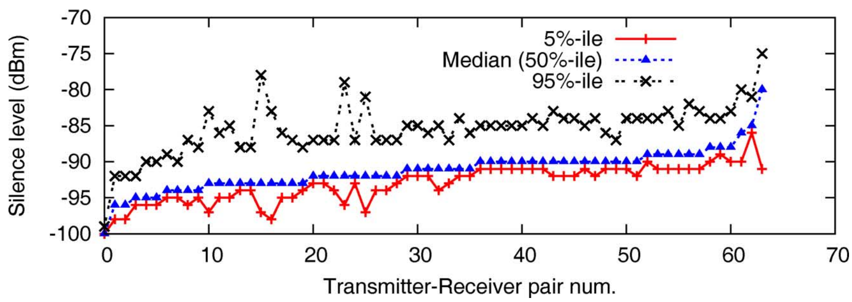

Fig. 4. Noise levels in Roofnet data (1 Mbps).

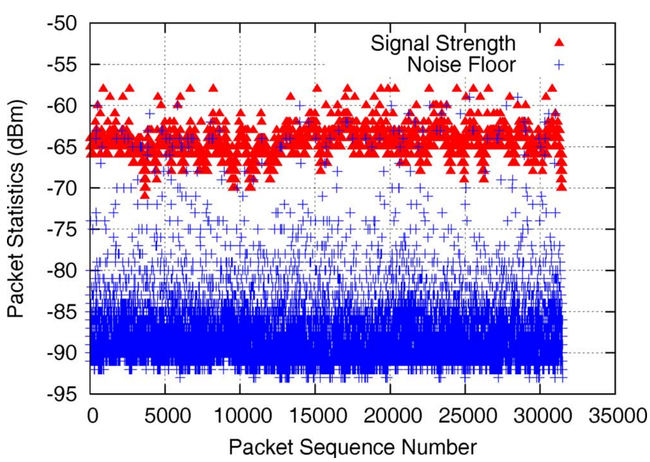

Fig. 5. Per-packet variation of RSSI and noise levels.

minus 5th percentile (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 −94 dBm, 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 −80 dBm (again, well above the card sensitivity of −88 dBm) and the error rate was between 20% and 80%. Once again, we see behavior 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 number 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 −90 dBm or more!

The only explanation for the high and variable noise levels is the presence of several 2.4-GHz sources. We believe it is likely

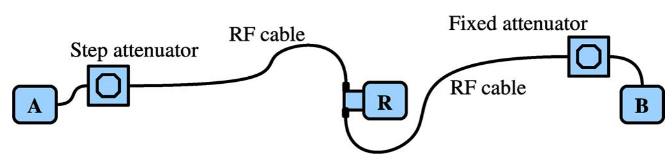

Fig. 6. Setup for interference experiment.

that 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 interpacket interval of 2 ms. Each experiment ran for about 2 min.

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 −75 dBm. We varied the attenuator at A across

TABLE II
CONTROLLED INTERFERENCE EXPERIMENT RESULTS

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

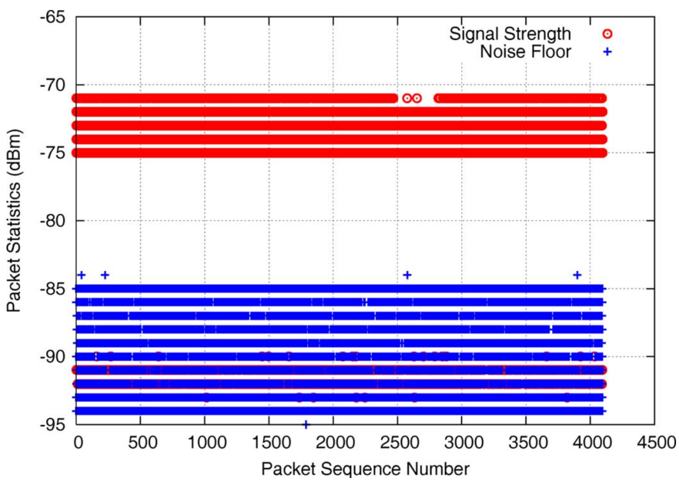

Fig. 7. Per-packet variation of RSSI and noise in the controlled interference experiment.

four experiments such that we had mean RSSI values from A to be about −90, −85, −80, and −75 dBm.

From R's log, we calculated 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.

1) 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 Section III-C. For comparison with Fig. 5, we plot a similar graph. Fig. 7 shows the per-packet signal strength and noise level variations for packets from B. This is shown for experiment number 2 in Table II, where B's average RSSI is about −75 dBm and A's average RSSI is about −85 dBm.

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

2) P2: Again focusing on experiment number 2 in Table II, we see that so far as B's packets are concerned, there is a loss of 18.3% (Col-4). However, the number of A's packets received are very few. With a loss rate of 99.2% and a sending rate of 500 packets/s, A's packets are received at an average rate of only 4 packets/s. Yet, this is sufficient to cause an 18.3% loss rate, which amounts to 91.5 lost packets/s! 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 −85 dBm, much below the sensitivity as measured earlier (−81 dBm, see Section III-B).

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

3) 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. 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 [6, Sec. 8]. On each link, they first measure the average rate at which foreign packets are seen in a 90-s duration. Then, they measure the packet error rate seen on that link in the immediate next 90-s period. A scatter plot of these two shows no correlation [6, Fig. 18], 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

1) P4: We now look at the question of whether the card reported noise level should 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. Furthermore, in Table II, we compared 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 -85 dBm, whereas A's interference B we know was at -75 dBm.

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 [26]. 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.

2) 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.

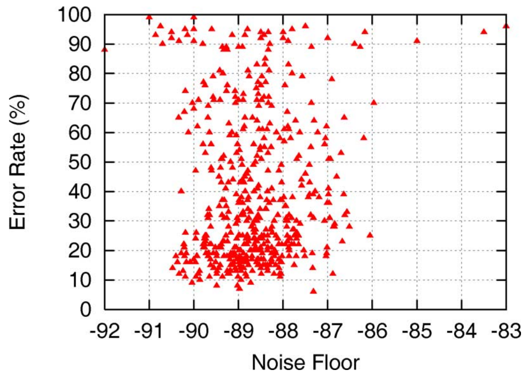

Fig. 8. Average noise level versus error rate, 100-packet bins.

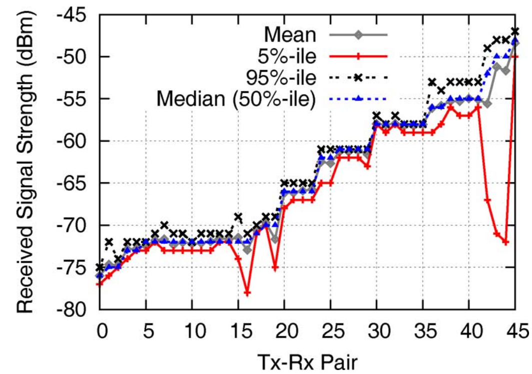

Fig. 9. Variation in the RSSI.

We compute this within 100-packet bins for the same transmitter–receiver pair in Roofnet as for Fig. 3. Fig. 8 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 behavior based on the average noise floor either.

In our controlled experiments, too, we can observe this lack of correlation between the average noise level and the error rate. To see this, note that from experiment number 1 to 2 and then to 3, B's average noise level increases, and so does the error rate. But in experiment 4, B's average noise level actually drops.[3]

IV. Stability Over Time

We saw above that link behavior 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 5th percentile, the 50th percentile (median), and the 95th percentile. Fig. 9 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.

The band between the 95th percentile and 5th percentile 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 numbers 16, 19, and 36 in Fig. 9, 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. 9 has three cases (pair numbers 42, 43, and 44) where the band was 18, 23, and 24 dB, respectively. A look at the RSSI variation with time revealed that there were several data points that 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 and receiving cards are connected via an RF cable. This hardware quirk has also been reported in other studies [7]. Except for these quirks, these data points also 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 in 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 95th percentile to 5th percentile band was within 5 dB for most links. There were a few links in the Roofnet data that showed larger bands (6–11 dB).

To understand RSSI variation over longer durations of time, we ran 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 h, while at *Apt2Dorm*, the experiment ran for 24 h.

Table III shows a summary of the results from the six transmitter–receiver pairs. We see again that in the LOS cases, the 95th percentile minus 5th percentile 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 exceeded 4 dB. In a few non-LOS cases, a few positions showed bands larger than 4 dB.

What does such variation mean? The short-term variation tells us when we *cannot* expect predictable behavior. Note that the steep portion of the error rate versus RSSI plot (Fig. 2) is

[3]The authors are grateful to the reviewer who pointed this out.

TABLE III
LONG-TERM RSSI VARIATION.

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

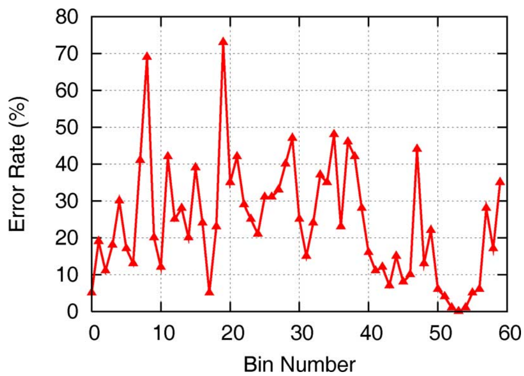

Fig. 10. Variation in the error rate.

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.

To illustrate the above aspect, we consider a receiver position in our data that 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 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. 10. 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. MEASUREMENTS IN OTHER WIRELESS TECHNOLOGIES

We have thus far presented measurements using 802.11b and compared our results with those from prior results on an 802.11b outdoor mesh network. We now look at the performance of two different PHY layers: 802.11g and 802.15.4. 802.11g uses an OFDM-based PHY layer, while 802.15.4 uses QPSK in the 2.4-GHz band. We consider 802.11g since it is the currently popular version of WiFi (successor to 802.11b). Prior measurement studies have either exclusively focused on 802.11b [7], [8], [23] or have considered higher layer measurements on 802.11g [27]; we undertake a link-level study of 802.11g.

802.15.4 is another wireless standard that is intended to be used in outdoor mesh network settings. Apart from traditional sensor network applications such as BriMon [18] or Volcano monitoring [19], we also envision the use of 802.15.4 mesh networks in rural settings for carrying low-volume voice traffic [28]. While there have been prior 802.15.4 studies in indoor conditions [29], ours explores the link-level behavior under a variety of outdoor settings. In comparison with prior work [29], we also consider much longer 802.15.4 links of a few tens of meters to a few hundred meters. Such link lengths are achieved through the use of various kinds of external antennas.

The next section (Section V-A) presents our 802.11g measurements, while the subsequent section describes the 802.15.4 results (Section V-B).

A. 802.11g Measurements

For logistical reasons, our 802.11g measurements were carried out at locations different from our earlier 802.11b measurements in Section III. The environment, however, was similar. We used another university campus, the Indian Institute of Technology Bombay. We conducted measurements on six different links, whose lengths were approximately 75, 150, 175, 450, 450, and 800 m, respectively. All six links had LOS. The various nodes were set up atop different buildings.

The software setup was similar to the 802.11b measurements, except that we used the open-source Madwifi driver v0.9.3 along with the Ubiquiti SR2 cards. Fig. 11 shows the error rate versus average RSSI for 100-packet bins, much like Fig. 2. We show different plots in the same graph for the six different links. The four graphs represent the data for four of the eight possible PHY rates in 802.11g. The plots for the remaining rates are not shown due to lack of space, but the results were similar for those, too.

We can see from Fig. 11 that there is a steep drop in the error rate with increasing RSSI, so there is a definite correlation of the error rate with RSSI. There is also a threshold RSSI beyond which the error rate is close to 0%. Therefore, as with 802.11b, we can conclude from this that multipath-induced packet errors have little role to play in both these settings as wells as for the 802.11g PHY.

Although we can see a steep drop, the plots in Fig. 11 show a band, similar to that seen for the 5.5-Mbps graph in Fig. 2. The same reasons mentioned in Section III-B apply here, too: There is inherent variation in RSSI, and at high error rates, the computed average RSSI could be an overestimate since we do not know the RSSI values for the lost packets. Additionally, for the 802.11g hardware (Ubiquiti SR2 cards), we observed higher RSSI variation even in controlled conditions, as compared to the 802.11b hardware (Prism2 cards). When having a transmitter connected via an RF cable to a receiver, we observed as much as 4–5-dB variation with the Ubiquiti SR2 cards, whereas this was only about 1–2 dB for the Prism2 cards. This explains the bands seen in Fig. 11.

In some cases in Fig. 11, especially for link-6, we have a higher error rate than for the other links, at the same RSSI. A deeper look into our logs explains this. Although the interference at each of the links is small, it is noticeable in our logs. We confirmed that the higher error rate was indeed due to external interference. (But the interference was not prominent enough for us to classify points in the plot as interference-free versus interference-prone. Anyway, we see a steep shape in the graphs despite the small amount of interference.)

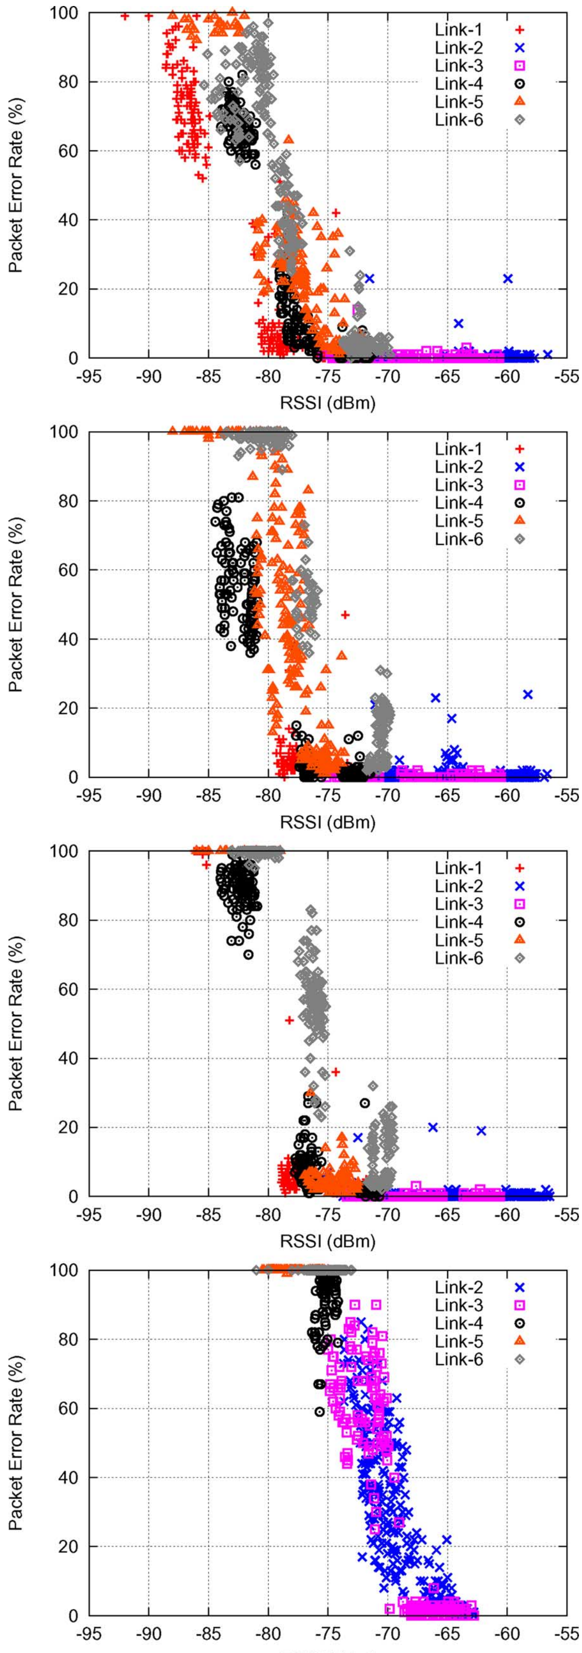

Fig. 11. Error rate versus RSSI (6, 18, 24, and 54 Mbps, respectively).

B. 802.15.4 Measurements

For our 802.15.4 measurements, we used the Tmote hardware (www.moteiv.com) which uses the CC2420 radio chip. We

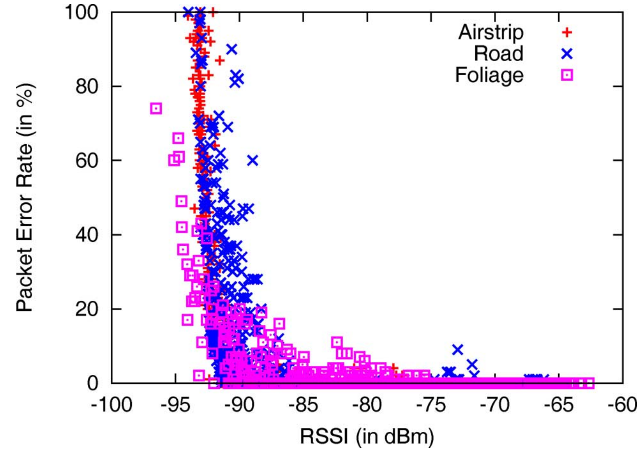

Fig. 12. Error rate vs. RSSI for 802.15.4 in 2.4 GHz.

used the TinyOS software (www.tinyos.net) for sending and receiving packets. The packet-error rate measurements were done in three different locations: 1) an "airstrip," where there is clear LOS and no trees or buildings in the vicinity; 2) a narrow "road" that has LOS and trees and some buildings on the side; and 3) a "foliage" environment, which has dense foliage between the sender and the receiver without any clear LOS. In all three environments, we used various kinds of external antennas: 3.1-dBi built-in, 8-dBi omni, 17-dBi sector, and a 24-dBi parabolic-grid. For the various measurements in each environment, we fixed the transmitter and moved the receiver to various positions until the maximum possible link length. The maximum link range in the three environments was about 800, 500, and 90 m, respectively. Our measurements included a total of 40 transmitter–receiver location pairs (further details are in [30]).

Fig. 12 shows the error-rate versus the average RSSI for 100-packet bins, much like our earlier plots. Here, too, we see that there is a strong correlation between the RSSI and the error rate. For the airstrip and road environments, there is a clear threshold RSSI of about -90 dBm, above which the error rate is close to 0%. The foliage environment's scatter-plot, however, shows a "spread" of the points. That is, even above -90 dBm, we have a few points with 10%–20% error rate. There was, however, no significant 2.4-GHz interference in the vicinity.

To examine the cause for this spread in the scatter-plot, we plot the per-packet RSSI as a function of the packet number. This is shown in Fig. 13. We find that there can be huge variability of about 7–10 dB in the per-packet RSSI even across time-scales as small as 1 s or less. In comparison, such variability for the other two environments (which had clear LOS) was only about 4–5 dB. This is quite like the higher RSSI variability we observed in non-LOS conditions for 802.11b in Section IV.

Now, a higher RSSI variability has the following effect: Even when a 100-packet bin has a higher average RSSI, there would be several packets whose RSSI would fall well below the threshold of -90 dBm and hence not be recorded in the log (we can calculate the average RSSI only based on what packets we receive, and it is therefore likely an overestimate). So, a higher RSSI variability explains the higher "spread" seen for the foliage environment in Fig. 13.

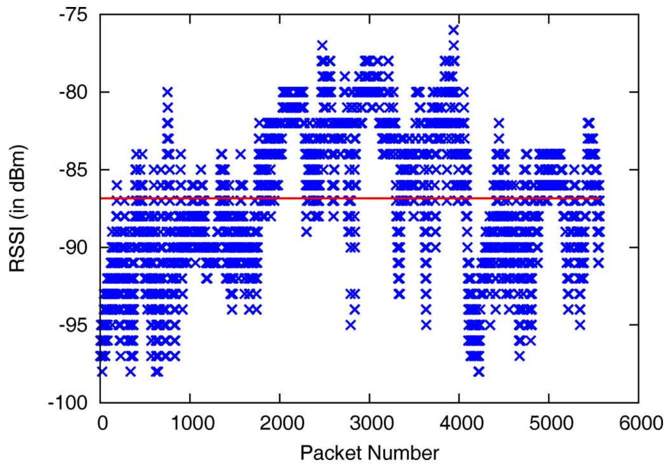

Fig. 13. RSSI variation in the foliage environment, 802.15.4.

As with 802.11b, we also examined the RSSI variability over long durations of time (12–24 h). Our results were similar for 802.15.4, too: The long-term RSSI variability was also within a few dB at most for LOS outdoor settings.

VI. 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.

Apart from rural locations, our results are likely applicable for many mesh networks in semi-urban residential communities as well since these too do not have tall buildings, similar to our measurement environment (e.g., see [14]). While our initial measurements focused on 802.11b, we have presented measurements on 802.11g and on 802.15.4 as well to show that the above conclusions apply broadly for these PHY layer technologies. 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 and 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 -79 dBm 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. Note that after this adjustment, the threshold value roughly matches the value of -75 dBm mentioned in the measurement study in [14].

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 behavior, 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 that expect predictable performance.

B. Implications on Routing

1) Routing Metrics: In the absence of the link abstraction, much work has focused on the design of appropriate routing metrics [11]–[13], [21]. These essentially seek to distinguish between links that 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. 10, such operation can lead to high variations in the error rate, which is unpredictable. This in turn would mean erratic behavior at the routing layer if we use metrics such as ETX [12] or WCETT [21].

2) Opportunistic Routing: In this technique, the approach to handle intermediate loss rates is to opportunistically use packet reception whenever it succeeds [16]. 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.

3) Interference Aware Routing: There are several prior efforts that have focused on interference-aware routing (e.g., [31], [32]). Most of these seek to mitigate *internal* interference, i.e., interference among the links of the wireless mesh itself. The work in [33] 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.

Furthermore, our measurements in Section III-E (**P4** and **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 [31]) 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 multihop mesh networks. It causes self-interference; i.e., 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 explained in Section 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, in interference-free environments, opens the door for TDMA-based MAC approaches. A TDMA-based MAC essentially controls *internal* interference; i.e., interference between the various links of a mesh network. Prior work [9], [34], [35] has already shown prototypes of TDMA-based MAC implementations on WiFi hardware. However, multihop 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 [20]. We now discuss one of the important aspects of a TDMA-based MAC.

D. Practicality of a Multihop TDMA MAC

One of the main challenges in the implementation of a TDMA MAC with scheduled slots is the time synchronization required. Whether fine-grained synchronization can be achieved using off-the-shelf 802.11 hardware over multiple hops is a significant question. In ongoing work, we have implemented a prototype multihop synchronization mechanism using the underlying 802.11 hardware time-stamping mechanism. We have been able to achieve synchronization errors of about a few tens of μs or less over multiple hops using off-the-shelf 802.11 hardware. It is worth noting that prior work on similar 802.11 hardware [34] reports about 25 μs error, but this is achieved over a single hop only. In 802.15.4 hardware also, in recent work, we have been able to achieve a few tens of μs error per hop [18].

We note that we do not foresee more than 5–6 hops in most settings in rural deployments of FRACTEL, since villages are anyway only about a few km^2 in area and link lengths of a few hundreds of meters are easy to achieve (see [20] for more details). Therefore, the increase of synchronization error with the number of hops is unlikely to be a major concern in practice.

VII. Discussion

A. Multipath and Delay Spread

In the environments in which we have tested, multipath-induced delay spread is clearly not a significant factor. This is consistent with the fact that suburban settings exhibit lower delay spread as compared to dense urban settings [36]. We have shown that in the Roofnet data, too, from a dense urban setting, external interference played a major role in causing error rates. However, 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 [24] and [25], which have measured delay spreads in such urban environments to be 1 μs. However, 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 behavior 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.

B. 802.11a

We have used 802.11b, 802.11g, and 802.15.4 in our measurements, all operating in the 2.4-GHz band. For 802.11a, like 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 [37]). However, two aspects will likely come into play for 802.11a. First, it is likely that multipath will have even less of an effect on 802.11a since the attenuation levels at 5 GHz are higher than at 2.4 GHz, but only actual measurements can confirm this. Second, and more importantly, 802.11a has more frequencies of operation—12 total, of which 8 are nonoverlapping. 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.

1) 802.16: The upcoming IEEE WiMAX standard [38] 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 [22]. 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.

VIII. 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 behavior 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 behavior. Fortunately, such interference is not an issue in rural settings where network connectivity is sparse or nonexistent 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 multihop TDMA scheduling are likely to gain more significance.

Acknowledgment

This work would not have been possible without the help of various student colleagues and project staff in conducting the experiments at both IIT–Kanpur, India and at IIT–Bombay, India. The 802.11g experiments were conducted by N. K. Dara at IIT–Bombay. The authors would like to thank all the reviewers for their in-depth and insightful comments.

References

[1] *The Working Group for Wireless LANs*, IEEE P802.11 [Online]. Available: http://grouper.ieee.org/groups/802/11/

[2] "Wireless community network list," 2009 [Online]. Available: http://www.toaster.net/wireless/community.html

[3] B. Raman and K. Chebrolu, "Experiences in using WiFi for rural Internet in India," *IEEE Commun. Mag.*, vol. 45, no. 1, pp. 104–110, Jan. 2007.

[4] "Wray community communications," 2009 [Online]. Available: 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," 2009 [Online]. Available: 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 *Proc. SIGCOMM*, Portland, OR, Aug. 30–Sep. 3,, 2004, pp. 121–132.

[7] K. Chebrolu, B. Raman, and S. Sen, "Long-distance 802.11b links: Performance measurements and experience," in *Proc. MOBICOM*, Los Angeles, CA, Sep. 23–29, 2006, pp. 74–85.

[8] A. Sheth, S. Nedevschi, R. Patra, S. Surana, E. Brewer, and L. Subramanian, "Packet loss characterization in WiFi-based long distance networks," in *Proc. INFOCOM*, Anchorage, AK, May 6–12 2007, pp. 312–320.

[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 *Proc. USENIX NSDI*, Apr. 11–13, 2007.

[10] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, "Architecture and evaluation of an unplanned 802.11b mesh network," in *Proc. MOBICOM*, Cologne, Germany, Aug. 28–Sep. 2, 2005, pp. 31–42.

[11] D. S. J. De Couto, D. Aguayo, B. A. Chambers, and R. Morris, "Performance of multihop wireless networks: Shortest path is not enough," presented at the HotNets-I, Princeton, NJ, Oct. 28–29, 2002.

[12] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A high throughput path metric for multihop wireless routing," in *Proc. MOBICOM*, San Diego, CA, Sep. 14–19, 2003, pp. 134–146.

[13] R. Draves, J. Padhye, and B. Zill, "Comparison of routing metrics for static multi-hop wireless networks," in *Proc. SIGCOMM*, Sep. 2004, pp. 133–144.

[14] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proc. MOBISYS*, Uppsala, Sweden, Jun. 19–22, 2006, pp. 96–109.

[15] "The experiment data used for the Roofnet 2004 Sigcomm paper," 2009 [Online]. Available: http://www.pdos.lcs.mit.edu/roofnet/roofnet-sigcomm04.tar.bz2

[16] S. Biswas and R. Morris, "Opportunistic routing in multi-hop wireless networks," in *Proc. SIGCOMM*, Philadelphia, PA, Aug. 22–26, 2005, pp. 133–144.

[17] "IEEE 802.15 WPAN Task Group 4 (TG4)," 2009 [Online]. Available: http://www.ieee802.org/15/pub/TG4.html

[18] K. Chebrolu, B. Raman, N. Mishra, P. K. Valiveti, and R. Kumar, "BriMon: A sensor network system for railway bridge monitoring," in *Proc. MobiSys*, Breckenridge, CO, Jun. 17–20, 2008, pp. 2–14.

[19] G. Werner-Allen, K. Lorincz, M. Welsh, O. Marcillo, J. Johnson, M. Ruiz, and J. Lees, "Deploying a wireless sensor network on an active volcano," *IEEE Internet Computing*, vol. 10, no. 2, pp. 18–25, Mar./Apr. 2006.

[20] K. Chebrolu and B. Raman, "FRACTEL: A fresh perspective on (rural) mesh networks," presented at the NSDR Workshop, SIGCOMM 2007, Sep. 2007.

[21] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proc. MOBICOM*, Philadelphia, PA, Sep. 26–Oct. 1, 2004, pp. 114–128.

[22] P. Bhagwat, B. Raman, and D. Sanghi, "Turning 802.11 inside-out," presented at the HotNets-II, Cambridge, MA, Nov. 20–21, 2003.

[23] "MIT Roofnet," 2009 [Online]. Available: http://pdos.csail.mit.edu/roofnet/doku.php

[24] D. C. Cox, "Delay Doppler characteristics of multipath propagation at 910 MHz in a suburban mobile radio environment," *IEEE Trans. Antennas Propag.*, vol. AP-20, no. 5, pp. 625–635, Sep. 1972.

[25] E. S. Sousa, V. M. Jovanovic, and C. Daigneault, "Delay spread measurements for the digital cellular channel in Toronto," *IEEE Trans. Veh. Technol.*, vol. 43, no. 4, pp. 1–11, Nov. 1994.

[26] Intersil Corp., "ISL3873: Wireless LAN integrated medium access controller with baseband processor," 2000, Appl. Note FN4868..

[27] T. Ireland, A. Nyzio, M. Zink, and J. Kurose, "802.11g long-distance measurements: Antenna placement and orientation," in *Proc. WiNMee*, Limassol, Cyprus, Apr. 20, 2007.

[28] B. Raman, K. Chebrolu, and P. Kulkarni, "Lo^3: **Lo**w-power, **Lo**w-Cost, **Lo**cal voice and messaging for rural regions," 2009 [Online]. Available: http://www.cse.iitb.ac.in/silmaril/br/doku.php?id=proj:lo3, to be published

[29] K. Srinivasan and P. Levis, "RSSI is under appreciated," in *Proc. EmNets*, Cambridge, MA, May 30–31, 2006.

[30] B. Raman, K. Chebrolu, N. Madabhushi, D. Y. Gokhale, P. K. Valiveti, and D. Jain, "Implications of link range and (in)stability on sensor network architecture," in *Proc. WiNTECH*, Los Angeles, CA, Sep. 29, 2006.

[31] A. P. Subramanian, M. M. Buddhikot, and S. Miller, "Interference aware routing in multi-radio wireless mesh networks," in *Proc. WiMesh*, Sep. 25–28, 2006, pp. 55–63.

[32] T. ElBatt and T. Andersen, "Cross-layer interference-aware routing for wireless multi-hop networks," in *Proc. IWCMC*, Vancouver, BC, Canada, Jul. 3–6, 2006, pp. 153–158.

[33] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *Proc. SIGCOMM*, Pisa, Italy, Sep. 11–15, 2006, p. 51–62.

[34] A. Sharma, M. Tiwari, and H. Zheng, "MadMAC: Building a reconfigurable radio testbed using commodity 802.11 hardware," in *Proc. WSDR'06*, Sep. 2006, pp. 78–83.

[35] B. Raman and K. Chebrolu, "Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks," in *Proc. MOBICOM*, Cologne, Germany, Aug. 28–Sep. 2, 2005, pp. 156–169.

[36] T. S. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 2003, ch. 5.

[37] Cisco, "Cisco Aironet 802.11 a/b/g cardBus wireless LAN client adaptor, data sheet," 2009 [Online]. Available: http://www.cisco.com/en/US/products/hw/wireless/ps4555/ps5818/

[38] R. B. Marks, "IEEE 802.16 WirelessMAN," 2009 [Online]. Available: http://www.ieee802.org/16/

Bhaskaran Raman received the B.Tech. in computer science and engineering from the Indian Institute of Technology, Madras, India, in May 1997, and the M.S. and Ph.D. degrees in computer science from the University of California, Berkeley, in 1999 and 2002, respectively.

He was a Faculty Member in the Department of Computer Science and Engineering at Indian Institute of Technology, Kanpur, India, from June 2003. Since July 2007, he has been a Faculty Member in the Department of Computer Science and Engineering at the Indian Institute of Technology, Bombay, India. His research interests and expertise are in communication networks, wireless/mobile networks, large-scale Internet-based systems, and Internet middleware services. His current focus and specific interest is in appropriate technology in the domain of communication and computing for rural use.

Kameswari Chebrolu received the B.E. degree in electronics and communication engineering from Andhra University in 1998 and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of California at San Diego in 2001 and 2004, respectively.

She was an Assistant Professor at the Department of Electrical Engineering at the Indian Institute of Technology, Kanpur, India, from February 2004 to July 2007. She is currently an Assistant Professor with the Department of Computer Science and Engineering at the Indian Institute of Technology, Bombay, India. Her research interests are in the areas of wireless network architecture, protocol design, and analysis. She is specifically interested in technology for developing regions.

Dattatraya Gokhale received the M.Tech. degree in computer science and engineering from the Indian Institute of Technology, Kanpur, India, in 2007.

He is currently serving in the Indian Navy.

Sayandeep Sen received the B.Tech. degree in computer science and engineering from the University of Kalyani, India, in 2004, and the M.Tech. degree in computer science and engineering from the Indian Institute of Technology, Kanpur, India, in 2006.

He is currently a second-year graduate student at the University of Wisconsin-Madison. His research interests are in the areas of wireless networks.