CHALMERS
UNIVERSITY OF TECHNOLOGY

Master’s Thesis report: Characterization of Radio Link Reliability

Master’s thesis done at Tele Radio AB

ANDREAS MAGNUSSON

Department of Electrical Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2017

Master's thesis EX094/2017

Characterization of Radio Link Reliability

ANDREAS MAGNUSSON

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Communication Systems
Chalmers University of Technology
Gothenburg, Sweden 2017

Characterization of Radio Link Reliability
ANDREAS MAGNUSSON

Supervisor: Jesper Ribbe, Tele Radio
Adviser: Andreas Buchberger, Electrical Engineering
Examiner: Erik Agrell, Electrical Engineering

Master's Thesis EX094/2017
Department of Electrical Engineering
Communication Systems
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Early model of a Hydra transmitter.

Typeset in LaTeX
Distributed electronically
Gothenburg, Sweden 2017

Abstract

Tele-Radio AB makes radio communication equipment for a wide range of industrial applications. The company's new high-end product line with its increased complexity has a lot of room for optimization since the effects of many parameter settings are not yet understood. In this project, the influence of different parameters is investigated. Specifically, the spatial orientation, transmit power and channel access algorithms such as the listen-before-talk (LBT) and clear channel assessment (CCA) are evaluated. For this, one transmitter/receiver link is set up in an office environment and their performance is measured in packet error rate (PER). Further, the effect of WiFi interference is investigated. In the presence of interference, a strong dependence of the PER on the orientation and the configuration of the channel access algorithms is observed offering the opportunity to increase the performance by reworking them. Without interference, the experimental setup does not achieve enough certainty to draw any conclusions about the influence of any parameter.

Acknowledgements

Thank you Jesper Ribbe and the R&D team at Tele Radio AB for the project. Thank you Erik Agrell for considering the thesis and Andreas Buchberger for great supervision.

Andreas Magnusson, Gothenburg, September 2017

Contents

List of Figures

1 Introduction

1.1 Purpose

Tele Radio AB produces remote control devices for industrial applications, such as big cranes that lift weights of several tons, mobile- and marine applications and many more. Regarding remote control of heavy machinery of this kind the system has many safety requirements due to the hazards which may exist in the operating environments. One would be an overhead crane lifting weights at around 30 metric tons. It has to be certain that it never makes any movements not ordered by the operator and if anything were to go wrong, it can stop instantly. Tele Radio's high-end products are all expected to meet SIL3 (safety integrity level 3) of the IEC (International Electrotechnical Commission) 61508 [2] standards and keeping this throughout the product line poses technical challenges.

In order to meet the required specifications of the safety certificates and stay competitive on the market it is important to explore new possibilities to improve the performance of the devices. There are many variables that affect the quality of the radio link between transmitters and receivers. It can be environmental, such as distance, orientation of the antennas (the signal is usually stronger in one direction), surrounding obstacles, temperature, humidity, etc. Some of these can be controlled to some extent but most of the time they have to be worked around. Parameters on a technical level include output power of the transmitted signal, frequency band, radio packet size and time interval between sending them and settings in software functionality such as jitter function and LBT (listen before talk algorithm).

Even though the knowledge level of radio communication systems is very substantial in these modern days, with computing power becoming cheaper and hence devices more complex, it is not always clear how a combination of variables will affect the overall performance of the system. In order to reach a sufficient level of understanding of the influence of individual parameters on the system's performance, a robust experimental setup is needed. Once this understanding is achieved, doors will open to find new ways of running the systems and gradually increase the performance without having to change the hardware or the amount of energy it uses.

1.2 Previous work

Tele Radio started off by buying existing products and adapting them for their customers. When they took the step to start developing their own systems they followed what was the industrial standard on the market at the time and kept their focus on making adaptations for each single customer's needs. In order to make a product that could be used all over the world they later took the step to have their new product line use the IEEE 802.15.4 protocol on the 2.4 GHz band (a frequency band that as opposed to the previously used 433MHz is permitted for commercial use in all countries [3]). The products were built and ran using parameter sets that are standard for the 802.15.4 protocol and for the hardware that was being used. Tele Radio has previously implemented the standard parameters from the protocols and in addition performed in house testing for specific needs doing trial and error until an acceptable system was found. They have then used the same knowledge on their new product line also here doing specific testing upon encountering an issue. This newer line is much more sophisticated than the previous set of products however and there has not yet been done any research thoroughly examining the behavior of the devices used for long periods of time. With the new high-end product line being an order of magnitude more complex than its predecessors there is a much larger interest to see what these devices can do thinking outside of the box.

1.3 Approach

In this project the new devices are being operated for long periods of time collecting data for analysis. The main point of interest is to investigate whether small changes to key radio protocol parameters have a significant impact on the performance of the system, and this way get a better understanding of the system as a whole. This way it can be possible to find better parameter sets that will increase the performance by potentially reducing overall packet loss or increasing the energy efficiency.

2

Theory

2.1 Radio protocol

In order for information to be transferred via electromagnetic waves, the signal needs to be modulated. This is the part of the communication that is covered by the physical layer of the OSI model. The Hydra devices use a low level protocol that is part of IEEE 802.15.4-2006, a standard which specifies the first and second layer of the OSI (Open Systems Interconnection) model. When speaking about radio communication or just communication very generally the OSI model is frequently mentioned. The model is a part of the OSI project at the ISO (International Organization for Standardization) under the identification ISO/IEC 7498-1 [4]. The model consists of seven layers where each represents a level of the communication hierarchy, where the two first ones are the Physical layer and the Data Link layer, see figure 2.1. Hydra however only uses definitions from the physical layer part of IEEE 802.15.4-2006.

2.1.1 Spreading

An overview of the full protocol can be seen in figure 2.2 showing the path a piece of data takes going from digital to a radio signal. The process in this protocol uses a spreading spectrum technique known as Direct-Sequence Spread Spectrum (DSSS). A piece of data is translated into a pre-determined bit sequence called a Pseudo Noise (PN) code. The reason for this is to reduce interference from other sources. The way it works is the PN pieces, called chips, are being transmitted at a very fast rate compared to the transmission rate of the radio signals. This allows for having the PN sequences quite long making it possible to be very distinguishable from each other, and also from sources of interference. IEEE 802.15.4-2006 defines 16 symbols (decimal values from 0 to 15) represented by 4 bits each. Of each data byte, the 4 least significant bits map to one symbol and the 4 most significant bits map to the next symbol. Each symbol is in turn mapped into a PN sequence of 32 chips ($c_0, ..., c_{31}$) that are almost orthogonal, hence making it very unlikely for a sequence ever to be mistaken for another. For details about the particular values in the pre-determined sequences for IEEE 802.15.4, refer to the standard document.

2.1.2 Modulation

Once the data is translated in to chip sequences by the spreading, it is to be transmitted by the radio (the right box of figure 2.2). The modulation scheme used in

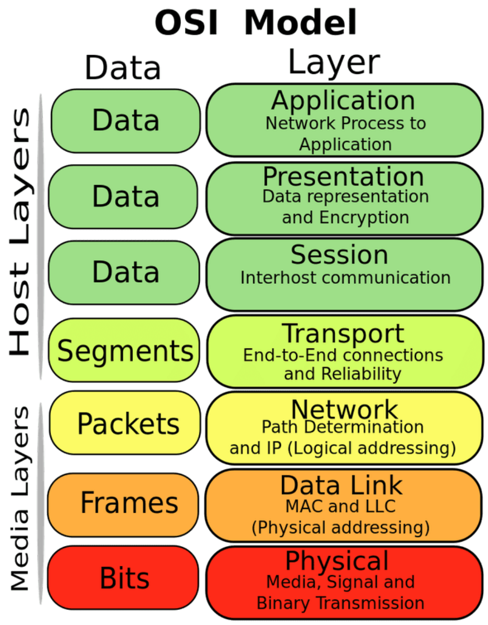

Figure 2.1: The Open Systems Interconnection (OSI) model. The devices used in this project use a part of IEEE 802.15.4-2006 which implements the Physical and the Data Link layers. Gorivero / Wikimedia Commons / CC-BY-SA-3.0 `https://creativecommons.org/licenses/by-sa/3.0/deed.en`

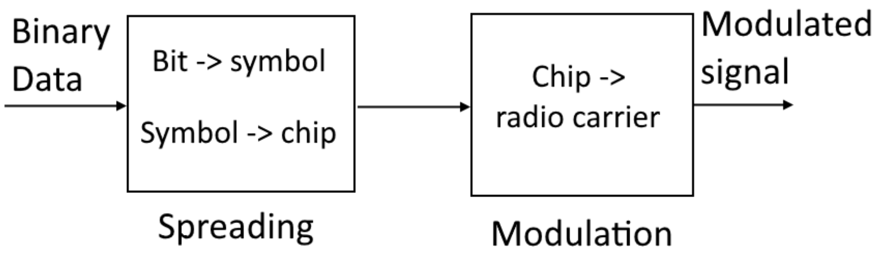

Figure 2.2: Diagram of the radio protocol.

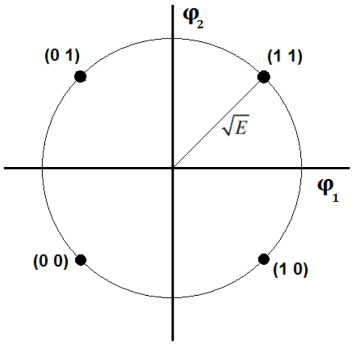

Bit pair	Phase	φ1	φ2
(1 0)	π / 4	$+\sqrt{E/2}$	$-\sqrt{E/2}$
(0 0)	3π / 4	$-\sqrt{E/2}$	$-\sqrt{E/2}$
(0 1)	5π / 4	$-\sqrt{E/2}$	$+\sqrt{E/2}$
(1 1)	7π / 4	$+\sqrt{E/2}$	$+\sqrt{E/2}$

Figure 2.3: The signal constellation of QPSK in ϕ_1/ϕ_2 (I/Q) space. Table summarizing the bit-phase translation of QPSK (right).

IEEE 802.15.4-2006 is a variant of Quadrature Phase Shift Key (QPSK). In regular QPSK the signal is modulated according to

$$s_i(t) = \sqrt{\frac{2E}{T}} \cos\left[2\pi f_c t + (2i-1)\frac{\pi}{4}\right]; \quad 0 \leq t \leq T; i = 1, 2, 3, 4 \tag{2.1}$$

Where E is the symbol energy, $T = 2T_b$ is the symbol duration where T_b is the bit duration and f_c is the carrier frequency. Note that a symbol in QPSK is not to be confused with a symbol in DSSS. As seen here, a symbol in QPSK is 2 bits long and hence covers 2 chips in a DSSS sequence. Therefore a time frame of $T = 2T_b$ always processes 2 chips so the relation between the chip duration in DSSS and bit duration in QPSK and simply becomes $T_c = T_b$. Through trigonometric expansion of equation (2.1) and by using the basis functions

$$\phi_1(t) = \sqrt{\frac{2}{T}} \cos 2\pi f_c t \quad \phi_2(t) = \sqrt{\frac{2}{T}} \sin 2\pi f_c t; \quad 0 \leq t \leq T \tag{2.2}$$

the modulation can be expressed as

$$s_i(t) \equiv I\phi_1(t) - Q\phi_2(t) \tag{2.3}$$

with

$$\begin{aligned} I &= \sqrt{E} \cos\left[(2i-1)\frac{\pi}{4}\right] \\ Q &= \sqrt{E} \sin\left[(2i-1)\frac{\pi}{4}\right] \end{aligned} \tag{2.4}$$

where I and Q are called the in-phase and quadrature axes. Figure 2.3 shows the resulting signal constellation in ϕ_1/ϕ_2 (I/Q) representation and a table with the bit symbols corresponding to phase. I/Q is used for practical reasons as for how the model is implemented in the hardware. A binary bit stream is split in two by a serial to parallel converter (demultiplexer) sending the odd bits to the I and the even to the Q. The two are then independently modulated according to the expressions in equation (2.2) and combined to a total QPSK signal. The receiver contains two

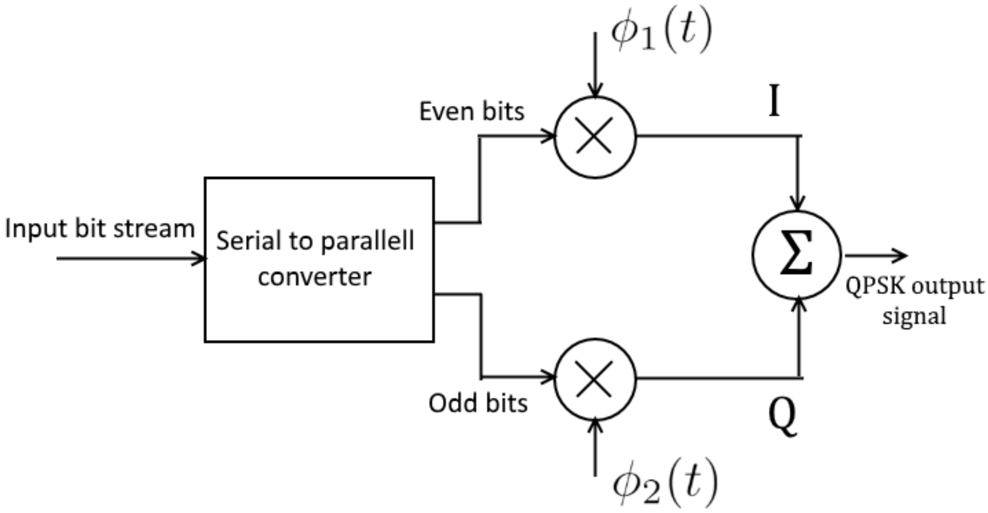

Figure 2.4: Schematic diagram of a QPSK modulator.

filters matching equation (2.2) followed by an integrator for each path. These are being sampled for one symbol period $T = 2T_b$ at a time after which the sampled signals are sent through decision makers who make a decision for each bit based on a threshold value. Finally the bits from each path are multiplexed in the original order. Modulation and demodulation diagrams can be seen in figures 2.4 and 2.5 respectively. The CC2520 radio chip used in the Hydra products implements a slightly modified version of the modulation called Offset QPSK (O-QPSK). It works exactly the same way as QPSK only the difference being the Q-phase is delayed by one bit duration T_b. In QPSK the phase can jump up to π at a time that can cause large amplitude fluctuations which are not desirable in a communication system. Adding the offset makes it so that the signal can only jump $\pi/2$ at a time, significantly decreasing the unwanted fluctuations.

A timing timing diagram is presented in figure 2.6 showing the offset delay of the Q-phase. An IEEE 802.15.4 symbol is translated to its chip sequence and the even chip numbers are sent to the I-phase while the odd go to the Q-phase.

2.1.3 Frame Format

Tele-Radio AB only implements the beginning of the frame format of 802.105.4-2006 and due to company secrecy the whole implementation for the Hydra devices cannot be included in this report. This does however not matter since the only part of interest in this project is the initiating part of the frame. The communication occurs in bursts with one burst being called a radio packet. The time between the packets is a predefined transmission rate which is tunable and can change for various reasons (more about this in section 2.1.4). The first piece of a radio packet frame is the synchronization header (SHR) whose purpose is to synchronize the transmission by indicating the start of a message. It consists of a preamble followed by the start of frame delimiter (SFD). The preamble is a range of 0x00 values (with 4 bytes being

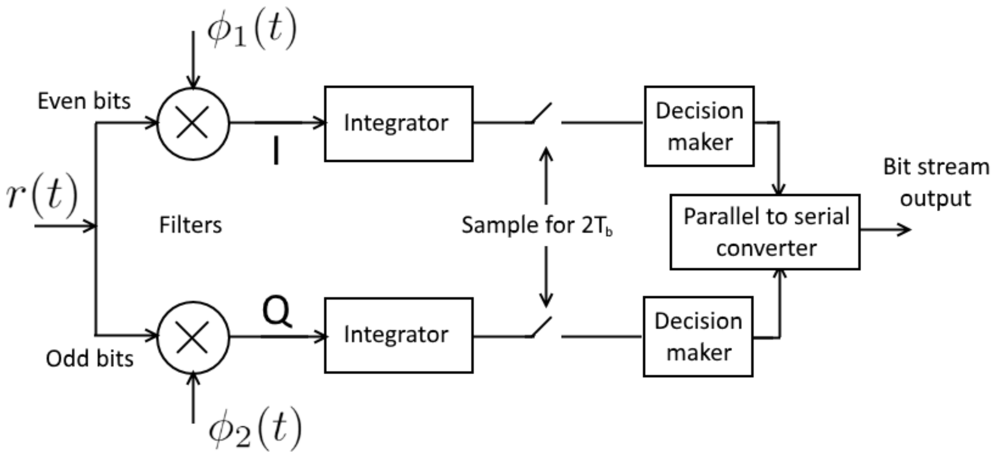

Figure 2.5: Schematic diagram of a QPSK demodulator.

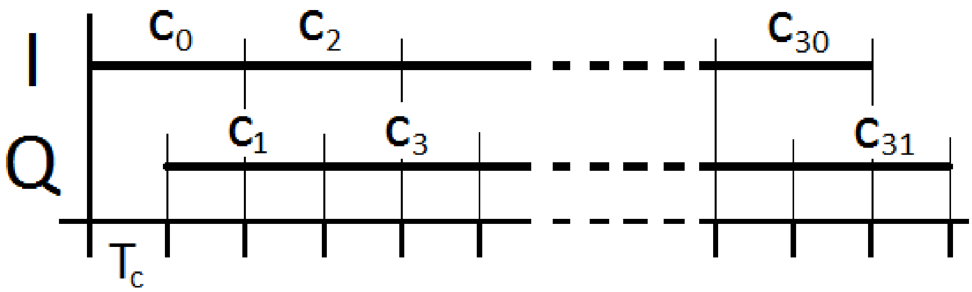

Figure 2.6: Timing diagram showing the offset of I-phase and Q-phase.

Preamble				SFD	Payload
0x00	0x00	0x00	0x00	0xA7	Tele Radio AB implementation

Figure 2.7: The frame format of a radio packet. Tele-Radio AB's implementation part cannot be included.

the default length) and the SFD is one byte with the hard coded value of 0xA7 (see figure 2.7). The length of the preamble is tunable and can be set by the developer implementing the standard.

2.1.4 Listen-before-talk (LBT)

If the transmission rate is constant it becomes very problematic having multiple devices operating on 2.4 GHz in the same area since there is a risk that the different signals interfere. The set transmission rate representing the time interval of when the communication bursts occur is a set value but what it actually represents is the average of time intervals between the bursts. The real time between every packet burst is set to differ and this can be solved in several different ways. The most common is a jitter function that randomizes the time interval for every packet, keeping it between a reasonable minimum and maximum value and the mean at the set transmission rate. This works very well for low range applications. The Hydra devices wish to achieve higher range however and therefore must transmit with a much higher output power. To do this there are rules established by ETSI [5] which require the devices to listen to other signals in their vicinity before transmitting not to cause interference for any other communication.

The procedure of checking whether the channel is free to use or not is called Clear Channel Assessment (CCA) which is a part of IEEE 802.15.4-2006. It measures the power, more precisely the RSSI (Received Signal Strength Indication, described later in the report) and determines based on a threshold whether it is clear to transmit or not. Additionally it will also only assess a clear channel if it is not already busy receiving a frame. It is possible to change the mode of the CCA behavior having it be clear all the time, determine based on the RSSI measurement only, only while receiving a frame or both the latter. The default mode and also the only mode used in this project is both. Only doing one CCA check for determining if a transmission should be performed or not is not enough. If for example a very intense other source would be close by, the radio link would simply die due to the CCA never indicating a clear channel. This is solved by performing several CCA checks with random time intervals between and if the channel every time is determined to be busy, the packet will be sent anyway but at a reduced output power. How this is done is described in the complete LBT algorithm below

1. Perform CCA checks with random times between each other until either

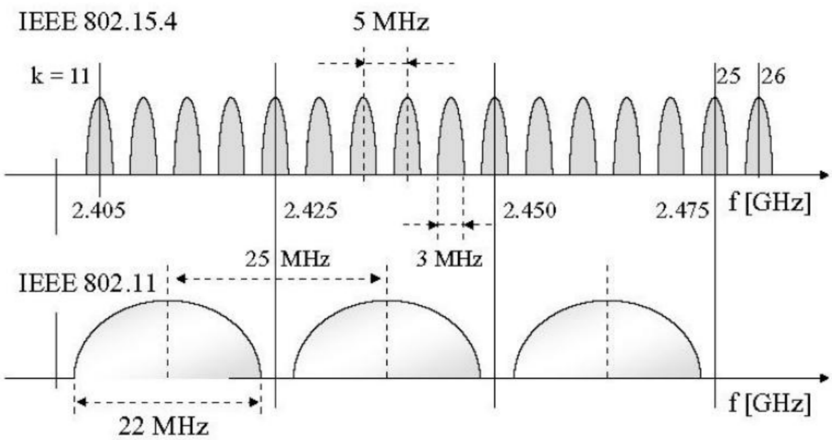

Figure 2.8: Comparison of channels between 802.15.4 and 802.11. Reprinted from Performance comparison of IEEE 802.11g and IEEE 802.11n in the presence of interference from 802.15.4 networks [11].

- the timer reaches a set max time or
- it determines a clear channel

2. Transmit the packet

   - with lowered output power if the CCA checks reached max timer
   - with full power if the CCA checks determined a clear channel

2.2 802.15.4 coexistence with 802.11 (WiFi)

The most common and most probable sources of interference for the Hydra devices is (besides other ones of its kind) are WiFi-networks, which is to be expected since they operate on the same frequency band of 2.4 GHz. The protocols work quite differently and giving a deep explanation of a typical WiFi protocol is out of the scope of this project. This section will address the most significant factors. The data transmission rate of a modern simple WiFi protocol is typically 54 Mbps as opposed to the 802.15.4-2006 which peaks at 250 kbps. Also WiFi is supposed to be able to operate at a range of 250 m and the Hydra specification say 150 m, this being in an outdoor like environment. This means that WiFi is allowed to operate at a higher output power and but above all uses much wider frequency bands.

In figure 2.8 is a layout of all the channels of 802.15.4 and channels 1, 6 and 11 of 802.11. The bandwidth of 22 MHz of a WiFi channel is substantially bigger than the 5 MHz of the 802.15.4 and only these three WiFi channels overlap most of the 802.15.4. The only channels of 802.11 shown here is 1, 6 and 11 because they are the least likely to overlap with 802.15.4 and will be the only ones considered in this project. Note that most 802.11 use extended protocol which means that the entire bandwidth of one extra channel is used. For example if channel 1 is used in extended mode, it uses a bandwidth spanning over channel 1 and channel 6. This leaves little room for 802.15.4 but if one is aware and can control which channels to use, inference can be avoided by for example using WiFi channel 1 and 802.15.4 channel 15. This is however rarely the case and there has been research to understand the problem of having these two protocols running in the same environment. In 2009, researchers at the University of Wisconsin ran hardware experiments on a network of 15 802.15.4 nodes having a 802.11n system in the same environment with controllable load ranging from 1 to 60 Mbps [6]. The first remark they made was that since 802.11 is operated at a much higher power any overlap will choke out the signals from 802.15.4 thus simply by this fact will make it difficult to run 802.15.4 close to WiFi. They concluded that overlap with the 802.11 extension channel can cause packet latency for 802.15.4 and it worsens increasing the traffic load of 802.11n, although the exact impact depends on the usage of the extension channel. Furthermore, overlapping with the control channel (that is the set operating channel) causes severe deterioration in both the loss rate and the packet latency for IEEE 802.15.4 traffic, much more serious than overlap with the extension channel. This also worsens when increasing the traffic load on 802.11n.

These conclusions fall back directly on the expected performance of the Hydra devices. If there is a channel overlap with WiFi, the transmitter will likely measure a high RSSI value and the CCA checks in the LBT will determine a busy channel causing a transmission with lowered output power after the timer exceeded the maximum timer. The signal with lower power will get choked by the WiFi signal with much higher power and the packet will be lost.

2.3 Significant parameters

This section explains the different parameters which affect the performance of the radio link and will be the basis of the experiments.

2.3.1 Output power

The biggest contributor to what determines the range of radio communication is the output power. Quite trivially the higher power the further the radio waves will reach but at the cost of energy and so battery life, a very critical component in applications of this kind. There are also legal restrictions and restrictions by the standard on average and maximum output power [5].

Regarding the relation between output power and range it is not linear so it is hard to think about it intuitively. It is estimated by [8]

$$\text{Distance (km)} = 10^{(\text{MPL}-32.44-20\log_{10}(f))/20} \tag{2.5}$$

where f is the frequency in MHz and MPL is the maximum path loss in dBm given by

$$\text{MPL} = \text{transmit power} - \text{receiver sensitivity} - \text{losses} - \text{fade margin} \tag{2.6}$$

where potential gains are included in the transmit power directly. (Both equations (2.5) and (2.6) come from a conversion of the Friis transmission equation [9]). Losses include unavoidable means of distortion to the radio signal, such as attenuation in the air. Fade margin is a term usually added by the manufacturer to while estimating the range of the device account for environmental conditions, such as obstacles and noise [8]. This means if the fade margin is set to zero the range value will only be valid in ideal line of sight conditions which is not practical.

2.3.2 Preamble length

In the IEEE 802.15.4-2006 standard a default value of 4 bytes is defined for the preamble however most manufacturers of radio chips implementing this standard recommend developers to perform tests of their own. The optimal value may differ depending on how the system is to run. A shorter preamble will make the reading and therefore also the synchronization faster and also leave room for more data to be sent with the packet. However it will be more likely for noise or other sources not part of the radio link to be processed. This will result in a packet loss since the correct packet may be missed while falsely processing an incorrect one. A longer preamble will make synchronizing more accurate and will make it more likely that a correct packet is accepted, but it also leaves less room for data. Additionally it will make the synchronization process slower which could also cause a packet loss, should it miss a correct packet arriving almost simultaneously.

2.3.3 LBT parameters

There are several different parameters that can be tuned in the LBT algorithm. Those included in this project are described below.

2.3.3.1 CCA threshold and hysteresis

To avoid the determination switching back and forth between clear and busy when the RSSI is close to the threshold value, a hysteresis is used. If the hysteresis parameter is set to be other than 0, the CCA check will give a clear channel if $\text{RSSI} < \text{Threshold} - \text{Hysteresis}$ or determine a busy channel if $\text{RSSI} \geq \text{Threshold}$. Should the RSSI reading fall in within the range $\text{Threshold} - \text{Hysteresis} < \text{RSSI} < \text{Threshold}$, the CCA will give the same result as the previous check. This means that if the hysteresis were to be set too high, the CCA would frequently determine a busy channel even though it would be OK to transmit.

2.3.3.2 Maximum time to perform CCA

Increasing the maximum time indicating how many times the CCA should be performed before the transmitter forces a transmission with lowered output power can have a big impact if the system is under interference from another source of the same frequency. If the other source outputs a stronger signal (like WiFi) and if the maximum time is set too short, it is very likely that many transmissions will be sent with a lower output power and therefore be choked by the other signal, or simply not reach the receiver due to the resulting shorter reach. If the timer was to be increased the transmission would be more likely to occur with full power but it may also take longer before the packet is transmitted, causing delays in the connection.

2.4 Environmental factors

For experimental purposes the environment of operation needs to be as free of variation as possible since upon changing a parameter and analyzing, it must be trusted that the potential effect was due to the parameter change and not to another variable. This section covers some information about variables that need to be taken into consideration when dealing with radio communication.

2.4.1 Air

This section should more generally be called propagation medium but the devices used in this project are always operated in air. Radio waves like any type of electromagnetic waves are affected by attenuation due to absorption or scattering of photons. These are photon molecule interactions and so the magnitude of the attenuation is determined by the amount of interactions which in turn depends on how many particles there are to interact with, meaning the thickness of the air. High pressure vs. low pressure will be of no significance since the difference in pressure and number of particles is very small. The main factor will be the humidity since the amount of water in the air can vary a lot and might have a big enough effect for experiments being run outdoors. In a normal office environment these effects can be neglected.

2.4.2 Surroundings and obstacles

Any kind of objects in the vicinity of the radio waves may have an effect on the performance, especially if located between the transmitter and receiver. This is also dependent what sort of material the objects are made of, their shape, etc. Since the system in this project uses 2.4 GHz the signal is also more prone to reflect off objects than to penetrate them as opposed to use a lower frequency band, e.g. 900 MHz. Any problems with obstacles can be avoided by simply having the air between the devices clean and off limits for people. However having the experiment setup indoors, the problem of the radio waves bouncing on the walls cannot be avoided. There are special rooms to fix this particular problem when performing measurements of this kind but the funds required for such an installation is completely out of the budget for this project. The only way to deal with this problem is to accept its presence, keep the environment as sterile as possible and collect enough data to be able to reduce the effects caused by it.

2.4.3 Orientation

The orientation of the devices can in many cases have a severe impact on the quality of the radio link. The most common reason for this is the location of the antennas on the device. To find the ideal position is a challenge because the device should be as compact as possible while the antennas should be as far away from metal as possible. Most of the material on the devices is made of plastic which does not cause any issue since it is almost entirely transparent to radio waves but all the electrical components as well as the batteries consist of mainly metals. As opposed to plastic most metals act as mirrors for the radio and the signal will bounce off the surface and thereby act as dampening in certain directions. In addition to that, on the receiver specifically, the output voltage and current can be up at 30 V and 3 A which can cause severe electromagnetic noise.

On the receiver the antennas are placed in a way that apart from the surrounding plastic cover have no material interference and more or less the same reception in all directions. The transmitter has many more components however, including buttons, levers, joysticks, more electrical components and batteries. The antennas are positioned on the front, meaning the side always pointing away from the operator and usually in the direction of the contraption being remote controlled. Due to this position, many components are located behind the antennas and the backwards transmitted signal gets a lot of material distortion. In total, this means that the orientation of the devices should have a large effect on the performance but as stated in section 2.4.2 the 2.4 GHz signal is very prone to bounce off the walls with the reflectance strong enough to keep the radio link up. This makes it very hard to theoretically predict how the signal direction will affect the quality.

2.5 Measurements

Since radio communication is transferring information in the form of packets the main thing to look at when determining the performance is usually the amount of received packets compared to the amount that was transmitted. There are also means to keep track of the strength of the signal to get an indication of the quality in real time.

2.5.1 Packet error rate

The main quantity for a quality measurement is the total packet loss or packet error rate (PER), which is simply given by

$$\text{PER (\%)} = (1 - \text{Received packets/Transmitted packets}) \cdot 100 \tag{2.7}$$

2.5.2 Packet error burst histogram

The total PER being a straight forward way of looking at the quality of the performance, does not show if a single packet was lost at one time or if several packets were lost in a row. Say there was an experiment with 90 received packets out of 100 expected (a PER of 10%), if all the packets were lost at different times the performance would still be good. If a packet is lost but the next is not the radio link model can compensate and the radio link can stay intact with the user potentially not even noticing any problems. Should all the 10 packets have been lost one after another however the radio link would most likely disconnect. To capture and quantify this dilemma the experiment setup captures in addition to received and expected packets a histogram with the bins representing packets lost in a row. As explained the lower packet error burst the better so measurement of performance is the shape of the histogram; the more right-skewed the better.

2.5.3 Orientation 2D histogram

Having the devices always pointing in the same direction throughout the project to avoid the variation of orientation is very challenging since just a small accidental push or misplacement can have an effect. To rule out this problem and also to investigate whether the performance is affected by the orientation or not, the experiment setup has the receiver and transmitter both rotating and their angles relative to a reference direction stored for every measurement point. This is presented as a two-dimensional histogram with the receiver and transmitter angles on each axis and a color scale representing the packet error rate for each transmitter/receiver angle pair.

2.5.4 Received Signal Strength Indication (RSSI)

The RSSI is a commonly used measurement of power in a received radio signal. It is not shown to the user but is very important for the underlying computations and decisions made by the algorithms in the software. The CC2520 radio chip provides an RSSI value which is the average received power over 8 symbol periods ($16\mu s \cdot 8 = 128\mu s$) given in decibel-milliwatts (dBm). With this value the actual power of the received signal will be

$$\mathrm{P} = \mathrm{RSSI} - \mathrm{OFFSET}\ [\mathrm{dBm}] \tag{2.8}$$

where the offset is an empirical value found during the characterization and approximately 76 dBm [10]. Although the RSSI is not used as a direct measurement of the radio performance in this project, it is a good indication of the quality of an experiment or to tell if the radio link at some point died during a run. Even if there is a weak or distorted signal, the RSSI value will still be measured. This means that one can tell the difference between the radio link being disconnected (if no RSSI is measured) or just having many packets lost in a row. That being said, the receiver is also configured to not accept too many lost packets in a row and simply disconnect the radio link if this occurs.

3
Methods

3.1 Experiment setup

The setup consists of one radio link using one transmitter and one receiver, both using firmware up to date as of January 2017. Although the firmware is currently under development and improvements are made on a weekly basis, the same is kept throughout the project. When running the WiFi interference experiments a regular for-private-use router is connected through a wireless network to a laptop, both standing right next to the radio transmitter and receiver.

3.1.1 Radio link setup

The devices are set up in an office environment in a room occasionally used for tests of different kinds. The receiver is placed on a platform which is mounted onto a DC motor that has a rotating speed of 1 rpm. The motor is operated by a micro controller which in turn has two switches connected to it that are located next to the rotating platform. If one switch is pressed the platform will stay still for 0.5 s (to avoid peak currents in the circuit) and then change the rotation to the opposite direction; vice versa for the other switch. Attached to the platform is an arm that will hit each of the switches on each direction of rotation as can be seen in figure 3.1. Having the receiver rotate in both directions is due to the cables that have to be attached to it for power supply and data collection. The transmitter is standing on an almost equal contraption but runs on batteries. Also since all radio data is collected from the receiver, the is no cable problem. This platform rotates clockwise only and hits a switch once every lap (see figure 3.2) which is also connected to the same micro controller as the receiver switches. The program doing the data acquisition saves a time stamp and in the case of the receiver switches also a stamp for the current rotation direction. This can then be used to compute the angles for the devices at any given data point.

The data collection is done by custom made Python-scripts using an in-house-made API that provides USB-communication to the products and the possibility to decide which data to read. It is optimized to be able to run one iteration of collection faster than the lowest packet transmission rate that is used for the radio communication, meaning faster than 20 ms. The speed at which it operates is always set to be the set transmission rate so that during a run it approximately collects as many data points as there should have been packets received. Acquiring faster than this will not be necessary since it would only save the same information multiple times, whilst acquiring any slower would cause loss of information. The data entries that

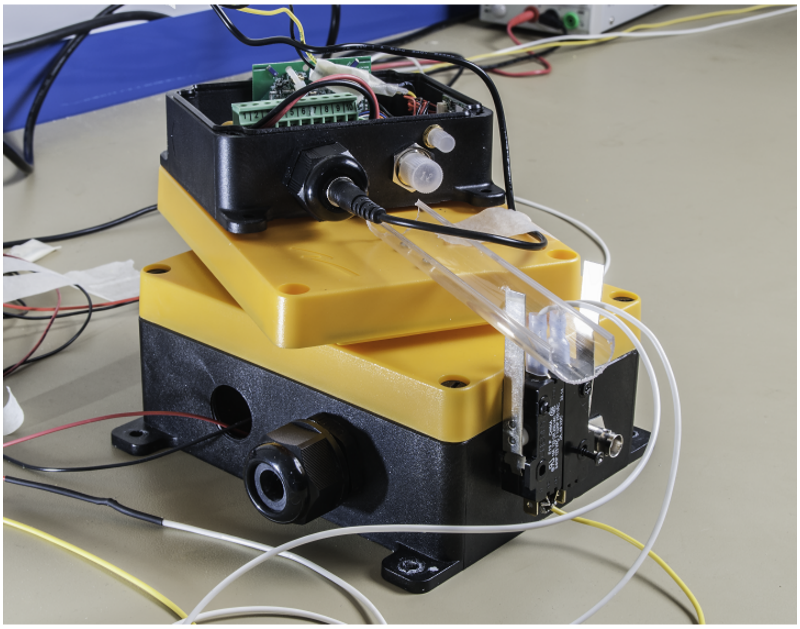

Figure 3.1: Hydra CAN receiver mounted on a rotating platform. An arm presses a switch on each side upon reaching it triggering a controller to swap the rotation direction and save a time stamp.

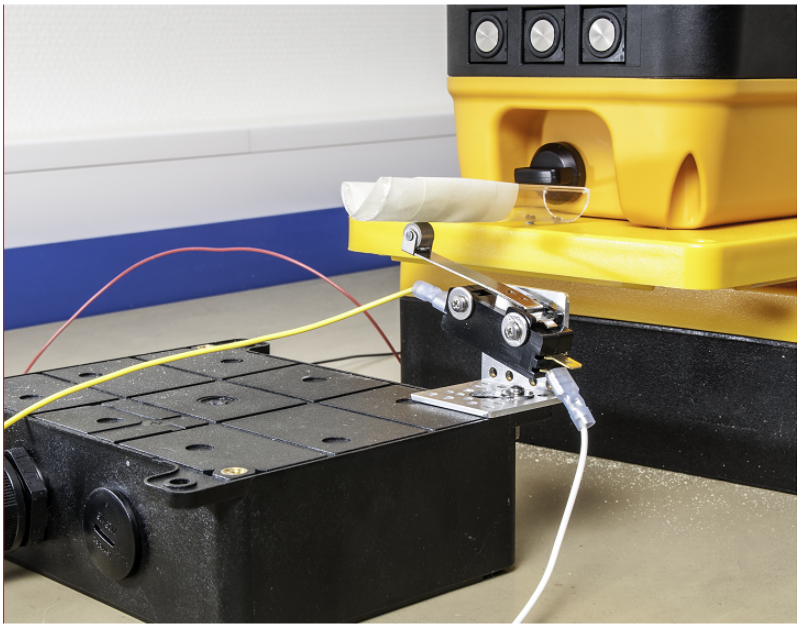

Figure 3.2: The Hydra transmitter (top right corner) mounted on a rotating platform. The arm presses a switch every lap triggering the controller to save a time stamp.

are captured in each iteration are listed below.

Ticks to last transmitter packet (TLP)
This presents the time since the last received packet with 0.1 millisecond precision. It is used to compute packet loss in a row (PLR) by counting how many multiples of the set packet transmission rate the value fits. Hence for a given transmission rate t_r the PLR becomes

$$\text{PLR} = \lfloor \text{TLP}/t_r \rfloor - 1 \tag{3.1}$$

where the subtraction with 1 is to account for the expected value with no packet loss is t_r.

Total packets
The total amount of packets expected to be received since the start of the radio session. If the acquisition script is working accordingly, this value should increase by 1 every iteration (since the script iterates with the same rate as the set radio transmission rate).

Received packets
The number of packets that have been received.

Missed histogram 1, 2, 3, 4, 5
The embedded software can count occurrences from 1 up to 5 missing packets in a row independently from the TLP. For most experiments the histogram will have a majority of its counts in this region which makes the data processing much quicker since these do not have to be computed of the TLP. It can nevertheless be beneficial to compute the lower values from TLP anyway and use these histogram counters to verify the data.

2D histogram
The histogram is made splitting the 360x360 degrees in to 5x5 degrees bins. For each bin, the number of packet losses occurring within that bin is divided by the total number of events (loss and successful transmissions) for that bin. This way the bins can be compared to one another better due to the receiver and transmitter platforms rotating at the same speed, making certain angle combinations occur more often than others.

RSSI
The received signal strength indication or input power measured by the receiver.

During a data acquisition a regular radio session is running meaning no buttons are pushed and the transmitter is simply sending one type of packet with the set transmission rate and packet size. Due to many things going on in the building during day time (WiFi usage, cell phones, people walking around, etc.) the time for running a reference experiment is overnight. However doing experiments during the day will still be of great value since the devices are usually operated in an environment with noise of this kind. Running the experiments with the controlled WiFi setup is fine either day or night since it will choke out most of the other noises.

3.1.2 WiFi setup

The router is running IEEE 802.11n with an extended protocol. It has the possibility to store data connecting a USB flash memory which can be reached by connecting to the router using FTP. The computer is running a script which is downloading a file stored on the router over and over again throughout the length of an experiment. The work load on the router is kept constant at the maximum of what it can handle, which peaks at around 65 Mbps. During an experiment, the total run time is being logged as well as the number of downloads that have completed. This is not used as experimental data but is of value if the downloading were to be interrupted (computer crashing, script lagging out or other). This way it is possible to determine the exact time when WiFi was disconnected and to take this into account when analysing the data. A schematic view of the entire setup can be seen in figure 3.3.

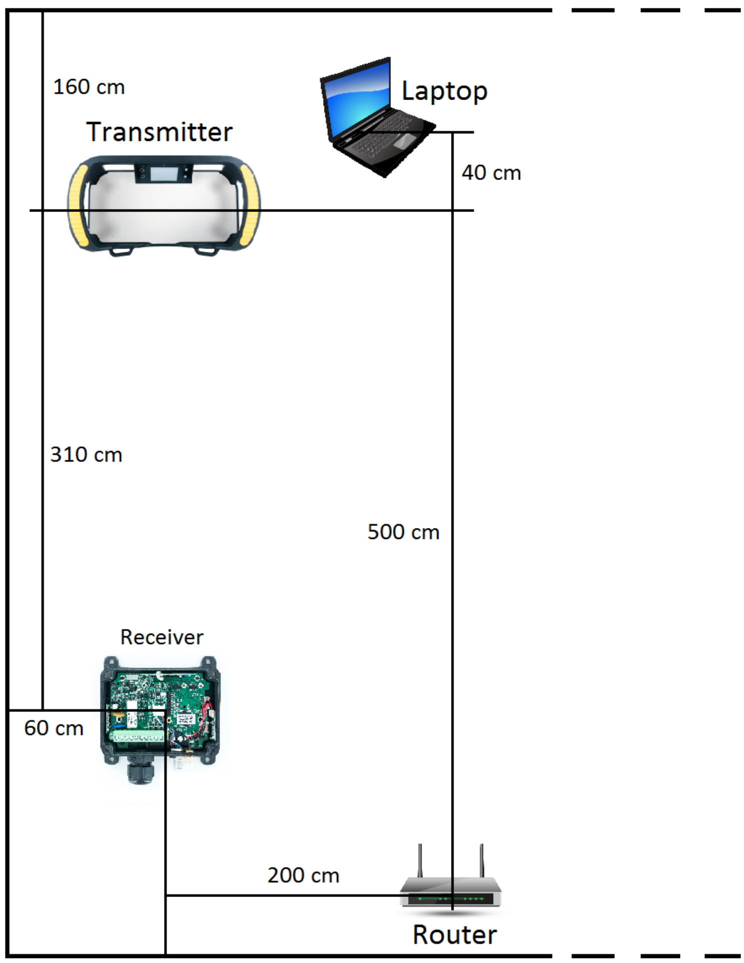

Figure 3.3: Schematic view of the full experiment setup. The transmitter and receiver are mounted on rotating platforms running the idle radio protocol. The laptop is connected to the router with a 802.11n WiFi.

4

Results

The results will be presented in two parts: without noise and with an 802.11n source producing interference. The default parameters referred to are presented in table 4.1.

4.1 WiFi disabled

4.1.1 Influence of the LBT

The devices were set to the default parameters apart from the output power being lowered to -10 dBm. The reason for this was to decrease the performance of the system intentionally to pick up potential variety more easily. During daytime, eleven runs each lasting between 8 and 12 hours with enabled and disabled LBT were carried out. Three runs of each set are presented in figure 4.1 (LBT enabled) and figure 4.2 (LBT disabled). As it can be seen, there is no difference in performance. Going from minimum to maximum total PER the histograms show that higher PER means more lost packets in a row aside from it being single packet loss events. To further investigate if disabling the LBT had any effect on the overall PER, a two sided t-test was perform on all of the 11+11 runs (details on how such a statistical test is performed can be found in appendix A.1). The test came out $P < 0.8$ meaning the probability is less than 80 % that any change in PER was due to disabling the LBT. The total PER values for all tests are plotted in figure 4.3.

4.1.2 Orientation

The 2D histograms as described in section 2.5.3 of the 11+11 runs (the same runs as in LBT vs. no LBT) are presented in figure 4.4 and 4.5. The color scale represents PER in % for each bin. There is an X-like pattern which can clearly be seen in the

Output power	18.23 dBm
Transmission rate	40 ms
CCA max time	20 ms
LBT	Enabled
Channel	12

Table 4.1: Default settings used for the experiments.

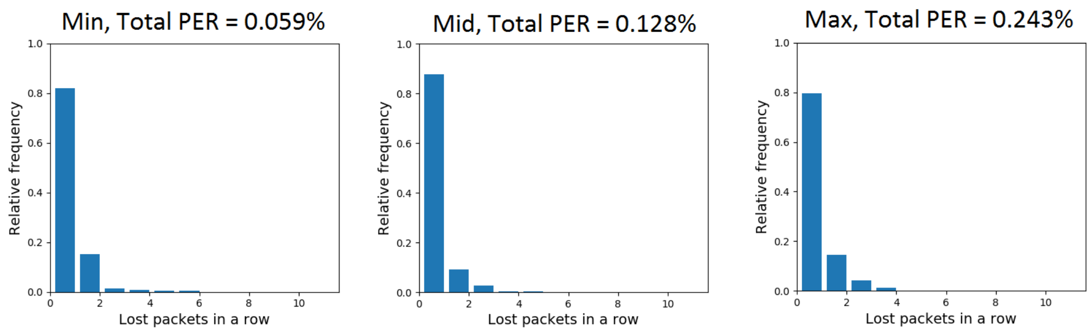

Figure 4.1: Default setting runs with minimum total PER (left), closest to average total PER (middle) and maximum total PER (right).

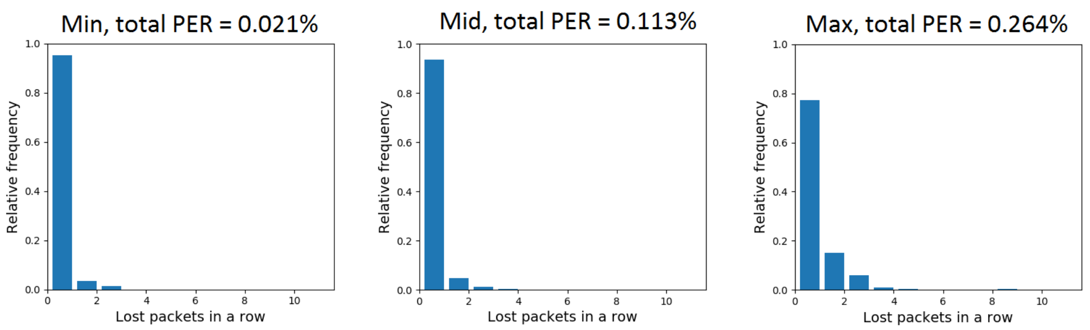

Figure 4.2: Runs with disabled LBT; minimum total PER (left), closest to average total PER (middle) and maximum total PER (right).

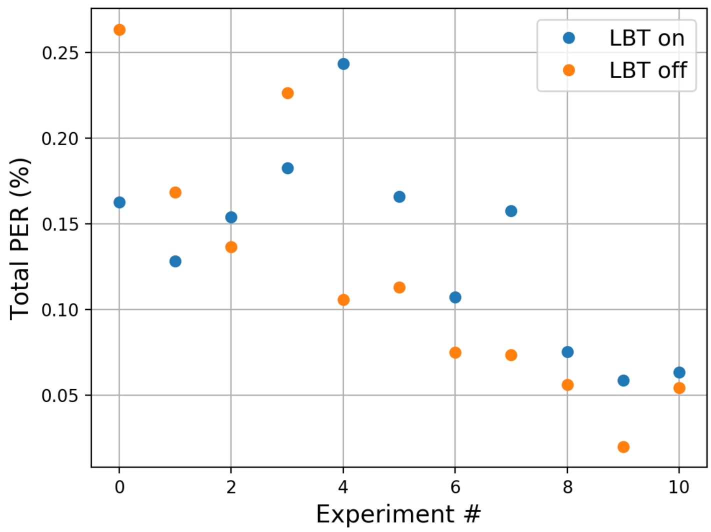

Figure 4.3: Total PER of default settings vs. disabled LBT. Experiments 9, 1 and 4 with LBT are shown in figure 4.1 and experiments 9, 5 and 0 without LBT in figure 4.2.

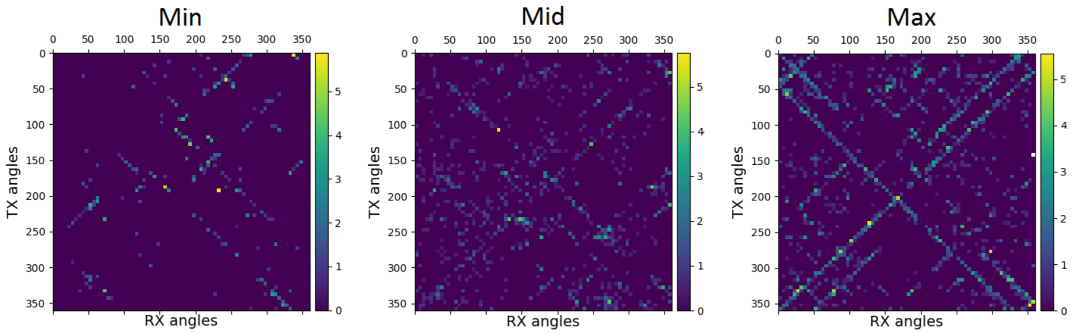

Figure 4.4: PER (color) vs. orientation for the experiments presented in figure 4.1; LBT enabled. The color translates to the PER in % for the angular coordinate.

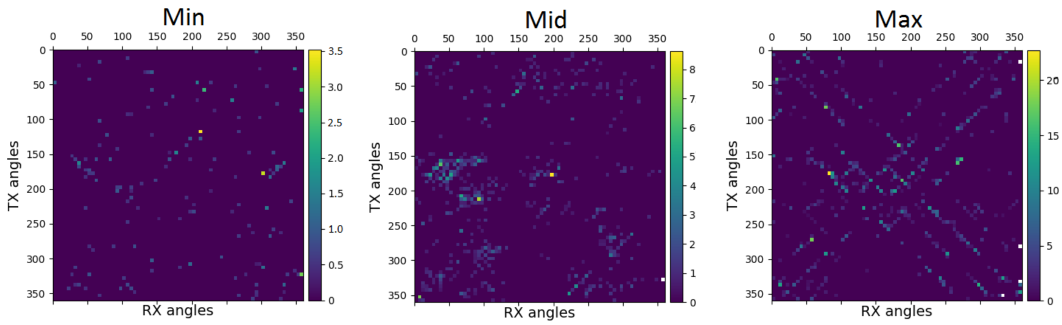

Figure 4.5: PER (color) vs. orientation for the experiments presented in figure 4.2; LBT disabled. The color translates to the PER in % for the angular coordinate. Note the different values on the color scale between these plots. These are kept to maintain proper visualization.

max in figure 4.4 but also appears slightly in the other plots. More packet loss occur along these lines and indicates a correlation between packet loss and orientation.

4.2 Under 802.11n interference

Much larger effects could be seen having the noise present. For all the experiments in this section the output power of the radio was set to its maximum possible value of 18.23 dBm and the lowered power which is used if the LBT times out was set to 10 dBm. The WiFi always used channel 1 and the output power of 23 dBm.

4.2.1 Influence of the LBT

Three samples from runs with LBT enabled collected the same way as the previous section and are presented in figure 4.6. Even using the maximum output power on the transmitter, the PER went up substantially adding the WiFi noise. The histograms change the same way here as without the noise showing more packet loss in a row for higher PER.

Disabling the LBT significantly increased the performance of the system in terms

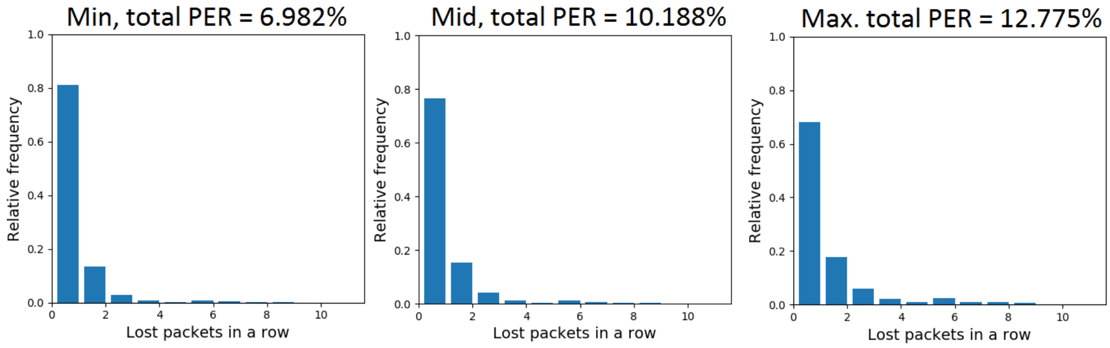

Figure 4.6: Runs with enabled LBT under the influence of WiFi noise.

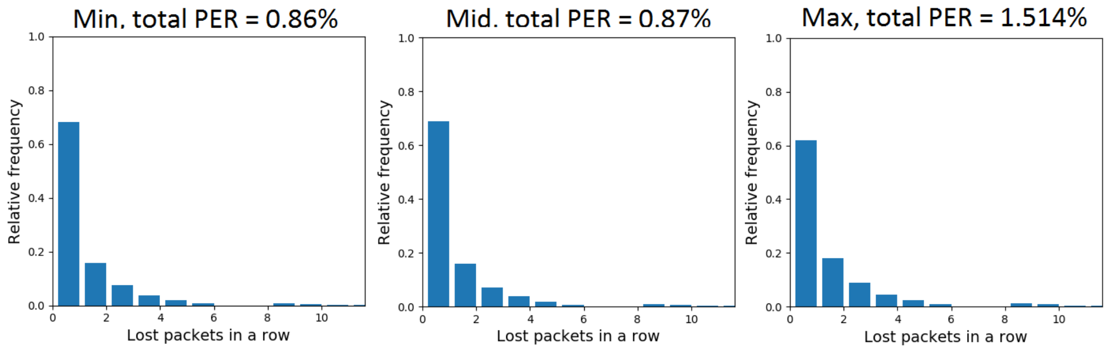

Figure 4.7: Runs with disabled LBT under the influence of WiFi noise.

of total PER, as can be seen in figure 4.7. The histograms changed shape slightly having more packet loss in a row relative to the total PER.

4.2.2 Increasing CCA maximum time

To evaluate if the CCA maximum time was set in a way that made it very hard for the radio to ever get a free channel from the WiFi due to mismatched timing with the WiFi protocol, two tests were run setting the CCA max time to 60 ms and two tests setting it to 100 ms, see the results in figure 4.8 and 4.9 respectively. An increased performance can be observed in both cases, especially for the longer CCA maximum time.

4.2.3 Free channel test

A simple test was run to confirm that a channel lying between the WiFi control and extension channels would be free of interference. Running the system on channel 15, which as explained in section 2.2, should be free of interference gave a total PER of 0.004%. Running an additional test on channel 14 which lies on the edge of WiFi channel 1 resulted in a total PER of 7.255%, hence these tests confirm the theory.

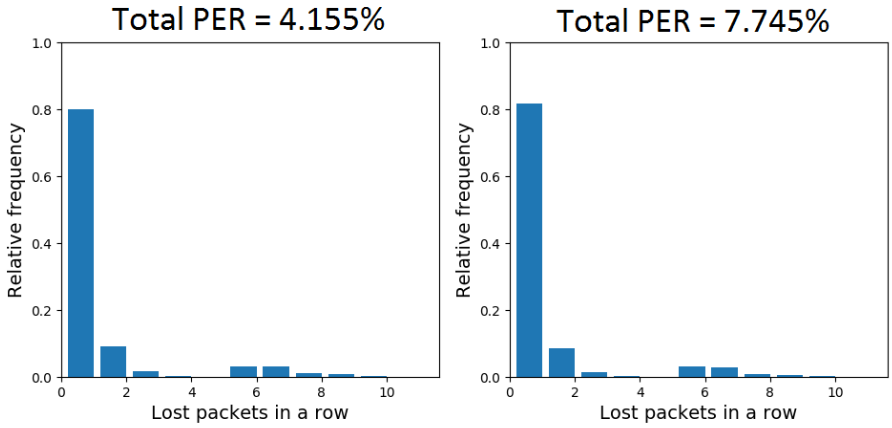

Figure 4.8: Runs with WiFi noise having the CCA max time set to 60 ms.

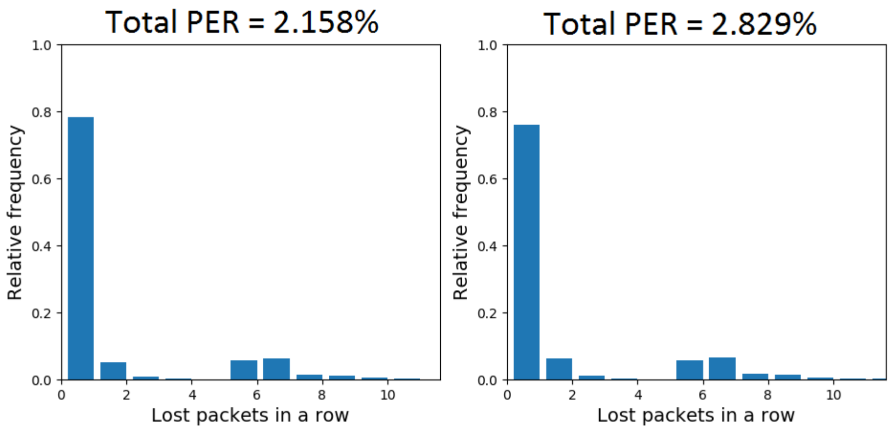

Figure 4.9: Runs with WiFi noise having the CCA max time set to 100 ms.

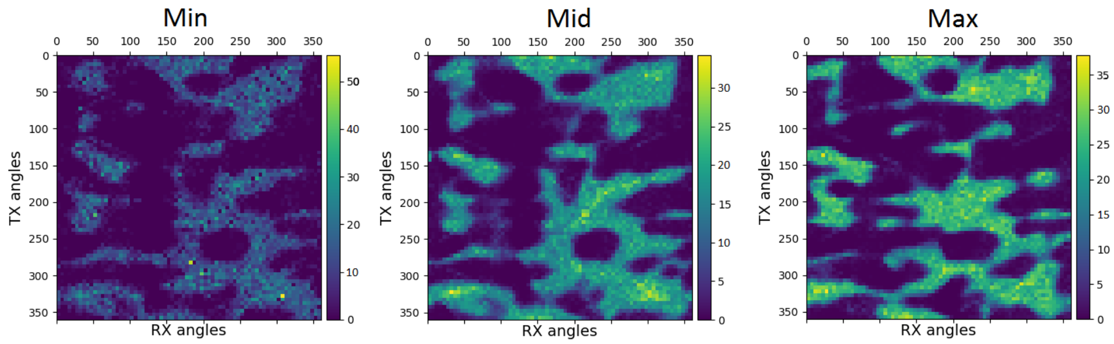

Figure 4.10: PER (color) vs. orientation for the three runs presented in figure 4.6. The color translates to the PER in % for the angular coordinate.

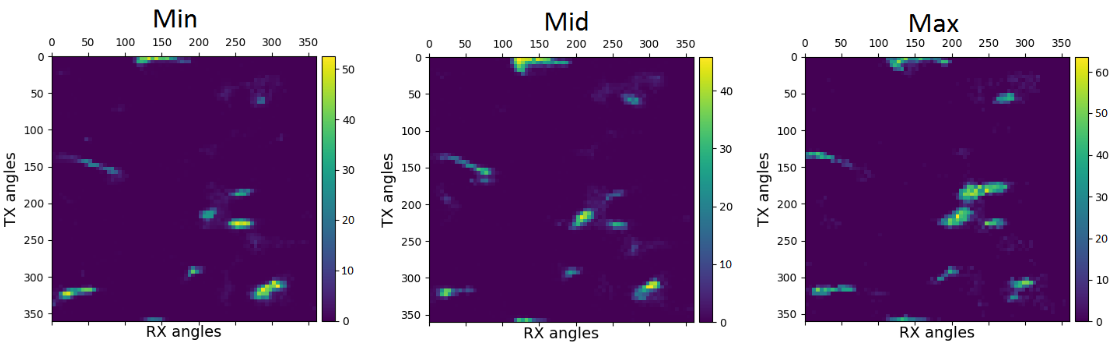

Figure 4.11: PER (color) vs. orientation for the three runs presented in figure 4.7. The color translates to the PER in % for the angular coordinate.

4.2.4 Orientation

Adding the WiFi noise made the system have a substantial dependency between direction and when packet loss occur. In figure 4.10 the total orientation 2D histograms from the same experiments presented in section 4.2.1 having the LBT enabled are drawn. First note that the X-like patterns from the results without the WiFi are gone. Although the PER differs and the 2D histograms are not perfectly alike, they are very similar.

In figure 4.11 the next three experiments where the LBT was disabled are shown. Here there are even more distinguishable areas in the TX/RX angle space where packet loss is more prone to occur. Note that in all the areas where there is packet loss for the disabled LBT case, there is also packet loss for the enabled case. For visualization of what the angles mean in terms of the setup orientation, see figure 4.12.

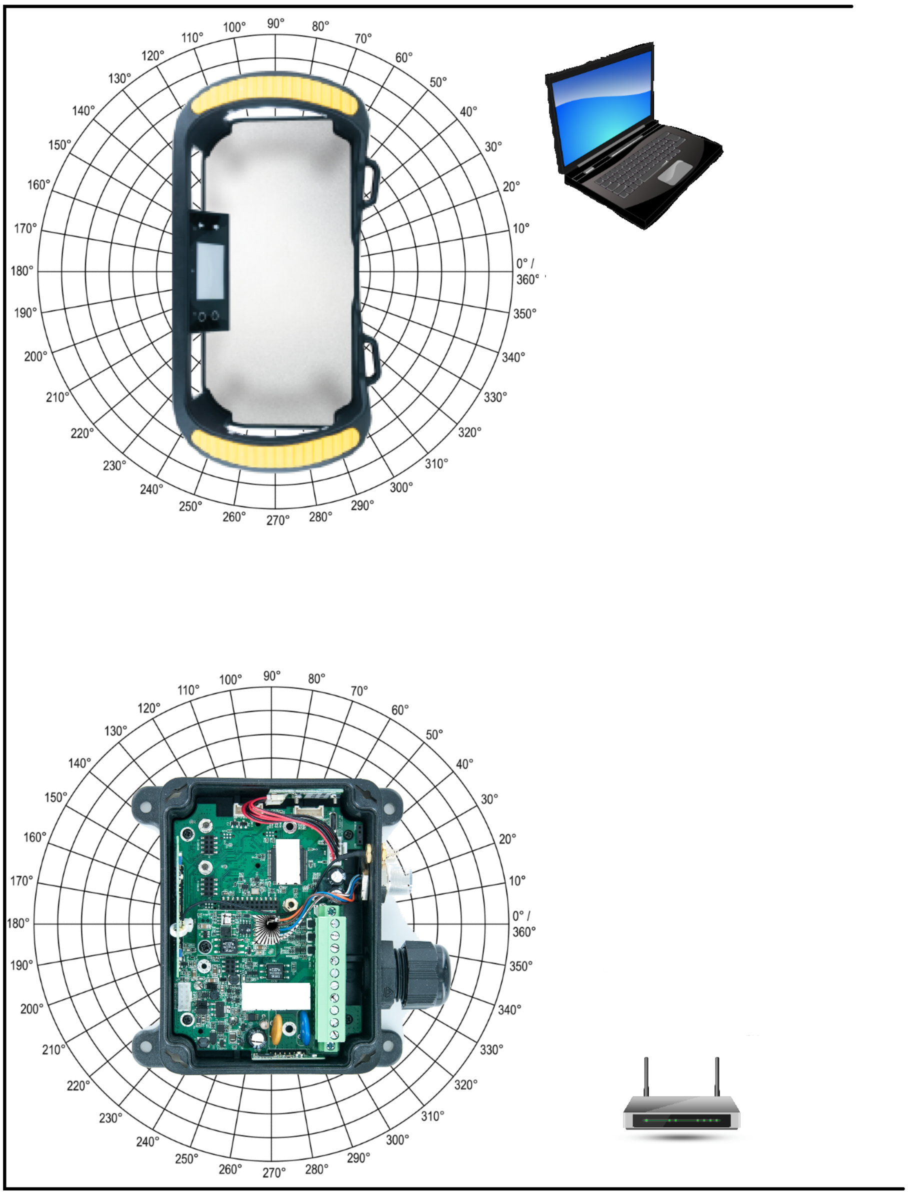

Figure 4.12: Modified view of the setup to show how the angles align in the room and for the devices. TX angles are for top and RX angles for bottom.

5
Discussion

Performing radio communication tests in an environment of this kind is usually not reliable. Most companies that do testing on their devices to see effects of small changes in the settings have special rooms built for this purpose. The walls in these rooms consist of material where a large portion of the incoming radio waves are absorbed and not reflected. Additionally the outside walls are thick making it unlikely that any signal from unknown external sources can enter the room causing noise. This changes the whole way one can reliably draw conclusions from the results of the experiments. Experimental environments of this kind however cost millions of Swedish crowns to build and was completely out of budget for this project. The main argument to still make an attempt to perform experiments in this kind of environment was that the devices are generally used in these conditions so for commercial purposes it will still be of value. The idea was to compromise for the uncertainties by always doing big changes in the parameter that was tested and run the tests for a very long time. For the tests made in the noiseless environment however no tweaking ever made the performance norm break out of its own variance and nothing could be concluded. Even as seen in the results the quite radical change of disabling the listen-before-talk algorithm showed very little change in performance; the total PER being the same only the shape of the histogram changed slightly for the lower PER examples. Other parameters than those presented in the results, more specifically preamble length and CCA threshold, were also tested for some experiments but chosen to be left out since they did not show any sign of performance change at all.

Regarding the observations in the orientation plots in the noise free environment, the X-shape which can clearly be seen in the case of enabled LBT is much more faint when disabling the LBT. Also the packet loss occurrences are spread in the case of enabled LBT and much more concentrated to a less number of bins which can be seen due to the color scale reaching much higher values for disabled LBT. A first thought to why this pattern appears is that since the motors have the same speed, every time they rotate in the same direction the same relative angle between the devices is kept and is represented by a line in the plots. However since the PER in each bin is computed independently from all the others to compensate for this problem, the effect has to be real. The most probable reason for this is the surrounding material inside the transmitter and the receiver. Since for 2.4GHz the wavelength is $\lambda = c/(2.4 \cdot 10^9 s^{-1}) = 0.125$ m, there could be that there is a piece of metal located a quarter wave length from the antenna making the reflection cause destructive interference. This will cause a particular direction from the device to have a much weaker signal and if this exists in both devices, the packet loss will be more severe when these directions align. This explanation makes sense fore the positive slope of the X, but is not sufficient for the negative slope. In this case the transmitter and receiver rotate with opposite directions hence their relative angle changes at a constant pace. What could occur here is that from each of the radio chips point of view, the chips are perfectly mirrored to one another throughout the span of that rotation. This means that certain spots on the antennas on the receiver will always be hit by signal coming from the same spot on the transmitter antennas. If an alignment of this kind exist, where having these transmitter/receiver spot pairs where the signal is clearly non-optimal, this could be an explanation for why the negative slope appears.

Upon adding WiFi to the environment one could suddenly observe big effects from the settings modifications. The output power had to be set to the maximum for the radio link to stay active throughout a run showing the significance of the output power. The results from disabling the LBT further proves this since the most probable reason for the severe PER having the LBT enabled, is that the CCA checks very often determine a busy channel and the transmission is performed with lower power and is choked by the WiFi. Increasing the CCA max time seems to make the system able to find a gap in the WiFi transmission time spectrum where the channel is left clear. This could either be that simply having more CCA checks would give system more chances to "squeeze in" a transmission when the WiFi (which is most likely running a similar random time based algorithm for transmission) happen to have a longer downtime. It could also be that WiFi has a longer downtime in its protocol to allow for other transmissions and this is captured by the increased CCA max time. Note also the appearance difference of the histograms between different experiment sets. Here single packet loss is more common which would indicate a better performance on its own but take note, in the case of 100 ms especially, that also 5 or more packets in a row are more common. This could be a direct effect of having the CCA max time this high, since the receiver is still expecting packets with the default transmission rate of 40 ms. The delay caused by performing CCA checks for this long will be counted as a lost packet.

The reasons for the sudden dependence on orientation adding the WiFi noise could be many. The phenomena causing the X-like pattern could still be there but completely overshadowed by much stronger effects. It is very likely that the certain relative angles between the receiver and the transmitter represented by the areas in the angular space were more prone to packet loss even without the noise source but due to the lack of events, it was not possible to see the effects. It is also important to remember that the output power was increased and this could lead to radio waves interactions with the surrounding material changing. It surely made more signal bounce off the walls which can cause both positive and negative interference. The actual presence of the WiFi could make certain directions much more noisy than others simply due to the location of the laptop and the router since these signals will also bounce (even more because of the major output power) and have areas with constructive and destructive interference. Regarding the relative angle between the devices, the fact that orientation matters packet loss wise is a very big finding. The antennas are located away from the other electronics as much as possible and mostly have plastic surrounding them. The hope is that both transmission and reception should be as good as possible in every direction.

The work done with the WiFi noise is yet somewhat insufficient to draw any definite conclusion regarding a way to optimize the system. The reason is that the WiFi experiments were not considered from the start but began as a last resort with one month to go. This is unfortunate since only running them for this short period of time has shown very critical aspects of the behavior of the devices. This has laid the foundation for mainly two parts that should be followed up on:

The LBT algorithm

Since WiFi is and will most likely continue to be a main source of interference for radio systems of this kind, reworking the LBT algorithm to specifically deal with WiFi can be of great value. First the transmitter (and/or the receiver) would have to be able to determine whether there is WiFi in its vicinity. This could be done by RSSI measurement; the values are often much higher than that from radio signals and they persist longer in time (concluded from the increase in CCA check time). The WiFi channels also span wide frequency bands meaning measurements could be done on a span of frequencies covering these bands and upon reaching certain RSSI values on certain channels for a certain time, it can accurately be determined if there is WiFi. The LBT can then adapt and for example have a longer CCA check time or the protocol could run on a higher transmission rate, as long as it is not long enough to cause noticeable delay for the user. It may also be possible to have the algorithm work around the WiFi protocol but this will be very difficult to achieve. There are very many WiFi protocols operating and they work very differently from the protocol in the Hydra radio devices. Very often the rate of which they transmit for instance is based on the current demand of the network, e.g. how many peers that are connected and how much data that is to be transferred. The Hydra devices operate on quite small CPUs so there is limited computational power to be used and algorithms with too much complexity cannot be run properly.

Specific orientation experiments

To begin with one has to determine which of the transmitter and receiver has the biggest direction dependence. Under the assumption that it is limited to the same experimental setup as in this project, the receiver should be standing still having only the transmitter rotating. This should be done both without and with the WiFi noise since the effect may emerge even without removing the uncertainties of having the receiver rotating. The receiver then has to be strapped down very securely or the problem of its position moving slightly between experiments will be present which was the idea to be rid of having the rotations in the first place. The transmitter and receiver will then swap places, and additionally several different positions for both should be used and same goes for the laptop and the router running the WiFi. Doing a set of experiments of this kind could give angle spans where the transmitter and/or receiver operates more poorly which can be very useful for future designs of the hardware.

6

Conclusion

Running experiments on Tele Radio AB's high-end devices on low output power in an office environment disabling the listen-before-talk algorithm has no effect on the total packet error rate, although they showed indication of slightly better performance in terms of packet loss in a row occurrences. The orientation of the devices did not have any correlation of when packets are lost. Adding WiFi noise to the experiment setup the output power had to be set on maximum for the radio link to stay active but even so the packet error rate was severely increased by the noise. Increasing the LBT CCA check timer resulted in a much lower packet error rate and disabling LBT completely even more so. The influence of WiFi noise also made occurrences of packet loss very dependent on the orientation of the devices. This demonstrates that the overall performance can be increased by implementing a more complex LBT algorithm. Regarding the orientation, more specific tests need to be performed to reach an understanding of why this behavior emerged upon adding WiFi noise.

Bibliography

[1] Frisk, D. (2016) A Chalmers University of Technology Master's thesis template for LaTeX. Unpublished.

[2] International Electrotechnical Commission, (2017). Functional Safety and IEC 61508. [online] Available at: `http://www.iec.ch/functionalsafety/` [Accessed 20 03. 2017]

[3] Eroglu, K. (1998). The worldwide approval status for 900 MHz and 2.4 GHz spread spectrum radio products. *1998 IEEE EMC Symposium. International Symposium on Electromagnetic Compatibility. Symposium Record (Cat. No.98CH36253)*, Denver, CO, pp. 1131-1135 vol.2.

[4] International Organization for Standardization/International Electrotechnical Commission, 1994, ISO/IEC 7498-1:1994(E) *Open Systems Interconnection Model*, Genève, ISO/IEC Copyright Office

[5] European Telecommunications Standards Institute, *Wideband transmission systems; Data transmission equipment operating in the 2,4 GHz band and using wide modulation techniques*, ETSI EN 300 328 V2.1.1 (2016-11). Available from: ETSI. [Accessed 01 09. 2017].

[6] Polepalli, B., Xie, W., Thangaraja, D., Goyal, M. and Hosseini, H. (2009) 'Impact of IEEE 802.11n Operation on IEEE 802.15.4 Operation', *2009 International Conference on Advanced Information Networking and Applications Workshops* DOI 10.1109/WAINA.2009.102

[7] *Module 5: Carrier Modulation, Lesson 25: Quaternary Phase Shift Keying (QPSK) Modulation* Version 2 ECE IIT, Kharagpur

[8] Downey, C (2013). Understanding Wireless Range Calculations. [online] Electronic Design. Available at: `http://electronicdesign.com/communications/understanding-wireless-range-calculations` [Accessed 13 April 2017].

[9] H. T. Friis (1946). *A Note on a Simple Transmission Formula, in Proceedings of the IRE*, vol. 34, no. 5, pp. 254-256, doi: 10.1109/JRPROC.1946.234568 Available at: `http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1697062&isnumber=35759`

[10] Texas Instruments, "CC2520 datasheet 2.4GHz IEEE 802.15.4/ZIGBEE RF Tranceiver", Dec. 2007.

[11] Haani Masood, S.(2013) *Performance comparison of IEEE 802.11g and IEEE 802.11n in the presence of interference from 802.15.4 networks*, Department of Electrical Engineering, McGill University

[12] Peters, C. A., (2001) 'Statistics for Analysis of Experimental Data', *Environmental Engineering Processes Laboratory Manual*

A

Appendix

A.1 Two sided t-test

The procedure made for statistical comparison between two measurements in this project were performed according to [12]. We have the two experimental conditions A and B, let x be the observation variable and η_A and η_B denote the true means of x. The assumption is that the errors in the measurements are normally distributed and the null hypothesis that the difference between η_A and η_B is zero. Since $\bar{x}$ is a random variable, the difference between the two observed sample averages $\bar{x}_A - \bar{x}_B$ is also a random variable and the difference in error has a t-distribution,

$$(\bar{x}_A - \bar{x}_B) - (\eta_A - \eta_B) \sim t_\nu s_{(\bar{x}_A - \bar{x}_B)}, \tag{A.1}$$

where t_ν is the scaling factor with ν being the degrees of freedom. Since the difference $\eta_A - \eta_B = 0$ according to the null hypothesis, equation (A.1) can be written as

$$(\bar{x}_A - \bar{x}_B) \sim t_\nu s_{(\bar{x}_A - \bar{x}_B)}. \tag{A.2}$$

The variance of a difference between two variances is the sum, hence,

$$\sigma^2_{(\bar{x}_A - \bar{x}_B)} = \sigma^2_{\bar{x}_A} + \sigma^2_{\bar{x}_B}. \tag{A.3}$$

Having the sample variances s_A^2 and s_B^2 and the sample sizes n_A and n_B, the standard error in the difference is

$$s_{(\bar{x}_A - \bar{x}_B)} = \sqrt{\frac{s_A^2}{n_A} + \frac{s_B^2}{n_B}}. \tag{A.4}$$

Assuming that the error magnitude in the measurements are approximately equal, the degrees of freedom is the sum the degrees of freedom in each of $\bar{x}_A$ and $\bar{x}_B$, meaning $\nu = n_A + n_B - 2$. An estimated weighted average of the variance in the two samples known as the common variance is

$$s^2 = \frac{(n_A - 1)s_A^2 + (n_B - 1)s_B^2}{n_A + n_B - 2}, \tag{A.5}$$

which leads to the standard error in difference can be written as

$$s_{(\bar{x}_A - \bar{x}_B)} = s\sqrt{\frac{1}{n_A} + \frac{1}{n_B}}. \tag{A.6}$$

The two sided t-test is carried out with the probability

$$\mathrm{P}\left\{ -\left|\frac{\bar{x}_A - \bar{x}_B}{s_{(\bar{x}_A - \bar{x}_B)}}\right| \leq t_\nu \leq \left|\frac{\bar{x}_A - \bar{x}_B}{s_{(\bar{x}_A - \bar{x}_B)}}\right| \right\} = 1 - \alpha \tag{A.7}$$

where α is looked up in a two sided t-test table. For example, the test presented in section 4.1.1 in this report gave $\left|\frac{\bar{x}_A - \bar{x}_B}{s_{(\bar{x}_A - \bar{x}_B)}}\right| = 0.6588$ which is out of range for most existing t-tables, resulting in $\mathrm{P} < 0.8$. This means the probability of rejecting the null hypothesis is less than 80%. In other words it concludes there is no statistical significance of the difference in mean value being due to the difference of experimental settings.