

Link Layer Measurements in Sensor Networks

Niels Reijers* Gertjan Halkes Koen Langendoen
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology, The Netherlands

E-mail: {N.Reijers,G.P.Halkes,K.G.Langendoen}@ewi.tudelft.nl

Abstract

Key issues in wireless sensor networks such as data aggregation, localisation, MAC protocols and routing, all have to do with communication at some level. At a low level, these are influenced by the link layer performance between two nodes. The lack of accurate sensor network specific radio models, and the limited experimental data on actual link behaviour, warrant additional investigation in this area.

In this paper we present the results from extensive experiments, exploring several factors that are relevant for the link layer performance. These include (i) the effect of interference from simultaneous transmissions, which has not been looked into before, (ii) the degree of symmetry in the links between nodes, and (iii) the use of calibrated RSSI measurements. Finally, we present some guidelines on how to use the results for effective protocol design.

1. Introduction

Wireless sensor networks promise many new applications through the use of small and cheap wireless sensing devices that can run on battery power for several months or years. Nodes will have a very tight energy budget, few processing resources, little memory, and limited communication capabilities. This combination of strict constraints had not been previously addressed in research on wireless networks and mobile computing, and opens up new challenges. Together with the promise of new and exciting applications, this has given rise to the development of new algorithms for wireless sensor networks on various topics like localisation, routing, medium access control, etc.

Since wireless communication is expensive in terms of energy consumption, managing it efficiently is at the heart of many sensor network specific protocols. However, the

mental model that protocol designers often use during development is not much more sophisticated than a simple circular model. Although it is well known that this is an oversimplification, empirical data on the behaviour of wireless links in realistic environments using cheap hardware, has so far been limited; we are only aware of a few publications that include detailed measurements [2, 8, 9]. Having a better understanding of the actual behaviour of wireless links can benefit protocol design by drawing attention to potential problems otherwise hidden by a too simple model. For example, geographic routing fails miserably in the presence of bad links [10].

Sensor networks are expected to be deployed at a scale of potentially hundreds or more nodes. Real experiments at this scale can be difficult, so simulations are often used. Currently, popular simulators like the ns-2 network simulator [14] and GloMoSim [12] gloss over link layer details and use antenna, propagation, and interference models that are inaccurate. The OPNET simulator [6], on the other hand, includes very detailed models, but is expensive and computationally complex. In order to develop adequate models that capture the important effects, and nothing more, we need detailed information about actual link behaviour.

In this paper we present the results of extensive link layer measurements with prototype hardware (51 nodes) in different environments (indoor, outdoor, open space). We focus on aspects that are important for protocol design: the influence of the environment, directionality in the nodes' antennas, the degree of symmetry in the communication links, the effect of interfering transmissions, and the possible use of the received signal strength indication. Finally, we present some guidelines on how to use the results for effective protocol design.

2. Test method

For our tests, we used nodes from the European EYES research project [11], built around 5 MHz Texas Instruments MSP430F149 processors equipped with 2 KB RAM (data) and 60 KB ROM (code). The nodes have a

* Supported by NWO (Dutch National Science Foundation) in the CONSENSUS project.




Figure 1. EYES node with embedded antenna.

256 KB EEPROM memory. The radio used is an RFM TR1001 [7], similar to the radio used on the popular MICA motes [13]. It uses amplitude shift keying (ASK), and operates at 868.35 MHz and 115 Kbps. The radio is connected to a dipole antenna, which is embedded in the printed circuit board (PCB) for robustness. The node, and a closeup of the antenna are shown in Figure 1. Using a digital potentiometer, we can control the transmit power in 32 linear steps, where 0 is the lowest and 31 the highest transmit power setting.

Packets are sent and received using the MSP430's internal UART. All packets start with a preamble, consisting of a training sequence of three 0x55 bytes, followed by a 0 byte and 11 high bits to synchronise the receiver's UART to the correct start bit (tests were also done with a more balanced preamble where the UART was manually synchronised, but this led to significantly worse performance). Then a start byte, 0x50, is sent to signal the beginning of the packet, followed by the header and payload in a 4b8b encoding that ensures a DC-balanced signal and allows for limited error correction. In contrast to ordinary Manchester encoding, which takes only care of DC balancing, we use an encoding that increases the Hamming distance between most codewords. This enables us to correct bit errors in some, but not all cases.

We performed experiments in three different environments. The hardest one was in a corridor in our building, where there are a lot of reflections. The second one was a tennis court, which is mostly featureless, except for aluminium net posts, lying flat on the ground, an iron wire fence surrounding the courts, and some iron lines around the concrete slabs that make up the courts. Finally we tested in the middle of an Astroturf field-hockey pitch, with no obvious sources of reflection within at least two radio ranges (over 20 metres to the nearest fence).

In all our tests, the nodes are aligned on a straight line. Nodes were placed on the ground, with the batteries touching the ground as shown in Figure 1. The node at the beginning of the line sends a packet every 50 ms, containing a payload of 10 bytes, which includes a sequence number. Each run lasts for 2048 packets, or 102.4 seconds. The other nodes are receivers and keep a log of which packets they receive and which they have missed. This data is written to

EEPROM, as well as various statistics such as received signal strength indication (RSSI), and the number of bit errors, which we can detect as long as the synchronisation is not lost because the receiving nodes know the contents of the messages that will be sent. In practice, we discovered that messages are either received and decoded correctly without any errors, or the start symbol is not correctly detected at all; we rarely encountered a checksum failure. We conjecture that our preamble and start symbol detection only succeeds when the Signal-to-Noise Ratio (SNR) is at a rather high level, which warrants the correct reception of the subsequent data bits. Note that this binary behaviour is unique to our hardware/software setup as others do report frequent bit errors in noisy conditions [9].

3. Environmental effects

In our first set of experiments we will look at the influence of the environment on packet reception. While this is generally ignored in simulations because it is too hard to model, we found it to have a significant impact on our results. We placed 51 nodes in a straight line in the corridor. The first receiver was placed at five metres from the sender, and nodes were placed one metre apart up to 15 metres from the sender. From there they were placed half a metre apart. We tuned the transmit power such that the whole radio range fits in this area, resulting in power setting 23 (out of 31).

Figure 2 shows the result of our first run, which we will use as a base to compare other data to. The figure clearly shows an area, up to about 16 metres, where packet reception is almost perfect. Beyond this we would expect a sudden drop to 0 percent reception as the signal-to-noise ratio drops below the required minimum. However, instead we see a very erratic picture with nodes relatively close to the sender performing poorly, and nodes farther away receiving almost all packets. This is consistent with the *gray area* shown in [9]. For reference, we will indicate the gray area in our results by shading along the x-axis (cf. Figure 2). The gray area is important. For example, in Figure 2, the gray area covers about half of the radio range, and assuming a circular model and uniform node distribution, about 75% of the neighbours are involved.

3.1. Link quality

We classify links as *good*, *medium* or *bad*. Good links are links with 85% reception or more. These links are useful for communication since the packet loss can be handled by retransmissions. Bad links have a reception of 15% or less, and are links that are mostly dead. Medium links do not deliver enough messages to be useful, but do deliver enough so that they can be a problem for many protocols. For example, routing protocols need to consciously avoid them be-




Figure 2. Packet reception in the corridor (base run).

cause many packets will be lost, even though enough may get through to setup the route. Similarly, for MAC protocols, these links will cause collisions and interference, but the performance may not be good enough to complete a whole RTS/CTS/DATA/ACK exchange.

Looking at Figure 2 again, we observe that most links are either good or bad. Of the 37 nodes in the gray area, only 7 have medium performance links.

3.2. Factors contributing to the reception pattern

There are many different factors that can influence the number of packets a node receives:

1. Multipath effects
2. Human activity
3. Receiver sensitivity
4. Sender characteristics
5. Node orientation (Section 4)
6. Interference from other nodes (Section 6)
7. Background noise / temperature / humidity

For each node the performance is determined by the sum of all these factors. We believe the pattern we see in the gray area is caused by complex multipath effects in the corridor that degrade the signal in some spots, and amplify it in others (even increasing the radio's range). To verify this, we need to rule out possible other causes. Only the second and third factor could potentially be responsible for the gray area – the others will affect packet reception, but cannot cause the gray area by themselves.

Temporal behaviour There are several effects that are time dependant and that can influence packet reception. Some of those may cause local disturbances that contribute to the gray area effect. The activity of people within the building changes during the day, and the equipment they operate may affect our measurements. Also, the temperature changes




Figure 3. Base run setup on a different day.

during the day, which may affect our radios and batteries. Although the measurements were done during daytime, human activity was very limited and the experiments were done without any people in the corridor.

So one of the first things we want to know about the pattern we see in Figure 2 is how stable it is over time. To determine the 'short' term behaviour, in the order of minutes, we did three identical runs of our base scenario immediately after each other. The difference between them was minimal. Incidentally, an individual node may perform better or worse, but the overall picture was identical. This is confirmed by the average reception of all nodes. The three runs scored 53.54%, 55.53%, and 55.49%.

We also tested long term stability of the behaviour, by doing more runs on different days. The run that showed most deviation from the base run is shown in Figure 3. There is clearly a big difference between the base run and this one. The average reception over all nodes has gone up from 53.54% to 67.62%. But still the pattern of highs and lows is the same, except for some exceptions around 22 metres. Nodes that performed well in the base run still perform well, and nodes that performed badly still perform badly. Similar fluctuations were found between identical runs a few hours apart.

This is consistent with the idea that different factors influence the performance, but that the pattern is caused by reflections in the corridor. Clearly some external factor has changed compared to the base run, and has improved the performance of all nodes. This could for instance be temperature, since this test was done at a different time of day. However, the positions of the nodes, and the geometry of the corridor have not changed. Therefore, the reflection pattern is still the same, although all nodes are receiving more packets.

Individual node performance Another possible candidate for causing the reception pattern is differences in the sensitivity of individual receivers. To determine how this influences our measurements, we shifted all nodes 3 positions




Figure 4. Shifted receivers (a,b) and different sender (c,d).

forward, wrapping the last three around. The first receiver is now at position 4, and the last receiver has come around to position 3. The result is shown in Figure 4(a).

If receiver sensitivity was causing the gray area, the entire pattern should have shifted three positions to the right. However, this has not happened. Figure 4(b) shows the difference between Figure 4(a), and the base experiment. If the pattern had shifted, we would see many more long bars. Instead all but a few show very little change. We see individual nodes performing better or worse. This is something we see throughout all our tests. It is caused by various sources of noise that are beyond our control. However, when we look at the characteristic peaks and drops, we see that they are still in the same location.

We also replaced the sending node, and again the basic picture remained the same, shown in Figures 4(c) and 4(d). This indicates that the difference in the characteristics of individual nodes is not large enough to have a significant impact on our measurements.

Other environments Because all the features in an office corridor are bound to cause many reflections, we repeated these experiments in two different environments. In Figure 5(a) we see the reception for the same experiment on a tennis court. The gray area effect is indeed much less pronounced than in our initial tests, since there are less things for the signal to reflect off. However, even in this clean environment, there is still a gray area of significant size. It is most likely caused by to the fence around the courts and the aluminium net posts. Also note that the radio range was reduced significantly, and we had to do this experiment at the




Figure 5. Packet reception on the tennis court (a) and hockey pitch (b).

maximum power level. This is possibly due to the fact that the reflections in the corridor caused the signal to propagate farther than in a completely open space.

We did a third set of experiments on a hockey pitch. In this environment, there was not anything within the radio range that could interfere, and as a result, the gray area is minimal.

The last two experiments show that the gray area is not just a result of the hard conditions in the corridor, but that it will occur even in quite benign environments.

4. Directionality

Next, we turn our attention to the directionality of the nodes' antennas. For practical reasons we have limited ourselves to one axis in these tests, because in many real deployment scenarios it will be feasible to ensure nodes land, or are placed with the proper side up, but it will be more difficult to control their orientation along the vertical axis.

Directionality can be split into two factors: directionality in the sender's output power, and directionality in the receiver's sensitivity. We will examine both separately.

It needs to be said that the design of the antenna obviously has a large impact on the directionality. Therefore, the results presented in this section are more specific to our type of nodes than the rest of this paper. Having said that, we feel that if nodes are to be mass produced at a very low cost, and are expected to be handled roughly during deployment, the EYES node's embedded dipole antenna may be an attractive option, both from a cost and robustness perspective.

4.1. Sender directionality

We performed several different experiments to determine the sender's directionality. The difference in output power is the same regardless of the environment, but the presence of the strong gray area effect in the corridor makes it difficult to clearly visualise this. Therefore Figure 6(a) shows the re-

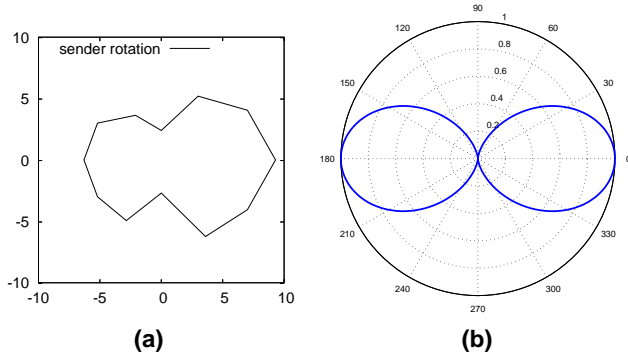


Figure 6. Observed sender directionality (a), and theoretical radiation pattern (b); hockey pitch.




Figure 7. Receiver directionality; hockey pitch.

sults from a test in the cleanest environment: the hockey pitch. As in the previous section, we placed the nodes in a straight line. We then rotated the sender in steps of 30 degrees. For each run, we determined the first node where reception dropped below 85%. The previous node’s location is used as the radio range for that angle.

The resulting 8-like shape matches reasonably well with the theoretical radiation pattern, shown in Figure 6(b). Observe that the actual signal is stronger in one direction than in the other. A likely cause of this phenomenon is the presence of track segments on the PCB, parallel to the antenna, that act as parasitic directors for the electromagnetic field.

4.2. Receiver directionality

We examine receiver directionality by rotating the odd-numbered nodes, and determining the range as for the sender directionality. The even numbered nodes are used to verify the results. These nodes should show similar performance in all experiments.

According to theory, the pattern for receiver sensitivity should be the same as for the transmitted power [4]. The results shown in Figure 7 show a very similar pattern compared to the graph shown for the sender, although slightly




Figure 8. Directionality in the corridor.

less pronounced. We can also see that the static nodes, which were not rotated, show a perfectly circular graph, indicating that conditions did not change significantly during the test.

4.3. Directionality and environmental effects

Both the sender and receiver directionality were also tested in the corridor environment. The results for this environment are, especially for the receiver directionality test, mangled by reflections. Figure 8 shows the results where we took the average of each node and its four closest neighbours before determining the radio range to prevent individual nodes from cutting off the range too quickly and making the picture completely unrecognisable.

The reflections in the corridor environment have a stronger effect on receiver directionality than on sender directionality. One explanation for this, is that because of the many reflections, receivers may receive quite strong signals from directions that are at an angle to the direct line to the sender. The directions of the reflected signals may be different for each node, and the nodes’ most sensitive side is facing another direction each rotation. Therefore the receivers are affected by the rotation differently.

For the sender rotations the receiving nodes face the sender with their most sensitive side, so the direct signal from the sender is the dominant signal.

5. Symmetry

One important question for protocol designers is the degree of symmetry in communication links: If I receive a message from my neighbour, what are the chances that he will receive my reply?

This is especially important for the gray area, where we see individual links performing much better or worse than one would expect judging by distance alone. These good links may be very useful, especially since the gray area can cover up to 75% of the neighbours. But they are only of

	Good	Medium	Bad
Good (1319, 53%)	1162	117	40
Medium (432, 18%)	117	212	103
Bad (699, 29%)	40	103	556

Table 1. Link classification: full dataset, 2450 links.

	Reverse link		
	Good	Medium	Bad
Good	88%	9%	3%
Medium	27%	49%	24%
Bad	6%	15%	79%

Table 2. Conditional probabilities for the reverse link quality with full dataset, 2450 links.

use if the reverse link is good as well. If not, they may even cause problems. For example, at the MAC layer, if one node’s transmissions are interfering with another node, but that node cannot communicate with the sender to prevent collisions.

If we believe the gray area to be the result of multipath reflections, we would expect the conditions to be reasonably symmetric. Although it is possible to construct scenarios where the following does not hold, in most environments including our corridor, a radio wave travelling from A to B over one or more reflections should be able to follow the same path of reflections from B to A.

To determine the degree of link symmetry we conducted the following experiment: 50 nodes were placed in a line, spaced 50 centimetres apart. We then did 50 separate experiments similar to the base run, with a different node sending in each run. This gives us a full 50×50 matrix of individual links. The whole experiment lasted for about two hours.

The 50 nodes resulted in 2450 individual links (one way, i.e. two entries for every pair of nodes), which we then qualify into good, medium, and bad links, as in Section 3.1. The results are shown in Table 1. The left column shows the total number of good, medium and bad links. In these experiments, we found more than half of the links to be good, but of course this depends on the locations of the nodes, so the figure in itself does not say much.

The other three columns show the performance of the link in the other direction. Of the 1319 good links, the reverse link was good as well in 1162 cases, and the reverse link was medium or bad in 117 and 40 cases. Next we calculated the percentages for each link type. Table 2 shows that if we have a good link one way, there is a 88% chance that the reverse link is also good, and only a 3% chance that it is bad. The picture for bad links is similar; the reverse link




Figure 9. The gray area in the symmetry test setup.

	Reverse link		
	Good	Medium	Bad
Good (273, 30%)	82%	15%	3%
Medium (160, 17%)	26%	38%	36%
Bad (497, 53%)	1%	12%	87%

Table 3. Symmetry results: gray area, 930 links.

is also likely to be bad. For medium links, the picture is a bit fuzzier, but these only constitute 18% of the total links. These numbers confirm our earlier visual impression from Section 3. Most links are either good or bad, which is good news. However, the number of medium quality links is significant enough that they should be taken into account during protocol design.

Although the links seem quite symmetric, many of the good/good links are in the area before the gray area, where reception is almost always good, and thus likely to be symmetric. To look more closely at the gray area, we exclude all links that are less than 10 metres apart. The value of 10 metres was chosen by visually inspecting the results from the first of the 50 symmetry test runs. This run, where the left-most node was sending, is shown in Figure 9 (the node at 10.5 metres was defective).

The results for this limited set (see Table 3) are surprisingly similar to the results of the whole dataset. Of course the percentage of good links has dropped, from 53% to 30%, and the bad links have increased. But the symmetry is still high. The chance of a good link having a good return link as well has dropped only from 88% to 82%, which is not very significant given the limited number of experiments and high degree of noise.

Most other probabilities have changed by similar amounts, indicating that link behaviour is quite symmetrical for links both in and outside of the gray area. Also, it is interesting to note that although we have more bad links and less good links in the gray area, the percentage of medium links has not changed much.




Figure 10. Spread of asymmetry ($A \rightarrow B - B \rightarrow A$).

5.1. Behaviour of asymmetric links

Now let us look at the asymmetric links. Exactly how asymmetric are they? In Figure 10 we plot the difference between the reception rate of all link pairs, except those that are both good, or both bad. These links are uninteresting because they both have very similar performance. From Table 1 it follows that we are looking at the least symmetrical 472 (432+40) of the 2450 links. The histogram is normalised to have a surface area of 1.

The first interesting thing we see in Figure 10 is that the histogram is quite flat. If the performance of the links was completely random, the slopes of the histogram would be almost perfectly linear, as shown by the line in the histogram (the line is not completely linear, because we excluded the cases where both links are good or bad).

This is easily explained. As we can see in Table 2, more than half of the link pairs we consider here have at least one good or bad link. This means that for all those pairs with at least one link close to 0% or 100%, there is a larger than average chance that the other link will have a very different performance, resulting in a flatter histogram.

5.2. Effect of directionality on symmetry

The second thing we notice about Figure 10 is that it is quite symmetrical. In Section 4 we saw that the nodes transmit stronger in one direction than in the other, and also that they are more sensitive to reception in that same direction (in this test we only look at the reception rate along the strongest axis of the node). However it is unclear what this means for symmetry.

Figure 11 shows two nodes aligned as they were in the symmetry tests, and also the pattern we found in our sender directionality tests. The pattern we found for receiver directionality was similar, at least so far as that it was stronger in the same direction (to the left in this figure). Of course there were more nodes in between, but these are not important for this section. The question arises whether B will be able to receive A's messages.




Figure 11. (A)symmetry due to directionality; can A hear B?

If we only look at sender directionality, we would expect an asymmetrical link from B to A, because A's transmission will not reach B. If we only look at receiver directionality, we would expect an asymmetrical link in the reverse direction because A's sensitivity will not be high enough to hear B's message, but B will be sensitive enough to hear A.

If either of these effects clearly outweighs the other, we would see this as a bias in Figure 10. The bias would be towards negative values if sender directionality is more important, or positive values if receiver directionality is more important. In fact, there is only a very small bias: the mean of the histogram is 0.038. The fact that we hardly see any bias suggests that both factors keep each other in balance, which is good news for symmetry.

A bias in Figure 10 could also be caused by changing external conditions. Although our earlier results suggested that during a time window of about two hours changes in the external conditions are limited, we did a simple check to be sure. We split the 50 nodes up into a left and right half, and calculated the average reception rate of all links *within* each of those two halves. Links from one half to the other need to be excluded because they could be influenced by the effect of directionality. Within each half each link is counted in both directions, so even if there is such an effect, it will be the same for both halves.

The resulting averages are 76.9% for the left half, and 76.1% for the right. This difference is small enough to be confident that external conditions did not affect Figure 10 by much. Further, the fact that the left half performed slightly better pushed Figure 10 towards a positive bias, so without the slight change in external conditions, the bias would even be smaller.

This experiment was limited to one axis only, where the difference in power/sensitivity was significant, but not as large as compared to nodes at a 90° angle. More tests are needed to confirm this initial result.

6. Interference

Previous results on measuring link layer performance in sensor networks have dealt with a single transmitter. However, for certain problems, especially MAC layers, it is im-




Figure 12. Positions of interference node.

portant to know the actual behaviour when a receiver is in range of more than one transmitter.

We measured the effect of interference from other nodes on the hockey pitch, in a setup where nodes were uniformly spaced 30 centimetres apart. We used one node to act as an interferer. It continuously sends the interference byte: 0x55. Including start and stop bits, this translates to a ‘01010101’ binary pattern. The interference node was placed at different positions along the line of nodes, as shown in Figure 12. During each run, the sender was aimed with its strongest side pointing towards the receivers. The interferer had its strongest side pointing towards the sender. Because of the limited range on the hockey pitch, transmission power was set to the maximum.

6.1. Radio range

When analysing the effect of interference, we need a metric for the radio range. We used the position of the last node that received 85% of the packets or more as the radio range. This means we include the gray area. We will look at what happens within the radio range in Section 6.2. The two most commonly used models for interference are the circular model, where no communication is possible when there is a collision, and SNR based models, where one message may be received correctly as long as its SNR is high enough. To examine which of these is the most accurate, Figure 13 shows the established range for each experiment, as well as the range under the circular model, and the range under a very simple SNR model: packets are received if Equation (2) holds, where the various parameters were tuned to best fit the measurements.

$$\text{Signal}(\text{distance}) = \frac{\text{TX power}}{1 + \text{distance}^\beta} \quad (1)$$

$$\frac{\text{Signal}(\text{distance to sender})}{\text{Signal}(\text{distance to interferer}) + \text{noise}} > \text{required SNR} \quad (2)$$

Halfway between the sender and interferer, both signals should be about the same strength, so we expect the actual range to fall slightly short of that point. For the circular model, we selected a radio range of 9 metres, which is the sender’s range in its strongest direction under ideal circumstances (open space, no interference).

The results in Figure 13 are clear. The circular model does not work. For nodes in between the sender and interferer it is much too pessimistic. When we move the in-




Figure 13. Radio range in the presence of interference.




Figure 14. Reception within the reduced radio range; interferer located at 12 m (a) and at 8 m (b).

terferer beyond the sender, we get a situation where the nodes are all closer to the sender, but still receive interference from a node beyond the sender. In that case the circular model would predict that nodes close to the sender do not receive the packets, but nodes at a large distance from both the sender and interferer do! Clearly, this does not match reality. In Figure 13 see the radio range increase, starting with the nodes near the sender.

This is exactly what the SNR model predicts: All nodes now receive a stronger signal from the sender than from the interferer. The radio range is not fully restored immediately, because nodes at larger distance from the sender do not reach the required SNR yet. But as the interferer moves away, its noise level drops quickly, and the range is already 66% restored when the interferer is moved 1 metre beyond the sender.

6.2. Performance within the reduced radio range

Having established some limits on which nodes can still receive packets in the presence of interference, we will now look at the performance within the reduced radio range. Looking at Figure 14, we see that not only the range is reduced, but also the quality of the links that are within the reduced radio range.




Figure 15. Radio range in the corridor in the presence of interference.

This happens mostly when the interferer is at a reasonable distance (in these tests from 9 to 15 metres) from the sender. When the interferer is closer to the sender, both signals are strong, and diminish quickly when the distance is increased. Therefore it is very clear at what point the signal becomes too weak.

Surprisingly, we did not find that many more bit errors in the medium quality links. We suspect the problem to be our UART losing synchronisation, or incorrect preamble detection, but further investigation is necessary to test this hypothesis.

6.3. Environmental effects

As before, the results in the clean environment of the hockey pitch are quite different from what we see in the harsh environment of the corridor. Figure 15 shows the radio range in the corridor. The range is reduced even more by the many reflections adding extra noise. In this scenario, neither model works very well, although SNR is still better than the circular model.

It should be noted that this graph is the result of a single set of measurements, in a difficult environment. It is reasonable to assume that the same experiment in a different environment, or even at a different location in the corridor, would produce different results. However, the results do show that the commonly used interference models break down in a difficult environment like this.

7. RSSI

As mentioned before, the nodes recorded a histogram and true average of the received signal strength indicator values (RSSI). These values may be of use for localisation algorithms and for determining link quality. Figure 17(a) shows the measured RSSI values on the tennis court, an environment with few obstacles. Even in this situation the RSSI-vs-distance curve is erratic. For reference, the fig-




Figure 16. Example of RSSI calibration.

ure includes a sample signal decay based on the free space model [4]. The measured values deviate considerably from this model. One cause is that the nodes are not calibrated, and different nodes report different values at one and the same location.

We performed a simple calibration procedure based on some experiments in the corridor, in which we performed six circular shifts of all receiving nodes. This gives us RSSI values at six adjacent positions for all nodes. We chose one node as the reference node, and shifted the other nodes' RSSI graphs to minimise the mean absolute error of the node measurements with respect to the reference node and/or the other calibrated nodes (Figure 16). The calculated offsets were on average ± 40 , and ± 100 at most.

This calibration allows us to compare the results from different nodes more accurately. Figure 17(b) illustrates the effects of this calibration. The calibrated values are less erratic and resemble a sample theoretical signal strength decay much more closely. The dips in the RSSI values are most likely due to environmental circumstances since we verified that the involved nodes worked fine in other settings. In the more hostile environment of our office corridor, we find that even with calibration the RSSI values still vary wildly. This renders (calibrated) RSSI readings of little use for localisation.

To determine whether or not RSSI can be used to estimate link quality, we studied the correlation between the RSSI value and the reception rate. Figure 17(c) plots the average RSSI value against the average reception rate on the tennis court. Note that from an RSSI value of 1750 downwards the RSSI value and reception rate seem hardly correlated at all. In the office corridor, see Figure 17(e), we find similar results with a threshold of about 1850. With the calibration procedure we can improve the correlation between RSSI value and reception rate somewhat, see figures 17(d) and (f). However, the improvement is less clear compared to the RSSI/distance case.

When we look at using the RSSI values for link quality estimation, we see that there is a threshold above which reception is consistently good. Below the threshold, the RSSI value says very little about the reception rate. Unfortunately, the threshold depends on the environment, so if we want to

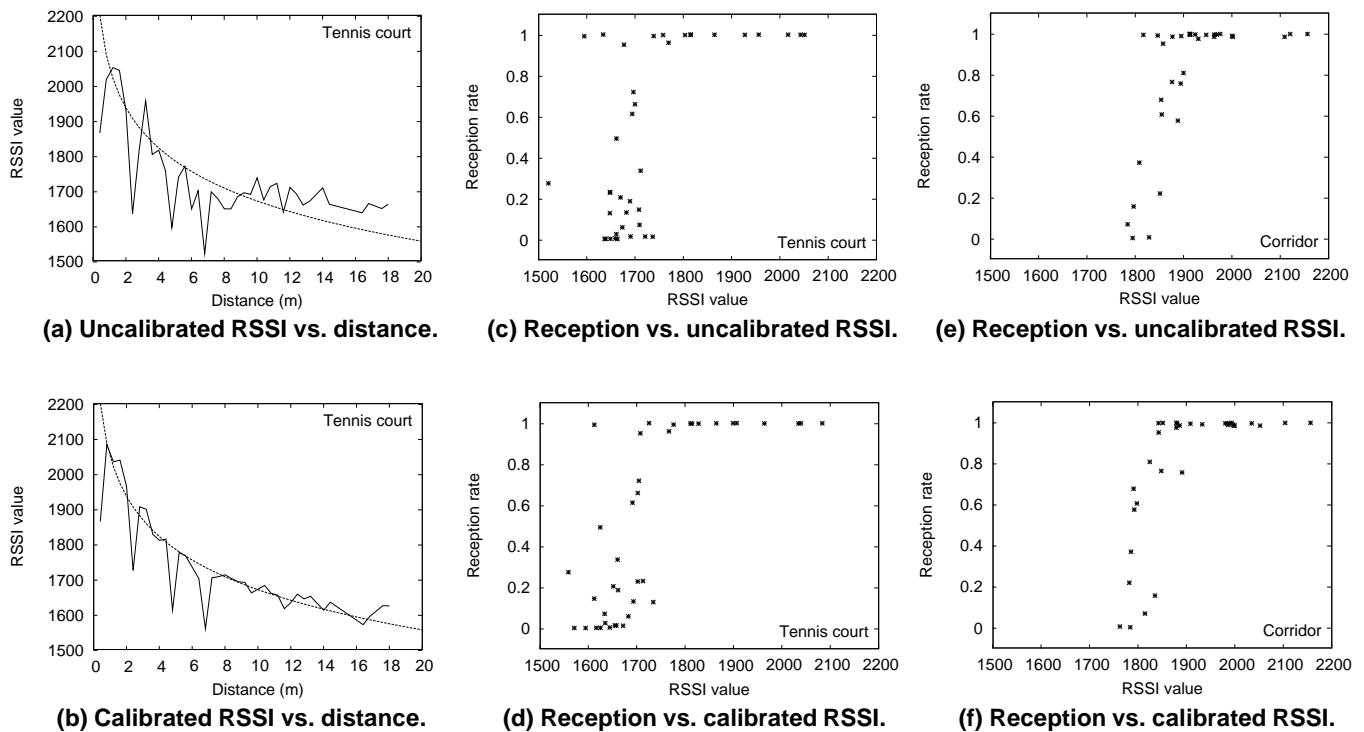


Figure 17. RSSI graphs for the tennis court and corridor.

use RSSI to classify links as 'good', it either has to be chosen conservatively, or a priori knowledge of the deployment area is necessary.

8. Discussion

We will now discuss some guidelines for protocol design and simulations that follow from our measurements.

Localisation protocols that reason about a node's location based on which nodes it can hear need to be aware of the gray area: even if we cannot communicate with a certain node, it may be much closer than another node with which we have a good link. Also, the directionality in the antennas means that a hop in one direction may be significantly shorter than a hop in another direction. This could cause problems for algorithms that use hop counts to determine distances, like DV-hop [5]. Finally, our data on calibrated RSSI values indicates that its use to determine distances is limited. It seems the only possible use may be to determine which of two neighbours is closer by, which may be used by some localisation algorithms [3].

For routing protocols, the good news is that most links are either good or bad. This means it will be possible to use the long links in the gray area, but it is important to monitor the link quality to filter out the occasional medium or bad quality link. Also, links tend to be symmetrical, which

makes setting up routes easier because we can reasonably assume we can send data back to the node we received a request from. Again this is true in most cases, but there is a small fraction of asymmetric links.

Even for the bad links, a (very small) number of messages still gets through. This means we should not consider every node we hear from a neighbour, but we need to think about when a link is reliable enough to call the other node a neighbour. Of course the problem is how to get this information without exchanging a large number of messages just to determine the link quality. The RSSI value may be of help to us here. There is a threshold beyond which reception is consistently good. By discarding links with a lower RSSI (or at least marking them as suspect), only a limited number of good links is lost, and it may be possible to recover them at a later stage. The threshold changes depending on the environment. The threshold could be determined by keeping track of the highest RSSI value we find for poor links.

The interference tests show that more messages will be received correctly than predicted by the circular model. It will be interesting to see if this can be exploited by MAC protocols to increase spatial reuse. Finally, the directionality we observed shows that a node's neighbours do not always form a nice circle. This may be important for routing algorithms or data aggregation.

9. Related work

Since practical experience with (large-scale) sensor networks is limited, we are only aware of a few other publications that report on, or include some, link layer measurements [2, 8, 9]. The work by Zhao et al. [9] is most closely related to our measurements since we basically use the same radio (ASK modulation) and experimental setup (straight line configuration). We extend their results in a number of ways, including the use of a dipole antenna instead of a whip antenna, and a study of the impact of interference by a neighbouring node. Like Zhao et al. we observe a rather large gray area, in which multi-path effects determine the large variances in packet reception rate between adjacent locations. In particular, reception is either good (30%) or bad (53%), but rarely in between (17%). This binary distribution in the gray area is also observed in grid topologies [8].

An important new result from our analysis is that the reception rates at both ends of a link are highly correlated, so in practice about 25% of the links in the gray area can be used for bidirectional communication. We demonstrated that calibrated RSSI provides a reasonable indication for good links, which obviates the need for time consuming link estimation as, for example, proposed by Woo et al. [8].

Determining the impact of asymmetric links on MAC, routing and data gathering/distribution protocols in sensor networks is an important problem that has received some attention lately [2, 8, 10]. Simulation is an attractive option, but validity is a major concern even if the models are based on actual measurements like the work by Zhou et al. [10]. Their RAM model, for example, does include sender directionality and interference, but neither accounts for the gray area effect, nor for the correlated reception rates at both link ends. Since our results show that both effects are significant at the link layer, it remains to be seen what the real consequences are for the upper layers.

10. Conclusions

In this paper we presented the results of extensive link layer measurements with prototype sensor nodes, which include a simple radio (ASK modulation) and an embedded dipole antenna, in three different environments (corridor, tennis court, hockey pitch). We confirmed the existence of grey areas caused by multi-path reflections as previously reported by others [8, 9]. An important new finding is that links within a gray area are symmetric, mostly good-good and bad-bad, so a considerable fraction of long links exist that can be exploited, for example, by routing algorithms. We also showed that many of these long links can be identified by (calibrated) RSSI readings.

Our study of the effects of interfering transmissions showed that simple circular and SNR-based models are inaccurate, in particular, in the gray area. In future work therefore, we would like to use our empirical data to develop practical models that capture the important link layer effects, and nothing more, such that we can study the behaviour of various algorithms (MAC, routing, data aggregation) in large-scale sensor networks.

The raw data for the experiments presented in this paper will be made available on our website [1].

11. Acknowledgements

We like to thank Ivaylo Haratcherev for explaining the electrical principles of radio communications, our ‘native speaker’ Tom Parker for proof reading the draft version of this paper, and the reviewers for their useful comments.

References

- [1] Consensus project website. <http://www.consensus.tudelft.nl>.
- [2] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical Report CSD-TR 02-0013, UCLA, February 2002.
- [3] T. He, C. Huang, B. Blum, J. Stankovic, and T. Abdelzaher. Range-free localization schemes in large scale sensor networks. In *MOBICOM*, 2003.
- [4] Ian Poole. *Newnes guide to radio and communications technology*.
- [5] Dragos Niculescu and Badri Nath. Ad hoc positioning system (APS). *Global Telecommunications Conference, 2001. GLOBECOM '01. IEEE*, 5:2926 – 2931, November 2001.
- [6] OPNET Technologies Inc. *OPNET Wireless Module User Guide*, 2003.
- [7] RF Monolithics (www.rfm.com). *RFM TR1001*. <http://www.rfm.com/products/data/tr1001.pdf>.
- [8] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *1st ACM Conf. on Embedded Networked Sensor Systems (SenSys 2003)*, pages 14–27, Los Angeles, CA, November 2003.
- [9] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *1st ACM Conf. on Embedded Networked Sensor Systems (SenSys 2003)*, pages 1–13, Los Angeles, CA, November 2003.
- [10] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic. Impact of Radio Asymmetry on Wireless Sensor Networks. In *MobiSys*, Boston, MA, June 2004.
- [11] IST project EYES. <http://eyes.eu.org/>.
- [12] GloMoSim. <http://pcl.cs.ucla.edu/projects/gloMosim/>.
- [13] Motes design. <http://webs.cs.berkeley.edu/tos/hardware/>.
- [14] The ns-2 Network Simulator. <http://www.isi.edu/nsnam/ns/>.