

# Implications of Link Range and (In)Stability on Sensor Network Architecture

Bhaskaran Raman

Kameswari Chebrolu

Naveen Madabhushi

Dattatraya Y. Gokhale

Phani K. Valiveti

Dheeraj Jain

IIT Kanpur, INDIA 208016  
{braman,chebrolu,mnaveen,goks,phani,dgolchha} AT iitk.ac.in

## ABSTRACT

In this work, we consider sensor networks used for scientific instrumentation, where we have a set of nodes collecting data and relaying the same to a central base node. From the point of view of deploying such networks, we re-look at the design choices available for the network architecture and protocol design. A design choice which has not received in-depth attention thus far is the use of external antennas for improving the communication range. We present extensive measurements to quantify the use of external antennas. We show that this is a simple yet effective mechanism, in many cases allowing the use of just a single-hop network architecture. Such an approach of course, also greatly simplifies protocol design. Related to the range studies, we also look at the time variability of RSSI and the packet error rate. We find variability at time scales as small as a single packet and also at several hours. Given this, we argue that dynamic metric based routing becomes highly questionable, and a centralized protocol design becomes a serious alternative.

**Categories and Subject Descriptors:** C.2.1 Network Architecture and Design: Wireless communication

**General Terms:** Measurement, Performance, Experimentation.

**Keywords:** 802.15.4 link characteristics, 802.15.4 link range, Wireless Sensor Networks, Network architecture

## 1. INTRODUCTION

Wireless Sensor Networks (WSNs) have been thus far used in several scientific instrumentation efforts [7, 11, 12, 6]. A few others are in the process of being deployed [8, 5]. The typical use of these networks is in gathering data from each *mote* and collecting them at a central *base node*. There are several important considerations while deploying such networks. What should be the network architecture? What is the radio communication range? What is the expected num-

ber of hops from/to the base node? Does dynamic metric-based routing make sense? At what time-scale?

In this paper, we take a fresh look at the design choices available for answering these questions. First, we study in-depth the possibility of using external antennas to improve communication range. Using the CC2420 radio-based *tmotes* [4], we have conducted extensive measurements with external antennae in different environments: an air-strip, a narrow straight road, and dense foliage. In the first two environments, we have line-of-sight, and are able to achieve communication ranges as long as 500m-800m. Even in dense foliage, we are comfortably able to achieve ranges of 70m-90m with the use of a high gain sector/parabolic-grid antenna at only one end of the link (the base node).

Our range measurements suggest that in many application scenarios, we can potentially eliminate multi-hop scenarios and have just a single-hop network from the base node. This would greatly simplify the network architecture and protocol design/implementation. In other applications, we can greatly increase the number of nodes which are one-hop away from the base. Since the one-hop nodes from the base are burdened with relaying data from nodes further away, they are typically the bottleneck in terms of communication energy consumption, and thus also in terms of the network lifetime. In intuitive terms, it is thus beneficial to increase the number of one-hop nodes, i.e. the coverage of the base node. This would provide the opportunity for distributing the data relaying functionality over a larger set of nodes and thus improve network lifetime.

The study of link range leads us next to look at the temporal stability of links' error rate and RSSI characteristics. We find, like in [9] (and with a wider range of data) that average RSSI is correlated strongly with the packet error rate. However, much unlike in [9], we observe high variability even at per-packet granularity. This means that a packet's RSSI sample cannot be taken as an indication of the average RSSI or the average error rate.

Further, in different environments, we also observe RSSI variability at various time scales: 2 s, 20 s, as well as several hours to a day. We quantify such variability through extensive studies. We also observe similar variability in the LQI metric.

Our data on RSSI and error rate variability suggests that trying to dynamically predict a link's quality for routing (or other) purposes may not be meaningful. However, based on our data, we do suggest a way of distinguishing between

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WiNTECH'06, September 29, 2006, Los Angeles, California, USA.  
Copyright 2006 ACM 1-59593-538-X/06/0009 ...\$5.00.

stable links and potentially unstable links. This can be done in the early stages of deployment, and/or with a very low periodicity.

Suppose that we use only the links with stable (and low) error rate, and given that sensor network applications anyway employ a centralized, relatively powerful *base node* to collect data, we argue that a centralized routing protocol becomes a serious design alternative. This would not only greatly simplify protocol design and implementation, but also provide much better opportunities for load-balancing than metric-based dynamic routing. Such load-balancing may eventually do much better to improve network lifetime than does dynamic routing.

In sum, based on link range and stability measurements, we propose a set of simple yet effective design choices which we believe have not received adequate attention thus far. The rest of the paper is organized as follows. Section 2 provides a brief description of the various applications under consideration. Section 3 details our studies on the communication range with the use of external antennae. Subsequently, we present the link-quality stability studies in Section 4. We present further discussion points in Section 5 and conclude in Section 6.

## 2. BACKGROUND

The network architecture is highly dependent on the nature of the application. Hence, before our discussion of any design choices, we provide a brief overview of the various applications which have been designed/prototyped thus far. While the above list of applications is clearly not exhaustive, it is a representative list. We shall refer back to these applications frequently as we proceed with our discussion.

(1) *Habitat monitoring* [7]: motes sense temperature & humidity in a particular habitat of sea birds. (2) *Redwood micro-climate study* [11]: up to 30 sensor motes deployed between heights of 15-70 m on a redwood tree. (3) *Volcano monitoring* [12]: up to 16 motes with seismometer sensors to gather data on volcanic activity. (4) *Industrial monitoring* [6]: motes with accelerometers to periodically measure the vibration pattern of motors. (5) *SenSlide* [8]: motes with strain gauges inserted into the earth, to predict land-slides; several nodes are expected to be deployed in a geographical region. (6) *BriMon* [5]: intended to periodically monitor the structural health of a railway bridge with accelerometer sensors deployed in each pier of the bridge.

## 3. EXTERNAL ANTENNA COMMUNICATION RANGE

Fundamental to the sensor network architecture is the achievable communication range. Longer reach of a link effectively means 1) smaller number of hops to the base node and 2) larger number of nodes that are one hop from the base node. As mentioned earlier, this could greatly simplify protocol design, improve performance and extend network lifetime.

In this section, we study the possibility of the use of external antennas to improve the communication range. We concern ourselves with applications in outdoor environments since distances are expected to be longer in these settings.

Some possible concerns with the use of external antennas could be the relatively large form-factor involved, cost and symmetry of communication (high gain antenna at the base

node and low gain (internal) antenna at the sensor mote). For 2.4GHz operation, high gain parabolic grid antennas [2], sector antennas [1], or even omni-directional antennas [3] can be 0.5-1m in length and weigh 0.5-5kg. They can cost between 50\$ - 120\$ depending on the gain and functionality. In all the outdoor applications listed in Section 2, we note that the base node does not particularly have any form-factor or cost constraints. In fact, in the case of Volcano monitoring, SenSlide, and BriMon, the individual nodes also do not particularly have form-factor constraints.

Note that while external antennas for 2.4 GHz operation are manageable, those for lower frequencies (433 MHz, 900 MHz) are likely to be much larger due to the larger wave-length. Commercially available external antennas for lower frequencies are also much more expensive (for e.g. see [www.hyperlinktech.com](http://www.hyperlinktech.com)). In this paper, we focus on 2.4GHz operation only.

An important, sometimes non-intuitive point to note is that the *symmetry* of a wireless link is *not* affected by the use of any kind of antenna at either end. Basic antenna theory says that for any antenna, the transmit gain is the same as its receive gain. So, in either direction, the overall gain (sum of transmit gain of transmitter and receive gain of the receiver) is constant. Note however that the transmit power or the receive sensitivity may have an effect on the symmetry of the link. So simply boosting the transmit power at the base node will not improve the bidirectional communication range. Also note that the asymmetry observed by researchers in sensor networks [14] is orthogonal to our proposed use of external antennas. Below, we first describe the experimental setup, followed by the results.

### 3.1 Experimental setup

We used Moteiv Tmote Sky motes [4] for our experiments. These motes come with a Chipcon CC2420 radio chip that is compliant with the 802.15.4 standard. To connect the external antennas to the motes, we soldered an SMA (Sub-Miniature ver-A) connector while also disconnecting the internal antenna. Note that the Tmote Sky comes with a 3.1 dBi internal antenna. For the external antenna, we experiment with three different types: a 24 dBi parabolic grid with a beam width of 8°, a 17 dBi 90° sector antenna, and an 8 dBi omni-directional antenna. We denote these henceforth as *grid*, *sector*, and *omni* respectively.

We try several combinations of antennas at the two ends of the wireless link: (1) internal-internal, (2) omni-internal, (3) sector-internal, (4) grid-internal, (5) omni-omni, (6) sector-omni, and (7) grid-omni.

One of the motes was designated as the transmitter and was mounted and clamped along with its antenna on a mast, whose height details are furnished in the individual experiments. The transmitter was programmed to continuously transmit a configurable number of packets with 20ms inter packet arrival. All packets were broadcast by the transmitter. The transmit power was fixed at 0 dBm, the maximum allowed by the CC2420 radio. Each packet contained a 10 byte header and 14 bytes of data that included a sequence number to help calculate packet loss rate.

The receiver consisted of a mote mounted along with its antenna on a tripod at a height of 1.5m to 1.7m. The receiver was programmed with TOSBase running in listening mode, and it forwarded the captured packets to a laptop along with the individual RSSI and LQI values.



Figure 1: Dense foliage, Narrow road environments

We performed the range studies in two different environments: dense foliage and a narrow straight road with foliage in the vicinity. The former was meant to capture the environment in which applications such as the Redwood study are deployed, while the latter was meant to be representative of applications such as Volcano monitoring, and Bri-Mon, where we have line-of-sight mostly. Fig. 1 shows views of the foliage environment with the omni antenna and the narrow road environment with the grid antenna respectively.

### 3.2 Results

#### Dense foliage environment

In this environment, we experimented with just the internal antenna at the receiver, though the transmitter’s antenna was varied between internal, omni, sector and grid. The transmitter antenna was fixed at a distance of 1.5m above the ground. For each choice of the transmitter antenna, we took readings at several receiver locations. We stopped at a distance where the received signal strength fell to about  $-85dBm$  or worse. The experiments in this setup used 6000 transmitted packets: about 2 minutes, with one packet sent every 20 ms.

Tx Antenna-Dist.	Avg. Pkt. Error % (Std. Dev)	Avg. RSSI (dBm) (Std. Dev)
Internal-20m	0.03 (0.26)	-74.98 (2.37)
Internal-25m	0.15 (0.63)	-74.41 (3.64)
Internal-30m	0.08 (0.53)	-79.44 (2.35)
<b>Internal-35m</b>	<b>0.3 (1.25)</b>	<b>-78.79 (3.43)</b>
Omni-20m	0.2 (1.1)	-72.97 (3.43)
Omni-30m	0.18 (0.81)	-76.75 (3.94)
<b>Omni-40m</b>	<b>0 (0)</b>	<b>-79.42 (2.35)</b>
Omni-50m	7.22 (16.5)	-86.68 (4.57)
Sector-30m	0.03 (0.26)	-67.76 (3.15)
Sector-40m	0.07 (0.52)	-69.33 (2.92)
Sector-50m	2.27 (4.55)	-82.76 (3.7)
<b>Sector-60m</b>	<b>0.53 (2.6)</b>	<b>-80.77 (3.55)</b>
Sector-70m	13.01 (14.37)	-90.01 (3.91)
Grid-70m	1.5 (3.61)	-79.79 (5.03)
Grid-80m	0.28 (0.99)	-77.07 (3.11)
<b>Grid-90m</b>	<b>1.6 (4.08)</b>	<b>-85.05 (4.19)</b>

Table 1: Range meas. in the dense foliage env.

In Table 1, we show statistics on the packet error rate and the RSSI for various configurations. The average and the

standard deviation of these metrics are calculated as follows. We first partition the transmitted 6000 packet sequence into bins of 100 packets each. We compute the error-rate and average RSSI within each bin. We then compute the average as well as the standard deviation of the resulting 60 values. We ignore bins which do not have any received packets, in which case, the computation of the average RSSI/LQI within the bin would not make sense.

In the table, we have marked those rows with bold font which seem to indicate the approximate range (in our particular setting) of the particular transmitter antenna in use. This range happens to be about 35m for the internal antenna (we did not receive any packets at 40m distance). It is slightly better, about 40-45m for the omni antenna. The range is about 60m for the sector antenna and is about 90m for the grid antenna. Thus the use of a grid/sector antenna in these settings can nearly halve the number of hops to the base node.

Given the high gain of the sector and parabolic grid antennae (17 dBi and 24 dBi), we had expected better ranges than 60m or 90m respectively. However, it is worth noting that in our environment the foliage was so dense that beyond about 35m, the people standing near the transmitter and the receiver could not see each other. Beyond 60m, we had to use cell-phones to communicate between the two ends.

#### Narrow road environment

Similar to the foliage environment, we conducted various range measurements on a narrow road. The road was about 4m wide and 500m in length. In this environment, we experimented with both internal and omni-directional antenna at the receiver. The transmitting antenna was placed at a height of 1.5m when using internal and omni. However we had to increase the height to 3.8m for the sector and grid antenna as the ground reflections at the lower height of 1.5m introduced losses in our measurements even at short distances<sup>1</sup>.

We realized towards the end of our study that the length of the road was limiting our range measurements. Hence, we continued further range studies in a nearby air-strip which was about 1 km in length. The measurements over 500 m in Table 2 are on the air-strip.

Table 2 shows our various range measurements in this environment. We note that the line-of-sight range measurements are much higher than in the case of our foliage measurements. This is true even in the case where we only used the internal antenna at the receiver. There was an increase of about 410 m in range, going from a foliage environment to direct line-of-sight environment, when using a grid antenna at the transmitter and internal antenna at the receiver. In the line-of-sight environment, replacing the internal 3.1dBi antenna to an 8dBi omni at the receiver, increased the range further by 300m.

We have also verified during the course of some of these experiments if the communication is indeed symmetric. We did not observe any significant asymmetry to report.

### 3.3 Implications

The range improvement due to the use of high gain directional antennas is relatively small in the foliage environment.

<sup>1</sup>Increasing height of the transmitting antenna in the foliage environment however did not seem to improve the range

Tx Antenna-Dist.	Avg. Pkt. Error (%) (Std. Dev)	Avg. RSSI (dBm) (Std. Dev)
<b>Internal antenna at receiver</b>		
Internal-60m	0.18 (1.03)	-81.11 (2.97)
<b>Internal-75m</b>	<b>1.37 (4.34)</b>	<b>-83.74 (3.61)</b>
Omni-60m	0 (0)	-77.45 (2.17)
<b>Omni-75m</b>	<b>0 (0)</b>	<b>-80.64 (2.47)</b>
Omni-90m	35.92 (33.42)	-94.91 (1.6)
<b>Sector-210m</b>	<b>0 (0)</b>	<b>-81.92 (0.49)</b>
Sector-310m	1.02 (4.3)	-91.85 (0.81)
Sector-400m	0.62 (2.24)	-92.33 (1.03)
Sector-500m	0 (0)	-90.12 (0.5)
Grid-90m	0 (0)	-75.35 (1.36)
Grid-210m	0.03 (0.18)	-75.82 (2.37)
Grid-300m	0 (0)	-80.42 (1)
Grid-400m	0 (0)	-82.21 (0.9)
<b>Grid-500m</b>	<b>0 (0)</b>	<b>-85.67 (0.94)</b>
<b>Omni antenna at receiver</b>		
<b>Omni-90m</b>	<b>0.04 (0.33)</b>	<b>-80.92 (0.88)</b>
Omni-150m	7.63 (12.46)	-90.86 (0.64)
<b>Sector-500m</b>	<b>0.13 (0.68)</b>	<b>-82.16 (0.37)</b>
Sector-600m	0.07 (0.25)	-89.48 (0.35)
Sector-700m	0.5 (1.05)	-91.22 (0.34)
Sector-800m	3.42 (4.83)	-91.58 (0.41)
Grid-500m	0.12 (0.49)	-75.25 (0.07)
Grid-600m	0.07 (0.25)	-79.85 (0.24)
Grid-700m	0.15 (0.61)	-82.07 (0.2)
<b>Grid-800m</b>	<b>0.13 (.39)</b>	<b>-85.76 (0.31)</b>

**Table 2: Range measurements in line-of-sight env.**

This is likely due to the heavy path loss which 2.4 GHz transmissions suffer in this environment. However, we feel that a range of 60-90m using the high gain antennae is still quite useful for many applications. For instance, the Redwood study [11] could likely use a sector/grid antenna at the base, facing up, to communicate with all of the motes (up to a height of about 70m) in a single hop.

The Volcano monitoring deployment [12] reports a range of 200-400 m using omni antennas at either end. In our experiments, we were able to achieve only about 150 m reliably (less than 10% error rate, across several days). It is not clear as to what the long-term error-rates are in [12].

Replacing the omni with a grid antenna at the base node, we are able to achieve a range as high as 800 m. This is about four times the range used in the current Volcano monitoring deployment. Note that since the nodes in the deployment are arranged in a line, the fact that the parabolic grid antenna has only a beam width of  $8^\circ$  should not be an issue. Our measurements thus have direct implications for a future deployment.

The implications are similar for a deployment of BriMon as well, since there too we expect to have nodes in a straight line, with form-factor of the individual motes' antennas not being an issue. In fact, for railway bridges which are about 800 m long (a significant fraction, in India), we could simply use a single-hop architecture.

Our measurements with the sector antenna are likely to be useful for deployments where the nodes will not necessarily be in a straight line. We are able to achieve a range of at least 200 m with the internal antenna at the receiver, and over 500 m with an external antenna at the receiver. Such ranges would be useful in application deployments such as Habitat monitoring and SenSlide, were they to use 2.4 GHz radios. Although the Habitat monitoring deployment [10] considers the use of external antennas, they do not use high-gain antennas for the communication between each sensor mote and the base node. This design decision is likely because of the use of 433 MHz and 900 MHz for communication, in which cases, the high gain external antennas are cumbersome as well as expensive.

We wish to caution that the ranges we have reported should be taken as providing a rough idea, rather than exact measurements. The exact values are likely to vary depending on the environment.

In summary, the consideration of high gain external antennas and the implications in terms of potential (a) increase in the number of one-hop nodes, and (b) reduction in the number of hops are novel aspects of our measurement study. These have direct implications for improving network lifetime significantly. In cases where a single-hop architecture is feasible, the network protocol design, implementation, debugging, management, etc. become far simpler.

## 4. STABILITY OF LINK-QUALITY

The previous section had looked at the communication range. Link range implicitly means communication at a certain tolerable loss rate. The important questions in this regard are: what is the behavior of the link over time? And what does this behaviour mean for dynamic metric-based routing? Before we answer these questions for realistic environments, we undertake a study in a controlled environment to isolate the radio behavior from environmental effects. We present this next.

### 4.1 A controlled calibration

We used a setup as shown in Fig. 2 for the calibration of the CC2420 radio. A 50-ft RF-cable with a step-attenuator in-between was used to connect the transmitter to the receiver. We ensured sufficient physical separation between the transmitter and the receiver so that we could safely discount any RF leakages from the connectors.

We varied the attenuation from 0 dB to 93 dB in steps, and in each instance, conducted an experiment where the transmitter sent a series of 5000 packets to the receiver, much as in our earlier experiments outdoors. We partition the packets into bins of 100 transmitted packets. For each such bin, we compute the average RSSI as well as the average packet error rate. A scatter plot of this data is presented in Fig. 3.

This plot is close to what we expect theoretically, with the packet error rate climbing from close to 0% to 100% within just about 5-6 dB. While this is not really surprising, what we observe next is. With the attenuator setting fixed at 90 dB, we look at the various 100-packet bins. Fig. 4 plots the observed error rate in each bin, against the bin number (i.e. time).

We see that there is a large temporal variation in the error rate, in the time-scale of 2 sec (100 packets at 20 ms inter-packet gap). The average RSSI over all of the received

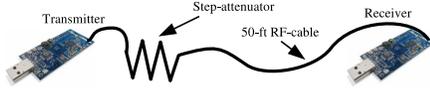


Figure 2: The setup for controlled calibration

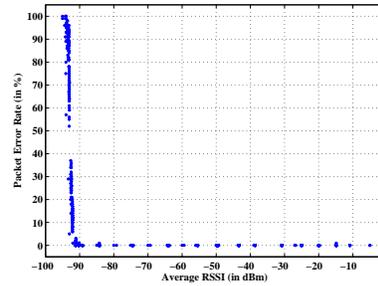


Figure 3: Error rate versus RSSI: using RF cable

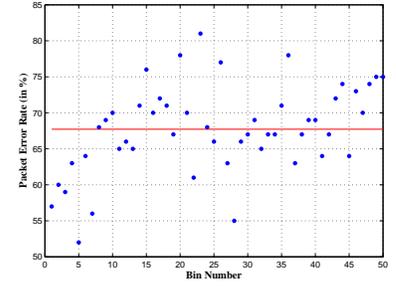


Figure 4: Error rate variation: using RF cable

packets was  $-94$  dBm in this experiment. Given that the environment is very controlled, we can attribute the variation to operating very close to the receive sensitivity of the radio. We observed similar large variations in the error rate in other cases too, where the overall error rate was neither close to 0% nor to 100%.

## 4.2 RSSI and error rate in other environments

Next, we look at plots similar to Fig. 3 for different real-life environments: (a) the air-strip, (b) the narrow road, and (c) foliage. Among these environments, we expect multipath to be minimum for the air-strip, maximum for the foliage and in-between for the road. For the air-strip and the road, we had conducted experiments at various distances, and on three different days around 6-8am in the morning (apart from the range study data presented in Sec. 3). We took care to mark the specific positions of the transmitter and the receiver so that the setup across the three days was as similar to each other as possible. All of these experiments were conducted using omni-antennas at either end. For the foliage environment, we consider here the same set of data as presented in Sec. 3. These measurements were done on a single day, but using different antenna combinations.

Fig. 5 presents the error-rate versus average RSSI plots for the three environments respectively. As earlier, the error-rate as well as the average RSSI were averaged over bins of 100 transmitted packets.

We see that in general there is a good correlation between the average RSSI and the packet error rate. This is similar to that observed in [9], except for the fact that we also see a larger “spread” among the points in the plot for the multipath prone environment (foliage). Our study spans several days, distances, and environments, and hence strengthens the conclusion in [9] that such correlation exists. While this conclusion is similar to that in [9], what we observe below is not only different, but also has starkly different implications; it also explains the “spread” we alluded to above.

## 4.3 Temporal variability in RSSI & error rate

The authors in [9] indicate that the RSSI variability across packets (within the 200 packet bin they considered) is not very high. And hence conclude that taking a single sample of RSSI would be indicative of the average RSSI. We now look at the per-packet variability of RSSI, and also the variability at larger time scales. (The study in [9] did not consider variability at larger time scales).

## Temporal variability in RSSI

To illustrate the RSSI variability, we consider a specific experimental run: the Omni-50m case of the foliage environment. We choose this run since it has an average packet error rate of 7.2%, i.e. neither close to 0% nor to 100%. Fig. 6 shows the CDF of the observed RSSI values. This is shown across the received packets of the 6000 transmitted packets (without any binning).

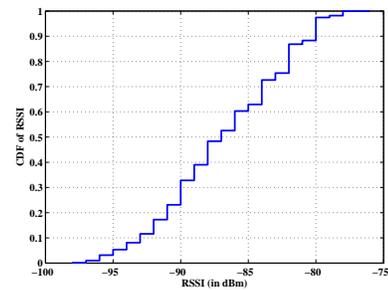


Figure 6: CDF of RSSI: Omni-50m case of Tab. 1

We see a large variation of the RSSI (about 15 dB) over the duration of the run. To further quantify the temporal variation of RSSI, we plot the observed RSSI against the received sequence number in Fig. 7. (Keep in mind that we can know the RSSI only for the successfully received packets.) From the figure, it is clear that we cannot use the RSSI in a single packet as an estimate of the average as suggested in [9]. We see that there is large variability in the RSSI (of 5-10 dB) even across single packets. There is no particular temporal pattern in the RSSI variability. To see if this variability is in fact subsumed by averaging within a time window, we average out the RSSI over bins of size 100. Such a plot is shown in Fig. 8. Fig. 9 shows a similar plot for a bin size of 1000 transmitted packets.

We observe that the time variability does not really go away by averaging it over 100 packets (2 sec) or even over 1000 packets (20 sec). This essentially means that we cannot really have a scheme where we measure the average RSSI over 100 (or 1000) packets and use it as an indication of the future. Of course such a scheme to average the RSSI would also beg the question: if the RSSI has to be averaged out over 100 or more packets, why not just measure the error rate directly?

We wish to stress that the RSSI variability was not really specific to this location or environment. We observed such

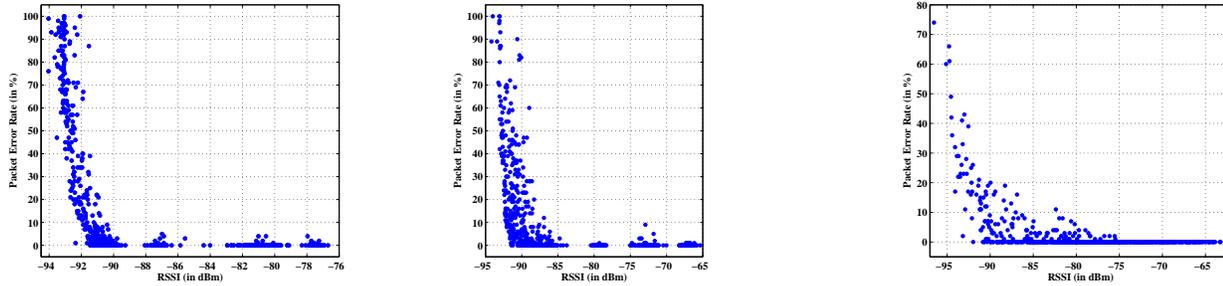


Figure 5: Error rate versus RSSI: (a) air-strip, (b) road, (c) foliage

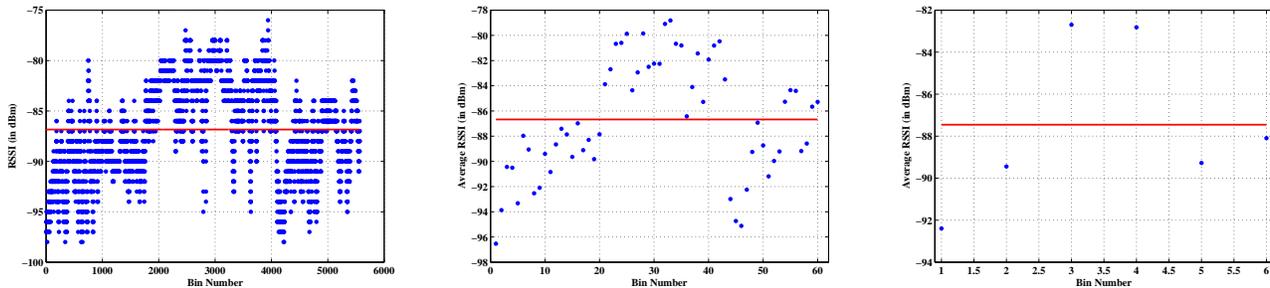


Figure 7: RSSI varn. with time: fo-  
liage, Omni-50m, binSz=1

Figure 8: RSSI varn. with time: fo-  
liage, Omni-50m, binSz=100

Figure 9: RSSI varn. with time: fo-  
liage, Omni-50m, binSz=1000

variability in several other cases too (including those with close to 0% error rate, as we shall summarize shortly). In some environments (air-strip, road), we had repeated the 6000-packet experimental runs after a gap of a few hours to a day. In such cases, we also observed variability in the average RSSI across different runs! Variability of about 4-5 dB across such periods was not uncommon. Thus there is variability at these large time scales too, of several hours to a day.

### Temporal variability of error rate

We have just shown that in a given environment, there could be large RSSI variability. This means that in the course of a run, we could have cases where the RSSI falls in the steep region of Fig. 3, although the average RSSI may well be above the threshold (as in our example above, the average RSSI was about  $-87$  dBm). Now, put together the observation from Fig. 4 that in the steep region of operation, there is an inherent variability in the observed error rate. We thus have the implication that we would observe variability in the error rate too (at the same time scales). This is illustrated in the plots in Fig. 10 and Fig. 11.

In case the overlap with the steep region is higher, the variability is even higher. Fig. 12 shows the error-rate variation for a bin size of 100, for an experimental run in the road environment when we had the mean RSSI to be about  $-89$  dBm. We see error-rate variation all the way from about 0% to about 90%!

In fact, throughout all our data, we observed high variability in error rate whenever the error rate was significant (say, over 5%). This observation essentially implies that a routing metric such as  $1/\text{PSR}$  (PSR is the Packet Success Rate) [13] would not really work. That is, by the time we end

up measuring the packet error rate meaningfully, it would have already changed.

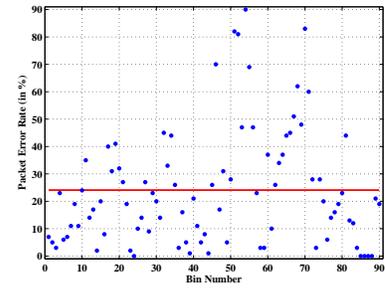
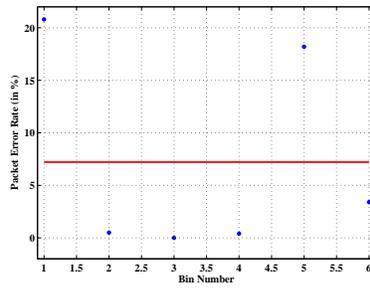
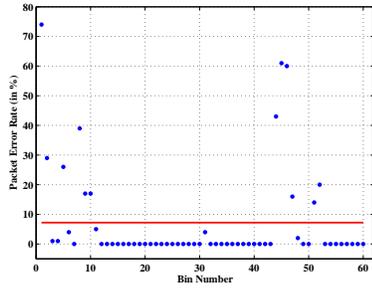
### RSSI variability in other environments

We further quantify the variability in the RSSI in various environments. We have considered two additional environments for this data. The first is a corridor in a student hostel, about 70 m in length, and expected to have significant multi-path due to the adjacent buildings, the floor, the ceiling, pillars, etc. The second environment is inside the structures-lab of our institute. This is an indoor environment which has not only a lot of multi-path, but is also a highly dynamic environment. During the course of our few-minute experimental runs, there was significant movement of personnel, machinery, material, etc. Both the environments are shown in Fig. 13.



Figure 13: Hostel-corridor, structures-lab env.

Recall that RSSI can be observed only for the packets successfully received. Hence, to capture the full extent of RSSI variability, we restrict ourselves to data points where we observed an overall error rate of less than 0.1%. In our



**Figure 10: Err. rate varn.: foliage, Figure 11: Err. rate varn.: foliage, Figure 12: Err. rate varn.: road**  
**Omni-50m, binSz=100                      Omni-50m, binSz=1000                      90m, binSz=100**

data collection, the number of such data points was very large; we present only a small, representative subset of such points.

For various environments, Table 3 captures the RSSI variability in terms of the difference between the 95th percentile and the 5th percentile in the range of RSSI values. This metric, instead of the standard deviation, removes the effect of outliers. In the table, the different distances indicated in the first column are the distances between the transmitter and the receiver. For the lab environment, we took measurements at four different locations. These locations were at different points in the lab, at distances of 10.3 m, 14.8 m, 21.5 m, and 18.4 m respectively.

Location	Tx ant.	Rx ant.	Day/time	5 <sup>th</sup> perc. (dBm)	95 <sup>th</sup> perc. (dBm)	Diff. (dB)
Airstrip-90m	Omni	Omni	Day 3	-84	-80	4
Foliage-20m	Internal	Internal	Day 1	-81	-70	11
Foliage-40m	Omni	Internal	Day 1	-86	-76	10
Foliage-40m	Sector	Internal	Day 1	-76	-66	10
Road-210m	Grid	Internal	Day 4	-80	-70	10
Corridor-60m	Internal	Internal	Day 1	-71	-68	3
Corridor-60m	Internal	Internal	Day 3	-76	-70	6
Corridor-30m	Internal	Internal	Day 5	-76	-67	9
Road-55m	Omni	Omni	Day 2	-69	-66	3
Road-90m	Omni	Omni	Day 3	-82	-79	3
StrLabLoc1	Internal	Internal	Day 1 15:45	-76	-66	10
StrLabLoc1	Internal	Internal	Day 1 23:30	-71	-66	5
StrLabLoc2	Internal	Internal	Day 1 15:45	-87	-78	9
StrLabLoc2	Internal	Internal	Day 1 23:30	-78	-74	4
StrLabLoc3	Internal	Internal	Day 1 15:45	-80	-73	7
StrLabLoc3	Internal	Internal	Day 1 23:30	-78	-76	2
StrLabLoc3	Internal	Internal	Day 2 11:15	-84	-76	8
StrLabLoc4	Internal	Internal	Day 1 23:30	-89	-82	7
StrLabLoc4	Internal	Internal	Day 2 11:15	-88	-78	10

**Table 3: RSSI variation (95th - 5th percentile)**

We see in the table that the variation can be as high as 11 dB in many cases. Again, we wish to stress that Table 3 is only a small sample of our overall data. In other instances too, we observed such high variability: cases where the variability was less than 4 dB were quite rare in comparison.

In our earlier experiments (air-strip/road), we had observed average RSSI variability at time scales of about a day. To see if there was variability at the time scale of a few hours, we had measured once in several hours in the lab environment. On close examination of the individual data points tabulated above, we observed that there was RSSI variation at this time scale too. In sum, we have thus observed RSSI variability at time scales of 2 s, 20 s, several hours, and a day as well.

## 4.4 Implications

Our measurements of RSSI and error-rate variability as quantified above have far reaching implications. Similar to RSSI variability, we observed variability in the LQI too, although we have not reported it. The LQI variability is to be expected, of course, since LQI is only a measure of chip-level error-rate (which is what manifests itself as packet error rate). Now, measuring the RSSI or LQI as an indicator for the error-rate makes sense only if such a measurement can be done with a smaller sample than for the error-rate. The variability we have observed effectively rules this out.

The fact that we have variability in the error rate or the LQI over scales of 100-1000 packets means that metrics such as 1/PSR or 1/LQI would likely not produce stable/optimal behaviour.

The variability at the time scales of several hours to a day has implications for applications such as BriMon or Industrial monitoring. In these scenarios, the nodes need to collect data for every few hours to a day, sleeping in the duration in-between. The variability means that we cannot reliably use a mechanism whereby we measure the RSSI (or the packet error rate) at one wake-up time, and use it for routing during the next wake-up time. The measurement done earlier would likely not be valid by the time of the next wake-up period.

We have observed RSSI variability at time scales of 2-20 sec, and also at time scales of several hours to a day. Now, there does remain the open question of whether at a given location, the error rate is stable when averaged over time scales larger than 20 sec, but smaller than several hours. We feel that this is unlikely, although not ruled out by our measurements as such.

## 5. DISCUSSION

We have argued above that dynamic link metric estimation based on RSSI, error-rate (or LQI) may not be meaningful/optimal due to variability. However, we hasten to point out that the instability in error rate is only when the RSSI variability window overlaps with the steep region of Fig. 3. That is, when the link sometimes operates close to the receive sensitivity, and the overall error rate is significant (over 5%). The immediate next question then is whether it is possible to say when a link has a stable error-rate.

The answer to this is suggested by the data in Table 3. We could take into account the RSSI variability to provide sufficient link margin to operate away from the steep region, proceeding as follows. We can make RSSI measurements for a certain duration. Such measurements can be done

passively in the background, or may be done before actual deployment. Given sufficient samples, the maximum of these RSSI measurements would fall beyond the 95th percentile of Tab. 3. We could then make a worst-case assumption of say 11 dB (as suggested by our data) for the RSSI variability. If the RSSI variability window does not overlap with the steep region, we could say for sure that the link would have stable and low error rate. In such an approach, we would be erring on the side of safety while determining a link to be stable; i.e., we could potentially miss out on some stable links, but would not classify unstable links as being stable.

This in essence provides a means to plan a sensor network deployment (such as the examples cited in Sec. 2). The inter-node distances and the network architecture can be designed to accommodate any RSSI variability. With such an approach, we would have links which have low error rate and are also stable over time.

Now, suppose that we only use such stable links, then routing only needs to take care of occasional failures and does not need to track dynamic link metric change. Such operation would be beneficial for predictable behaviour. It may even have lower overhead since we avoid the frequent message exchange potentially triggered by using unstable links.

We can take the above argument of using stable links even further by observing the following. Sensor network applications such as those in Sec. 2 anyway have a central base node which is more powerful in terms of CPU, memory, and power. A routing protocol (such as shortest path, or its variant) should not be too difficult to run on the central node even for say networks of size of a few thousand or more nodes<sup>2</sup>. For instance, Dijkstra's shortest path algorithm is only  $O(E \log(N))$  in time complexity and  $O(E)$  in space complexity for a network of  $E$  edges and  $N$  nodes.

Given these, we argue that centralized (as opposed to distributed) routing becomes a serious option to consider. Fault tolerance in such a centralized approach may not be a concern since the failure of the central node will anyway bring any data collection operation to halt. This would not only simplify protocol design, implementation, and management, but would also make possible other optimizations. For instance, load-balanced routing (when traffic patterns are known apriori) among one-hop nodes would be much easier to do with centralized routing. Recall that in applications such as those in Sec. 2, the one-hop nodes are the bottleneck in terms of the network lifetime.

Finally, given the variability we have observed in our experiments, we advocate a simple mechanism as a design option for indoor sensor network deployments such as [6]: simply increase the number of base nodes to cover all the other nodes within a single hop. This is because, power off-the-grid for several base nodes should not be an issue at least in most indoor environments. Such an approach is further supported by the fact that 2.4GHz coverage planning is already a well developed technology in the context of WiFi Access Points.

In sum, our measurements have thus opened up important design alternatives to consider and evaluate further.

<sup>2</sup>Even such scale is quite futuristic given that current deployments hardly have a few tens of nodes.

## 6. CONCLUSION

In this work, we have first presented and evaluated extensively a significant design choice for outdoor sensor network deployment, namely, the use of high gain external antennas. We have shown that ranges of up to 90 m are possible in dense foliage environments, while ranges up to 500-800 m are possible in line-of-sight environments with the use of high gain sector/grid antennas at the base node. This has significant implications for sensor network architecture: we can greatly reduce the number of communication hops, and/or increase the number of one-hop nodes. These would have significant impact on the network lifetime.

Next, we have evaluated the variability in RSSI and error-rate extensively. Our measurements indicate variability at time scales of 2-20 s as well as several hours. This raises serious doubts regarding the use of dynamic link metric estimation and its use for routing. We advocate a mechanism to plan links to operate in a region where we have stable (and close to zero) error rate. Given the use of stable links, we further suggest that design options such as centralized routing become a serious possibility to consider in the application scenarios prototyped thus far.

## Acknowledgment

We thank Akhilesh Bhaduria, Ranjeet Khas, Ramchandra, and A. K. Singh for having helped us with some of the experiment logistics.

## 7. REFERENCES

- [1] Hyperlink Technologies: HG2417P-090. [http://www.hyperlinktech.com/web/hg2417p\\_090.php](http://www.hyperlinktech.com/web/hg2417p_090.php).
- [2] Hyperlink Technologies: HG2424G. <http://www.hyperlinktech.com/web/hg2424g.php>.
- [3] Hyperlink Technologies: HGV-2409U. <http://www.hyperlinktech.com/web/hgv-2409u.php>.
- [4] Tmote sky. <http://www.moteiv.com/products-tmotesky.php>.
- [5] H. Haridas. BriMon: Design and Implementation of Railway Bridge Monitoring Application. Master's thesis, Indian Institute of Technology, Kanpur, May 2006.
- [6] L. Krishnamurthy, R. Adler, P. Buonadonna, J. Chhabra, M. Flanigan, N. Kushalnagar, L. Nachman, and M. Yarvis. Design and Deployment of Industrial Sensor Networks: Experiences from a Semiconductor Plant and the North Sea. In *SenSys*, Nov 2005.
- [7] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson. Wireless Sensor Networks for Habitat Monitoring. In *ACM International Workshop on Wireless Sensor Networks and Applications (WSNA)*, 2002.
- [8] A. Sheth, K. Tejaswi, P. Mehta, C. Parekh, R. Bansal, S. Merchant, T. Singh, T. C. A., and K. Toyama. SenSlide - A Sensor Network Based Landslide Prediction System. In *SenSys*, 2005. Poster.
- [9] K. Srinivasan and P. Levis. RSSI is Under Appreciated. In *EmNets*, May 2006.
- [10] R. Szewczyk, A. Mainwaring, J. Polastre, J. Anderson, and D. Culler. An Analysis of a Large Scale Habitat Monitoring Application. In *SenSys*, Nov 2004.
- [11] G. Tolle, J. Polastre, R. Szewczyk, D. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong. A Macroscopic in the Redwoods. In *SenSys*, Nov 2005.
- [12] 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*, Mar/Apr 2006.
- [13] A. Woo, T. Tong, and D. Culler. Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In *SenSys*, Nov 2003.
- [14] J. Zhao and R. Govindan. Understanding Packet Delivery Performance in Dense Wireless Sensor Networks. In *SenSys*, Nov 2003.