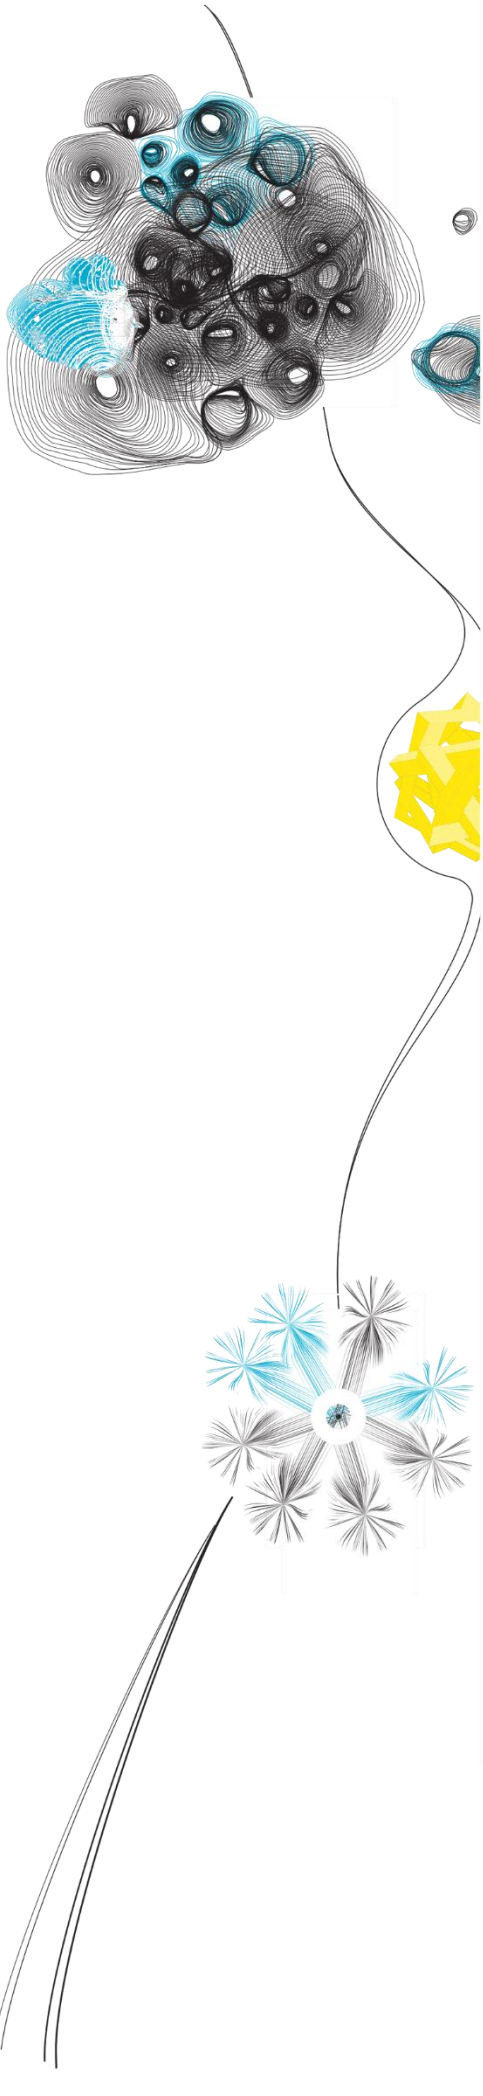

MASTER THESIS

EXPERIMENTAL EVALUATION OF LORA(WAN) IN INDOOR AND OUTDOOR ENVIRONMENTS

Chiel Hakkenberg

FACULTY OF ELECTRICAL ENGINEERING,
MATHEMATICS AND COMPUTER SCIENCE

PERVASIVE SYSTEMS RESEARCH GROUP

SUPERVISORS
dr. ir. N. Meratnia
dr. ir. K. Das

EXAMINATION COMMITTEE
dr. ir. N. Meratnia
dr .ir. K. Das
prof. dr. ing. P.J.M Havinga
ir. J. Scholten
dr. ir. M.J. Bentum
dr. ir. A.B.J. Kokkeler

AUGUST 2016

UNIVERSITY OF TWENTE.

UNIVERSITY OF TWENTE

Experimental Evaluation of LoRa(WAN) in Indoor and Outdoor Environments

By

Chiel Hakkenberg

Submitted to the Faculty of Electrical Engineering, Mathematics
and Computer Science on 17-08-2016,
in partial fulfilment of the requirements for the degree of
Master of Science

Acknowledgements

The completion of this thesis would not have been possible without the support of several people. First and foremost, I would like to gratefully thank my supervisor, Nirvana Meratnia, for her guidance and kind assistance throughout my research, as well as for providing useful suggestions for improving this thesis. I also would like to thank my second supervisor, Kallol Das for our insightful discussions and the friendly aid provided during my thesis. My gratitude is also extended to the other members of my (ever-growing) examination committee: Paul Havinga, Hans Scholten, Mark Bentum and André Kokkeler, for taking on the task to read and assess my work.

I would also like to thank Jan-Pieter for his help, time, and fruitful discussions on various parts of my thesis. Moreover, I really enjoyed the off-topic talks about almost anything and the times we drank a few beers. Most often, in these settings, we were accompanied by Jacob, whom I also would like to thank for the many hours of useful and not so useful discussions alongside the healthy lunch-walks.

Finally, my most warm thanks go to my girlfriend Ineke for her unconditional support and sometimes helping hand during these eventful times.

Abstract

Over the past years, Low Power Wide Area Networks (LPWANs) have become one of the fastest growing wireless technologies. These networks are ascribed a big role in the future of Internet of Things (IoT) applications and have gained a lot of interest, even though much practical details about them are still unknown. One such an LPWAN is LoRa, a proprietary spread spectrum modulation scheme. The spreading of the spectrum in LoRa modulation is achieved by generating signals in which the frequency linearly increases or decreases over time with a speed indicated as the spreading factor (SF).

This thesis provides a theoretical and experimental evaluation of LoRa for both indoor and outdoor environments in order to determine LoRa's suitability to provide a reliable solution for network connectivity. LoRa is proven an ideal technology for low energy consuming data transmissions over a long-range with low data rate requirements.

Extensive experiments performed to evaluate LoRa performance have shown that it performance iss influenced by its environment. Indoor experiments have shown that LoRa's performance is sensitive to both the location inside the building, as well as daily changes in the physical environment due to functional use of the office floor. Outdoor experiments have shown that mobility and the speed of the movement do not negatively influence performance. The distance and type of environment do play a role with less dense environments in terms of physical objects are performing the best. Empirical results have shown that LoRa's outdoor performance can be increased by lowering the transmission data rate.

Simultaneous LoRa transmissions on the same frequency and SF can cause collision-induced packet loss depending on the timing of the transmissions. During such an event, there are conditions under which the stronger of the two signals can be successfully received due to the Capture Effect. Most of the time, however, at least one of the signals is lost. The ability to provide reliable network connectivity in a dense IoT environment is examined by means of a simulator, whose behavior is based on the empirical results found in this thesis. Subsequent simulations of environments in which a large number of LoRa end-devices are present show that a few thousand end-devices can already induce performance degradation.

Several solutions to cope with collision-induced packet loss or to reduce the chance that it will happen are examined. A combination of (i) decreased base station coverage, (ii) smaller data packets, and (iii) more transmission channels is found to be the best solution to prevent network congestion. Ultimately, the reliability of a LoRa network depends on the density of LoRa devices in the neighborhood and the physical characteristics of the environment in which end-devices are placed.

Table of contents

1 INTRODUCTION

Wireless network technologies have become one of the most influential technologies of this time, pervading into our daily lives by connecting mobile devices to various cyber-physical systems. The explosive growth of wireless networks over the past decade has changed the way people communicate and interact with each other. Nowadays more users connect to the Internet through mobile devices such as smartphones and tablets than via fixed devices such as desktops [1]. Not only do we utilize well-known wireless standards such as Wi-Fi, 4G and Bluetooth to connect ourselves to the rest of the world, our devices use these technologies to continuously retrieve information for us, even when we are not actively using the device. A smartphone can for example transfer audio data to a nearby wireless headset and in the meantime receive notifications for various applications via Wi-Fi or 4G.

The trend of connecting more ever-smaller devices to the Internet has brought us to what is commonly believed as the start of the Internet of Things (IoT) era [2]. While the Internet itself revolutionized the interconnection between people, it is believed that the next step will be the revolution of the interconnection between objects (or 'things') embedded with intelligence [2][3]. One of the key enablers for the IoT is the growing pervasiveness, ubiquity and constant reduction in size and cost of (embedded) computing devices [4]. Another enabler is the innovation in wireless communication technology that has resulted in a large variety of ways to connect these devices to the Internet or each other, almost anywhere and anytime. While the Internet can be utilized for the communication between various independent systems, wireless technologies offer mobility and flexibility when connecting devices to the Internet. Many of these technologies are well known (e.g. Wi-Fi, Bluetooth, Zigbee, and 3G/4G) and each of these can enable devices to wirelessly communicate. However, each technology has its own pros and cons, depending on the application design and requirements. Factors such as range, data (or bit) rate, latency and energy consumption play important roles in the choice of a wireless technology.

Currently, one of the fastest developing technology in the field of wireless networks is the Low Power Wide Area Network (LPWAN) technology [5][11]. LPWAN networks are typically characterized by their low power consumption, long range, and low data rate. LPWAN technology is considered as an important factor in enabling the future development of smart metering, smart city and (smart) home automation applications [6][7]. Although LPWAN technologies are still considered to be "under construction" in terms of market adoption, maturity and standardization [2][8], a number of these technologies have gained significant attention from both the scientific world and the media [2], and some of them are already deployed in various applications [9][10].

LPWAN technology is best suited for communication between devices that only require small amounts of data transfer, a long lifetime and low implementation costs. These networks operate on unlicensed frequency bands and are characterized by their long range and star topologies, e.g. end-devices communicate directly to a gateway, which connects them to the Internet [12]. Figure 1 shows how various wireless technologies can be categorized depending on their range. The most

commonly known technologies such as Bluetooth and Wi-Fi belong to the Wireless Personal Area Network (WPAN) and Wireless Local Area Network (WLAN) categories, whereas cellular connectivity can be classified under Wireless Wide Area Network (WWAN) technology. LPWAN technologies complement and sometimes supersede cellular networks in terms of energy consumption and low implementation and deployment costs [13]. These features enable LPWAN technology to play a crucial role in the future of IoT [11], resulting in forecasts stating that this technology will connect billions of IoT devices in various sectors within the next decade [13][14].

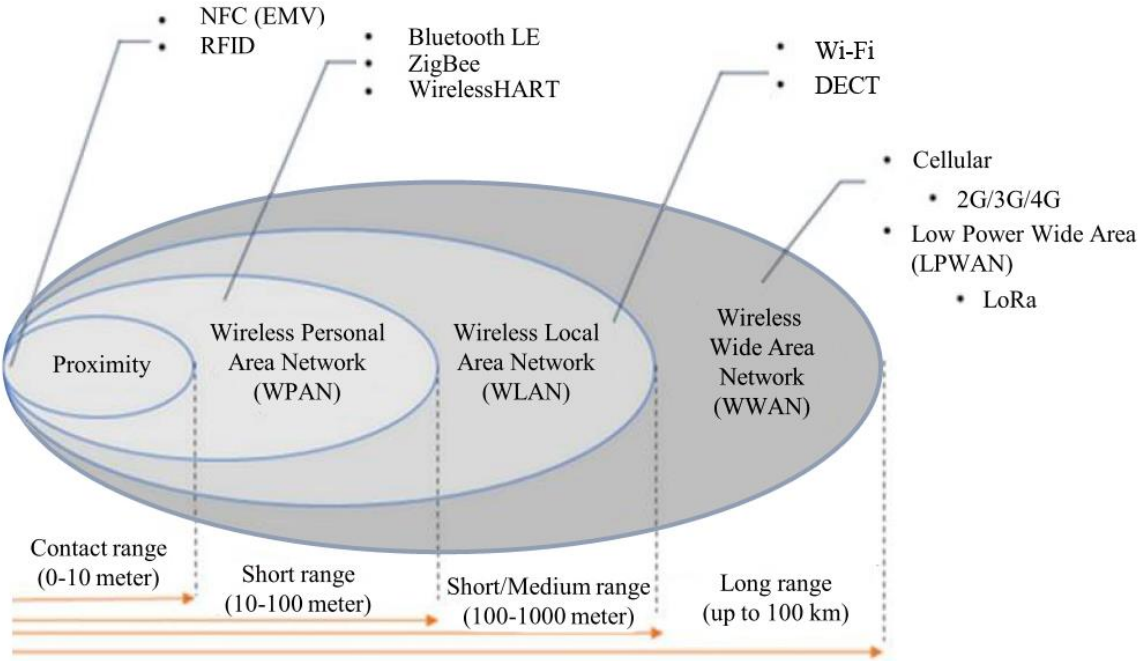

Figure 1. Classification of wireless technologies based on their range. Image redrawn from [2]

The incentive to research LoRa technology did not only come from the potential that LoRa has to facilitate a huge number of future IoT applications. The Pervasive Systems research group at the University of Twente, at which this thesis is conducted, is currently involved in the *Rhino project*, an anti-poaching initiative to prevent the poaching of rhino's in Africa. An LPWAN technology such as LoRa can play an important role in facilitating connectivity to different sensors placed throughout wildlife parks or on the animals themselves.

The purpose of this thesis is to examine the potential of LPWAN technologies by researching a new wireless technology called LoRa. LoRa, a name derived from 'Long Range', is created with the intention to facilitate low power and long range wireless communication links between devices. Networks utilizing LoRa technology can therefore be classified as a Low-Power Wide Area Network (LPWAN). Due to the novelty of LPWAN technology, it is difficult to determine what to expect from current implementations of LoRa. This thesis will therefore focus on the performance of LoRa technology in its current state, together with the accompanying open communication standard for LoRa called LoRaWAN, which is developed with the aim to standardize this technology. The main points of focus are on the performance of LoRa in terms of packet loss and signal reception caused by environmental influences, in both indoor and outdoor areas, as well as the energy consumption of a LoRa device. Moreover, this research aims to

validate claims made by the manufacturer on range and network capacity and uses the obtained empirical results to predict the behavior of an LPWAN network such as LoRa in a (future) dense IoT environment such as a smart city. The main research question is therefore: *Is LoRa technology a suitable candidate to provide a reliable solution for network connectivity in different IoT environments?*

To answer this question, we will focus on LoRa performance in indoor and outdoor environments in different configurations. The focus lies on examining the capability of LoRa to provide a good and reliable communication link for IoT applications. This is determined by examining the characteristics of LoRa, evaluating the behavior of LoRa end-devices, and conducting experiments under various circumstances to find potential weaknesses.

The rest of this thesis is structured as follows: A number of LPWANs, including LoRa, will be discussed in Chapter 2 to indicate the characteristics of LPWANs and gives an idea of what to expect from this technology. It provides an overview of the current LPWAN landscape and compares LoRa to similar competitive technologies. Chapter 3 gives a more in-depth view on LoRa and its proposed standard to obtain an overview of LoRa technology and to determine its capabilities, network structure, and theoretical performance characteristics. These characteristics are then evaluated in Chapter 4 by means of measurements to validate the documented behavior.

Next, the performance of LoRa technology is examined by means of empirical research in indoor and outdoor environments, whose results are described in Chapters 5 and 6 respectively. Additionally, these chapters describe the most important characteristics of an end-device using LoRa technology, as well as where the weaknesses are located. The results obtained in these chapters are used in Chapter 7 to build a simulator capable of replicating the behavior of a large-scale LoRa network. Such a simulation helps predicting how a LoRa network would behave in a scenario in which many devices are involved, what its weaknesses are, and provide a way to implement optimization measures to increase the performance of a LoRa network. Finally, concluding remarks are presented in Chapter 8.

2 Low-Power Wide Area Networks

LPWAN technologies are designed to offer a combination of features to complement and sometimes supersede conventional cellular and short-range wireless technologies in terms of energy consumption and low implementation and deployment costs [13]. This can also be seen in Figure 2, in which wireless communication technologies are compared in terms of range, bandwidth, energy consumption and implementation costs.

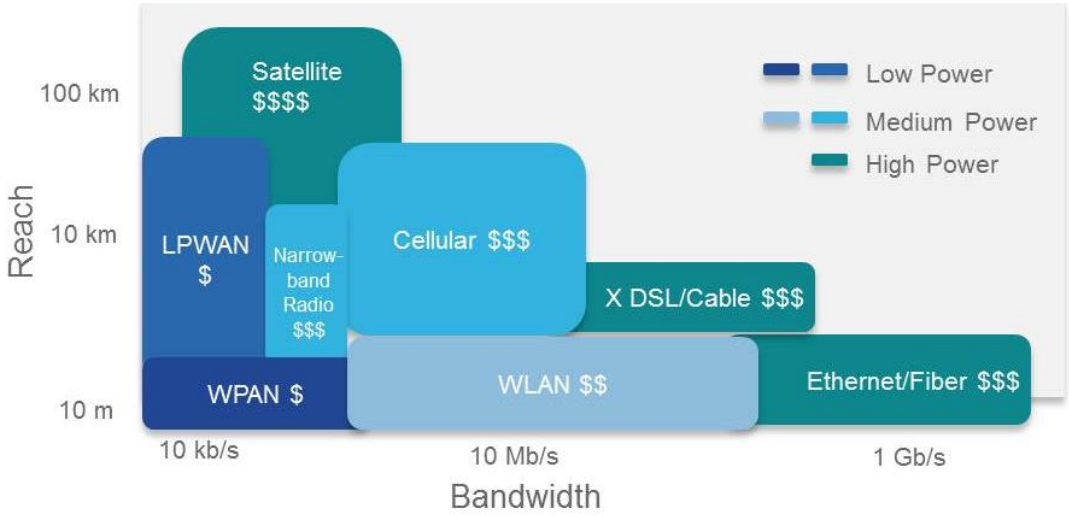

Figure 2. Comparison of wireless communication technologies in terms of their bandwidth, range, energy consumption and implementation costs. Image taken from [15]

The characteristics of LPWAN technology and the expected contribution to the rise of IoT communication has made this technology an interesting research topic, resulting in the rise of many implementations over the past few years [16]. The common characteristics for all of these technologies are the following [17]:

- *long range*: end-devices can be more than 10 kilometers away from the gateway
- *low data rate*: less than 5.000 bits per second, often only 20-256 bytes per message
- *low energy consumption*: enough to last between five and ten years
- *use of less congested sub-GHz frequency bands*: to decrease attenuation due to obstacles

This chapter gives a brief overview of the most prominent and developed LPWAN technologies apart from LoRa, which is discussed in more detail in Section 2.4.

2.1 SIGFOX

SIGFOX is developed by a French company founded in 2009 and is the first LPWAN technology on the market [8]. It uses a cellular style system and binary phase-shift keying (BPSK) modulation to transmit data using ultra-narrowband (UNB) technology. Data is transmitted in the 868MHz ISM frequency band and the spectrum is divided into channels of 100Hz wide. Messages are said to travel up to 50 km in rural and 10 km in urban areas. SIGFOX limits both the number of messages per day and the payload data size. For uplink transmissions, this means 140 messages with a maximum payload of 12 bytes are allowed per day. Downlink messages can only contain

up to 8 bytes and are limited to 8 messages per day [18][19]. According to SIGFOX, each gateway can handle up to a million connected devices. SIGFOX's network layer is not publicly available, which explains the lack of publicly available documentation for its internal workings.

2.2 Ingenu (On-Ramp Wireless)

Ingenu, formerly known as On-Ramp Wireless, has been founded in 2008. The company developed and patented its Random Phase Multiple Access (RPMA) technology, which it uses in uplink (device to network) communication [20]. Downlink messages are also supported but use conventional Code Division Multiple Access (CDMA) modulation. RPMA utilizes the unlicensed 2.4GHz band with 1 Mhz wide channels. Messages can contain between 6 bytes and 10 kilobytes of data and can range to over 500 km in case of an open line of sight [21].

2.3 Weightless

Weightless is a set of three open wireless technology standards developed by the non-profit organization Weightless Special Interest Group (SIG). The weightless protocols operate in sub-GHz frequency bands with each standard having its own characteristics. The *Weightless-N* and *Weightless-P* standards operate in license-free ISM bands whereas *Weightless-W* operates in the (TV whitespace) frequencies between 470MHz–790MHz that were previously allocated for analog TV broadcasts. The use of these frequencies could come with licensing requirements [22]. The packet size is flexible and has a minimum of 10 bytes for *Weightless-P* and *Weightless-W* with no defined maximum. The maximum packet size for *Weightless-N* is 20 bytes. The communication range in urban areas lies between 2 km (for *-P*) and 5 km (for *-W*) and the channel width is set to 200Hz, 12.5 KHz and 5 MHz for *Weightless-N*, *-P* and *–W,* respectively [19]. Little is known about the technology behind it and only weightless-N technology is currently being deployed.

2.4 LoRa

The LPWAN technologies described above give an indication of the current state of this technology. This section will give a more in-depth view of the characteristics of LoRa, i.e., a proprietary modulation scheme owned by semiconductor manufacturer Semtech [23].

2.4.1 Spread spectrum modulation

LoRa is a spread spectrum modulation scheme based on Chirp Spread Spectrum (CSS) modulation. Chirp Spread Spectrum originates from the 1940's and has since then been used for military and space communication due to its wide communication range and robustness to interference. In a spread spectrum system the transmitted signal is spread over a wide frequency band, which is much wider than the minimum bandwidth required to transmit the data. This type of modulation is robust against both narrowband and wideband disturbances due to the spreading of the signal over a large bandwidth [24]. Other benefits of CSS include the resistance against Doppler shift, multi-path fading, the low transmission power needed to transmit over a certain distance [25][26], and the ability to demodulate the spread signal from below the noise floor [24].

LoRa modulation itself can be used in many different network architectures. The spreading of the signal accros the spectrum in LoRa modulation is achieved by generating a chirp signal, i.e., a signal in which the frequency linearly increases (up-chirp) or decreases (down-chirp) over time. Figure 3 depicts examples of both an up-chirp and a down-chirp, and their representation in the time-frequency domain. The speed at which each chirp is varied in frequency is known as the spreading factor (SF), or chirp data rate. The bandwidth of these chirp symbols equals the spectral bandwidth of the signal. Each chirp symbol carries SF encoded bits ranging between 6 and 12. The SF influences the transmission time (also known as the Time on Air) and therefore the energy consumption, the transmission range, and the data rate. It therefore determines the amount of redundant data spread across the transmission. A high spreading factor means that more redundant information is transmitted, increasing the range but decreasing the data rate [27]. Since LoRa employs spreading factors orthogonal to each other, multiple spread signals can be transmitted at the same time and on the same channel without signal degradation [26].

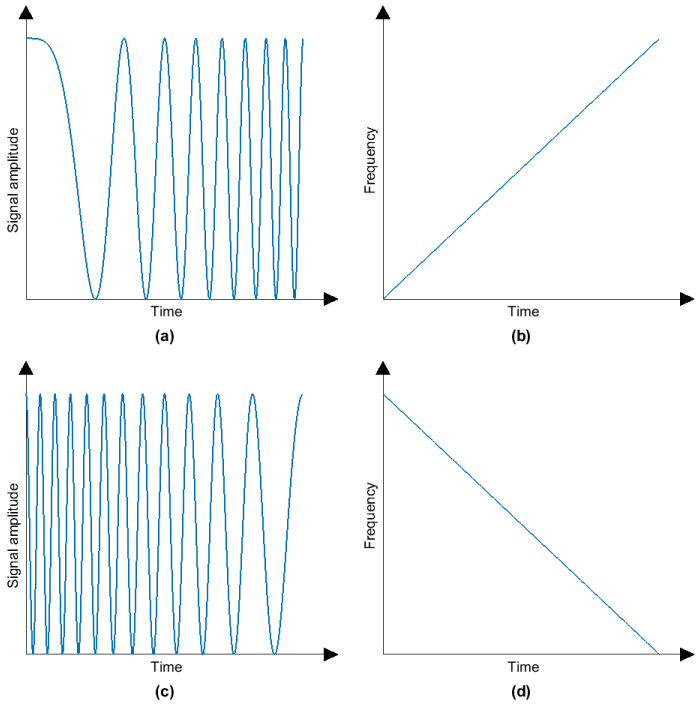

Figure 3. a) A linear up-chirp waveform b) The corresponding up-chirp in the time-frequency domain c) A linear down-chirp waveform d) The corresponding down-chirp in the time-frequency domain

2.4.2 ISM frequency band usage

LoRa operates in the license-free sub-1 GHz ISM bands such as 433 and 868 MHz (EU) and 169, 433 and 915 MHz (USA). The allocation of these license-free frequency bands as well as the regulations for these bands differ per region. LoRa usage specifications can therefore vary between regions. Currently, these specifications only exist for Europe and North America. This thesis focusses on LoRa usage in Europe, where it can be used on the LPD433 (low power device 433 MHz) and SRD860 (Short Range Devices 860MHz) bands. Since the LoRa network used for the experiments in this thesis uses the SRD860 band, the usage of LoRa in the LPD433 band will not be discussed further.

The radio spectrum allocation and standards in the ISM band for Europe are defined by the European Telecommunications Standards Institute (ETSI) [28]. The SRD860 band ranges between 863 and 870 MHz and is referred by ETSI as band G. In general, usage of this band is bound to regulations in terms of maximum transmission power (14dBm) and duty cycle (0.1%). There are, however, five specific sub-bands (G1, G2, G3, and G4) with their own separate regulations [29]. A summary of the current SRD regulation as defined in [29] can be seen in Table 1. Less strict duty cycle requirements apply when the end-device implements Listen Before Talk (LBT) and Adaptive Frequency Agility (AFA), meaning that the device is able to detect an unoccupied sub-band or channel prior to transmitting and is able to dynamically change the temporary operational channel if another transmission is detected. The devices in this thesis do not implement these features, as will be described in Section 3.1.2, and are therefore bounded to the restrictions described below.

Aside from when it operates in the G3 band, a LoRa end-device is not allowed to exceed an effective isotropic radiated power (EIRP) of 25mW (14 dBm). This means that an end-device should not radiate more than 25mW when an theoretically ideal isotropic antenna is used. ETSI has also defined duty cycle constraints for each band, meaning that a device cannot exceed a cumulative transmission time (Time on Air) on one carrier frequency, relative to a one-hour period. This restriction therefore results in a total transmission time of 36 seconds per hour when the maximum duty cycle is 1%. Although it is not mentioned in [29], there are several sources [30][31] stating that the maximum continuous transmission time is much lower than the maximum transmission time in a one-hour period. For example, an end-device is not allowed to transmit 36 seconds and then be quiet for an hour.

Band name	Range	Maximum transmission power (dBm)	Maximum duty cycle
G	863,0 – 870,0 MHz	25 mW EIRP (14 dBm)	0,1%
G1	868,0 – 868,6 MHz	25 mW EIRP (14 dBm)	1%
G2	868,7 – 869,2 MHz	25 mW EIRP (14 dBm)	0,1%
G3	869,4 – 869,65 MHz	500 mW EIRP (27 dBm)	10%
G4	869,7 – 870,0 MHz	25 mW EIRP (14 dBm)	1%

Table 1. SRD860 sub-band regulations.

The implementation of these transmission regulations, the communication mechanism, and network architecture can be done in many ways. Guiding such design choices into a standard would help with the market adoption of the technology and creates interoperability between LoRa enabled systems. An open communication standard for LoRa called LoRaWAN is currently under development and is promoted by the LoRa Alliance [32]. LoRaWAN will be discussed in the next chapter.

3 THEORETICAL CHARACTERISTICS OF LORA(WAN)

This chapter describes the most relevant characteristics of a typical LoRaWAN compliant end-device and the related theoretical performance of such a device. First, we look at the communication characteristics defined in the proposed LoRaWAN standard including the network architecture, usage of the frequency spectrum, and the implementation of transmission data rates. Then, the LoRa packet structure is discussed to reveal LoRa's data structure and to see what LoRaWAN adds to that. Next, a LoRa packet's Time on Air resulting from this packet structure is discussed. This provides an indication of LoRaWAN transmission times, which in turn affect the maximum transmission interval of an end-device. Finally, the described characteristics are used to get an indication of some performance characteristics of an end-device running LoRaWAN.

3.1 LORAWAN

Whereas LoRa is a proprietary modulation scheme, e.g. the physical layer implementation, the LoRa Alliance is working to design a MAC protocol and network architecture for LoRa in an open global standard called LoRaWAN. With this standard, the LoRa Alliance aims to provide interoperability among LoRa devices and to promote standardization of LPWAN technologies. Their aim is to offer a predictable battery lifetime, network capacity and to allow a large degree of freedom for users and businesses to develop IoT applications [33][34]. This thesis will follow the specifications proposed by the LoRa Alliance [35].

3.1.1 Network architecture

Most LPWAN networks use a star network topology. A star network architecture is characterized by a central node or a gateway to which all devices connect. In contrast to a mesh network where individual end-devices (nodes) forward data to (sometimes all) other end-devices to increase the communication range. The LoRaWAN network architecture is laid out as a star-of-stars topology, which means that gateways act as a message forwarder between end-devices and a central network server. A visual comparison between the above-mentioned topologies can be seen in Figure 4.

A long-range star or star-of-stars network preserves battery life and reduces complexity since nodes do not have to receive and forward data coming from other nodes. LoRaWAN specifies a network in which nodes are not associated with one specific gateway but instead transmit their data to every gateway in reach [33]. Every gateway that receives a message will forward it to a network server through an Internet connection (e.g. via Ethernet, Wi-Fi or a cellular connection). The intelligence of the network is located at this server, which is responsible for filtering duplicate packets, performing security checks, and scheduling acknowledgements through the optimal gateway.

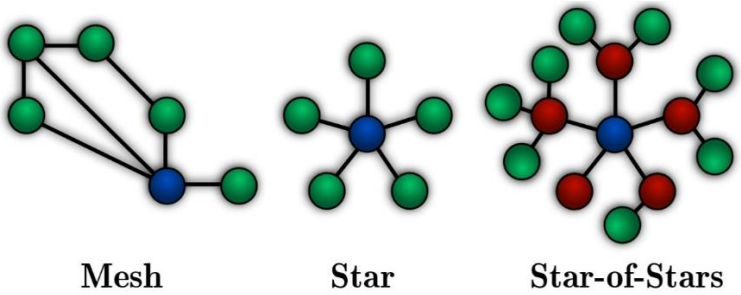

Figure 4. Different network topologies. LoRaWAN implements a Star-of-Stars topology. Green spheres indicate end-devices, red spheres the gateways, and blue spheres the network server

3.1.2 Channels

End-devices in LoRaWAN communicate on different frequency channels created within an ISM band, such as the SRD860 band (indicated by [35] as EU868MHz). LoRaWAN requires all devices to implement a set of default channels, which differ per used frequency band. Devices operating in the EU868MHz band should implement at least the three channels as listed in Table 2.

Channel	Channel frequency	Channel bandwidth	Maximum duty cycle	Data Rate range
0	868.10 MHz	125 KHz	1%	0-5
1	868.30 MHz	125 KHz	1%	0-5
2	868.50 MHz	125 KHz	1%	0-5

Table 2. Default channels for EU868MHz LoRaWAN devices in the SRD860 band

The ETSI regulations described in Section 2.4.2 allow the choice between duty cycle limited transmission or a Listen Before Talk and Adaptive Frequency Agility transmission scheme. LoRaWAN currently only uses duty cycle limited transmissions and implements this by enforcing a per sub-band duty cycle limitation [35]. This means that the transmission time, also known as the Time on Air (T_{onAir}), is used to restrict transmissions on the same sub-band for a time T_{off}, as given by the following formula:

$$T_{off} = \frac{T_{onAir}}{DutyCycle_{subband}} - T_{onAir} \qquad (3.1)$$

The entire sub-band is unavailable for any transmissions during T_{off}, meaning that the device (both gateways and end-devices) should transmit on another sub-band or have to wait before transmitting new messages. Therefore, a device that only implements the three default channels in the G1 sub-band has to wait at least 99 seconds before it can transmit a new packet on the G1 sub-band after transmitting a packet with a T_{onAir} of 1 second. The duty cycle limitation therefore decreases the maximum data rate of the device, i.e. in this case by a hundredfold.

3.1.3 Data Rates

LoRaWAN defines a specific set of Data Rates (DRs) to be used for transmissions. A LoRaWAN DR indicates a combination of SF and bandwidth used for a transmission. These two factors determine (i) how many bits (equal to SF) are encoded per chirp, (ii) the speed at which these chirps are sent (i.e. the speed of the frequency increase or decrease), as well as (iii) the start and end frequency of the chirp, as also described in Section 2.4.

Table 3 shows the relation between the LoRaWAN DRs and the corresponding LoRa transmission characteristics. It is worth noting that DR7 uses FSK and not LoRa modulation. Additionally, the default 3-channel LoRaWAN implementation does not utilize DR 6 and 7 and LoRaWAN does not use SF 6.

Data Rate no.	Spreading Factor	Bandwidth (kHz)
0	12	125
1	11	125
2	10	125
3	9	125
4	8	125
5	7	125
6	7	250
7	-	-

Table 3. LoRaWAN Data Rates and their corresponding LoRa transmission characteristics

3.1.3.1 *Maximum payload*

The maximum length of the LoRaWAN payload segment is region-specific and also depends on the used data rate. For LoRaWAN usage in Europe, this means that the maximum payload size for each data rate is defined based on specifications of Table 4. This table shows the eight data rates as defined by LoRaWAN, together with their corresponding configuration and maximum payload sizes. As mentioned in 3.1.3, DR 7 uses FSK modulation instead of LoRa modulation, and the default 3-channel implementation does not utilize DR 6 and 7.

Data Rate No.	Spreading Factor	Bandwidth (kHz)	Maximum payload size (bytes)
0	12	125	64
1	11	125	64
2	10	125	64
3	9	125	128
4	8	125	235
5	7	125	235
6	7	250	235
7	-	-	235

Table 4. Maximum payload size per LoRaWAN Data Rate [35]

3.1.4 Adaptive Data Rate

LoRaWAN allows a network to employ an Adaptive Data Rate (ADR) scheme, which enables the network infrastructure to manage the data rate and transmission power of end-devices. End-devices allowing ADR can receive downlink messages from a gateway containing the data rate and transmission power they should use. This enables the network to increase packet delivery and network capacity as well as the battery life of end-devices [36]. End-devices close to a gateway can be configured to use a high data rate and a low transmission power whereas end-devices at the edge of the network use the maximum transmission power and the lowest data rate to maximize their reach.

3.1.5 LoRaWAN classes

LoRaWAN has three different device classes to support an optimal implementation for a wide range of LoRa based applications. The difference between these classes relate to the timing of the so-called receive-windows. These windows determine the moment at which the end-device listens for downlink messages from the gateway in order to receive them. LoRaWAN defines the following classes:

- Class A is the default mode of operation for an end-device in which it opens two short receive-windows after it has sent an uplink message. These receive-windows allow the reception of downlink messages from the network. The first receive-window is opened on the same channel and uses the same data rate as the initial uplink transmission. The second window uses a fixed configurable frequency and data rate to increase resilience to channel fluctuations [35]. By default, the data rate and channel frequency of the second receive-window are set to 0 and 869.525 MHz, respectively. This frequency is part of the G3 sub-band, which allows higher power transmissions (500mW) with an increased duty-cycle (10% instead of 1%). The configurable delay between the end of the uplink transmission and the opening of the first window is 1 second by default whereas the second receive-window will always be opened 1 second after the first. Due to this implementation, class A end-devices can only receive messages directly after an uplink transmission. This implementation allows the end-device to be very energy-efficient. Class A must be supported by all LoRa devices.
- Class B end-devices have extra receive-windows at scheduled times. These devices listen at scheduled times for beacon packets from the gateway and open their receive-window once such a packet is received. This mechanism makes sure that the server can synchronize the downlink packet with the receive-window of the end-device. Applications that require data transmission on request or remote control (e.g. actuators) can benefit from downlink message reception independent from uplink transmission.
- Class C end-devices only close their receive-windows when they are transmitting themselves. This constant listening for messages therefore makes class C the most energy consuming class of end-devices. Devices implementing class C functionality should preferably not be battery-powered.

3.1.6 Joining a LoRa network

There are two methods to connect to a LoRa network, i.e., using Over-The-Air-Activation (OTAA) or Activation by personalization (ABP). Networks allowing ABP do not require the end-device to register before they start their transmissions. The end-device is simply recognized by its personal device address and the network security keys to encrypt the data are predefined in the end-device. OTAA requires a LoRa end-device to perform a join procedure by issuing a join request. During this procedure, a dynamic device address and network session keys are obtained by the end-device.

Depending on which method is used, some parameters need to be set before communication with a LoRa network can take place. When joining a network based on personalization, the end-device needs to be configured manually with (i) a *device address* to identify the end-device, (ii) a *network session key* to verify the Message Integrity Code of a packet to ensure data integrity, and (iii) an *application session key* to encrypt and decrypt the payload. Using OTAA, a device needs to specify (i) a global *device EUI* (Extended Unique Identifier) that uniquely identifies the end-device, (ii) an *application EUI* to uniquely identify the application provider, and (iii) an *application key* that is used to independently generate the encryption *network session* and *application session* keys. The *application key* is therefore never transmitted over the air. To join the network to conduct the measurements described in this thesis, OTAA is required.

3.1.7 Acknowledgements

End-devices can transmit either *unconfirmed* or *confirmed* data messages. Unconfirmed messages are not acknowledged by the network, so an end-device will never know if the message was received properly. If the application requires that the end-device knows whether transmission was successful, it should send its data using a confirmed message. If correctly received, the network will respond with an acknowledgement during one of the end-device's receive windows. When either the initial transmission or the acknowledgement fails, the end-device will retransmit the (same) data, if necessary.

3.2 LoRa Packet Composition

As depicted in Figure 5, each LoRa packet is composed of three elements, (i) a preamble, an (ii) optional header, and (iii) the payload. The header and payload of each packet can be checked for integrity with the use of cyclic redundancy checks (CRCs). These CRCs are appended to the header and payload of the packet, respectively. A CRC is a commonly used error detection mechanism based on a predefined number of check bits, often called a checksum. CRC enables the detection of small changes in the accompanying data based on the remainder of a polynomial division of the data's content [37]. When the CRC fails, the packet is regarded as corrupted and the receiver discards it. To improve the robustness of LoRa packets, the payload of each packet is encoded with cyclic error coding to allow forward error correction. The encoding rate can be set to a minimum of 4/5, meaning that for each four bits of information, one bit is added. The maximum coding rate is 4/8, doubling the amount of payload bits. The maximum packet length is 256 bytes [36].

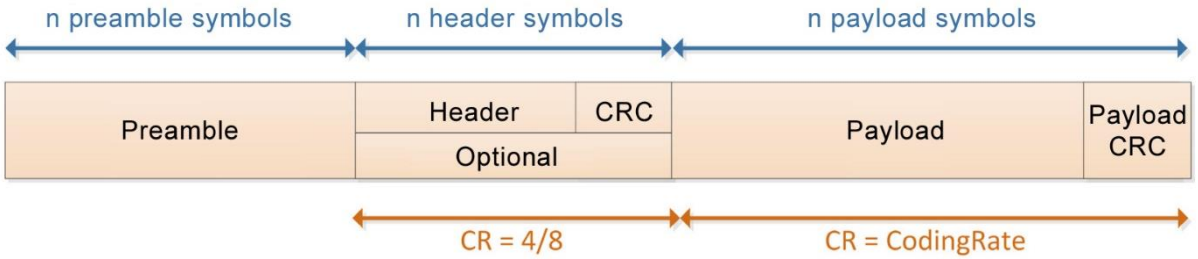

Figure 5. Composition of a LoRa packet. The header, header CRC and payload CRC are optional.

3.2.1 Preamble

The LoRa preamble signals a receiver that the end-device is about to transmit its data. The receiver uses the preamble to synchronize with the incoming data and detects the start of the header or the payload if no header is used. The programmable length of the preamble can range between 6 and 65535 (chirp) symbols. Together with the fixed overhead of 4.25 symbols [38], this yields to a minimum preamble of 10.25 symbols. In LoRaWAN, the preamble length is defined as 8 (+4.25) symbols.

3.2.2 Header

The optional header contains information about the content of each packet. This information includes the payload length in bytes, the forward error correction coding rate, and whether the payload CRC at the end of the packet is present. As can be seen in Figure 5, the header itself is encoded with the maximum coding rate of 4/8. The length of the header is not specified in the documentation. The use of a header is optional and if no header is used, which is called the Implicit Header Mode, the payload length, coding rate, and CRC presence are fixed and known (i.e. manually specified) at both sides of the radio link. The default mode of operation is, also for LoRaWAN, the Explicit Header Mode, in which the header is present in each packet. Only when SF 6 is used (which does not apply to LoRaWAN), the use of Implicit Header Mode is mandatory [38].

3.2.3 Payload

The payload segment contains the actual data, encoded with the coding rate specified either in the header, or manually at the receiving side of the radio link. A 16-bit payload CRC as shown in Figure 5 is optional for default LoRa end-devices but is always used in the LoRaWAN protocol.

3.2.3.1 *LoRaWAN payload fields*

In LoRaWAN, the payload segment itself, as depicted in Figure 6, is divided into a few additional fields. In LoRaWAN, the payload segment is called PHYPayload and comprises of three fields, i.e., a 1-byte MAC header (MHDR), a MAC payload (MACPayload), and a 4-byte message integrity code (MIC). The MACPayload field (called data frame) contains the (optional) actual data (FRMPayload) preceded by a frame header (FHDR) and an optional port field (FPort). This 1-byte port field is application specific and can only be left out if the frame payload field is empty. If set to 0, the port field indicates that the payload consists of only MAC commands. These commands are used for network administration purposes. The frame header consists of four fields, i.e., the 4-byte address of the end-device (DevAddr), a frame control byte (FCtrl), a 2-byte frame counter (FCnt), and a maximum of 15 bytes in size frame options (FOpts) field for communicating MAC commands.

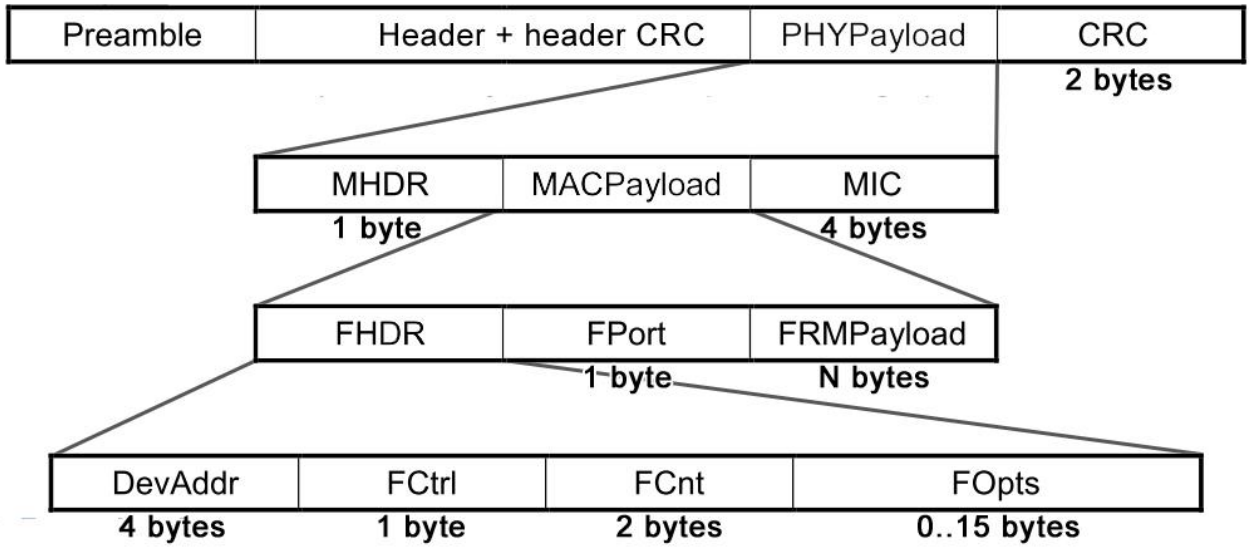

Figure 6. LoRaWAN packet composition and field sizes. Image redrawn from [35]

Assuming that a device implementing the LoRaWAN protocol is utilized to transmit a 1-byte payload without any MAC commands, the size of the PHYPayload would equal:

$$\underbrace{\text{PHYPayload}}_{14\text{ bytes}} = \underbrace{\text{MHDR}}_{1\text{ byte}} + \underbrace{\text{DevAddr}}_{4\text{ bytes}} + \underbrace{\text{FCtrl}}_{1\text{ byte}} + \underbrace{\text{FCnt}}_{2\text{ bytes}} + \underbrace{\text{FOpts}}_{0\text{ bytes}} + \underbrace{\text{FPort}}_{1\text{ byte}} + \underbrace{\text{FRMPayload}}_{1\text{ byte}} + \underbrace{\text{MIC}}_{4\text{ byte}}$$

In other words, the LoRaWAN protocol adds an additional 13 bytes to the data that needs to be transmitted. The additional fields mentioned above result in the minimal size of 14 bytes for the payload segment shown in Figure 5. Moreover, the additional payload data reduces the maximum amount of application data an end-device can transmit per data rate, as described in Section 3.1.3.1, by 13 bytes. Since all measurements conducted in this thesis use the LoRaWAN protocol, the payload of a packet is considered to include these extra 13 bytes.

3.3 Time on Air

According to the LoRa specification, the transmission time of a packet (Time on Air) is predominantly determined by the payload of the packet, the used spreading factor, the coding rate of the payload, and the channel bandwidth. The Time on Air of a LoRa packet is based on the duration of each individual chirp symbol. This symbol time T_s is defined as [26]:

$$T_s = \frac{2^{SF}}{BW} \text{ seconds} \tag{3.2}$$

The total Time on Air is defined as the sum of the preamble's Time on Air and the Time on Air of the actual packet. As described in the previous section, a fixed overhead of 4.25 symbols is added to the total symbol length of the preamble. The number of packet symbols (including the symbols for the header and CRC, if present) is given by the following formula [38]:

$$n_{payload} = 8 + \max\left(\text{ceil}\left[\frac{8PL - 4SF + 28 + 16CRC - 20IH}{4 \times \ SF - 2DE}\right] \ CR + 4 \ , 0\right) \tag{3.3}$$

This formula depends on the following denotations:

- PL: number of payload bytes (the length of the PHYPayload segment)
- SF: spreading factor
- CRC: indicates if the payload CRC is present. CRC = 0 if not, otherwise CRC = 1.
- IH: indicates if a header is present. IH = 1 if not, otherwise IH = 0.
- DE: indicates Low Data Rate Optimization usage. DE = 0 if not, otherwise DE = 1.
- CR: indicates the used coding rate. This value can range from 1 to 4.

The total amount of symbols n_{total} and corresponding Time on Air for each packet T_{packet} can then be calculated with the following equations [38]:

$$n_{total} = (n_{preamble} + 4.25) + \ n_{payload} \tag{3.4}$$

$$T_{packet} = n_{total} \times T_s = \left((n_{preamble} + 4.25) + \ n_{payload}\right) T_s \tag{3.5}$$

$$T_{packet} = \left((n_{preamble} + 12.25) + \max\left(\text{ceil}\left[\frac{8PL - 4SF + 28 + 16CRC - 20IH}{4 \times \ SF - 2DE}\right] \ CR + 4 \ , 0\right)\right) T_s \tag{3.6}$$

As mentioned in Section 3.2.3, the LoRaWAN stack adds 13 bytes to the total payload that needs to be transmitted. Equations (3.3) and (3.6) can therefore be updated to suit transmissions using the LoRaWAN protocol:

$$n_{payload,\ LoRaWAN} = 8 + \max\left(\text{ceil}\left[\frac{8(13 + PL) - 4SF + 28 + 16CRC - 20IH}{4 \times \ SF - 2DE}\right] \ CR + 4 \ , 0\right) \tag{3.7}$$

$$T_{packet,LoRWAN} = \left((n_{preamble} + 12.25) + \max\left(\text{ceil}\left[\frac{8(13 + PL) - 4SF + 28 + 16CRC - 20IH}{4 \times \ SF - 2DE}\right] \ CR + 4 \ , 0\right)\right) T_s \tag{3.8}$$

The Low Data Rate Optimization feature (denoted by DE) is a feature that aims to improve the robustness of the transmission to frequency variations during the transmission of the packet on high spreading factors. The use of this optimization is mandated by Semtech for symbol durations exceeding 16ms, and is therefore implemented in LoRaWAN. This leads to the reduction of the amount of encoded bits in each chirp symbol by 2 bits, so the encoded bits per chirp becomes SF minus 2. Low Data Rate Optimization is enabled when T_s becomes more than 16ms [38], which is the case on spreading factors 11 and 12 (and therefore on LoRaWAN data rates 0 and 1).

The LoRaWAN specifications define that both up- and downlink messages include the LoRa header (and header CRC) and that the Low Data Rate Optimization is to be switched on for spreading factors 11 and 12 [35]. Uplink messages also require the use of the payload CRC, whereas downlink messages do not implement a payload integrity check to keep these messages as short as possible. We therefore assume that the header and CRC payload are present and the Low Data Rate Optimization is switched on for spreading factors 11 and 12 when calculating the Time on Air of uplink packet's. The same assumptions are made for downlink packets, with the difference that the payload CRC is assumed not to be used. Default values for the parameters needed to calculate the total amount of payload symbols are shown in Table 5.

Denotation	Value
CRC	1 (uplink messages) or 0 (downlink messages)
IH	0
DE	0 (data rates < 11) or 1 (data rates > 10)

Table 5. Default values for calculating the number of LoRaWAN payload symbols

The Time on Air and other packet characteristics such as energy consumption can also be calculated with the LoRa Modem Calculator Tool provided by Semtech [39]. The graphical interface of the tool is shown in Figure 7.

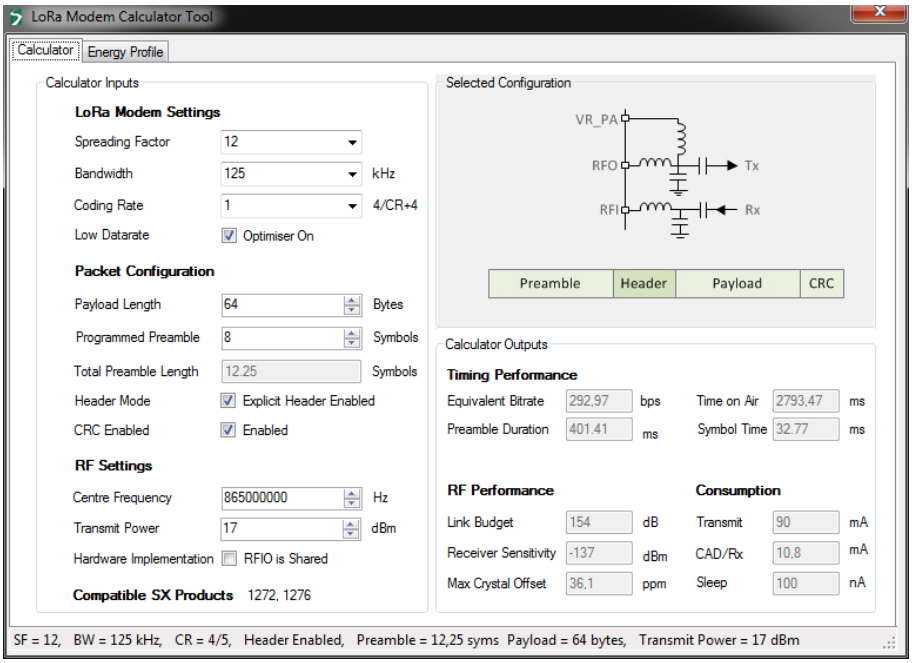

Figure 7. The LoRa Modem Calculator Tool, downloaded from [39]

3.4 TRANSMISSION POWER

LoRaWAN utilizes a default radiated transmit output power of 14 dBm. As specified by ETSI, this equals the maximum radiated power of 25mW for radio devices in Europe. The output power of a device may be lowered with steps of 3 dBm. The LoRaWAN specification defines six output power configurations of which five are available for operations in Europe, as shown in Table 6. The sixth and the strongest output power (20dBm) is not supported due to the maximum transmission power restriction defined by ETSI.

LoRaWAN TXPower setting	Corresponding TX power (dBm)
0	20 (not supported)
1	14
2	11
3	8
4	5
5	2

Table 6. Transmit powers as defined in the LoRaWAN specification.

3.5 THEORETICAL PERFORMANCE ANALYSIS

As a starting point, we take an end-device equipped with a (transceiver) LoRa modem implementing the default Class A LoRaWAN protocol as described in Section 3.1.5.

Based on Equation (3.8), we can calculate the Time on Air for a variation of payloads. This then gives an indication for the average time taken between sent packets based on the duty cycle restrictions specified in Section 2.4.2. It also indicated the maximum achievable throughput for different transmission configurations (based on data rate and payload size).

For this calculation, we assume that the transmitting device sends a packet as often as possible. The graphs shown in Figure 8 and Figure 9 show the calculated Time on Air and corresponding throughput as a function of payload for five different spreading factors. Since the number of payload symbols does not necessarily increase for each additional payload byte, an increase in payload bytes does not have to lead to a higher Time on Air and therefore throughput, as it can be seen from Figure 9.

These graphs are obtained by using the default LoRaWAN values, as presented in Table 5. The coding rate is equal to 4/5 (so CR equals 1 in the equation for T_{packet}) and the channel bandwidth is equated to 125 kHz. The payload size is varied between 14 and 64 bytes since a minimum data (FRMPayload) size of 1 byte results in a total payload of 14 bytes as described in Section 3.2.3, and the maximum payload for the highest three spreading factors is 64 (51+13) bytes. Table 7 shows the maximum transmittable data per packet for each LoRaWAN data rate and the corresponding Time on Air and maximum throughput. It is worth noting that when the channel bandwidth is doubled, the Time on Air will be halved and therefore the throughput will be doubled as well.

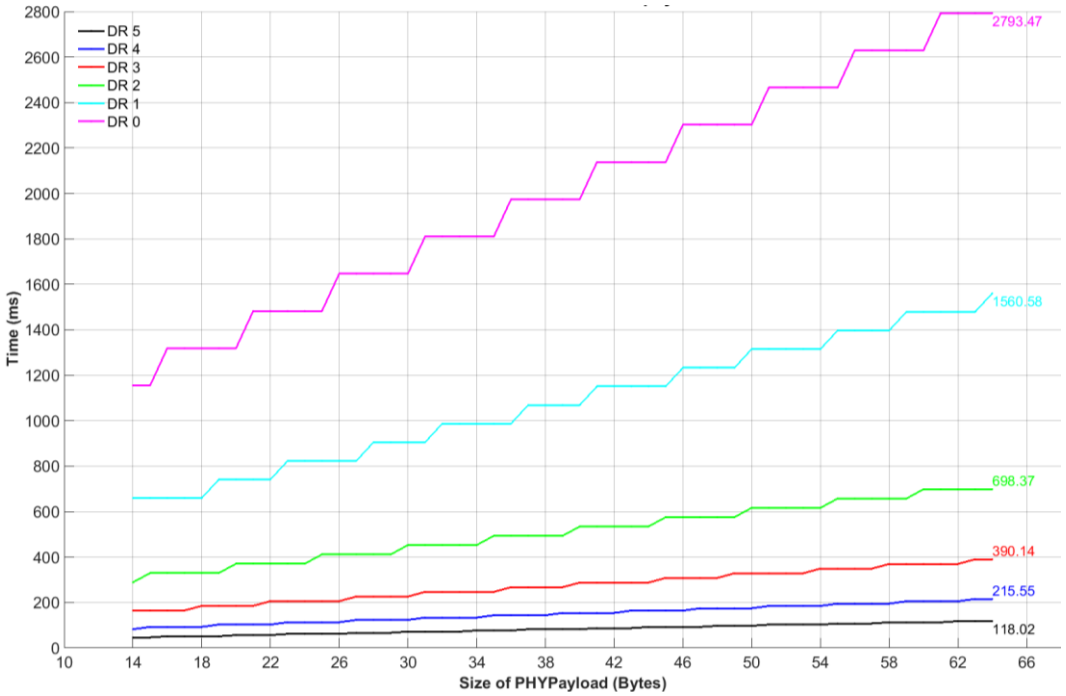

Figure 8. Time on Air as a function of payload for different Data Rates.

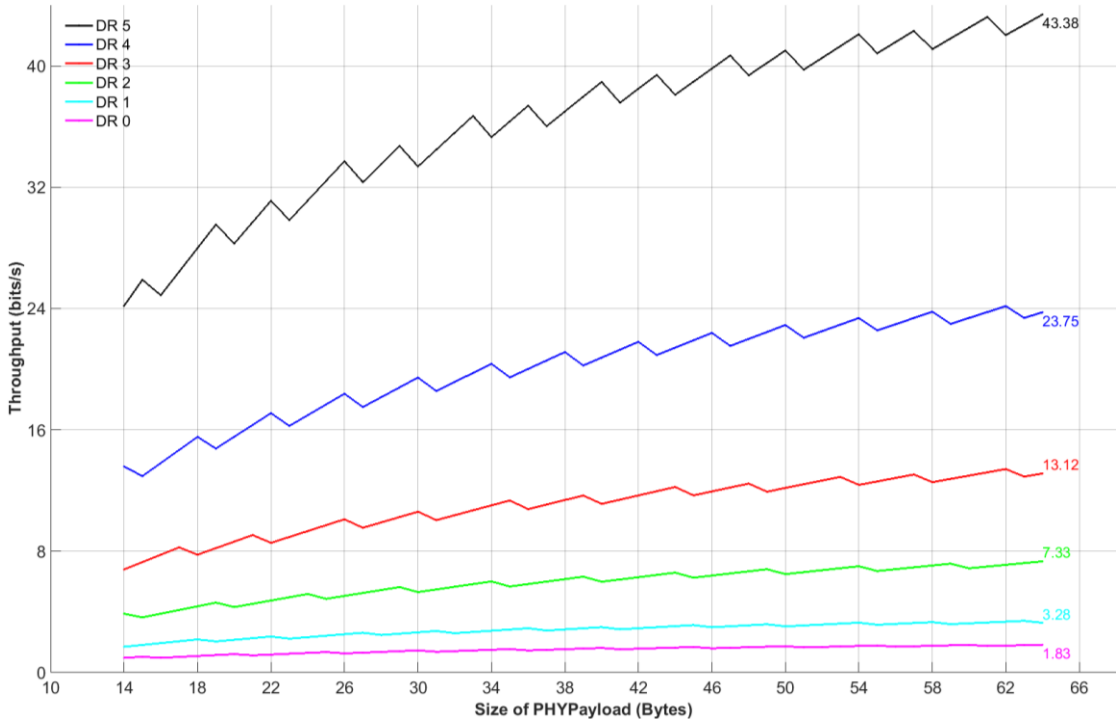

Figure 9. Maximum throughput per payload based on the 1% duty cycle restriction and a transmitter trying to send as often as possible.

Data Rate No.	Maximum PHYPayload size (bytes)	Time on Air (ms)	Maximum throughput (bits/s)
0	64	2793.47	1.83
1	64	1560.58	3.28
2	64	698.37	7.33
3	128	676.86	15.13
4	235	655.87	28.66
5	235	368.90	50.96

Table 7. Time on Air of LoRa packets with a maximum (LoRaWAN specified) payload and the corresponding throughput based on the 1% duty cycle limit.

The analysis above shows how various payloads influence the Time on Air for each of the five default LoRaWAN data rates and the maximum achievable throughput using a LoRaWAN end-device without implementing additional channels. Without the 1% duty cycle restriction, the maximum throughput would be hundredfold of the values calculated above. It is, however, worth noting that even though the regulations for the LPD433 band allow a higher duty cycle [28], LoRaWAN still limits the duty cycle to 1% to avoid network congestion [35]. The results obtained above are analyzed and verified in the next chapter.

Evaluating the Time on Air for different data rates and payloads has yielded throughput figures, which can be used to determine transmission schemes for LoRa end-devices. Small (one-byte payload) packets can approximately be transmitted 2.5 times more frequently than packets carrying a 64-byte payload. The trade-off between transmission frequency and payload size implies that end-devices requiring frequent data updates should keep their messages as short as possible to decrease the time they have to wait before a new transmission is allowed.

As long as they are able to reach the gateway, end-devices should use the fastest (lowest numerical) possible DR when Time on Air reduction is required. This is because each increment in DR almost doubles the Time on Air. This has, however, a downside in terms of diversity of used data rates. Devices requiring either frequent transmissions or saving as much battery power as possible are enticed to only use the fastest data rate. This causes an increased chance of collisions when many of such devices operate in the same area.

In general, the amount of data a device following LoRaWAN specifications can transmit is very limited. The maximum payload of 235 bytes on DR 4 and DR 5 still allows transmitting of a reasonable amount of data However the maximum throughput of 50.96 bits/s (on DR 5, with a payload size of 235 bytes), makes fast transmission of large amounts of data impossible. Creating additional channels with higher bandwidths or higher duty cycles would make it possible, for example, to increase the maximum throughput. A channel in the G3 sub-band would for example increase the duty cycle limitation to 10%, whereas the G1 sub-band supports the creation of a 250 KHz wide channel.

4 VERIFICATION OF THEORETICAL ANALYSIS OF LORA

This chapter first describes the physical components utilized to create and evaluate LoRa end-devices. The LoRa end-device and set up of the LoRa network are discussed, together with some important aspects such as the communication interface for the used LoRa end-device, its energy consumption, and the gateway used for receiving uplink messages. The operational behavior of the created devices is empirically validated by transmitting LoRa packets and visualizing them using a software-defined radio (SDR) spectrum analyzer, verifying the theoretical characteristics described in the Chapter 3, and analyzing the energy consumption of the LoRa end-device.

4.1 USED HARDWARE

The starting point for a LoRa based application is an end-device equipped with a LoRa modem transceiver and a gateway capable of receiving LoRa modulated signals. First, the characteristics of the created end-device are given and then the properties of the used gateway are described.

4.1.1 The LoRaWAN end-device

The main component of the end-device is the LoRa modem and its ability to implement the LoRaWAN protocol stack. A combination of these two can be found in the Microchip RN2483 Transceiver Module. The block diagram of the RN2483 module can be seen in Figure 10.

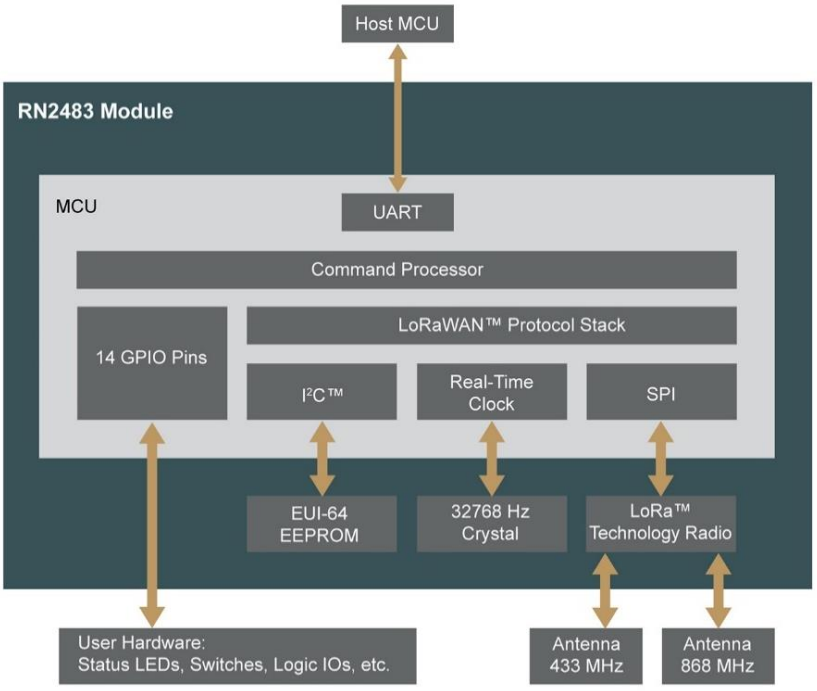

Figure 10. Block diagram of the RN2483 module. Image taken from [40]

This module is based on the Semtech SX1276 transceiver chip, which implements the LoRa technology. The block diagram shows that the module implements the LoRaWAN (Class A) protocol stack. It can work on both the 433 and 868 MHz frequency bands and can be operated with a microcontroller.

The RN2483 module has a text command/response interface in the form of a simple UART interface, which allows connectivity of the RN2483 to the host system. The available commands listed in [41] can be used to configure and control the device. There are three distinct types of commands when communicating with the RN2483:

- *mac* commands to configure and control the LoRaWAN stack
- *radio* commands to directly access or update the transceiver's setup
- *sys* commands to issue system level behavior actions

Although both *radio* and *mac* commands can be used to transmit LoRa packets, the use of *mac* commands is preferred to make sure the LoRaWAN specifications are honored when transmitting packets or altering transmission settings.

Combining the RN2483 module with an Arduino and adding a suitable antenna would yield an end-device capable of transmitting LoRaWAN compliant packets. Such an end-device is shown in Figure 11, in which an RN2483 module is mounted on top of a custom Arduino shield equipped with an antenna connector and accompanying antenna that is attached to the top of an Arduino Uno. This end-device has been used to conduct experiments described in this thesis.

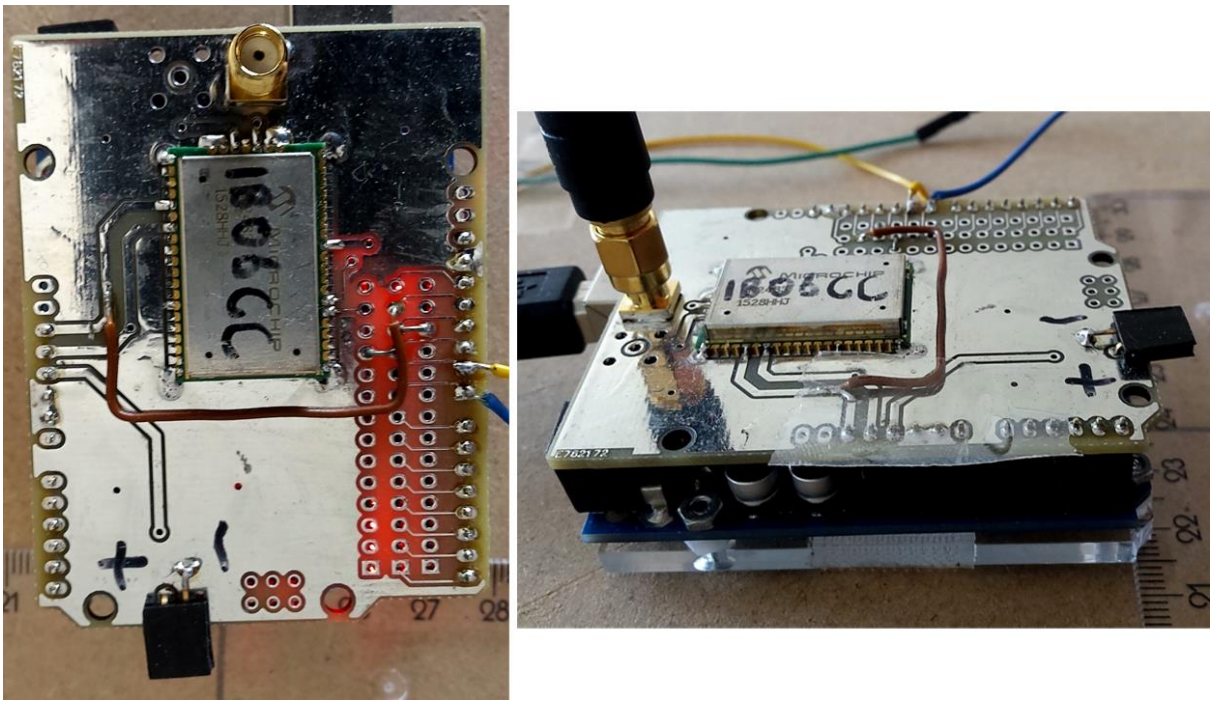

Figure 11. The RN2483 module mounted on a custom Arduino shield and attached to an Arduino.

As described in Section 3.4, an end-device that implements the LoRaWAN protocol has five TX power settings defined. The output power of the RN2483 module for each of these settings, according to its datasheet [40], is shown in Table 8 along with the corresponding current draw.

LoRaWAN TXPower setting	Corresponding TX power (dBm)	RN2483 Output power (dBm)	Supply current at 3V (mA)
1	14	13.5	38.0
2	11	10.4	33.7
3	8	6.9	30.0
4	5	3.6	26.1
5	2	0.4	22.3

Table 8. Theoretical LoRaWAN transmit power settings and corresponding output power and current draw of the RN2483 module

4.1.2 The gateway

The receiving node of the network is a MultiConnect Conduit MTCDT-H5-210L gateway manufactured by Multitech [42], which is shown in Figure 12. The gateway is designed to listen on the 868MHz band for incoming messages [43], and therefore requires all end-devices to support the three default channels (0, 1 and 2, i.e. 868.1, 868.3, 868.5 MHz) that must be implemented in every EU868MHz end-device to comply with the LoRaWAN specifications. No additional channels are created.

The gateway is capable of retrieving various information from each packet received by the network server. This includes:

- node's address or id;
- gateway's own address;
- channel (and corresponding frequency) at which the packet was transmitted;
- used Coding Rate
- payload (the contents of the FRMPayload segment as described in Section 3.2.3);
- used Data Rate
- Signal-to-Noise Ratio (SNR) of the received packet
- Received Signal Strength Indicator (RSSI) of the received packet
- sequence number (a packet counter starting from 0 after each network join)
- timestamp of when the packet was received by the gateway

The parameters listed above provide the information needed to identify packets, the end-device that transmitted them, the conditions under which they were sent (data rate, coding rate, and frequency), and to determine how well the radio signal was received by the gateway (RSSI, SNR). Due to the nature of LoRa modulation, which is a spread spectrum modulation as described in Section 2.4, it is possible to receive signals below the noise floor. Therefore, it is possible to obtain negative SNR values [38].

Figure 12. Front view of the Multitech conduit gateway

4.2 SPECTRUM ANALYSIS

As described in Section 2.4, LoRa uses Chirp Spread Spectrum modulation to encode the transmission data. The best way to visualize this technology is to use a waterfall plot that is centered on one of the three channels on which LoRa operates. With the use of a Realtek DVB-T Dongle [44] it is possible to listen to (parts of) the radio spectrum. When combining the dongle with signal processing software, a software defined radio (SDR) spectrum analyzer can be created.

To analyze a LoRa channel visually, we create a waterfall plot using GNU Radio Companion (GRC), i.e., a graphical application based on an open-source development toolkit to create SDRs and signal processing systems with the use of basic signal processing blocks [45]. In this case, a relative simple GRC application is needed to create an SDR that tunes to the first LoRaWAN channel (868.1 MHz, 125 KHz wide) and generates a waterfall plot as shown in Figure 13.

In this waterfall plot, the first symbols (the preamble) of a LoRa packet transmitted on spreading factor 12 are visible. The time on the right side scale makes it possible to calculate and verify the symbol time T_s, denoted below. Based on the length of the time scale and the length of each chirp visible in the waterfall plot, the symbol time should approximately be 32.75 milliseconds. This value matches the symbol time calculated with the definition of symbol time as described in 3.3:

$$T_s = \frac{2^{SF}}{BW} = \frac{2^{12}}{125 \times 10^3} = 32.77 \text{ milliseconds} \quad (4.1)$$

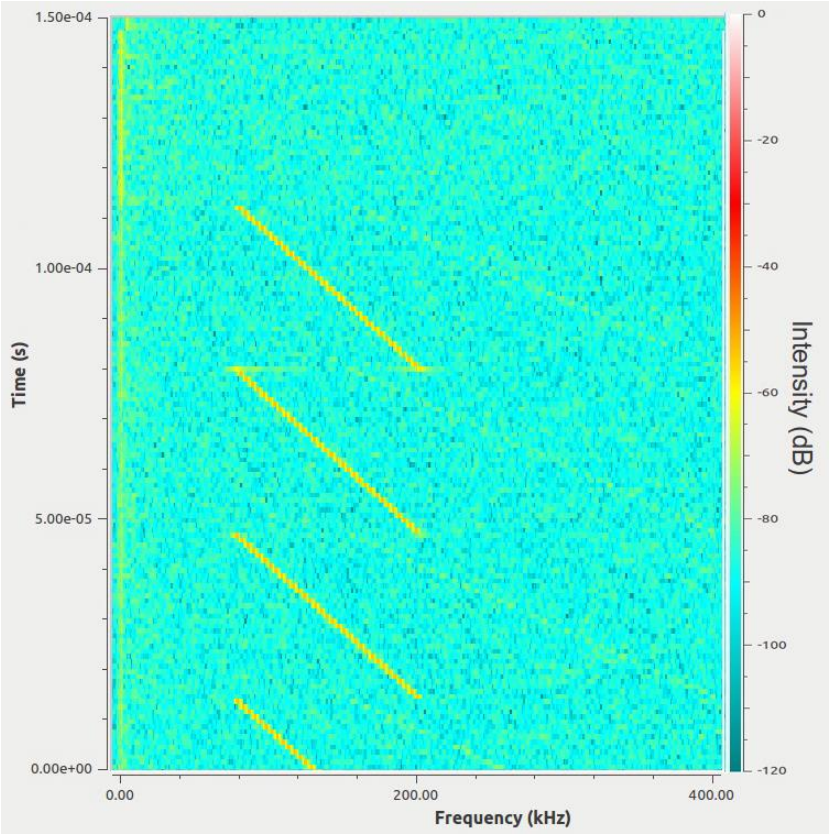

Figure 13. Waterfall plot of the spectrum around 868.1 MHz

The full LoRa packet, of which the beginning is shown in the image above, can be seen in Appendix A. The payload added to this packet is one byte and the used coding rate is 4/5. Due to software limitations, it is not possible to fit all the symbols of one packet in one frame of the waterfall plot. The images in Appendix A have therefore been constructed from a sequence of snapshots from the waterfall plot. The total number of symbols that can be distinguished in this image is 35.25. This corresponds to the value that can be calculated with Equations (3.4) and (3.7), and the default LoRaWAN values for the CRC, IH and DE parameters defined in 3.3. When substituting PL for 1, SF for 12 and CR for 1, the following calculations are obtained:

$$n_{\text{payload, LoRaWAN}} = 8 + \max\left(\text{ceil}\left[\frac{8\ 13 + \text{PL}\ \ - 4\text{SF} + 28 + 16\text{CRC} - 20\text{IH}}{4 \times\ \ \text{SF} - 2\text{DE}}\right]\ \text{CR} + 4\ \ , 0\right) \tag{4.2}$$

$$n_{\text{payload, LoRaWAN}} = \ \ 8 + \max\left(\text{ceil}\left[\frac{8\ 13 + 1\ \ - 4 \times 12 + 28 + 16 \times 1 - 20 \times 0}{4 \times\ \ 12 - 2 \times 1}\right]\ 1 + 4\ \ , 0\right) \tag{4.3}$$

$$n_{\text{payload, LoRaWAN}} = 8 + \max\left(\text{ceil}\left[\frac{108}{40}\right] \times 5, 0\right) = 8 + \max\ 3 \times 5, 0\ \ = 8 + 15 = 23 \tag{4.4}$$

$$n_{\text{total}} = \left(n_{\text{preamble}} + 4.25\right) +\ \ n_{\text{payload, LoRaWAN}} = \ \ \ 8 + 4.25\ \ + 23 = 35.25 \tag{4.5}$$

The packet shown in Appendix A also shows that the preamble consists of ten up-chirps and ends with two and a quarter down chirps. The subsequent symbols represent the header, payload data, and payload CRC respectively. These are all up-chirps, starting at an arbitrary frequency within the channel's bounds, which continue at the channel's lower bound once they reach the upper bound.

4.2.1 Time on Air for acknowledgement packets

When an end-device transmits a *confirmed* message, it has to be acknowledged by the network if the message is correctly received. In this case, our gateway responds with a downlink acknowledgement message. Since the gateway is also restricted by duty cycle regulations, it is interesting to see the size of such an acknowledgement. To determine this, the same approach as described above can be used to determine the acknowledgement packet's length. A full captured acknowledgement is also shown in Appendix A. This packet was received during the second-receive-window and is therefore sent on the channel around 869.525 MHz and at spreading factor 12. The acknowledgement consists of 35.25 symbols, which does not change if the length of the uplink message changes. This could be expected, as an acknowledgement is a packet with zero payload. By calculating the total symbols of a transmission of a packet with zero payload using Equations (3.4) and (3.7), a result of 35.25 symbols is obtained.

An interesting difference between the uplink data message from the end-device and the downlink acknowledgement from the gateway is that the latter consists of down-chirps (except the last 2.25 symbols of the preamble) instead of up-chirps. This mirrored behavior between up and downlink messages is only briefly mentioned as 'IQ inversion' by Semtech in its recommendations for EU868 LoRaWAN network operation settings [46]. A possible explanation for this implementation might be that this allows gateways and end-devices to distinct between up and downlink messages. Gateways will not try to receive messages from other gateways and end-devices can ignore transmissions from other end-device during their receive-windows.

The scarce information from the LoRaWAN specification does not provide exact information about the contents of an acknowledgement packet, except that the acknowledgement bit is set in the FCtrl field shown in Figure 6, and that downlink packets do not carry a payload CRC. Based on this information and the length of the received acknowledgement packet, such a packet most likely consists of a preamble, a header (+header CRC) and a payload containing no additional data (so FRMPayload is of zero length) and no payload CRC.

With this observation, the length of the acknowledgement packet transmitted on other data rates can be determined together with the corresponding Time on Air. This information provides an indication of the maximum number of acknowledgement packets a gateway can transmit within the duty cycle regulations for the two sub-bands in which the default communication channels are specified. Table 9 shows these maxima for each data rate, based on a 1% duty cycle limitation.

Data Rate No.	Symbols	Time on Air (ms)	Maximum acknowledgements per hour
5	40.25	41.2	873
4	40.25	82.4	436
3	35.25	144.4	249
2	35.25	288.8	124
1	35.25	577.5	62
0	35.25	1155.1	31

Table 9. Time on Air for acknowledgement packets and corresponding maximum packets per hour, per data rate

The additional channel in the G3 sub-band, which allows a 10% duty-cycle, enables the gateway to acknowledge an additional 311 packets in total, regardless of the received packet's data rate. When only these default channels are implemented, a gateway can acknowledge at most 1184 (873+311) packets per hour, and in the worst-case scenario only 342 (31+311). In a more fair scenario in which nodes transmit at equally distributed data rates, over 400 packets (almost 16 on each DR plus the additional 331 on the G3 sub-band) can be acknowledged per hour.

4.3 DUTY CYCLE RESTRICTION IMPLEMENTATION

According to the RN2483's Command Reference User's Guide [41], the module sets a duty cycle limit for the LoRaWAN protocol layer on each enabled channel individually. The default LoRaWAN setup uses three channels in the ISM G1 band to transmit data, which therefore have a duty cycle limitation of 0.33% each. This means that, to comply with the duty cycle regulations, the three default channels (0-2) are set with a default duty cycle of 0.33%. If new channels are created, all channels must be updated by the user in terms of duty cycle to comply with the ETSI regulations, including the default ones. The RN2483 module calculates if a transmission complies with the duty cycle restriction based on the current duty cycle settings for each channel's duty cycle. If this check fails, the transmission is rejected and the module responds with a *no free channel* notification via the UART command interface.

It is worth noting that although the duty cycle limitation applies to the whole G1 band, the RN2483 module calculates this restriction for each channel separately. This means that three packets can be transmitted directly after each other without violating the duty cycle limitation implemented on the RN2483 as long as these packets are transmitted on the three different channels. If the transmitter wants to send out a fourth packet, it has to wait for the first channel (with a duty cycle of 0.33%) to be available again. This behavior is, however, not compliant with the LoRaWAN per sub-band enforced duty cycle limitation as described in Section 3.1.2.

4.4 VERIFYING THEORETICAL PERFORMANCE INDICATORS

The following measurements are performed to verify the theoretical assumptions about the time between packets and the corresponding throughput as described in Section 3.5, as well as the implementation of the duty cycle limitation by the RN2483 module. On each spreading factor, 200 packets with a fixed payload length of 64 bytes and coding rate 4/5 are transmitted from an end-device to a nearby gateway.

4.4.1 Time between packets

The total time difference between the first and the last packet divided by the total number of transmitted packets in that time, gives the average minimum time between each packet for each spreading factor. Figure 14 shows the measured average time between the series of packets for each DR (in red), and the theoretical time between packets (in blue). The measured average times between packets lies very close to the calculated theoretical values as described in Section 3.5, as can be seen in the figure above.

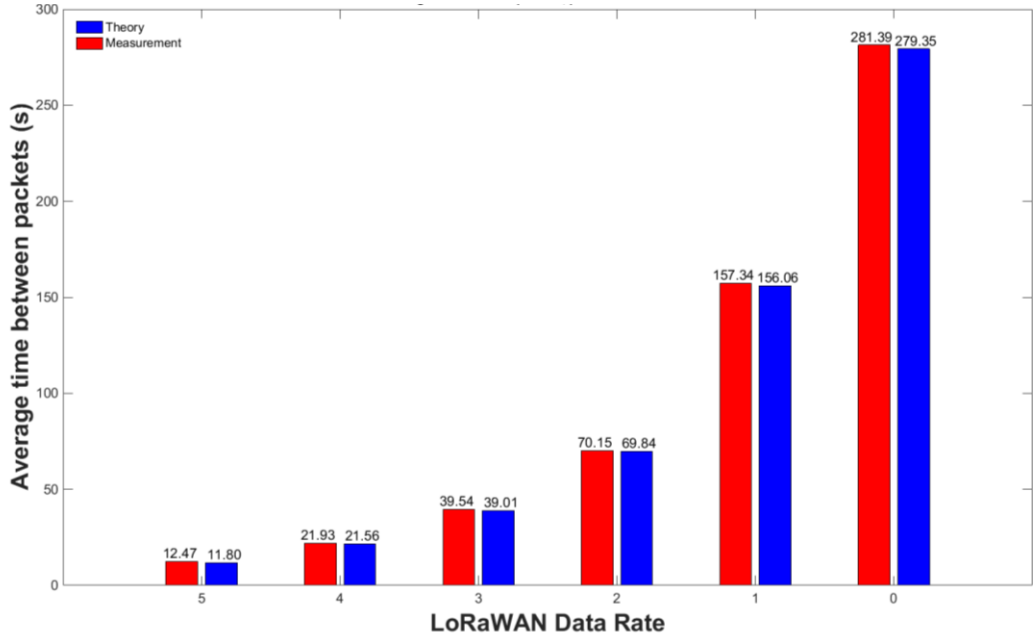

Figure 14. Theoretical and measured time between packets for each LoRaWAN data rate.

4.4.2 Transmission power and energy consumption

The power of a transmitted signal can be measured with the use of a spectrum analyzer[1] to check if the module is mounted properly on the shield and to verify the LoRaWAN transmission power settings. The transmission power is measured by connecting the spectrum analyzer to the antenna socket of the module's shield. The results are shown in Table 10 for each supported output power setting as defined in the LoRaWAN specification. The same measurement is repeated for the other end-devices to verify similar transmission powers between different end-devices. The other end-devices show similar results. The results from this measurement also show how the three channels are positioned in the spectrum, as can be seen in the images included in Appendix B. Each channel is somewhat centered around the frequency at which the channel is defined. This means that, in the case of a 125 KHz wide channel, the channel starts at 75 KHz before, and ends at 50 KHz after its defined channel frequency.

[1] A Hewlett Packard 8594E spectrum analyzer was used.

TXPower setting	Defined TX power (dBm)	RN2483 Output power (dBm)	Measured TX power (dBm)	Corrected for dipole gain (dBm)
1	14	13.5	11.42	13.57
2	11	10.4	8.33	10.48
3	8	6.9	4.75	6.90
4	5	3.6	1.17	3.32
5	2	0.4	-1.92	0.23

Table 10. Measurement results of the RN2483's transmission power when using LoRaWAN and the Arduino shield.

There are some differences between the measured values and the specified output power, which can be explained by how the ETSI regulations define the maximum radiated power for the frequency bands. The transmission power should not exceed an EIRP of 25mW (14 dBm), as described in Section 2.4.2. To obtain a maximum transmission power for the case in which a dipole antenna is used, the EIRP limitation should be converted to an Effective Radiated Power (ERP) measure. ERP refers to the radiation of a half wave tuned dipole instead of an isotropic antenna. Since a half-wave dipole is more representative of realistic simple antennas than an isotropic radiator, its gain (2.15dBm) has to be taken into consideration when supplying a signal to an antenna [47]. The transmission powers defined in the LoRaWAN specification therefore include the antenna gain for a dipole antenna, which makes it likely that the RN2483 module outputs 2.15 dBm less. When corrected for this, we obtain the values shown in the last column in Table 10. These values approximate the values stated in the RN2483's datasheet.

Although this thesis is not focused on optimizing the energy consumption of an end-device to make it last as long as possible, it is meaningful to get an indication of its power consumption. The energy consumption pattern of an end-device can be useful in future when end-devices need to consume as little power as possible. The energy pattern can help determine the lifetime of the end-device's battery or find an optimal transmission sequence to obtain the desired battery lifetime. To obtain the current consumption of an end-device when it is powered with a 3.7V Li-ion polymer battery, an Agilent 34401A digital multimeter is connected between the end-device and the battery. The end-device is programmed to join the LoRa network with OTAA and, once accepted into the network, the device will transmit an unconfirmed packet[2] on each possible LoRaWAN transmission power and will switch to a deep sleep mode for 10 seconds afterwards.

The current consumption of the end-device executing the transmission pattern described above is shown in Figure 15. This graph shows the typical behavioral pattern of a LoRaWAN class A device when joining a LoRa network and transmitting packets. After each uplink transmission, the two downlink receive-windows which characterize a LoRaWAN class A device are opened. The second receive-window would not have been opened if the end-device received a downstream packet during its first receive-window. This can be seen after the join request, where only one

[2] The packet's payload is 64 bytes, the coding rate is set to 4/8 and the transmission is done on DR 0, to maximize the Time on Air.

receive-window is present. This means the end-device received a packet from the gateway confirming the join request, which makes opening a second receive-window redundant. The average current draw when transmitting can be seen in Table 11. This table also shows the transmission current draw subtracted by the average current draw of the Arduino and the RN2483 module when it is not transmitting in the last column. In addition, the join procedure consists of a transmission at the maximum output power. Lastly, the current draw during a receive-window is not influenced by the transmission power. On average, the total current draw equals 45.5 mA during receive-windows, and 11.4 mA when removing the current consumption of the Arduino.

The measured values are slightly higher then could be expected based on the values stated in the RN2483 its datasheet. This is most likely due to the higher supply voltage (3.3 instead of 3.0V) we use to power the RN2483 module. Knowledge of the energy consumption can help determine the battery life based on transmission schemes of end-devices and the available battery power. Optimization based on these results can mostly be achieved by using a less consuming device to mount the LoRa module on, i.e. replace the Arduino with a less energy consuming (dedicated) host device, and by putting the LoRa module in sleep mode when no transmissions are needed.

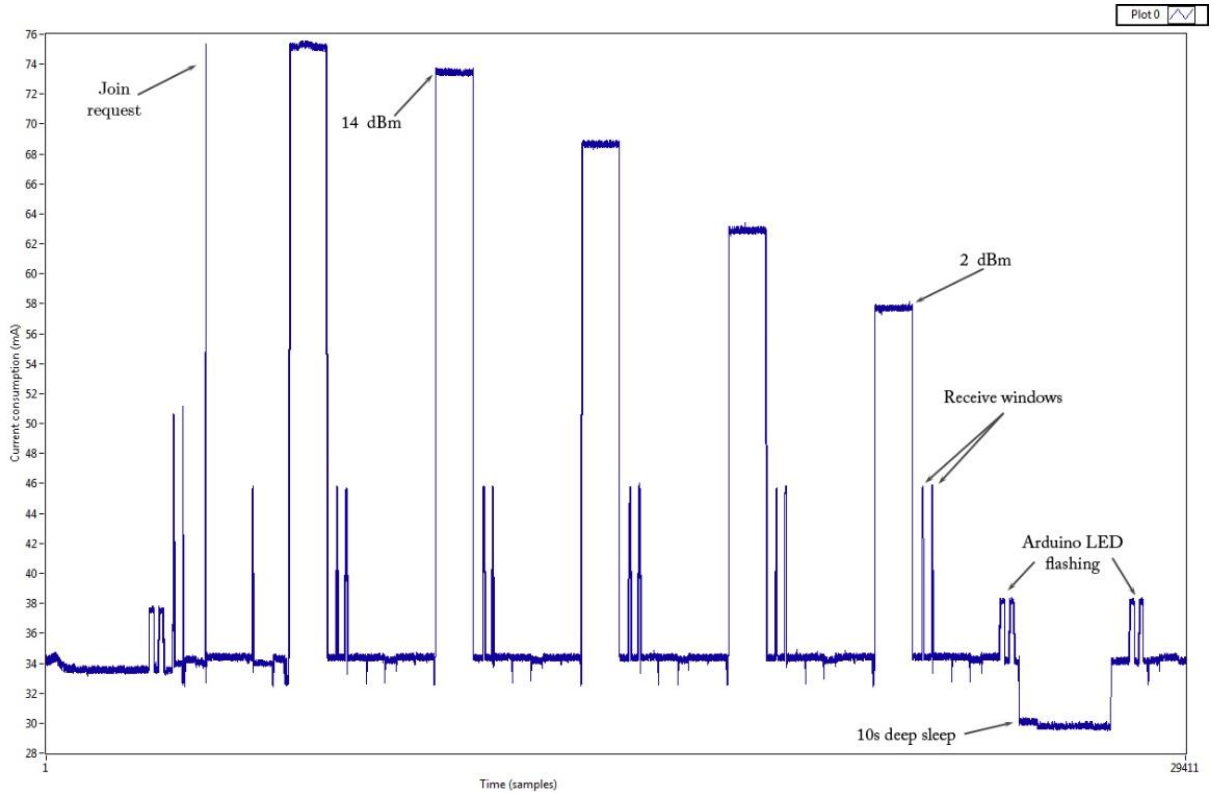

Figure 15. Current usage of an end-device when joining a network and transmitting 5 packets

LoRaWAN TXPower setting	Current from datasheet at 3V (mA)	Measured current consumption, RN2483 + Arduino (mA)	Measured current consumption RN2483 (mA)
1	38.0	75.2	41.1
2	33.7	73.4	39.3
3	30.0	68.7	34.5
4	26.1	62.9	28.8
5	22.3	57.7	23.6

Table 11. RN2483 theoretical current consumption and corresponding measured current consumption

The obtained energy consumption profile for this end-device can serve as a basis for getting an indication of the battery life of an end-device when its transmission power, data rate, packet characteristics and transmission interval are known. Due to the energy consumption of the Arduino itself, the end-device created in this thesis will at most last a few days on batteries. If required however, there are ways to decrease its energy consumption to a very minimum (in the order of tens of micro-amperes) by building a custom Arduino and putting it in sleep mode between transmissions [48][49]. Such a device would be able to last several years on only two alkaline AA batteries in a scenario in which it (i) transmits a packet once every ten minutes, (ii) with a 64-byte payload (iii) on DR 5, and (iv) with a transmission power of 14 dBm. In such a scenario it is assumed that the device only consumes 40 µA when in sleep mode and no additional (power consuming) tasks need to be performed. Devices that require more frequent transmissions or have to read and process data from a sensor before a transmission will of course have a shorter lifetime.

The measurements conducted in this section verify both the specified transmission power and energy consumption for the RN2483 module. Moreover, the image in Figure 15 reveals the typical energy consumption profile of a class A LoRaWAN end-device. The measured values approximate the values stated in the RN2483's datasheet. By calculating the Time on Air of packets a device transmits, an indication of the energy consumption for each transmission can be obtained for each output power setting. The energy consumption of the RN2483 during receive-windows was not mentioned in its datasheet. Results found in this section therefore allow a more precise calculation of the energy consumption during and after a transmission.

5 PERFORMANCE ANALYSIS OF LORA IN INDOOR ENVIRONMENTS

LPWAN technologies such as LoRa are designed to facilitate a large variety of IoT applications. A good number of these applications are situated in an indoor environment, e.g. in the fields of home automation, flex offices, and health monitoring. It is therefore important to know the performance of LoRa in indoor environments.

5.1 EXPERIMENT SETUP

To get an indication of LoRa performance inside a building, we place a LoRaWAN gateway on the top floor of an office building located at the University of Twente. An end-device, as described in Section 4.1.1, is then positioned on a number of different locations on the same floor to evaluate the communication link performance in terms of packet loss and SNR by varying the transmission power and the data rate. The location of the gateway inside the building as well as the location of the five end-devices are shown in Figure 16. The direct distance from the gateway to the end-devices is approximately 8, 12, 30, 70 and 102 meters for the locations marked as Corridor 1, Office, Staircase 1, Corridor 2, and Staircase 2 respectively. The total length of the corridor is 103 meters.

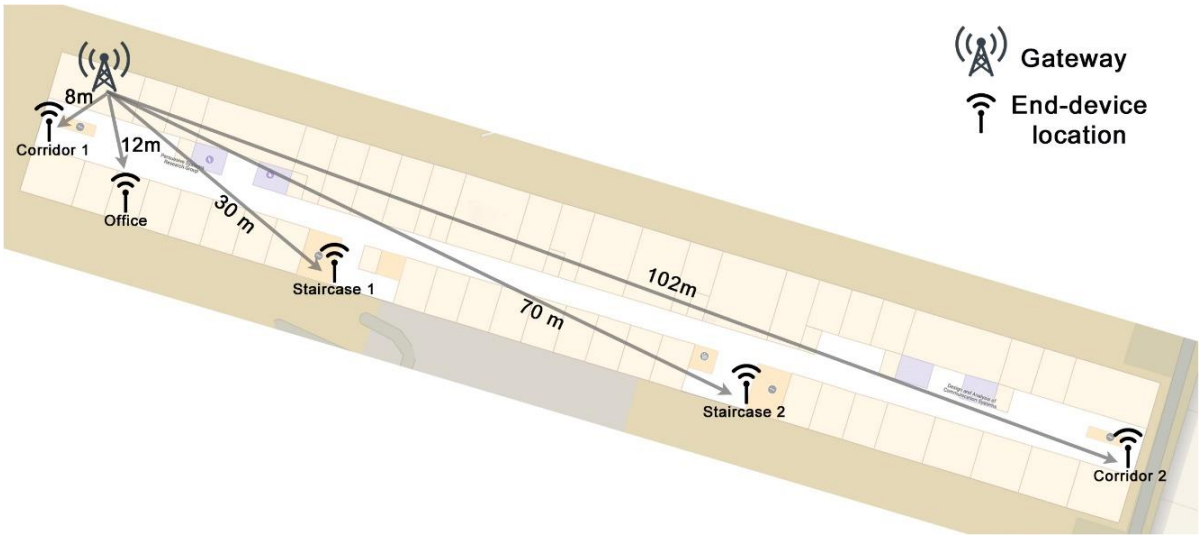

Figure 16. Measurement locations inside an office building.

5.2 VARYING LOCATION, TRANSMISSION POWER AND CODING RATE

In the first experiment, we transmit 15 packets on each coding rate, transmission power and three LoRaWAN data rates: DR 5, DR 3 and DR 1 at each transmission round. After every 15 packets, the coding rate is increased until the maximum is reached. Then the coding rate is set back to the minimum and the transmission power is decreased by one. In other words, every set of 15 packets is transmitted with the same configuration and every set of 60 packets is sent on the same transmission power with increasing coding rate increases. The end-devices were placed at the two corridor and two staircase locations on the map shown in Figure 16. The end-devices were programmed to start at DR 5, coding rate 4/5, and transmission power of 14 dBm. After each 15 packets, the coding rate was increased until the maximum of 4/8 was reached. Then the

transmission power was lowered one step and the same process was repeated. The same was done for DR 3 and DR 1. To avoid interference, simultaneous transmission of two (or more) end-devices was not allowed throughout the entire experiment.

The packet loss for each combination of coding rate, transmission power, and DR per location can be found in Appendix C. These results show that the location inside a building has a huge difference on performance of a LoRa end-device. Looking at these results, we see that the most notable results are obtained from the end-device located at the second staircase. Where all the other locations show a worst-case packet loss of 2 packets on any combination of variables, the packet loss on the second staircase location is much higher. The packet loss from the end-device in the second staircase transmitting on DR 5 is shown in Figure 17. This figure shows how varying the coding rate and the transmission power influence the throughput.

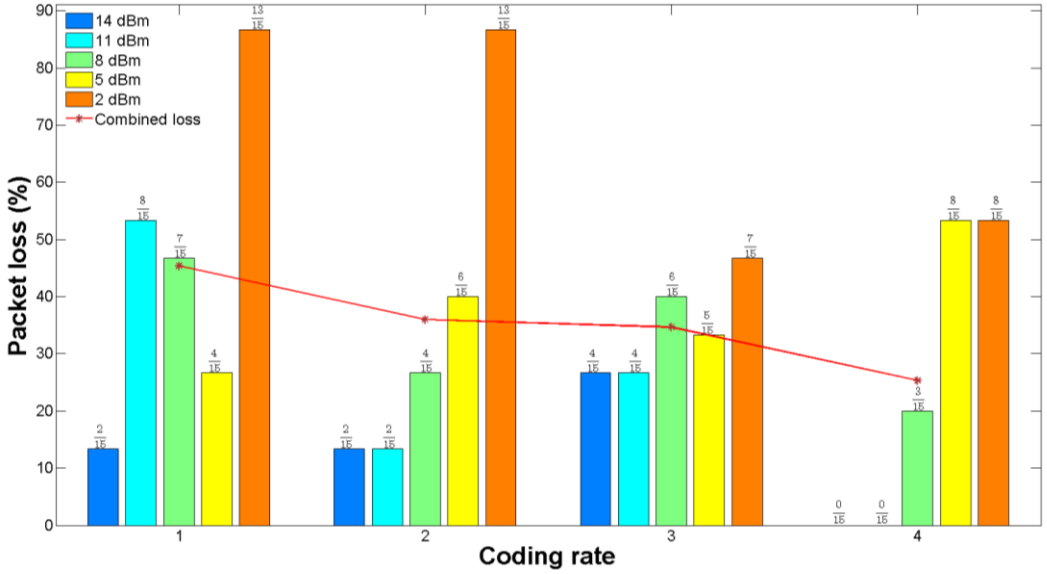

Figure 17. Packet loss when varying the transmission power and coding rate

Transmitting packets with a higher transmission power or with higher coding rates increases the chance the packet will be successfully received. When the data rate is decreased (resulting in an increase of the spreading factor), the packet loss decreases as well. This behavior is expected from spread spectrum technology as signals with higher spreading factors are more likely to reach a gateway.

There are, however, some notable differences in packet loss between the end-devices, as well as in the SNR of the received packets. Some end-devices show a slightly higher packet loss at higher coding rates or transmission powers, which contradicts the expected lower packet loss at higher coding rate and transmission powers. The numerical differences between the lost packets on different transmission settings are, however, small and might have been caused by changes in the environment. This is also observed when we look at the SNR values of the individual packets in Figure 18, which shows all packets transmitted on DR 3. The graph starts with packets sent with the minimum coding rate and the maximum transmission power. Figure 18 shows that the average SNR decreases when the transmission power decreases. This is as expected. Figure 18 also shows

many SNR fluctuations, even when identical configurations are used. This is most likely caused by the changing environment throughout the day in an office building, e.g. people moving around and doors opening and closing. We will look more closely into these influences in the next section.

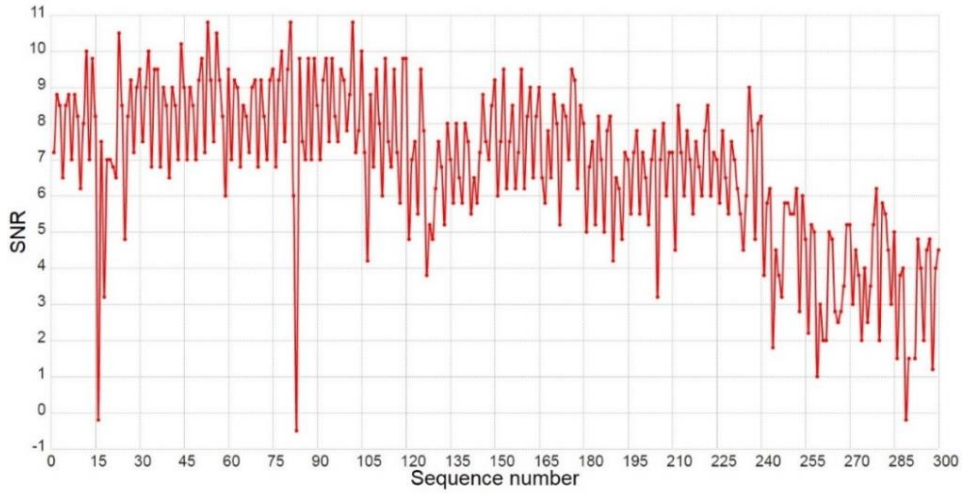

Figure 18. SNR values of individual packets. Every set of 15 packets is transmitted with the same coding rate and transmission power.

5.3 ENVIRONMENTAL INFLUENCES

The experiment described above led to the presumption that dynamics of an indoor environment can influence the performance the LoRa network. To investigate this further, an end-device located at the office location as shown in Figure 16 is configured to transmit a packet on DR 3 and with a payload of 64 bytes every 60 second. This measurement started at 4:26PM on a Friday afternoon when there were still people present in the building and lasted for five hours. As it can be seen in Figure 19, the SNR of the transmitted packets stabilizes after most people leave their offices. This behavior becomes even clearer by comparing the SNR of packets transmitted during the evening and night, and those transmitted during the day. Figure 20 shows this comparison. The left graph shows SNR values of packets transmitted between 6:30PM and 1:10AM whereas the right graph is based on packets sent between 7:50AM and 14:30PM. We conclude that the performance of LoRa technology is affected significantly by changes in the environment. Due to this fact, we perform the other indoor measurements during the evening and night or during the weekends to reduce these changing environmental influences.

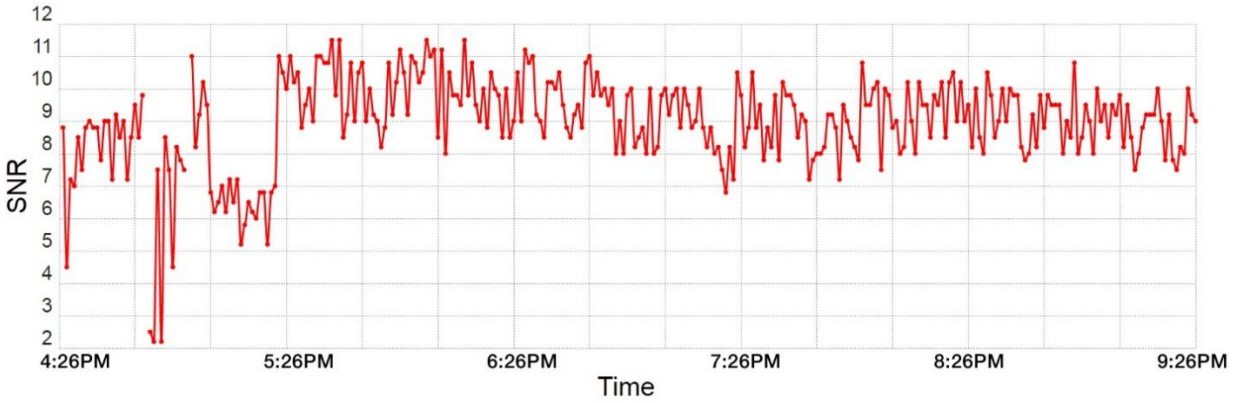

Figure 19. SNR of each packets received over time.

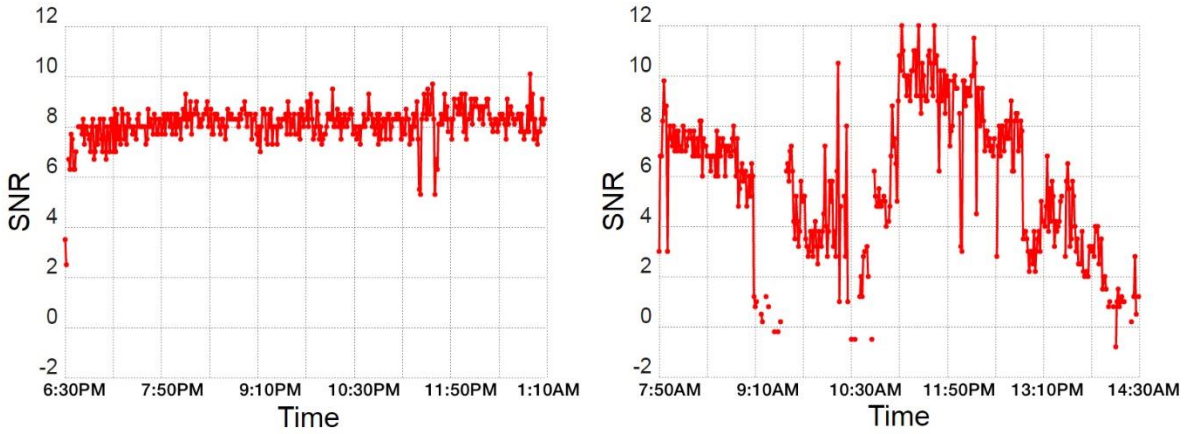

Figure 20. Left) SNR of packets tranmitted during the evening. Right) SNR of packets tranmitted in the morning

5.4 Interference from other LoRa end-devices

Aside from environmental influences, the performance of end-devices might also be influenced by other end-devices. LoRa is designed to be able to receive signals when they are transmitted at the same time but at different spreading factors. This ability verified by synchronizing three end-devices and letting them transmit packets at the same time and channel. Two of the three default LoRaWAN channels (channel 1 and channel 2) are switched off on the RN2483 module to force transmission on the same channel (i.e. channel 0). The experimental setup can be seen in Figure 21. The end-devices are connected to an Arduino, which requests a transmission from the devices every minute. The resulting transmission profile is shown in Figure 22. The end-devices transmit packets with a length of 64 bytes, which take 118, 216, and 390 milliseconds, respectively to transmit depending on the data rate. In this configuration, all packets are received which means that end-devices do not interfere with each other and simultaneous transmission of LoRa packets on different data rates is possible. The influence of the other devices can be noticed however by looking at the SNR graph of the three different transmissions in Figure 23. Since other signals at different spreading factors appear as noise to the receiver, the SNR values deteriorate in the presence of other LoRa modulated signals. The results from the described experiment have confirmed the behavior that could be expected based on the characteristics of LoRa modulation described in Section 2.4.

Another interesting configuration to evaluate is the performance of LoRa when two devices are transmitting at the same data rate (i.e. spreading factor) and at the same time, while transmission powers and timing differ between the devices. The question will be if the gateway is still capable of receiving one or even both packets. Other interesting situations occur when the two end-devices transmit their packets with different transmission powers or when their transmission is not completely synchronous and one device starts transmitting before the other. If receiving both packets is not feasible for the gateway, it would be interesting to see if the gateway prefers the stronger signal over the weaker one, even if it started receiving the weaker signal earlier. If the gateway is capable of receiving at least one of the overlapping packets, i.e. it locks onto one of them, this is known as the capture effect. This phenomenon and the experiments conducted to examine its existence are explained and analyzed in the next section.

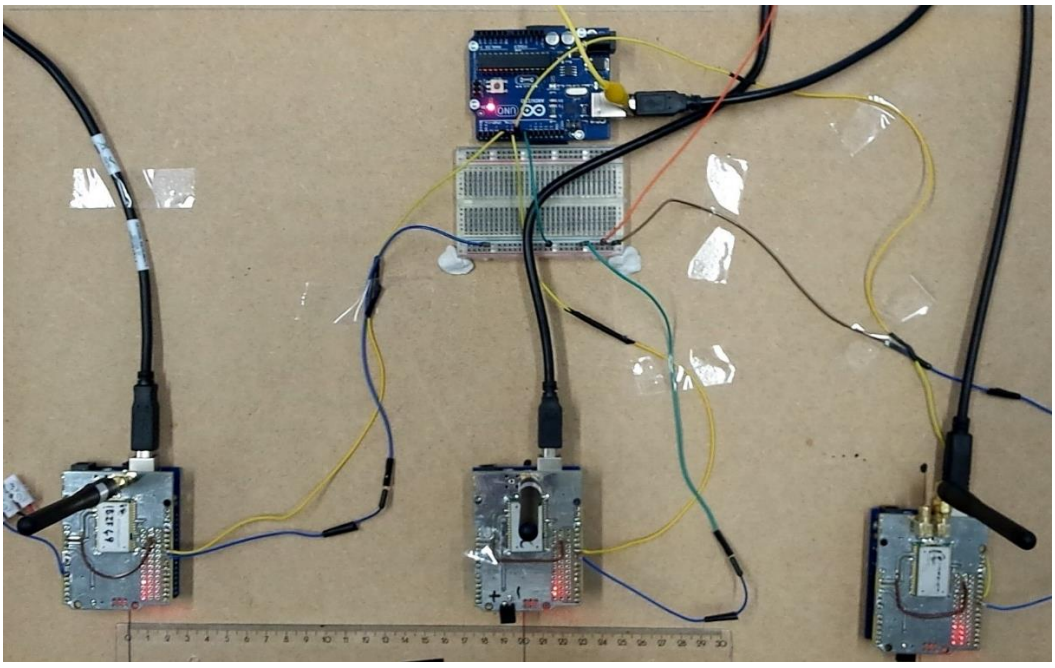

Figure 21. Measurement setup for three end-devices transmitting simultaneously

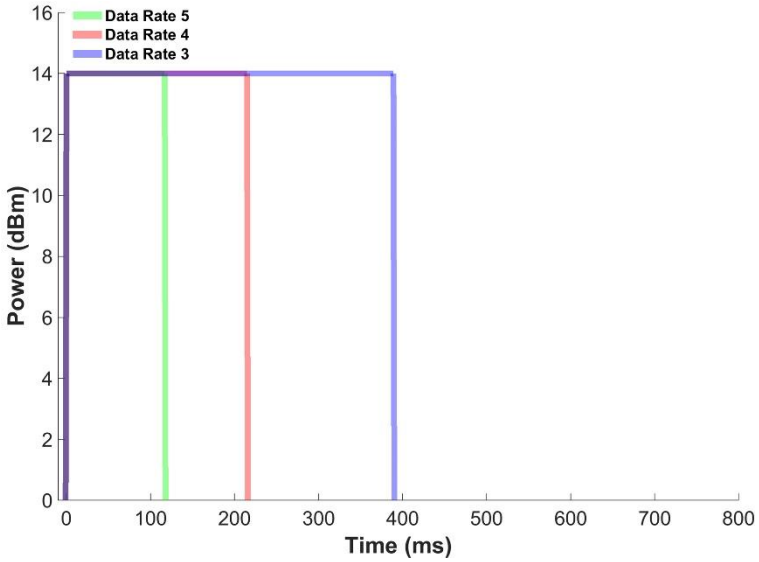

Figure 22. Transmission profile of three end-devices when transmitting on DR 5, 4 and 3

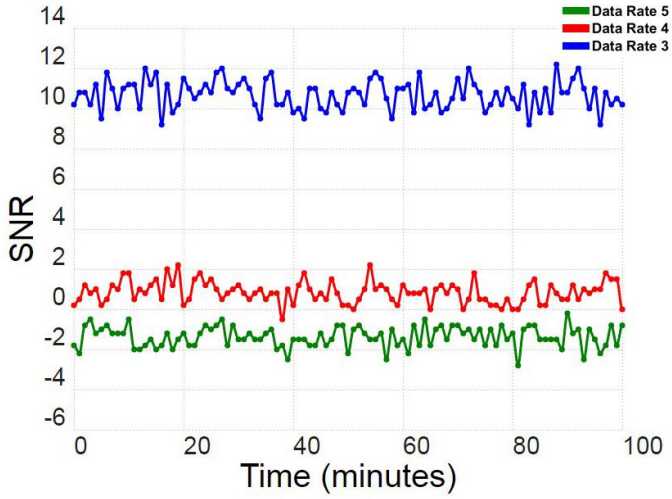

Figure 23. SNR of 100 packets from three end-devices transmitting at different data rates

5.5 THE CAPTURE EFFECT AND LORA

The capture effect is the ability of a radio to correctly demodulate a packet from at least one transmitter in the presence of overlapping packets from other transmitters, provided that the difference in signal power is sufficiently large [50] [51]. This means that the receiver is capable to lock onto one packet, which usually is either the first arriving signal or the strongest in the presence of interference. This packet has then captured the receiver [50]. Minimum signal power differences between the stronger and weaker signals required for a capture effect to occur differ per modulation type and range between 0.17dB and 3dB [52].

5.5.1 Experiment method

To determine if a capture effect can take place in a LoRa network, instead of transmitting at different data rates, all packets are sent at the same data rate (DR 3), which corresponds to spreading factor 9, and with a payload of 64 bytes and Coding Rate 1. Two end-devices have been used to transmit simultaneously, first at the same power (14 dBm) and then with a difference of 6 dBm between transmission powers. The experiment has been repeated more than ten times and with different end-devices, all showing similar results as shown in Table 12.

End-device	Tx power	Packet loss
1	14 dBm	0 %
2	14 dBm	99 %

End-device	Tx power	Packet loss
1	14 dBm	95 %
2	8 dBm	2 %

Table 12. Packet loss from two end-devices transmitting 1000 packets simultaneously every minute with left) the same transmission power and right) a 6 dBm difference in transmission power

The gateway is clearly only able to successfully receive one packet at a time and appears to prefer the stronger signal when two packets are transmitted simultaneously. To investigate the latter assumption more thoroughly, the experiment with end-devices transmitting at different transmission powers is repeated with the difference that one of the transmissions is delayed. This causes the gateway to receive one of the packets before the other. Cases in which the stronger signal arrives before the weaker are referred to as a *stronger-first* situation and cases in which a weaker signal arrives first are indicated as a *stronger-last* situation [52].

Both the weaker and the stronger transmitting devices are delayed for 100, 200, and 300ms relatively from each other. This is shown by the transmission profiles of the two nodes for this configuration in Figure 24 and Figure 25. An additional 400ms delay is introduced to verify a correct transmission and reception of both packets when they do not overlap (this is not depicted in the figures since this situation does not create an overlap between packets). If a capturing effect takes place at the gateway, the expected result would be a low packet loss for the stronger signal and a high packet loss for the weaker signal in the stronger-first situations. For stronger-last situations, a high packet loss is expected for packets from both end-devices.

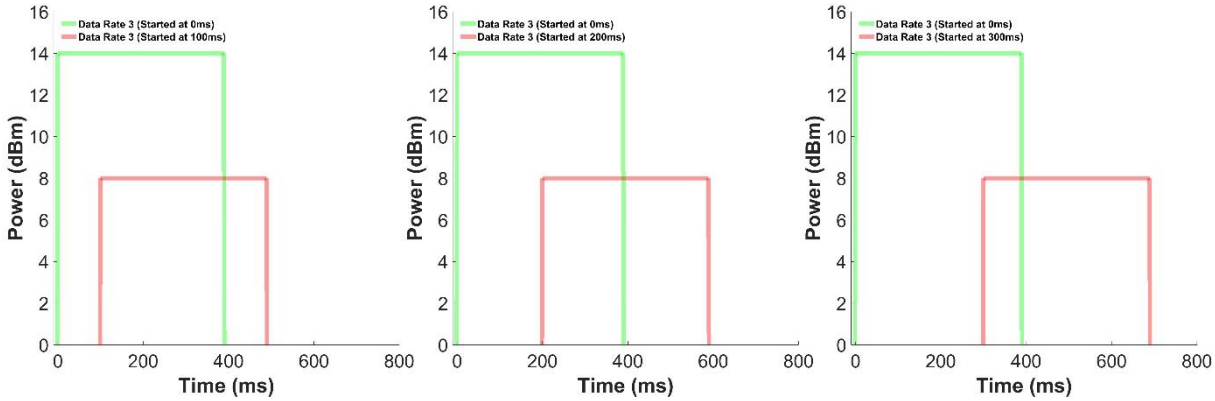

Figure 24. Transmission profiles to create stronger-first scenarios by delaying the weaker signal 100ms (left), 200ms (middle), 300ms (right)

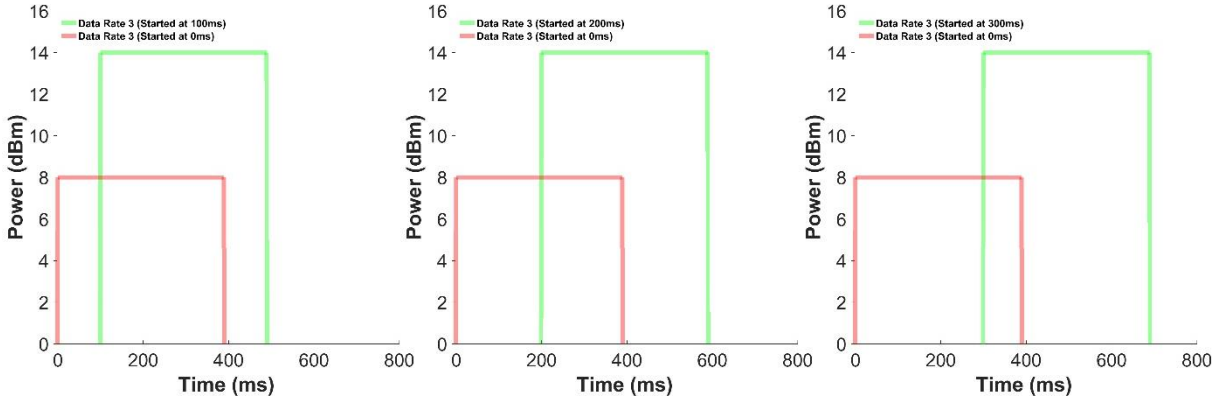

Figure 25. Transmission profiles to create stronger-last scenarios by delaying the stronger signal 100ms (left), 200ms (middle), 300ms (right)

The results of this experiment shown in Figure 26 indicate the presence of a capture effect. Packets originating from the end-device transmitting at the higher power are only received successfully in stronger-first conditions. In the stronger-last case, both packets are lost. This behavior does not change when the configuration between the two end-devices is interchanged or when a different data rate is used.

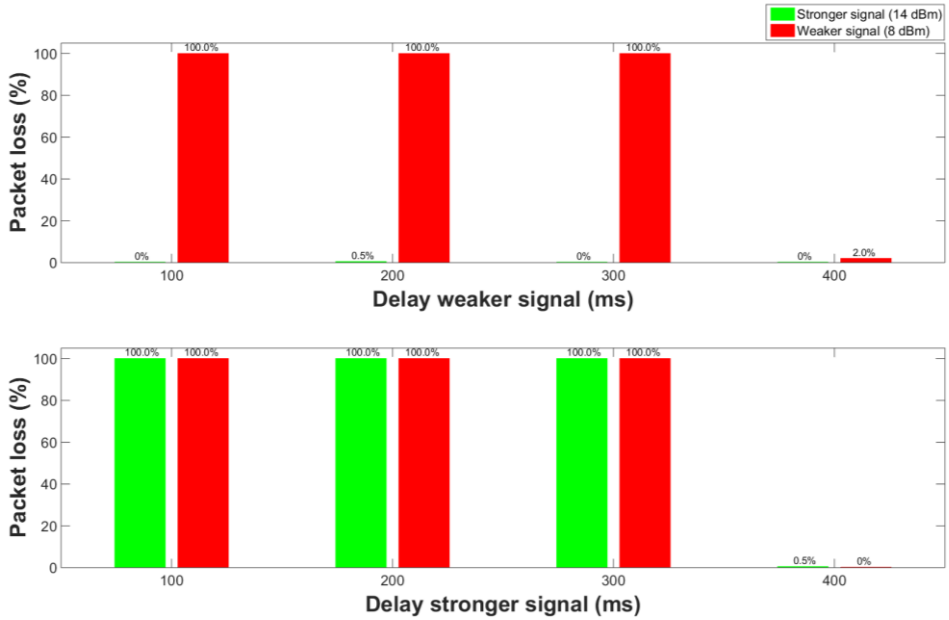

Figure 26. Packet loss in the a) stronger-first scenario and b) stronger-last scenario

5.5.2 A closer look at the timing of a collision

Now that the presence of the capture effect in LoRa is established, we will take a closer look at what happens to LoRa packets when the second transmission arrives during different segments of the first packet. By recalling the LoRa packet composition described in Section 3.2, we see that each LoRa packet consists of a preamble, header (+CRC) and payload (+CRC). In the previous experiment, all collisions occurred during the payload of the first arriving packet. The results of this experiment may change if the collision occurs during other segments of the packet, as will be described below.

The first interesting situation occurs when the collision takes place at the beginning or end of a packet, since there may be a bound for which one or both packets can still be successfully received even though there is an overlap between them. The situations in which this can occur are when (i) the stronger signal is slightly delayed compared to the weaker signal, (ii) the stronger signal is delayed until the end of the weaker transmission, and (iii) the weaker signal is received at the end of the stronger transmission. In a situation in which the weaker signal is slightly delayed with respect to the stronger signal, it is most likely that the stronger signal is received and the weaker one is lost.

A second scenario worth looking at is when the collision occurs during transmission of the header of the first arriving packet in a stronger last scenario. As we have seen so far, transmissions experiencing a stronger competitor are corrupted and cannot be successfully decoded. A gateway could therefore choose to stop listening to a signal that experiences too much interference from a stronger signal from a different end-device. Although we have seen in the previous experiment that this does not occur, i.e. both the stronger and the weaker signals are lost, the experiment conducted in Section 5.5.1 did not cover the situation in which the stronger signal starts during transmission of the weaker signal's header. The header contains information about the length of the payload, the presence of the payload CRC and the used coding rate. This indicates that the gateway is able to decode the header separately from the rest of the packet, since knowledge about the used coding rate is necessary to decode the payload. If the gateway is unable to retrieve content of the header and therefore the length of the packet and the used coding rate, e.g. due to interference from another end-device, it would make no sense to keep trying to receive the packet.

To determine behavior of a LoRa network in the scenarios outlined above, the setup as used for the previous experiment was utilized again to create the described *stronger-first* and *stronger-last* situations. This time, however, the delay between the packets was increased with steps of 1ms to increase the precision of the measurement, which is also depicted in Figure 27.

The focus of this measurement lied at the stronger-last scenario, since any collision occurring in a stronger-first scenario is most likely to result in the successful reception of the stronger signal, and the loss of the weaker signal. Moreover, collisions during the payload segment of a packet were not considered in this experiment since the results shown in Figure 26 have already shown

what happens in this type of situation, and there is no reason to expect any other behavior for collisions occurring at that moment. For each configuration regarding the timing of the stronger and weaker signals, 100 packets were sent.

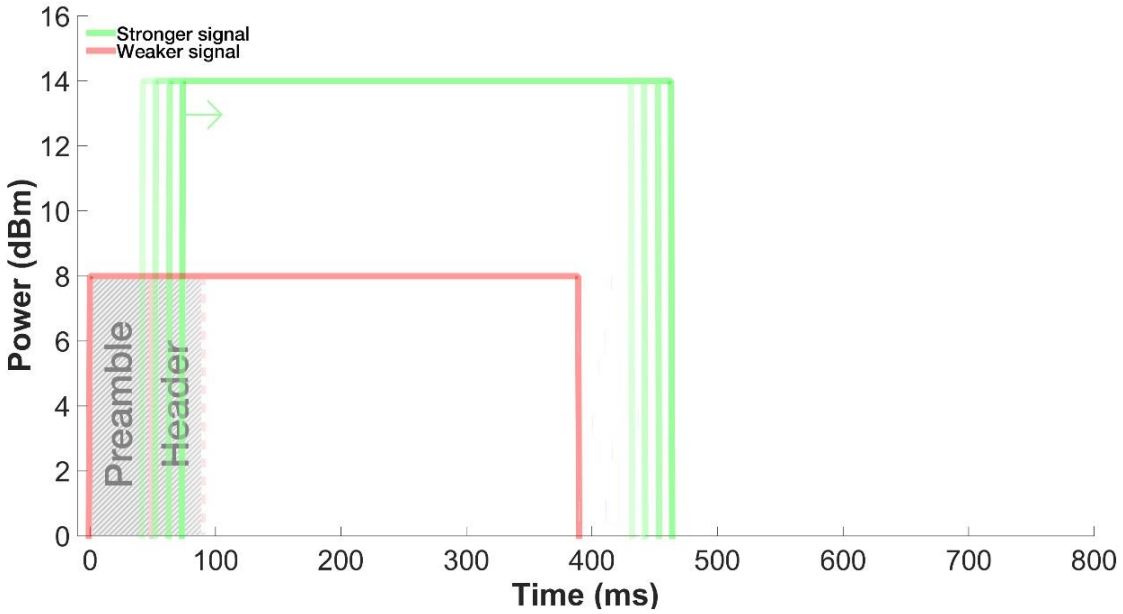

Figure 27. Timing of the stronger signal as compared to the weaker signal, during the experiment

5.5.3 Experimental results

The experimental results described below show that the packet loss is indeed affected by the timing of the collision. Figure 28 illustrates the percentage of received packets in a stronger-last scenario for every offset between the two transmissions. In what follows, we will look first into what happens when the edges of two packets overlap, and then into what happens in a stronger-last scenario during transmission of the header of the weaker signal. It is worth noting that the packet loss is assumed to be 100% between a delay of 90ms and 350ms, based on the previous experiment and the results that will be presented in Section 5.5.3.2. No experiment was conducted for delays between 90ms and 350ms as we have already seen from Figure 26 that the packet loss for both signals is 100% for delays of 100ms, 200ms and 300ms.

5.5.3.1 *Overlapping edges*

Figure 28 shows that there is a transition window at the edges of transmissions for both weak and strong signals in the stronger-last scenario. One can see that the gateway is able to receive more than 90% of all stronger signals, when they are delayed no more than 12ms. This roughly corresponds to 3 symbols (i.e. the symbol time is 4.096ms). After that, the packet loss increases noticeable up until the stronger signal is delayed 20ms, (i.e. 5 symbol times). From that point on, no stronger signals are received anymore. Similar behavior occurs at the other end of the weaker signal. Here, the stronger signal arrives at the end of the weaker signal and although the gateway was trying to receive the weaker signal, it is able to detect the preamble of the stronger signal after the weaker signal is fully transmitted. In this case, over 90% of the stronger signals are received when they overlap at most 6 symbols with the end of the weaker signals.

The weaker signal, on the other hand, can only be successfully received if there is almost no overlap. This is most likely caused by the fact that the stronger signal corrupts the CRC segment of the weaker signal, causing the gateway to discard the weaker signal. In this scenario, any overlap results in a high chance of packet loss for transmissions from the weaker node and only when there is less than 3ms overlap, more than 90% of the weaker signals are received.

The same experiments conducted for the stronger-first scenario show that the expected loss for all weaker signals and no loss for the stronger signals. Receiving a weaker signal in case of a stronger-first scenario will only be feasible if the weaker signal starts at the very end of the stronger signal so that the gateway only misses part of the weaker signal's preamble. This is the same behavior as we have seen when the stronger signal arrives at the end of the weaker signal. The difference is that in the stronger-first case, both packets can be received successfully since the CRC of the first (stronger) signal is not corrupted by the later arriving (weaker) signal.

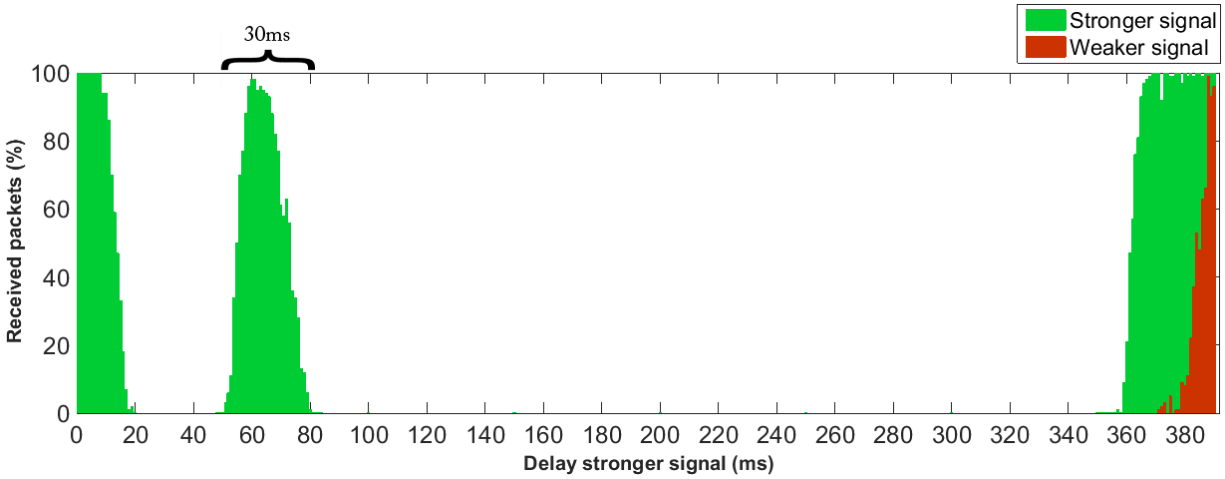

Figure 28. Received packets in a stronger-last situation based on a) empirical results

5.5.3.2 *Receiving stronger signals during a packet's header*

The same experiment also shows that the gateway does indeed switch to the stronger signal when its preamble arrives during the weaker signal's header. As can be seen in Figure 28, it is clearly visible that the gateway is able to switch to a stronger signal right after the preamble of the weaker signal ends. The total time window in which this behavior can occur is 30ms long, starting right after the weaker signal's preamble (which ends after 50ms). For delays between 59ms and 66ms, over 90% of the packets sent from the stronger transmitting end-device were received.

5.5.4 Discussion

The results above indicate the presence of a capture effect in LoRa. As soon as the gateway is able to detect a preamble, it will synchronize and lock to this signal. Stronger signals arriving on the same channel and data rate are ignored unless they arrive almost directly after either the start of weaker signal's preamble or its header. In that case, the gateway switches its focus towards receiving the stronger signal. In those situations, correct reception of the stronger signal is still feasible. Any stronger signal arriving after the header of a weaker signal will cause packet loss for both packet. This is because the gateway is unable to recover the weaker signal due to

the interference of the stronger one. Only when part of the stronger signal's preamble arrives after the end of the weaker signal, a correct reception of the stronger signal is possible despite the (small) overlap. The same holds for receiving the weaker signal if it arrives at the end of the stronger transmission. In that case, both packets can be received. It is worth noting that gateways cannot differentiate between LoRa packets originating from other LoRa networks. Each LoRa packet transmits the same preamble, which will make any gateway in its vicinity try to receive it regardless of the network it was intended for. A gateway can only determine if the packet was indeed addressed to the network it belongs to once it has fully received it.

End-devices deployed in systems that require a reliable data throughput could utilize acknowledgements to increase chances for the correct arrival of a packet. This can, however, be counterproductive as will be described in Section 7.4.1, were various options to decrease the chance of packet loss due to collisions are described.

6 Performance analysis of LoRa in outdoor and mobile environments

Wireless networks are widely used in various outdoor environments. Examples include the fields of environmental parameter monitoring (e.g. air pollution, water quality, or weather conditions), security and surveillance applications, mobile communication, localization (GPS), and wildfire detection [53][54]. The choice for a certain wireless technology will have a big impact on the overall performance and behavior of the system. It is therefore important to have prior knowledge about the expected behavior of the technology in different environments. Where indoor environments can be controlled up to a certain point, outdoor environments are more unpredictable due to, among other things, weather conditions and interference from physical structures like trees and buildings, potentially causing signal attenuation. The performance of a wireless sensor system can especially be affected if (sensor) node are not static and experience changing environments due to mobility. This chapter discusses outdoor and mobile measurements with LoRa end-devices in different configurations to give an insight into its outdoor performance.

6.1 Experiment setup

To conduct outdoor experiments, the same type of gateway as described in Section 4.1.2 was used. The gateway was placed on one of the highest buildings (over 50m) in the neighborhood, i.e. the 'Horsttoren' located at the university campus, to increase coverage and decrease the chance of signal obstruction by physical objects. To verify packet reception to this gateway, 100 packets were transmitted (on each data rate) from a grass field. From this field, at a distance of approximately 400 meters, there is a direct line of sight to the building. The received signal strength averages for each data rate of this measurement is shown in Table 13. It is interesting to see that particularly DR 3 and DR 1 stand out compared to the average values on the other data rates. DR 3 shows a much lower average RSSI value, whereas DR 1 has a higher average RSSI value. These deviations do not change when the experiment is repeated.

Data Rate No.	Average RSSI (dBm)
5	-94.55
4	-94.90
3	-100.52
2	-94.42
1	-91.29
0	-93.52

Table 13. Average RSSI for each Data Rate approximately 400 meters from the gateway.

Environmental influences on packet reception and signal strength caused by mobility were studied by transmitting LoRa packets while following a path in the neighborhood around the location of the gateway. This path is shown in Figure 29. The path contains areas with many physical objects such as buildings (residential and industrial zones) and vegetation (a park on the east side of the path and many trees around the campus), as well as areas with a clear line of sight to the building

at which the gateway is placed. The path has been traversed at two different speeds: walking speed (±6 km/h) and cycling speed (±15 km/h), for different configurations of data rates (DR 5 and DR 3) and transmission powers (14 dBm and 8 dBm). Four segments of the path are highlighted. These segments mark the following environments:

1. A wooded area closest to the gateway with some spots that have line of sight to the building. The segment is between 210 and 340 meters away from the gateway
2. A partly industrial partly residential area with a low building density and mostly low-rise buildings. This segment is between 820 and 1250 meters away from the gateway
3. A more densely built residential area with on average higher (three story) buildings. This segment is between 860 and 1230 meters away from the gateway
4. A heavily forested park. This segment is between 650 and 970 meters away from the gateway

We will use packets received at each segment of the path to compare the impact of different transmission configurations on performance, as well as the influence of different environments. Segments 2 and 3are roughly at the same distance from the gateway to eliminate differences in signal strength caused by path loss, i.e., attenuation of radio-waves over distance.

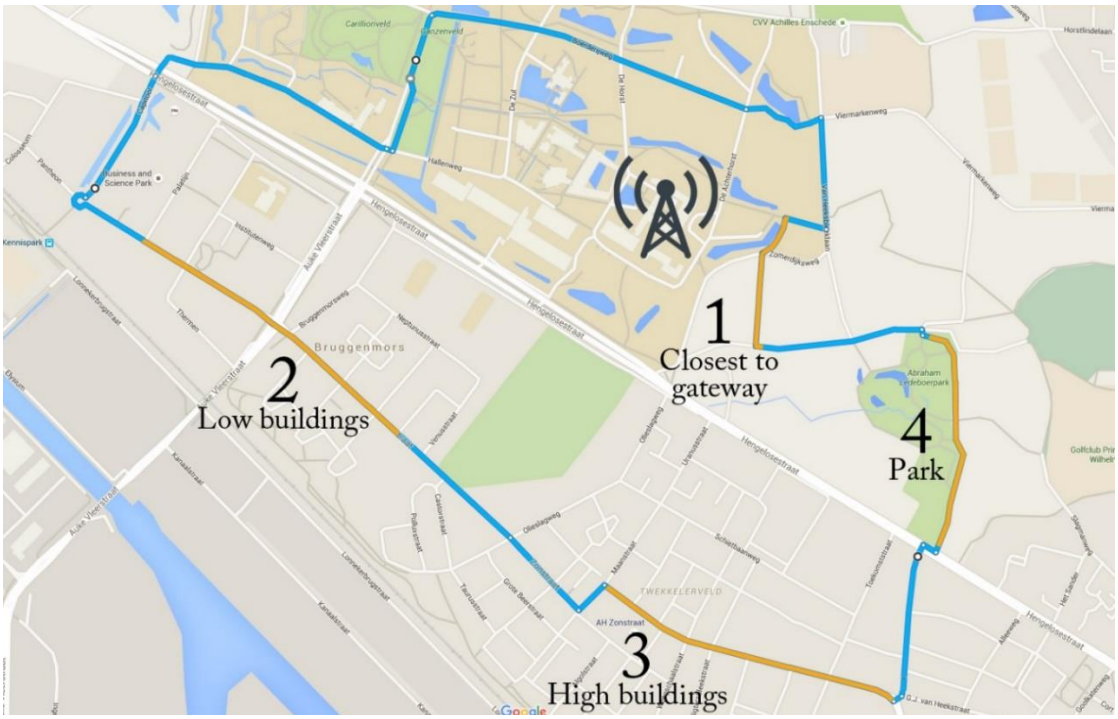

Figure 29. Path along which mobility measurements took place with the four segments highlighted in orange

To determine the packet loss and signal strength on places located on the path described above, an Adafruit Ultimate GPS Logger Shield [55] mounted on top of the end-device, as shown in Figure 30 was used. The GPS shield also comes with a SD card slot to allow storage of data. This setup makes it possible for an end-device to retrieve its position, transmit a LoRa packet, and store the retrieved GPS coordinates together with the ID of the transmitted packet. The latter makes it possible to link the location data to the sequence number of the received packets. The device is powered using a 5V USB powerbank.

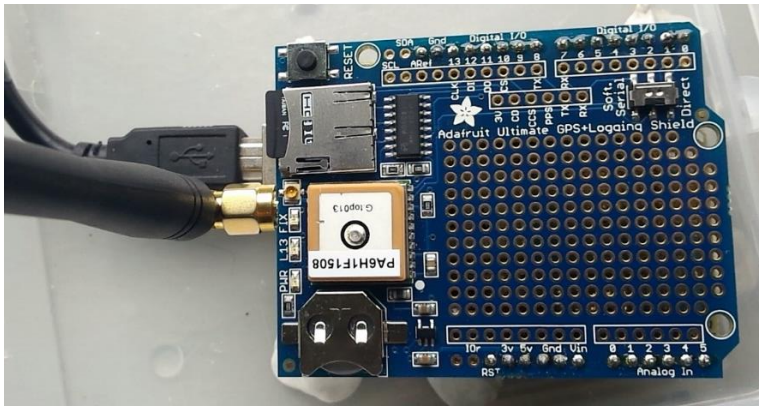

Figure 30. The Adafruit Ultimate GPS Logger Shield on top of an end-device as described in Section 4.1.1

6.2 EXPERIMENTAL RESULTS

This section presents the emperical results. First, the results for RSSI values along the pathway are shown and discussed for a combination of data rates and transmission power measured at walking speed. Next, the results at cycling speed are discussed and compared to the results obtained at walking speed.

6.2.1 Walking speed

Figure 31 shows the RSSI values retrieved from the gateway combined with the stored GPS data from the GPS shield. The results are obtained from an end-device moving at walking speed while transmitting packets at the fastest data rate (DR 5) and highest transmission power (14 dBm). Each waypoint on the map indicates a packet transmission. The better the received signal strength, the higher (redder colored) the RSSI value. Figure 32 shows the results for the same experiment for packets transmitted on DR 3. The increased distance between waypoints on DR 3 is caused by a higher transmission time resulting in a higher interval time between subsequent packets due to the duty cycle restriction. Figures showing the results for experiments with a lower transmission power (8 dBm) can be found in Appendix D.

The expected decline in RSSI values when the distance to the gateway increases can be derived from the highest values being measured at the closest distances to the gateway. However, some locations show a better or worse signal strength than their adjacent locations even when the distance to the gateway does not change much. These fluctuations are most likely caused by physical objects such as buildings and trees obstructing the signal's path. This does not change between different transmission configurations, as to be discussed next.

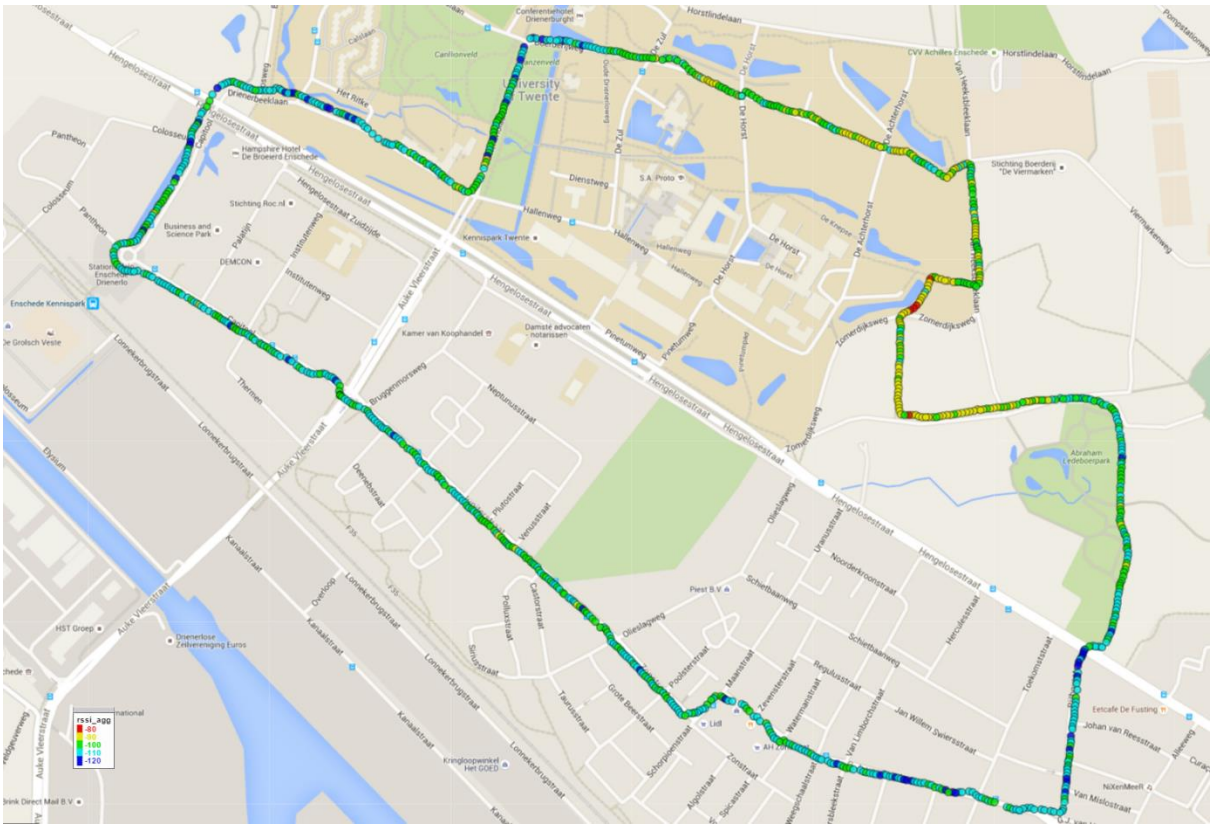

Figure 31. RSSI values measured while transmitting packets at DR 5 and TX power 1 (14 dBm), at walking speed

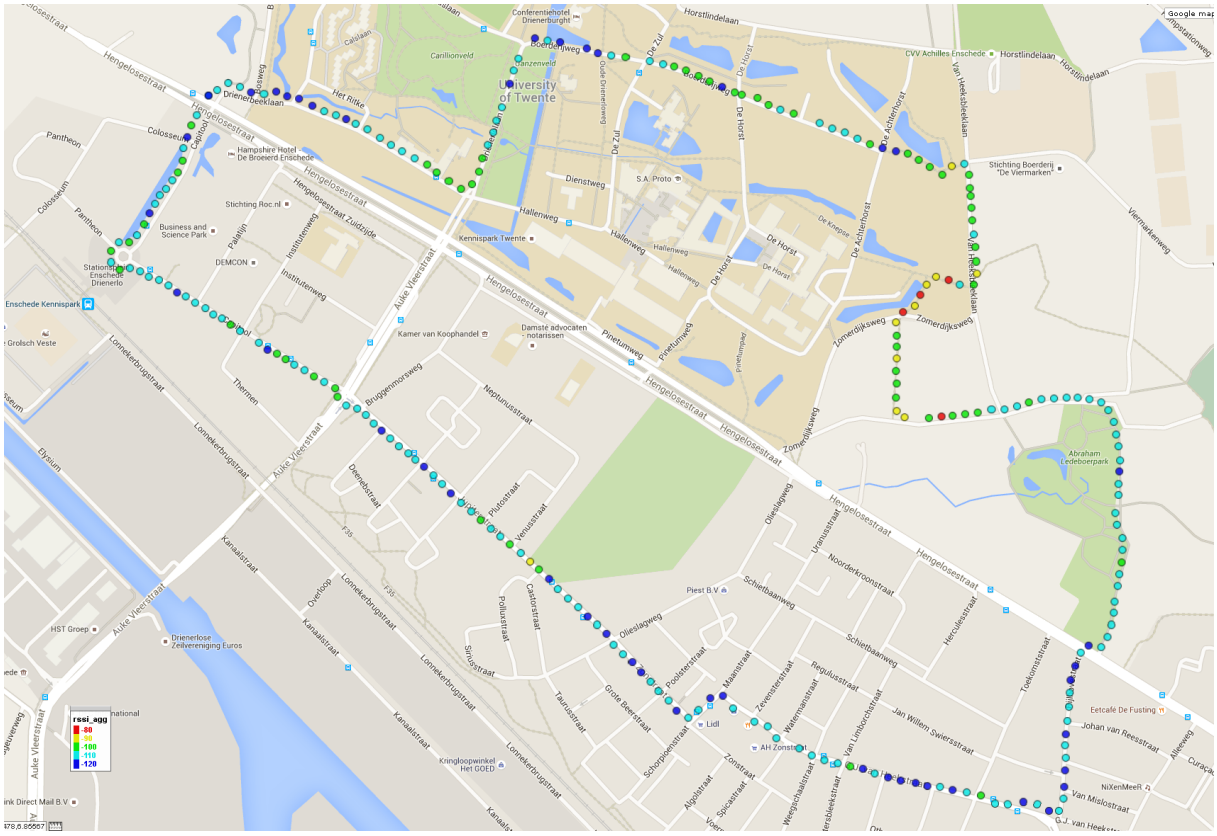

Figure 32. RSSI values measured while transmitting packets at DR 3 and TX power 1 (14 dBm), at walking speed

As it can be seen from Figure 31 and Figure 32, most packets arrive on both transmission configurations. When the output power of the end-device is lowered, the packet loss increases noticeably on both the individual segments and the entire path. This is reflected in Figure 33 and Table 14, where the packet loss and the average RSSI values for each segment for each transmission configuration are presented, respectively.

These results do not only show the increased packet loss on lower transmission powers but also the influence of the data rate. Packets transmitted on DR 3 show a higher probability of arriving than those transmitted on DR 5, especially those transmitted on Segment 2 and 3. The latter segment also shows the most packet loss of all regardless of the transmission configuration.

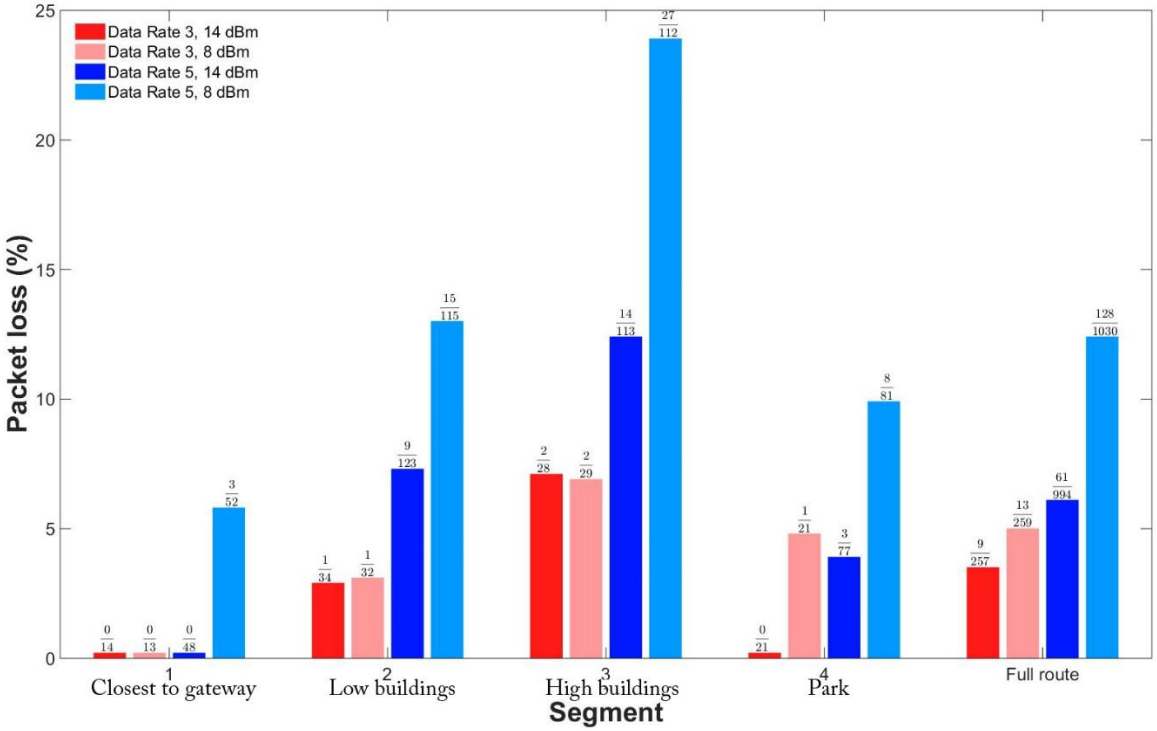

Figure 33. Packet loss at walking speed on each segment of the route, for each data rate and output power

Segment	Data Rate 5	Data Rate 3
1	-88.1	-87.1
3	-101.7	-103.5
4	-104.0	-108.0
5	-100.2	-106.0
Route	-99.5	-103.2

Segment	Data Rate 5	Data Rate 3
1	-91.7	-96.9
3	-107.5	-111.8
4	-110.0	-112.9
5	-107.3	-110.3
Route	-104.5	-108.9

Table 14. Average RSSI for each segment of the path for each data rate. Left: for a transmission power of 14 dBm Right for a transmission power of 8 dBm

Differences between individual segments are reflected in the increased packet loss on segments further from the gateway and in the lower average RSSI values. Segments 2 and 3 (both about equally far from the gateway) show the worst RSSI values and in general yield the most packet loss. The influence of buildings comes forward when we compare Segments 2 and 3. Both average

RSSI and packet loss are better on Segment 2, indicating that the higher building density on Segment 3 negatively influences performance. The influence of the vegetation on Segment 4 (the park) is less clear. Segment 4 lies closer to the gateway which has a positive effect on packet loss and RRSI values compared to other segments. When we look into the transition from Segment 4 to the residential area south to Segment 4 in Figure 31 and Figure 32, a noticeable decrease in RSSI values occurs. This indicates that nearby buildings have more influence on signal strength then vegetation and distance.

The influence of distance is best reflected when looking at the performance on Segment 1 (the path closest to the gateway). As expected, this segment performs best due to its close proximity to the gateway. Packet loss only occurs on the worst combination of transmission power and data rate. The RSSI values retrieved from packets transmitted at this segment are much higher than those on other segments.

It is worth noting that on DR 3, the number of transmitted packets on some segments is quite low due to the duty cycle restrictions. One additional lost packet can therefore influence the packet loss percentage quite a lot. It is, however, still clear that both a higher transmission power and a lower data rate affect the overall packet loss in a positive way. This is best visible if we look at the overall packet loss on the whole path. A higher data rate and a lower transmission power result in more packet loss. In this case, the data rate has a bigger impact on performance than the transmission power. It seems that the opposite holds when we compare the average RSSI values between DR 3 and DR 5 shown in Table 14. However, the differences in average RSSI values between the two data rates originate from a constant difference in RSSI values between DR 3 and the other DRs, as earlier shown in Table 13.

6.2.2 Cycling speed

To determine if the speed at which a mobile node travels has an influence on performance, the same experiment as described above has been carried out while traveling at a higher (cycling) speed. The results in terms of packet loss and average RSSI values are shown in Figure 34 and Table 15, respectively. Due to the increased speed at which the path is traveled, less packets are collected during each lap, making the results more susceptible to volatility. An example of this can be seen from the packet loss on Segment 3 and 4 in Figure 34, where one or two lost packets (on DR 3 and output power 14 dBm) have a huge influence on the results of different transmission settings. Figure 35 shows that the gap between each transmission from the measurement on DR 3 and an output power 8 dBm is increased considerably compared to Figure 32. Transmissions on DR 5, however, still clearly show the influence of the transmission power on packet loss. The overall packet loss for the whole path indicates that transmission power and data rate have similar effects on performance when compared with results from Section 6.2.1.

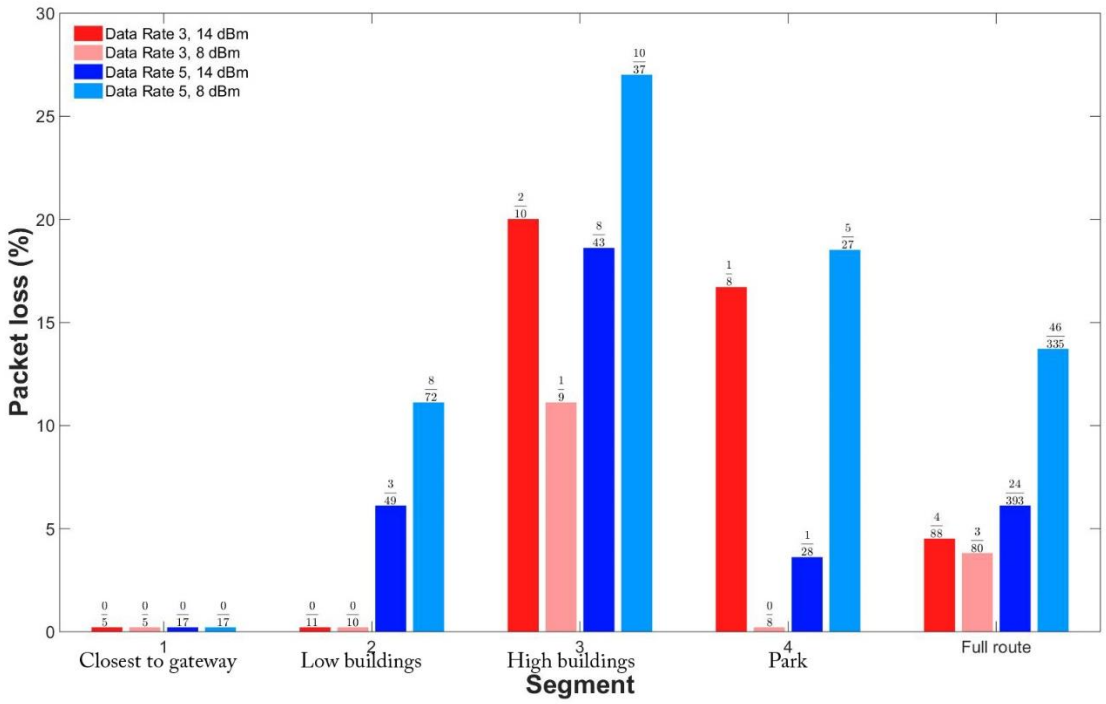

Figure 34. Packet loss at cycling speed on each segment of the route, for each data rate and output power

Segment	Data Rate 5	Data Rate 3
1	-90.8	-89.9
3	-101.6	-106.8
4	-103.7	-109.4
5	-102.4	-105.1
Route	-100.2	-104.6

Segment	Data Rate 5	Data Rate 3
1	-85.5	-97.2
3	-102.9	-110.0
4	-105.6	-113.9
5	-107.4	-112.0
Route	-101.8	-108.8

Table 15. Average RSSI for each segment of the route for each data rate. Left: for a transmission power of 14 dBm Right: for a transmission power of 8 dBm

The results for individual segments are comparable with the results described in Section 6.2.1. Segment 1 performs the best in terms of both packet loss (zero on all transmissions) and average RSSI. Segment 2 shows a better performance of DR 3 compared to DR 5. Segment 3 shows the most packet loss and the worst RSSI values. Segment 4's performance does not differ much from that of Segment 2. Figure 35 shows the distance between each transmission from the measurement on DR 3 and an output power of 8 dBm. These results show that the mobility experiments at low speed are more reliable because these experiments have yielded more data. When conducting experiments aimed at getting an indication of the coverage and packet loss in a certain area it is therefore advisable to conduct them at for a long period of time or on low speed, if the speed itself is not of great importance.

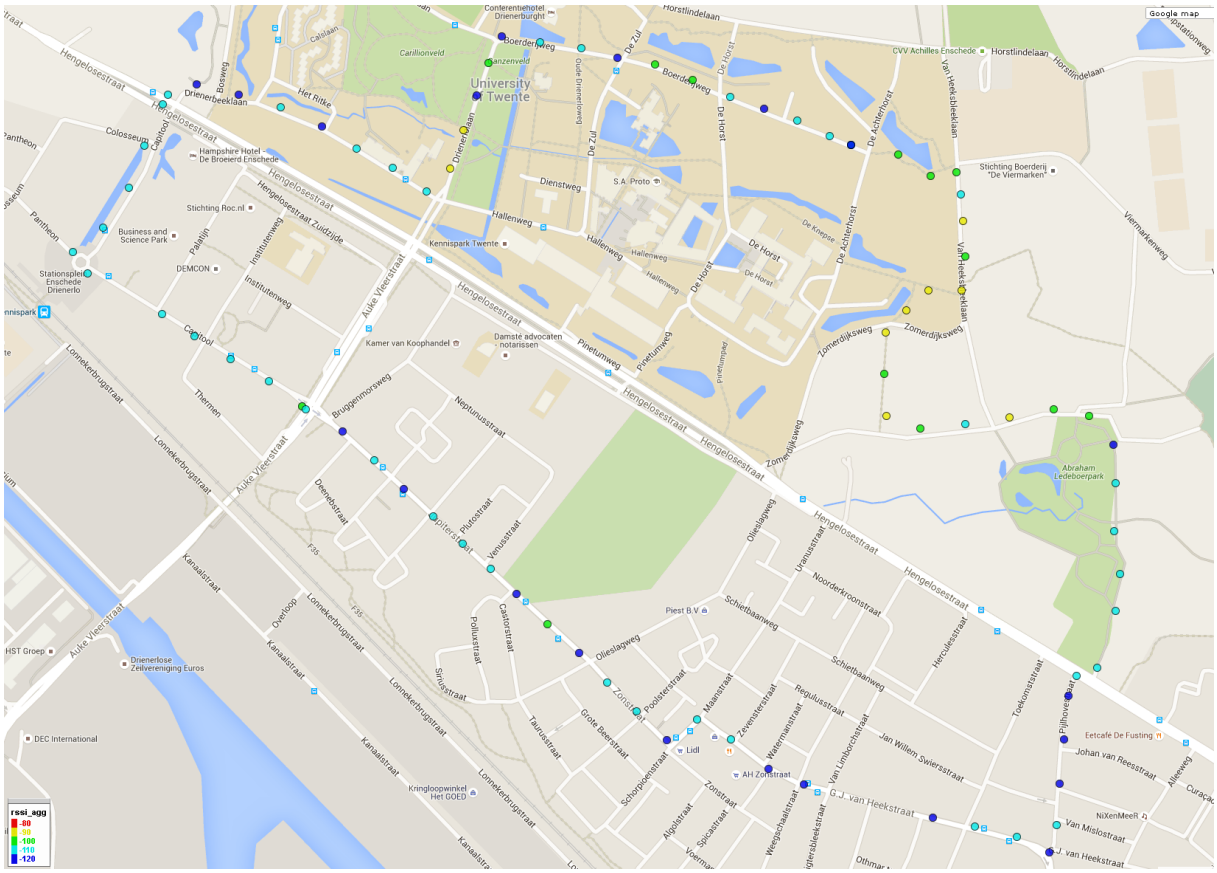

Figure 35. RSSI values measured while transmitting packets at DR 3 and TX power 1 (14 dBm), at cycling speed

6.2.3 Discussion

The most interesting comparison between the results described above can be made between the overall performance and speed. Looking at the packet loss on the full path for both measurement speeds (Figure 33 and Figure 34) indicates that there is no significant influence of speed on packet loss. The results are combined in Table 16 for easy comparison.

Speed	Data Rate 5	Data Rate 3
Walking	6.1 %	3.5 %
Cycling	6.1 %	4.5 %

Speed	Data Rate 5	Data Rate 3
Walking	12.4 %	5.0 %
Cycling	13.7 %	3.8 %

Table 16. Overall packet loss for each data rate. Left: for a transmission power of 14 dBm Right: for a transmission power of 8 dBm

The influence of speed is not present in these results and indicates that mobility itself does not negatively influence the performance of a LoRa network. Additional measurements on lower data rates could enhance these results, but would require a much longer route or measurement time to make them reliable.

Environmental changes caused by mobility are, however, of great influence on performance. When deploying a LoRa network with mobile end-devices, knowledge about the environment can help determining the best implementation in terms of transmission power and data rate. When high reliability (low packet loss) is required, most performance gain can be obtained from decreasing the data rate.

7 PERFORMANCE ANALYSIS OF LORA IN A DENSE IOT ENVIRONMENT

The previous chapters have shown an analysis of LoRa behavior under various circumstances. All of the conducted experiments took place in environments with little to no other LoRa traffic. The experiments with two transmissions overlapping in time and on the same data rate and channel (i.e. transmission frequency) have shown what happens when collisions between LoRa packets occur. When LoRa technology is going to be deployed in various IoT applications on a wide scale, these collisions are likely to take place. This is not a problem when they happen occasionally but if the number of end-devices increases, they can cause significant packet loss due to network congestion. In this chapter we will investigate the behavior and performance of LoRa in a realistic setting in which many end-devices are active by means of a LoRa simulator created in Matlab. First, the layout of the simulator is discussed along with some of its key components such as the implementation of the gateway and an end-device. The simulator is then used to recreate the scenario where two packets interfere with each other, as described in Section 5.5. By comparing the outcome of this simulation with the empirical results obtained in Section 5.5, we validate the correct behavior of the simulator. The simulator will then be used to create a dense IoT environment to analyze performance of the network in presence of many LoRa end-devices. A smart city is used as a reference scenario for such a dense IoT environment.

7.1 CONSTRUCTION OF THE SIMULATOR

The main purpose of the simulator is to mimic an environment in which (many) LoRa end-devices are present and transmit packets. This is achieved by (i) defining the characteristics of the environment, (ii) initializing the simulation environment and (iii) running the simulation for this environment. The simulator is written in Matlab since it offers an easy to learn and yet powerful environment commonly used in the scientific research community. The implementation strategy taken is to create the simulator in an Object-Oriented structure. The three main objects are the ones modeling the behavior of the gateway, an end-device, and the characteristics of a LoRa packet. To avoid significant memory usage and computation time, an overall system manager object is used to collect results for later analysis. This ensures that other objects only have to store a minimum set of variables and do not have to keep track of any state information during the simulation. First, we will look into different requirements for each of these models individually. Then the way to connect all these separate components is disused. Finally, we look at how a physical environment can be created in which the gateway and end-devices are placed.

7.1.1 The LoRa end-Device (LRD)

The main purpose of the LoRa end-Device (LRD) is to transmit packets at predefined timeslots within the bounds of the simulated time. To do so, an LRD object is given a list of transmission times upon initialization, in addition to the DR, transmission power, preamble and payload length, coding rate, and available channels. Every time the LRD should start transmitting a packet, it checks if the duty cycle restrictions allow it to do so and, if so, it creates a packet object. Each created packet is collected by the system manager to allow statistical analysis (e.g.

packet loss) after the simulation finishes. Then, for each time instance during the transmission, the LRD checks which part of the packet it is currently transmitting (e.g. the preamble, header or payload), and if it transmitted all the symbols of the packet so that it can stop the transmission. The LRD determines which part of the packet it is transmitting based on the number of transmitted symbols, to be described in the next section. The transmission status of the LRD is used by the gateway to determine how it should react to this transmission, which will be described in Section 7.1.3.

7.1.2 The LoRa packet

Packet objects are created by LRDs and contain information about the used DR and channel, the amount of transmitted symbols, and the strength of the signal as received by the gateway. When collisions occur, a gateway is only capable of detecting a new packet if it is transmitting its preamble during the preamble or header of the other packet and if this signal is stronger than other transmissions. Therefore, the packet model also contains the length of the preamble and the header.

Because there are no details about the exact length of the header, the length of the header is set to the total time for which there is a chance that a stronger signal is detected by a gateway. This time is derived from the empirical results obtained in Section 5.5.3.2. These results show that the total time during which a stronger signal can be detected during a weaker signal's header, is 30ms for transmissions on DR 3. In addition, the probability of a gateway detecting a stronger signal during a weaker signal's header for other data rates is also derived from the results obtained in Section 5.5.3.2. This is done using Matlab to calculate a probability distribution that fit these results. This distribution is then used to determine the header length and detection probabilities for other DRs, based on the fact that the symbol time between two adjacent DRs differs by factor 2 and that the assumption that the same behavior occurs on other DRs. The header of packets transmitted on DR 4, for example, is therefore assumed twice as short as the header of a packet transmitted on DR 3. Figure 36 depicts the detection probabilities during transmission of the header of a packet transmitted at DR 3, as described in Section 5.5.3.2, together with the Matlab generated probability distribution for this DR, alongside the distributions for DRs 2 and 4.

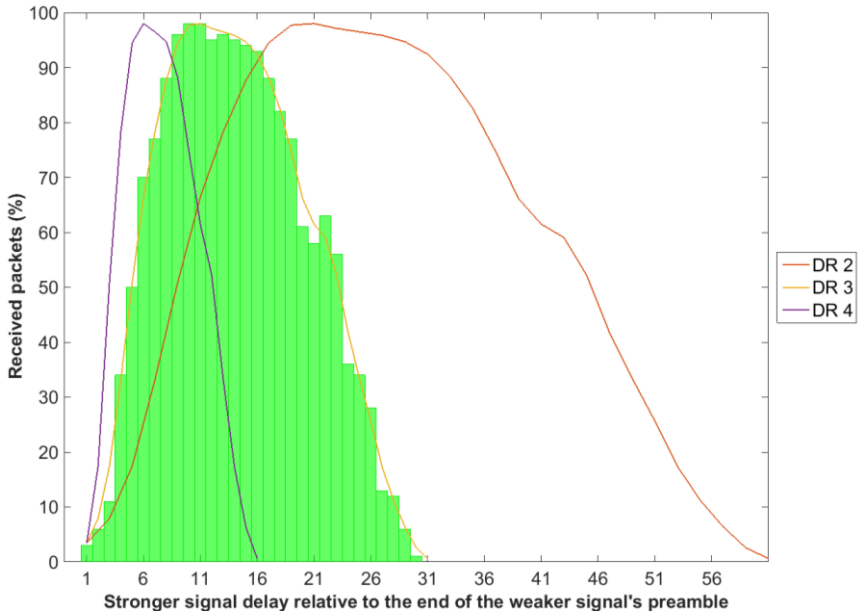

Figure 36. Measured detection probabilities of a stronger signal arriving during the header of a weaker signal together with the derived probability distribution for DR 2, DR 3 and DR 4

The same approach, as described above, is taken to determine the detection probability of a stronger signal's preamble (in a stronger-last situation) during the preamble or the end of the weaker signal, as well as the chance of detecting the weaker signal's preamble in a stronger-first scenario. These values are used by the gateway to determine packet detection probabilities for these situations, as will also be described in Section 7.1.3.

As mentioned earlier, a packet object will also contain the signal power per time instance from the gateway's point of view. Although it makes more sense to calculate and store this information in the gateway model, it saves a large amount of computation time if these values are calculated in advance and not during each iteration. A gateway can then simply compare the current signal power of each packet and make decisions based on these values without first having to calculate them. In addition to the values for signal power, the packet object is also used to store interference values for each time instance during the packet's transmission. The interference is calculated by the gateway and is passed on to the packet object during the simulation. These values are calculated by combining the signal power of all packets transmitted at the same DR and channel and are used afterwards to determine if a packet has been successfully received by a gateway. A more detailed description of this implementation approach is provided in the next three sections.

7.1.3 The gateway

The main goal of the gateway model is to determine which packets are detected and, subsequently, are received. This is achieved by keeping track of the transmission DR and channel of each LRD currently transmitting. If the gateway receives a preamble, it tries to lock onto this transmission. The ID, transmission DR and transmission channel of the LRD are stored in the gateway object if the transmission's signal power is above the sensitivity of the gateway for that DR [26]. The same variable is used to determine if the gateway is not already locked to another packet transmitted on the same DR and channel, when it receives a preamble. If this is the case,

the gateway checks if the newer signal is stronger. If so, the current state of the packet the gateway is currently locked to is examined. If either the preamble or the header of that packet is being transmitted at this moment, the gateway makes a choice about whether it will switch to the stronger signal based on the detection probabilities discussed in the previous section.

For every packet the gateway tries to receive, it calculates the total interference on the DR and channel of the packet. It therefore knows the received signal power of each transmitted packet for the current time instance. To decrease computation time of the simulation, the gateway model only calculates these values for the packet it is locked to, and stores these values in the accompanying packet object. Once a packet is fully transmitted, the gateway transfers this packet to the system manager and is able to detect any new transmission on the DR and channel combination used by the received packet. The system manager will determine if this packet was successfully received by comparing the signal power with the interference (both stored in the packet object) during the transmission. This is done after the simulation is finished to save computation time, and only for packets that the gateway was locked to until their transmission was finished.

7.1.4 The system manager

The system manager is an object created at the beginning of a simulation to store all created packets. It also stores the packets of which the preamble was detected by a gateway, and which the gateway tried to receive until they were fully transmitted. Packets present in the list of transmitted packets but not in the received packets list either never were detected by the gateway (e.g. due to stronger interfering signals, or due to a very low signal power), or experienced the capture effect. As mentioned earlier, each packet object contains its own signal power values as well as the strength of the interference for the duration of the transmission. After the simulation has finished, each received packet is analyzed to determine if there was any interference to corrupt the packet. If so, the packets are removed from the received packets list and transferred to a separate list for possible post-simulation analysis. A packet is considered lost if more than one symbol was corrupted due to interference (preamble symbols are excluded from this). The threshold for which packets are marked as corrupt is set to interference levels equal or higher than the packet's own signal power. In case of a stronger last situation where the stronger signal arrived at the very end of the weaker signal, there is a chance that the weaker signal is correctly received (as empirically shown in Section 5.5). This situation is taken into account based on the same approach described in Section 7.1.2 in which the reception probability of the weaker signal is determined by creating a probability distribution for this situation based on the results obtained in Section 5.5.3.

7.1.5 The path loss model

In a situation where two packets are transmitted at the same DR and channel, the behavior of the gateway is based upon the strength of each transmission. To determine the strength of a signal received by the gateway, knowledge about the attenuation of the signal while it travels to the gateway is needed. This attenuation is also known as path loss and is defined as the ratio of transmit power P_t to receive power P_r, usually expressed in dB [56]:

$$P_L = 10\log\left(\frac{P_t}{P_r}\right) \text{ dB} \tag{7.1}$$

Assuming transmit and receive signal powers are expressed in dBm, the received signal power can be expressed as:

$$P_r = P_t - P_L \text{ dBm} \tag{7.2}$$

Path loss can be predicted with a path loss model, which are widely used to predict a signal's propagation loss. These models take, inter alia, into account distance between a sender and a receiver, the frequency, and the type of environment to predict signal attenuation [56]. Past research has yielded many different path loss models, each suited for different situations [57].

Choosing a suitable path loss model for a given application can be a challenging task, which often requires empirical data to validate the model. The path loss model needed for the simulator only needs to give credible attenuation figures to differentiate between the signal power of each transmission, without high accuracy demands. Therefore, the requirement for the desired model is primarily its applicability for sub-GHz frequencies and (sub) urban environments. Based on this requirement, the COST-231 Hata Model is chosen for calculating the path loss in the simulator [56] [57]. This model is based on empirical data, widely used for predicting the path loss in wireless systems, and designed for outdoor transmissions on frequencies between 500 and 2000 MHz for both urban and suburban/rural environments [57].

The COST-231 Hata Model is formulated as follows [56]:

$$\begin{aligned} P_L\ d\ =\ 46.3\ +\ 33.9\log f\ -13.82\log h_b\ -a h_m\ f, h_r \\ +[44.9-6.55\log(h_b)]\log(d)\ +\ C_M \end{aligned} \tag{7.3}$$

, where $a h_m$ is defined for urban environments as:

$$a h_m\ f, h_r\ = 3.20[\log(11.75 h_r)]^2 - 4.97 \tag{7.4}$$

and for suburban/rural environments as:

$$a h_m\ f, h_r\ = [1.1 \log f\ - 0.7] h_r - [1.56 \log f\ - 0.8] \tag{7.5}$$

, where the distance d between an end-device and the base station is denoted in kilometers, f is the frequency in MHz, h_b is the height of the end-device's antenna in meters, and h_m the height of the base station's antenna in meters. C_M is defined as 3 dB for metropolitan areas and 0 dB for smaller cities, suburbs and rural areas. The COST-231 Hata model is implemented in the simulator with the values for each parameter as shown in Table 17.

Parameter	Value
f	868 MHz
h_b	1.5 meters
h_b	50 meters
C_M	0 dB
ah_m	urban

Table 17. Parameter values used in the COST-231 Hata model for calculating the path loss

7.1.6 Creating a simulation environment

Defining the simulation environment is done separately from the actual simulation, by generating and storing environment parameters in an environment file. This file can then be loaded when the simulator is started and allows the simulation of multiple environments in a row. Such an environment file contains the following set of parameters:

- Real-world time that needs to be simulated
- Size of the area in which LRDs can be located
- Number of gateways and LRDs in this area and the position of each of them
- Available DRs, channels, and transmission powers an LRD can transmit on
- Time at which each LRD should transmit its first packet
- Interval at which each LRD should transmit a packet
- Coding Rate, preamble, and payload length (in symbols and bytes respectively)
- Probabilities of the behavior of the gateway in case of each collision situation

Some of these parameters are fixed to save simulation time and to prevent complexity. For example, the Coding Rate, transmission power, payload length, available channels and number of gateways parameters are set to a default value and will not be altered during the simulation. It is however possible to generate environments in which (most of) these parameters are not fixed. Table 18 shows the default values for the above-mentioned parameters. The number of gateways is fixed to one to save a large amount of computation time and to decrease complexity. The available channels an LRD can transmit on is limited to one to speed up the simulation. When an LRD starts transmitting a packet, it randomly selects a transmission channel. By only allowing one out of the three default LoRaWAN channels, a collision is three times more likely to occur. This is the equivalent of allowing all three channels and deploying three times more LRDs.

Parameter	Value
Number of gateways	1
Transmission power	14 dBm
Coding Rate	4/5
Payload length	51 bytes
Preamble length	12.25 symbols

Table 18. Fixed values for some simulation parameters

The transmission DR and interval for each LRD are randomly selected from a list of available DRs and transmission intervals upon generation of the environment. We will discuss this, along with the placement of the gateway and LRDs, in more detail in Section 7.3.

7.1.7 Simulator process flow

Assuming an environment is created in the way described in the previous section, the simulation can proceed by using the environment parameters to generate the system manager and the objects modeling the gateway and LRDs. This initialization phase consists of the following steps:

1. Load an environment file
2. Initialize
    a. The gateway
    b. Every LRD
    c. The system manager
3. Create a list of every time instance at which an LRD should start a transmission

After the initialization phase, a simulation can start by iterating over each time instance until the total simulation time is reached. The time instance interval of the simulator is set to 1ms since the minimum symbol time (on DR 5) almost equals 1ms (i.e. 1.024ms), and to match the 1ms interval at which the measurements in Section 5.5 were conducted. The simulator will consecutively take the following steps for each time instance:

1. Determine which LRDs should start a transmission, and do the following for each LRD:
    a. Update its transmission status, make it create a packet, and add the created packet to the system manager
    b. Determine the signal power of each transmission and store it in the packet object
    c. Update the gateway with the signal power of each packet for this time instance
2. Determine which LRDs are already transmitting and do the following
    a. Determine the signal power of each transmission and update the gateway with this signal power for this time instance
    b. If an LRD is transmitting the last time instance of a packet, notify the gateway that it no longer has to be locked to this signal, and that it can transfer this packet to the system manager in case it was trying to receive this packet
3. Now that the signal power for each transmission is updated on the gateway, let the gateway determine which signal(s) it can detect and will listen to
4. If no LRDs are transmitting at the current simulation time, fast-forward the simulation to the next simulation time at which a new transmission starts (based on Step 3 in the initialization phase) to avoid unnecessary simulation overhead.

The process described above is depicted in Figure 37 to provide an overview of the simulator's process flow.

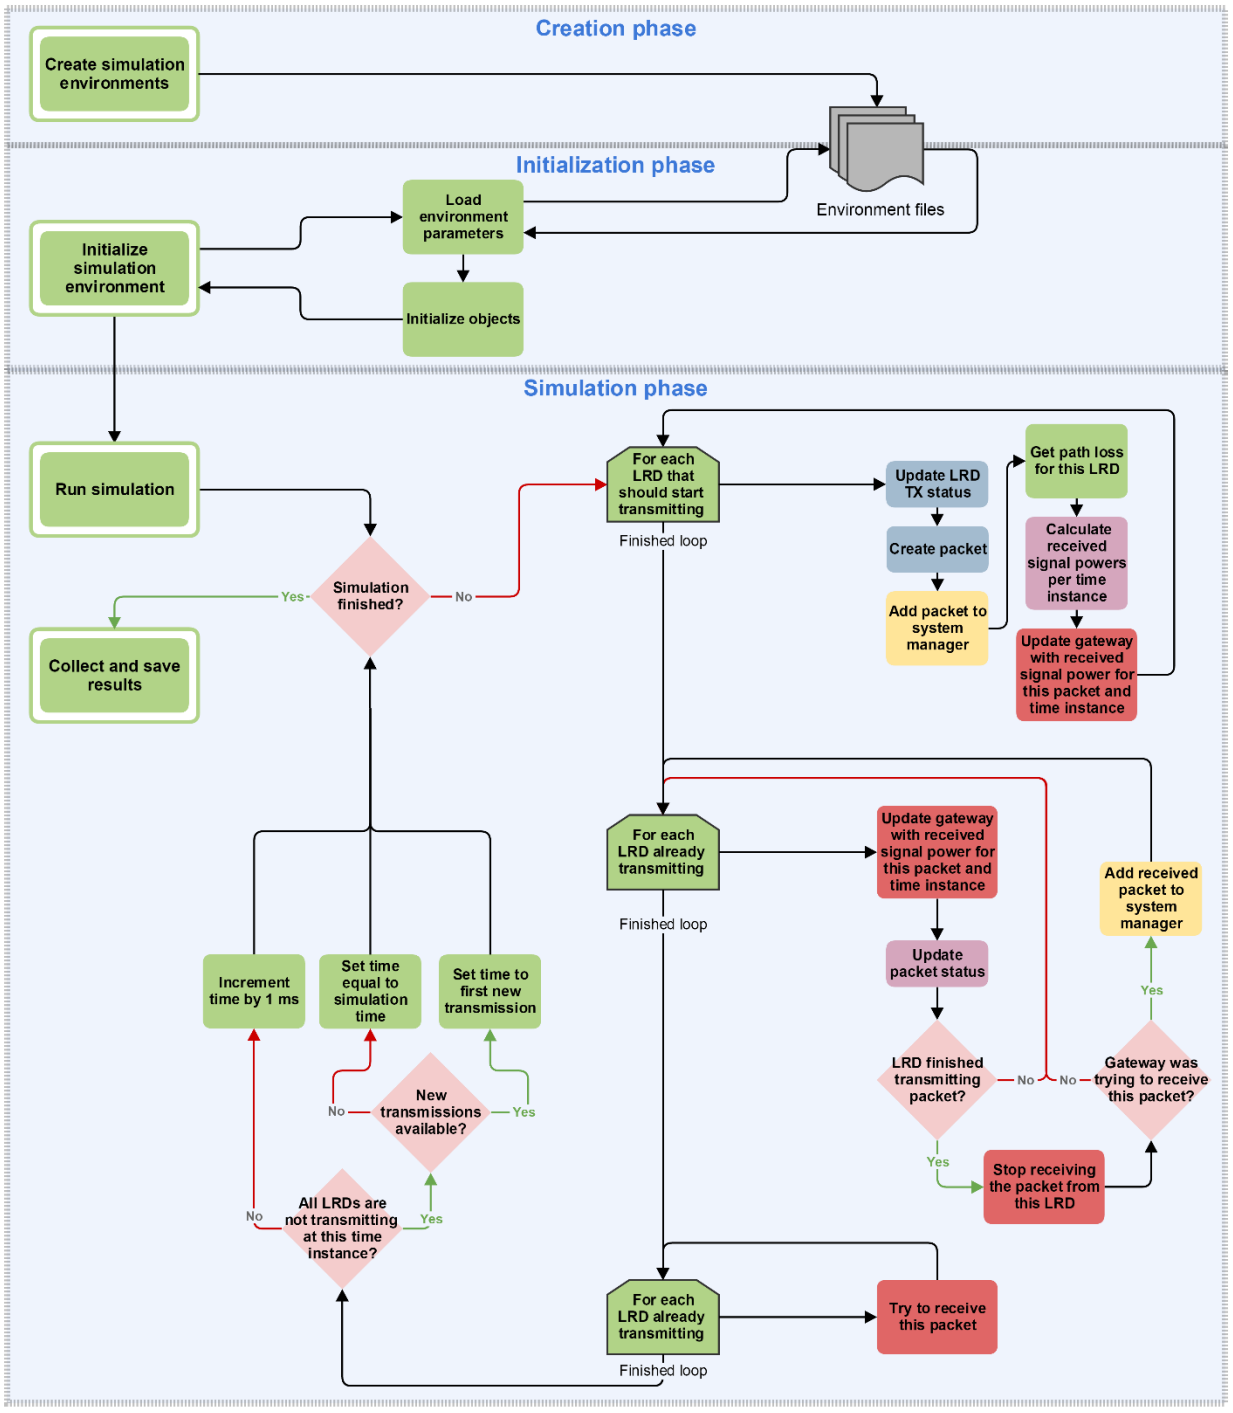

Figure 37. An overview of the simulator's process flow. Red blocks relate to the gateway object, blue blocks to a LoRa end-device, pink blocks to a packet object, and yellow blocks to the system manager. Green blocks relate to the simulator's own functions.

7.2 VALIDATION OF SIMULATOR BEHAVIOR

Now that we know how the simulator is implemented, a validation of its behavior is required. This is done by means of re-creation of the experiment described in Section 5.5. This experiment involved the forced collision of two packets to create a stronger-last scenario, and yielded the results shown in Figure 38a. The described verification simulation is carried out by repeating the experiment described in Section 5.5 and comparing those empirical results with the results obtained from the simulator, which are discussed below.

7.2.1 Simulation of the collision measurement

To reenact this scenario in the simulator, two end-devices (LRDs) at equal distance from the gateway, but with a 3 dB difference in transmission power, were created. The first environment created contains the two LRDs with equal transmission start times, a one minute interval between consecutive transmissions, and a total simulation time of 100 minutes. This yields to the transmission of 100 packets per LRD and should result in the loss of all the packets originating from the LRD transmitting at 3 dB less power. Figure 38b shows that this is indeed the case. This figure also shows the received packets per LRD obtained after simulating the same situation with a time delay (in steps of 1ms) between the two LRDs. Comparing these results with the empirical results shown in Figure 38a, we can conclude that the simulator is indeed capable of mimicking the behavior of a small LoRa network. The next step is to generate an environment with more than two end-devices and without fixing their position and transmission time.

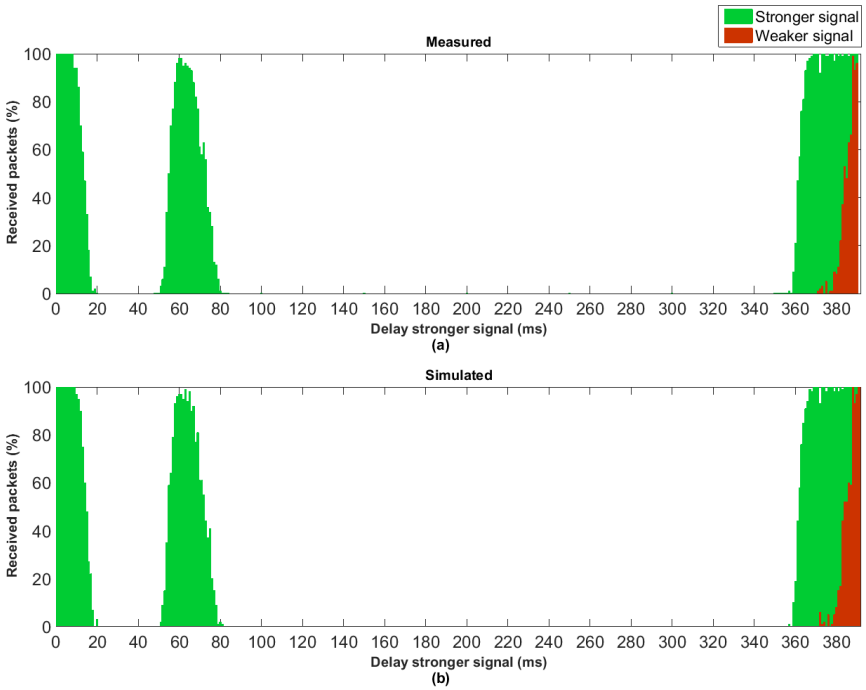

Figure 38. Received packets in a stronger-last situation based on a) empirical results b) simulation results

7.2.2 Simulating a simple environment with a few end-devices

Before simulating a situation with thousands of end-devices, a small-scale and short simulation is performed to obtain results that enable us to analyze effects of each transmission. This helps validating the ability of the simulator to generate environments with end-devices located randomly and transmitting at arbitrary intervals. To do so, an environment with 20 LoRa end-devices is created, which transmits packets on an arbitrary data rate but all on the same channel to increase the chance for a collision. This is equivalent to an environment in which there are 60 LRDs randomly selecting one of the three available channels to transmit their packet. The default parameters described in Section 7.1.6 are implemented, and the LRDs are randomly placed within a 2.5 km radius of the gateway. Each LRD will have an arbitrary starting time for its first transmission bounded by the simulated time, and tries to transmit a packet every 5 seconds. The simulated time is set to 30 seconds.

After running the simulation, five out of the 29 transmitted packets were lost due to collisions. The best way to analyze why some packets were lost is to look into the signal strength of each transmission in time, as shown in Figure 39. Here we can see the lost packets (indicated with 'lost'), and the circumstances that led to the loss of the packets. For example, both the first and the two last lost packets are caused by a stronger-last situation. In the first case, the gateway switched to the stronger signal resulting in the loss of only the weaker signal whereas in the latter collision both packets were lost.

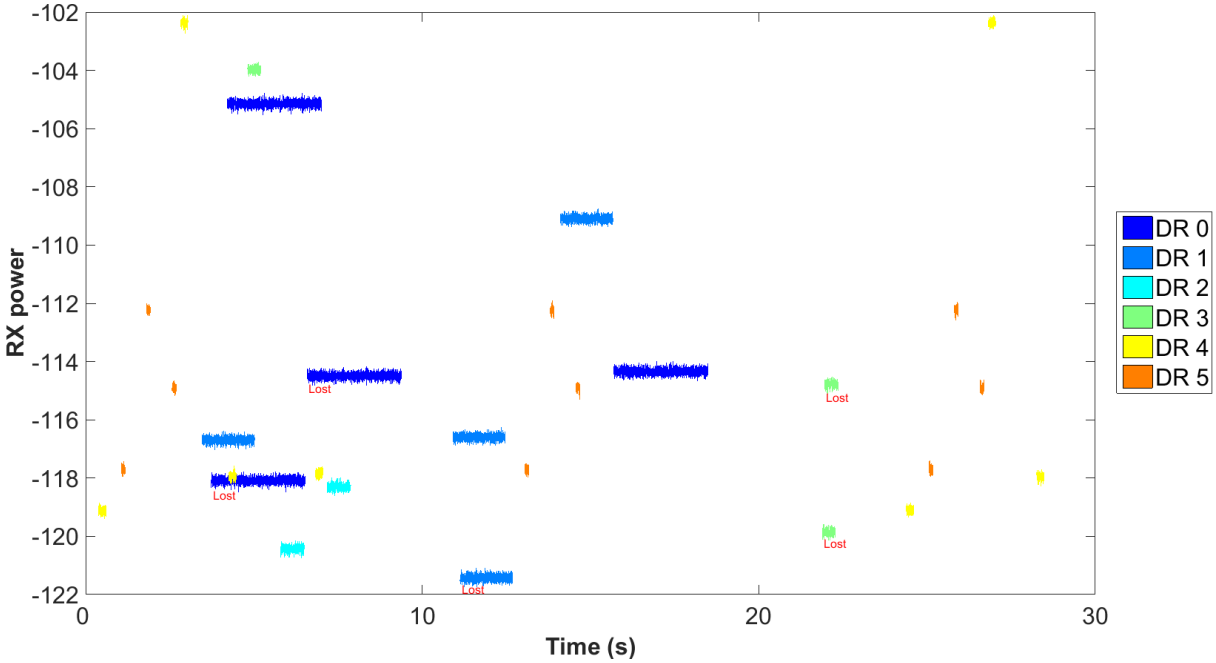

Figure 39. Received signal strength of each transmission in the time domain

Figure 40 shows the correlation between the received signal power at the gateway and the distance at which each transmitting LRD is located from gateway. This graph confirms that the implemented path loss model causes increased attenuation of signal power when the distance increases. These results, together with the results obtained from simulating the empirical experiment, prove that the simulator meets the requirements needed to simulate a dense IoT environment.

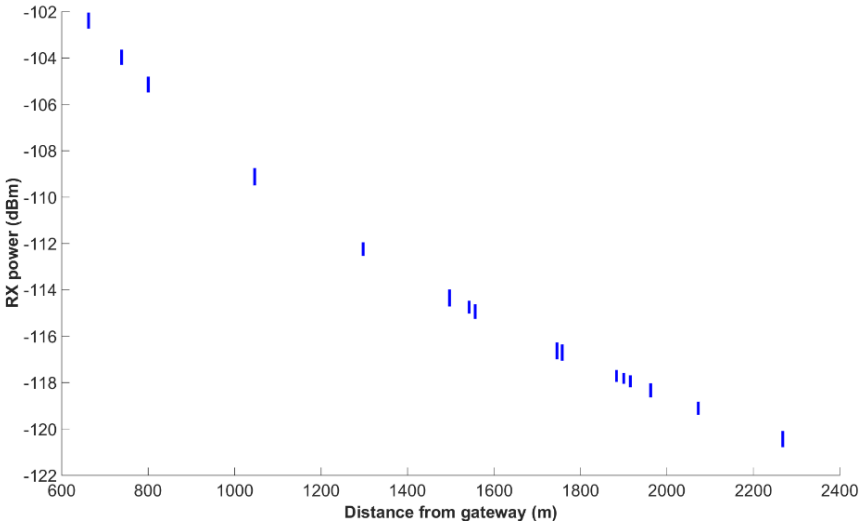

Figure 40. Received signal strength as function of the distance between gateway and transmitting LoRa device

7.3 SIMULATING A DENSE IOT ENVIRONMENT

In this section, a scenario will be created to simulate a large-scale IoT environment in which a considerable amount of end-devices are located. The environment created is based on a smart city as a reference scenario since smart cities are considered an important use case for LPWAN technology and are a good example of a situation in which a large amount of IoT devices are located in a relative close proximity to each other [14]. The purpose and characteristics of these devices depend on the application for which they are used. Examples of smart city IoT concepts include but are not limited to infrastructure & traffic monitoring, smart metering & home automation, structural health monitoring, waste management, smart healthcare & public safety services, and air quality & noise monitoring [58] [59] [60]. Many different IoT applications can be implemented within these categories, all with different characteristics. The one most relevant in a simulation of a smart city is the transmission interval of the end-devices. The transmission interval of an IoT device depends mostly on the purpose of the application, and can range anywhere between a transmission every few minutes (e.g. traffic and air quality sensing) and once a day (e.g. smart lighting). Some event driven applications do not transmit data periodically, such as wireless (fire) alarms to support public safety. In the smart city scenario for the simulator, it is assumed that all transmissions are periodic. In addition to this assumption, the following set of characteristics for the smart city environment is defined:

- The gateway is placed at a good location (e.g. a high building) and is able to cover an area with a radius of 2.5 km, which is equal to the largest part of a city like Enschede.
- Each LRD is located within the 2.5km radius of the gateway with a minimum distance of 50 meters. No LRDs are placed at the same location.
- Only one out of three LoRaWAN transmission channels is used to simulate three times the amount of actually created LRDs
- Upon the creation of an environment, each LRD is assigned a randomly selected DR from a set of available DRs. By default, there are six DRs.
- The transmission interval for each LRD is assigned randomly from a set of available transmission intervals. This set depends on the total simulated time.

This list of characteristics holds for all the simulations conducted whereas the number of LRDs, simulated time and available DRs are varied. The number of LRDs is varied between 100 and 10000 and the simulated time is set to either 1 hour or 24 hours. The associated transmission intervals from which one is selected for each LRD are defined as follows:

- For a simulated time of one hour, the available transmission intervals are 10, 30 and 60 minutes. A device can therefore either transmit 6, 2 or 1 packet(s).
- For a simulated time of 24 hours, the available transmission intervals are 10min, 30min, 1h, 2h, 6h, 12h and 24h. A device can therefore either transmit 144, 48, 24, 12, 4, 2 or 1 packet(s) during the simulated 24 hours.

The one-hour simulation is executed twice, once with the list of available transmission DRs containing all six possible DRs, and once with only DR 5. From the analysis in Section 3.5, we have seen that end-devices requiring low energy consumption are most likely configured to use

the fastest data rate (DR 5) in order to keep their transmission time short. The shorter transmission time on DR 5 will decrease the chance for collisions (i.e. packet loss), while on the other hand, the decreased diversity in data rates can lead to more packet loss. Simulations in which only DR 5 is available as transmission data rate are therefore included to see the effects of end-devices only using DR 5.

Table 19 shows the summarized characteristics of the simulations as described above.

Simulated time	Number of created LRDs	Number of simulated LRDs[3]	Data Rates	Transmission intervals (minutes)
1 hour	100-10000	300-30000	Either all DRs, or DR 5 only	10, 30 or 60
24 hours	100-10000	300-30000	All DRs	10,30,60,120,360,720,1440

Table 19. Summary of the smart city simulation characteristics

7.3.1 Simulation results

We start this section by looking into the simulation results of one individual environment, which allows us to zoom in at the transmission characteristics of a generated environment. For this analysis, the simulation results of an environment with 1000 LRDs in a 24-hour period are used. Figure 41 shows the distribution of DRs and transmission intervals over the LRDs (a and b), together with the total amount of transmitted and lost packets per DR (c and d). From these results, it is clear that the DRs and transmission intervals are more or less evenly distributed, which results in the distribution of transmitted packets per DR, as can be seen from Figure 41c. The packet loss per DR shown in Figure 41d indicates that the packets loss is the highest for packets transmitted on lower DRs. Packets transmitted on these DRs have an increased chance to collide due to their longer Time on Air, which results in loss of a large number of packets (30.76%) on DR 0 and only a few (0.13%) on DR 5 for a total packet loss of 9.67%. Transmitting every packet on DR 5, however, does not yield a big performance improvement. Repeating the simulation with the exact same environment but with each LRD transmitting on DR 5 yields to a packet loss of 8.87%, caused by the fact that each overlapping transmission is now a collision. A comparison between environments in which all DRs are used and ones where DR 5 is the only DR used is done below, when the results of the one-hour simulations are discussed.

[3] Due to the use of only 1 channel, instead of 3

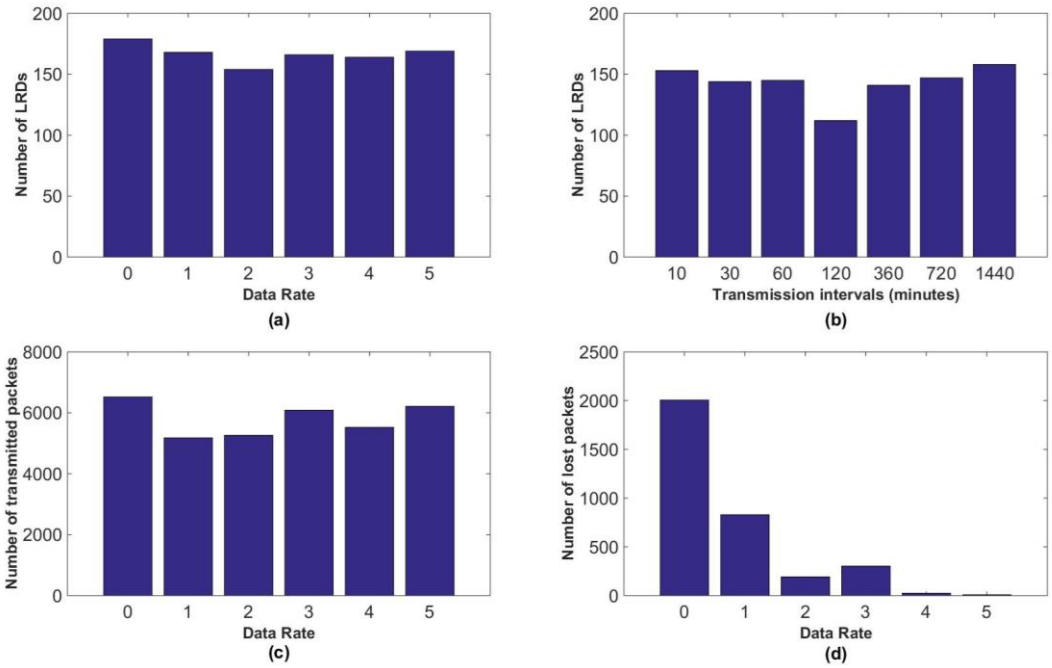

Figure 41. Simulation characteristics and results after simulation of 1000 LRDs in a 24-hour period. The top figures show a) the number of LRDs per DR and b) the number of LRDs per transmission interval, and the bottom two figures show c) the number of transmitted packets per DR and d) the number of lost packets per DR.

The above example has shown the uniform distribution of DRs and transmission intervals. A newly created environment with the same starting characteristics as the one described above can however yield to a different result. This is caused by the randomly division of the transmission DR, the transmission interval, and the first transmission time of each LRD. Although the differences are not particularly large, simulating only one environment could lead to either too pessimistic or optimistic results. The simulation of each situation has therefore been repeated ten times. In other words, for each set of environment characteristics, ten environments were created and simulated.

Figure 42 shows the average packet loss for the one-hour simulation for both situations in which each LRD is randomly assigned a DR and in which every LRD transmits on DR 5. As could be expected based on the analysis above, there is no big difference in packet loss between these situations. The packet loss is slightly better for small numbers of LRDs when only DR 5 is used, but this behavior inverts if the number of LRDs increased to above 10.000 devices, as can be seen in Figure 42. The bounds around the marker points indicate the lower and upper bound of each measurement and show how the performance for equivalent environments can fluctuate. For example, an environment with 300 LoRa Devices can result in a packet loss ranging between 0% and 5%. This indicates that it is possible to use 300 LRDs in the given environment without any packet loss by scheduling the moments of transmissions in a clever way. Looking at the magnitude of the packet loss we can notice that a few thousand devices can already induce the loss of more than 10% of the transmitted packets. This is, however, in a one-hour scenario were roughly a third of the LRDs is transmitting a packet every 10 minutes, which can quickly result in collisions between transmissions.

Figure 42. Average packet loss in a one-hour simulation for various amounts of LoRa end-devices

A similar graph as presented in Figure 42 is created for the 24-hour scenario, and shown in Figure 43. The shape of the curve equals the one obtained from the 1-hour simulation but, as expected, the packet loss for equal number of LRDs is lower due to the increased number of possible transmission intervals. The packet loss is, however, still substantial. Even for a small number of devices, the packet loss easily exceeds 10%. As already indicated by Figure 41d, this loss is not evenly distributed over different DRs due to the increased Time on Air on lower DRs. This is visible in a more clear way in Figure 44, where the average packet loss per DR for a few number of simulated LRDs is shown. From this figure, it is clear that the overall packet loss for an increased number of LRDs is initially caused by the transmissions on lower DRs.

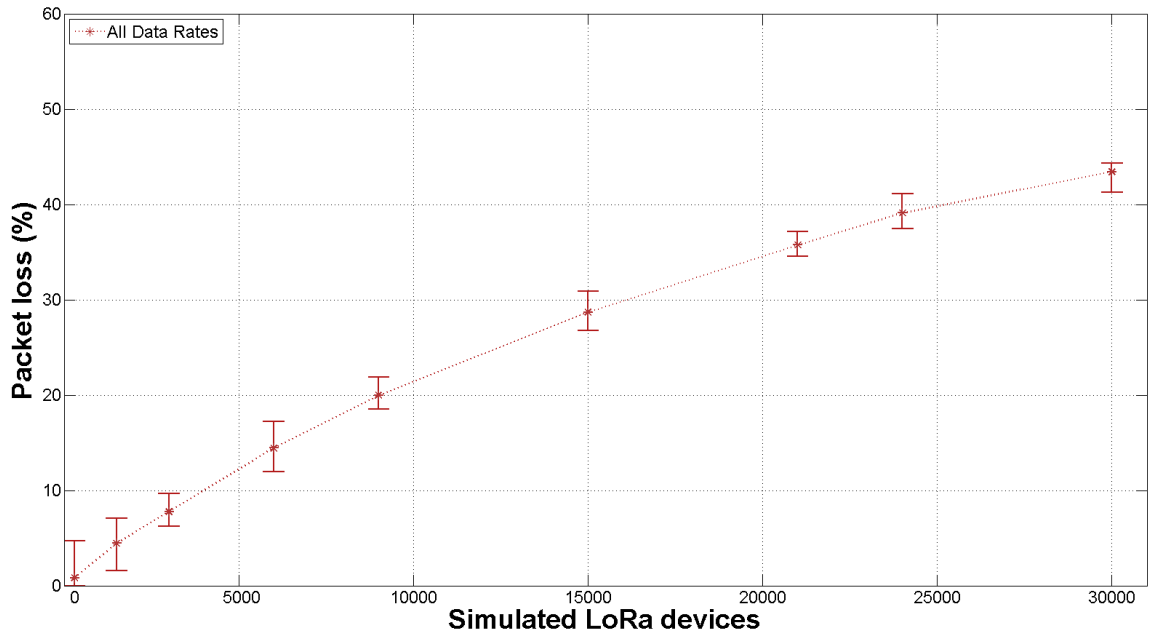

Figure 43. Average packet loss over a simulated period of 24-hour for various amounts of LoRa end-devices

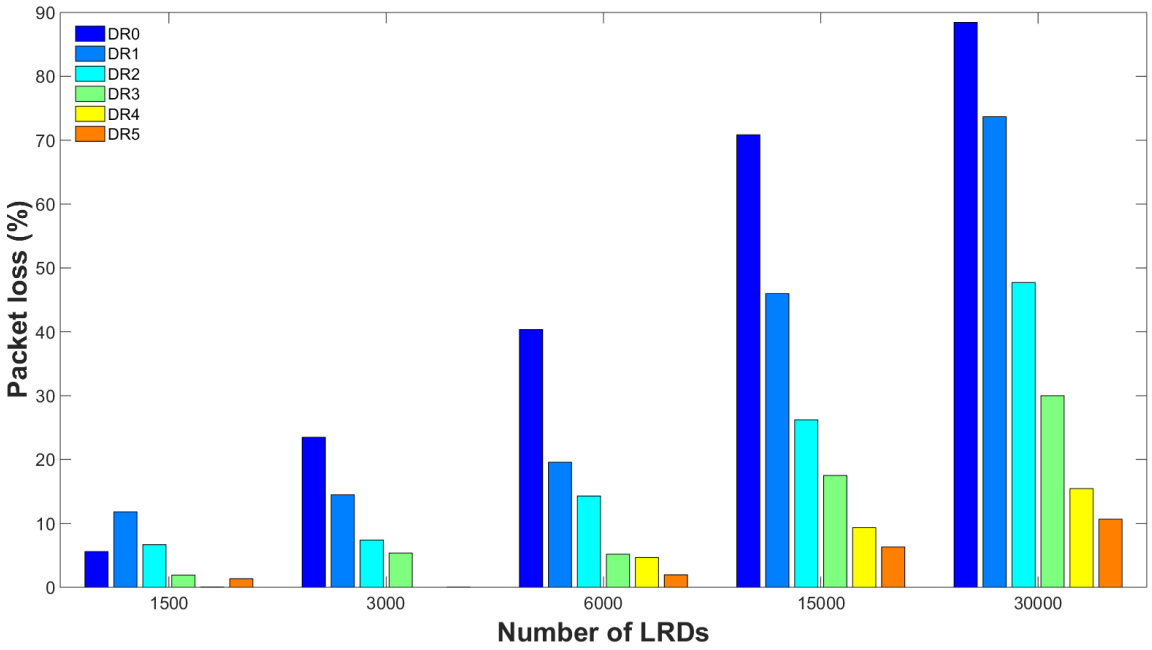

Figure 44. Average packet loss per LoRaWAN Data Rate for different numbers of LoRa end-devices

7.4 Discussion

From the simulation results obtained above it is clear that a noticeable number of packets are lost even with only a few thousand devices present. In what follows, we first discuss ways to resolve the decreased performance of a LoRa network and then we will discuss the performance of the simulator itself.

7.4.1 Dealing with collision-induced packet loss

The first solution that comes to mind to overcome the degraded LoRa performance when looking at Figure 44 is to chance the distribution of DRs among the LRDs.

7.4.1.1 *Distribution of data rates*

As already shown by Figure 42, using only the highest DR (DR 5) will not decrease the packet loss significantly. The figure shows that for a low number of end-devices the packet loss is slightly better when they all transmit on DR 5 but as the number of end-devices increases (above 10.000 in this case), the packet loss is worse than when all end-devices transmit on different DRs. The packet loss in a LoRa network can on the other hand be decreased by raising the use of higher data rates relative to the lower ones. However, this is only feasible in an environment where all end-devices can be controlled, e.g. the interference is not caused by other, external, LoRa networks. Moreover, even if there are no other LoRa networks present it would still be practically impossible to manage and maintain a network in which a fixed distribution of DRs is desirable.

7.4.1.2 *More channels*

One of the easiest solutions to cope with collision-induced packet loss is to implement additional channels alongside the three default LoRaWAN channels. It is worth noting that transmissions on additionally created channels can be bounded to stricter regulations, as can be seen in Section 2.4.2. In addition, more channels will only decrease the chance of a collision and will not prevent it.

7.4.1.3 *Listen Before Talk (LBT)*

Adding a Listen Before Talk mechanism on the end-device would make the device postpone its transmission until the occupied channel and the data rate are free. Such a mechanism would however, most likely negatively affect the energy consumption of the end-device and can, in densely populated IoT environments, lead to situations in which the end-device has to wait for a long period of time before it can start a transmission. Moreover, implementing LBT does not prevent every case of packet loss caused by a collision since the end-device is usually not as sensitive as the receiving base station and is therefore unable to accurately determine if it is the only device trying to reach a gateway at that moment.

7.4.1.4 *Acknowledgements*

Another way to ensure reliable data transfer is to request an acknowledgement. This will however cause more network traffic and will lead to even more packet loss in a dense IoT environment due to (i) overhead of the acknowledgements messages themselves, (ii) the effects of every lost packet being transmitted again, and (iii) the gateway not being able to listen for uplink messages while it transmits the downlink acknowledgement. Moreover, a gateway implementing only the three default LoRaWAN channels is only able to transmit just over a thousand acknowledgements per hour at most, as described in Section 4.2.1. This is insufficient to serve thousands of LoRa end-devices and is therefore only a solution in sparse IoT environments, in which end-devices seldom transmit data or when only a small number of end-devices requests an acknowledgement.

7.4.1.5 *Decrease gateway coverage*

One of the best solutions to cope with collision problems in a dense IoT environment is to increase the number of gateways and to decrease their coverage. This can prevent situations in which an end-device at a relatively long distance can capture the gateway, causing the loss of its packet due to a stronger nearby transmission on the same channel and data rate, and the likely loss of the nearby transmission. Having more gateways serving a smaller area ensures that the total number of end-devices reaching each gateway is small, decreasing the chance for collision-induced packet loss.

7.4.1.6 *Early collision detection*

Implementing a collision detection mechanism at the receiving end of the network can prevent gateways from trying to receive a signal that it will never be able to decode due to strong interference. This allows the gateway to timely stop demodulating the interfered message so that it is able to detect new messages. Collisions can, however, still occur and only transmissions starting in the period the gateway is now able to receive signals due to the early collision detection can be received by the gateway.

7.4.1.7 *Decrease transmission time*

Another good solution to decrease collision-induced packet loss is to decrease the Time on Air of packets. This decreases the chance of collisions, especially if there are many interfering signals such as in a smart city environment. Decreasing the Time on Air can be done by minimizing the payload of a packet. However, the ability to do so depends on the application and is not always possible.

7.4.1.8 *Use a LoRaWAN alternative*

A more drastic approach is to implement a LoRa network without using the LoRaWAN standard. This requires the implementation of a custom protocol but enables the user to (i) remove the 13 payload bytes added by LoRaWAN described in Section 3.2.3, (ii) remove the LoRa header described in Section 3.2.2 for uplink messages, and (iii) use SF 6, which is twice as fast as the fastest LoRaWAN data rate (DR 5 uses SF 7) and therefore decreases the Time on Air by a factor 2. The implicit packet format where no header is included is only possible when the payload length, coding rate, and the presence of the CRC at the end of the payload segment are manually configured at both sides of the radio link. Not using LoRaWAN makes implementing a LoRa network less straightforward but allows more flexibility in terms of the data packet composition and the use of the faster spreading factor (SF 6).

7.4.2 Performance of the simulator

The results obtained in the previous section hold for the smart city scenario as outlined in the beginning of this section, but also provide an indication for what to expect in different scenarios. Other scenarios can be easily examined by changing the parameters of the simulator, or adding new features like a Listen Before Talk mechanism to predict its influence on the network's performance. Another important addition would be the ability to simulate environments where multiple gateways are present to enable the simulation of more complex networks. The performance of the simulator itself can also be improved. For example, the computational time needed to simulate the 24-hour scenario with 10.000 LRDs currently takes more than fifteen hours on a desktop PC equipped with 8 GB of RAM and an Intel i7 3.4 GHz CPU. Simulating more complex networks (e.g. with more gateways) can drastically increase the time needed to finish a simulation. Optimizing the current simulator or even switching its implementation to a different programming language could therefore improve the performance of the simulator.

8 DISCUSSION AND CONCLUSION

Over the past years, Low Power Wide Area Networks (LPWANs) have become one of the fastest growing wireless networks. These networks are ascribed a big role in the future of Internet of Things (IoT) applications and have gained a lot of interest from both the scientific and the business world, even though much practical details about it are still unknown. An LPWAN frequently mentioned when these technologies are discussed, is LoRa. This thesis provides an analysis of LoRa and its communication standard called LoRaWAN to determine their suitability in different environments.

Our analysis of the (technical) characteristics of LoRa, presented in Chapter 3, has shown the technology is ideal for low energy consuming data transmissions over a long-range. Conversely, LoRa can only achieve a limited data throughput due to the limited payload size of each transmission in combination with the transmission regulations applied to the frequencies used by LoRa. After the theoretical analysis we created, as described in Chapter 4, a number of LoRa end-devices to verify the theoretical analysis and to conduct the experiments with. The verification has shown that our theoretical analysis was correct, and that the developed end-device is able to connect to a LoRa network. Analysis of the energy consumption has shown that our end-devices consume too much energy for long-term battery powered deployment. However, by optimizing the end-device for long-term deployment, it is possible to achieve a lifetime of several years with a LoRa end-device.

Extensive experiments were performed to evaluate LoRa performance under various circumstances. Our indoor experimental results, presented in Chapter 5, show that LoRa's performance is influenced by changes in the environment due to moving people and objects. This behavior can negatively influence the reliability of an end-device as its connectivity to the gateway becomes unpredictable. Experiments conducted in an office building show significant differences in performance depending on the location of the end-device in the building. This is caused by physical factors influencing the propagation of the radio signals. Moreover, a noticeable difference was observed in the stability and quality of signals while conducting experiments during day times and during night times (when almost no people were present). The daily activity in an office environment introduces significant fluctuations in the performance of LoRa.

Aside from the influences of a (changing) physical environment, the reliability of data transfer over a LoRa network is influenced by the presence of other LoRa communication signals. As presented in Chapter 5, simultaneous LoRa transmissions on the same channel and data rate will cause packet loss, depending on the timing of each transmission. Although, simultaneous packet transmissions are supposed to be lost due to collision, in some situations, the packets with stronger signal strengths can be received by the gateway. This phenomenon is known as the capture effect. Empirical results have shown that the stronger of the two signals is successfully received in case it arrives before the second, weaker, signal, and lost if the weaker signal is received first. In either case, the weaker signal is lost. We have shown that an exception to this principle occurs when a

stronger signals arrives during the header segment of the weaker LoRa transmission. A LoRa packet consists of a preamble, a header, and a payload. The header contains information regarding the payload segment. The receiver will no longer try to receive the rest of the LoRa packet when this information is lost due to interference from another LoRa transmission. This behavior can occur in any environment where two signals are received by the same base station. The behavior is, more importantly, independent of which LoRa network is used. Any transmission using LoRa can therefore influence the performance of any LoRa network in the neighborhood, as one of the main characteristics of LoRa technology is the ability to cover large distances (e.g. several kilometers).

The long-range capability of LoRa makes it appealing for outdoor IoT applications. For this reason, a series of experiments were conducted to determine the suitability of LoRa in mobile outdoor environments. Based on our observations in Chapter 6, we can conclude that the speed of an end-device does not influence the packet reception of applications relying on LoRa technology. The distance from its base station and the direct environment, on the other hand, do play a role. Denser environments with high buildings cause more packet loss. In such an environment, we can conclude that it is worthwhile to decrease the transmission data rate to influence the packet loss on that location in a positive way.

To determine if LoRa technology is able to provide a reliable solution for network connectivity in large-scale IoT environments, a simulator was built. The empirical results reported in Chapter 5 were used as a reference to implement the behavior of a LoRa network in the simulator. The simulator's ability to describe a LoRa environment accurately was validated and reported in Chapter 7 by recreating a measurement scenario and comparing the simulator's outcome with the results from the empirical measurements. From this, the conclusion was drawn that the LoRa simulator is suitable for testing the behavior of a LoRa network. The subsequent simulations of environments in which a large number of LoRa end-devices are present have shown that this can yield to a troublesome amount of packet loss. Even with only a few thousand devices present, a decline in successful packet reception is already observed. It is important to realize that these end-devices do not have to belong to the same LoRa network, as gateways can only determine if a LoRa packet was addressed to the network it belongs to once it has fully received it.

To decrease the chance for LoRa network congestion, one should think about implementing more than just the three default LoRaWAN transmission channels. One should bear in mind that additionally created channels in other parts of the spectrum could be bounded to more strict regulations in terms of transmission power or duty cycle. Using additional channels decreases the chance of packet loss but does not prevent it. The implementation of a Listen Before Talk mechanism on the end-devices, if their energy budget allows it, can prevent end-devices from transmitting packets on channels and data rates which are already used by another device at that moment. This can however lead to a situation in which some end-devices are never able to start a transmission in densely populated IoT environments and does not prevent every case of packet loss since the end-device is usually not as sensitive as a receiving base station.

The best approach to set up a LoRa network in a (future) dense IoT environment and to make it robust against possible interference from other LoRa networks is to implement a combination of solutions. For example, increasing the number of gateways in a city and decreasing their coverage prevents a situation in which end-devices at a relatively long distance can capture the gateway. This causes the loss of both the end-device's own message and the ones transmitted a fraction later on the same channel and data rate from end-devices closer to the gateway. Secondly, decreasing the transmission time decreases the chance for a collision to happen. Slimming down the payload of a packet is the easiest way to achieve this but is not always feasible. The length of a message can also be decreased by not using the LoRaWAN standard. This allows removal of the header in uplink messages and removes any overhead bytes present in the LoRa packet's payload segment. Moreover, LoRa's highest possible transmission data rate is not available in LoRaWAN. Although implementing a LoRa network without using LoRaWAN can be a bit more challenging, the availability of this fastest data rate will not only decrease transmission times, its use will also ensure no interference from any LoRa network implementing the proposed LoRaWAN standard.

Most of the other solutions presented in Chapter 7 only have a modest impact on performance. These include finding the best data rate to transmit on, detection of a possible collision at the base station, and the use of downlink acknowledgements. The latter can even cause more packet loss due to increased download traffic resulting in situations in which the base station is not able to receive any uplink message while it is transmitting the acknowledgements.

Based on the findings in this thesis we can conclude that the suitability of LoRa to provide a reliable solution for network connectivity in different IoT environments will largely depend on the environment of the application. In rural areas where only small number of LoRa end-devices are present, one or a few gateways can serve a large area without risking significant network congestion. In dense IoT environments such as smart cities, however, the high range of LoRa technology can work as a disadvantage and may contribute to an increased chance of packet loss. For such environments, it is therefore recommended to limit the coverage of base stations and to consider not using LoRaWAN so that the use of the fastest LoRa data rate is possible.

8.1 FUTURE WORK

The experiments described in this thesis only cover a limited set of transmission variables and environments and more work is needed to improve understanding of LoRa technology under various circumstances. Moreover, a number of measurements have been conducted over a short period of time, which sometimes has led to a limited data set to conclude from. Repeating and complementing the measurements on a large-scale, both in terms of time and geographical location, would possibly reveal more details about LoRa behavior and would certainly contribute to the validation of the conclusions drawn in this thesis.

More measurement data will also benefit the results of the simulator, since most of the simulator's implementation is currently based on measurements conducted on one data rate only. To improve the simulator, more empirical results can be added. By extending the abilities of the simulator to include, for example, the simulation of multiple gateways and additional transmission channels, more complex (IoT) scenarios can be predicted during future research. In addition, performance of a LoRa network or the impact of packet loss prevention mechanisms such as Listen Before Talk can be predicted by means of a simulation, before any implementation or deployment is needed. Lastly, the implementation of the simulator in another programming language than Matlab could be advantageous in terms of computational time, since the current simulator is primarily suitable for simulating LoRa environments with one gateway and less than a hundred thousand end-devices.

9 APPENDICES

Appendix A. Waterfall plot of LoRa up and downlink packet

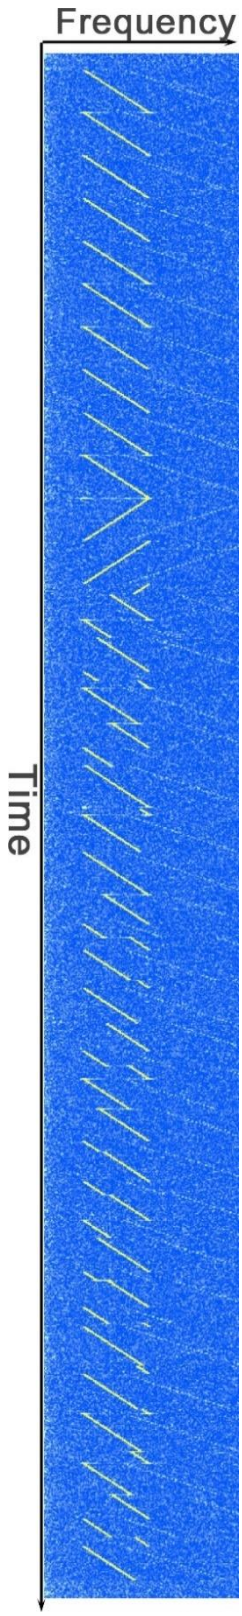

Figure A-1. Spectrogram of an uplink LoRa packet containing 35.25 symbols

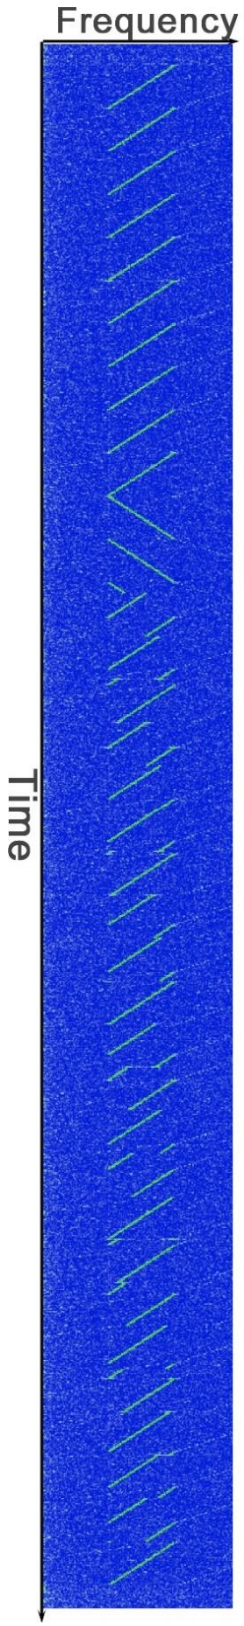

Figure A-2. Spectrogram of a downlink LoRa packet containing 35.25 symbols

Appendix B. Spectral analysis of transmission power

The image below shows how the three default LoRaWAN channels are utilized when a packet is transmitted on it. The channels are almost centered around the defined frequency of the channel; the 125 KHz wide channels, start 75 KHz before, and ends 50 KHz after their defined channel frequency.

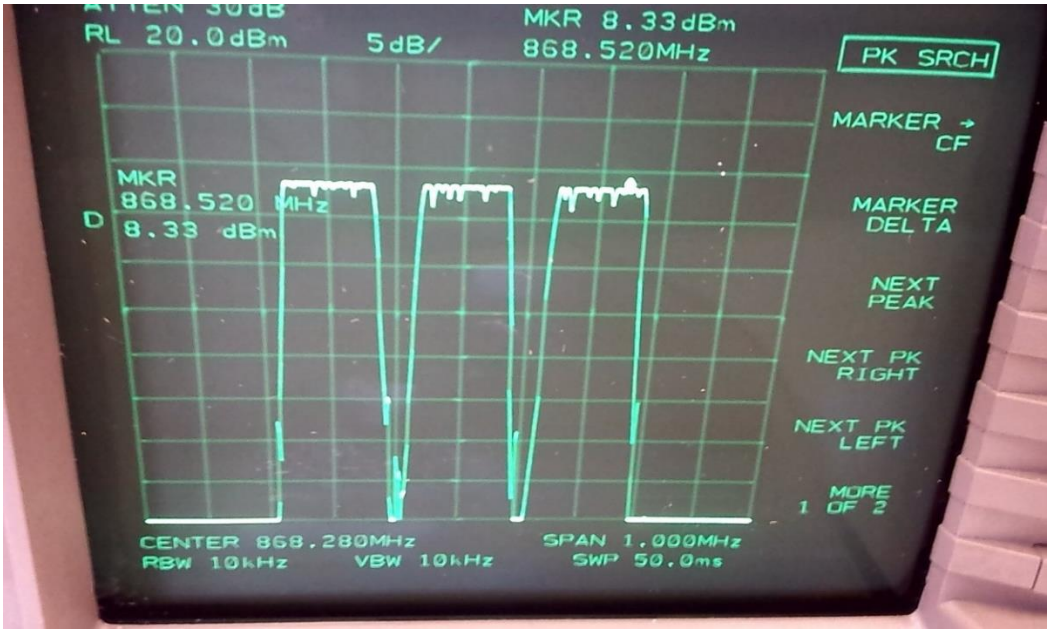

Figure B-1. The three default LoRaWAN channels on a spectrum analyzer

Appendix C. Packet loss on different locations and with different configurations

The three tables below show the results of the measurements described in Section 5.2, for the first corridor location.

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	0	0
2	0	0	0	0	0	0
3	0	0	0	0	0	0
4	1	0	0	0	0	1
Sum	1	0	0	0	0	1

Table C-1. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 5 from the first corridor location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	2	2
2	0	0	0	0	0	0
3	0	0	0	0	0	0
4	1	0	0	1	0	1
Sum	1	0	0	1	2	3

Table C-2. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 3 from the first corridor location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	2	0	0	0	0	2
2	1	0	2	0	0	3
3	0	0	0	0	2	2
4	2	1	2	0	2	7
Sum	5	1	4	0	4	14

Table C-3. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 1 from the first corridor location

The three tables below show the results of the measurements described in Section 5.2, for the first staircase location.

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	1	0	0	1	2
2	1	0	0	0	0	1
3	0	0	0	0	1	1
4	0	0	0	0	0	0
Sum	1	1	0	0	2	4

Table C-4. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 5 from the first staircase location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	0	0
2	0	0	0	0	0	0
3	0	0	0	0	0	0
4	0	0	0	0	1	1
Sum	0	0	0	0	1	1

Table C-5. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 3 from the first staircase location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	0	0
2	0	0	0	0	0	0
3	0	0	0	0	0	0
4	0	0	0	0	0	0
Sum	0	0	0	0	0	0

Table C-6. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 1 from the first staircase location

The three tables below show the results of the measurements described in Section 5.2, for the second staircase location.

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	2	8	7	4	13	34
2	2	2	4	6	13	27
3	4	4	6	5	7	26
4	0	0	3	8	8	19
Sum	8	14	20	23	41	106

Table C-7. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 5 from the second staircase location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	1	1	2	1	4
2	2	2	1	0	0	5
3	1	0	0	0	1	1
4	2	1	1	0	0	4
Sum	5	4	3	2	0	14

Table C-8. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 3 from the second staircase location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	0	0
2	0	0	0	0	1	1
3	0	0	0	0	2	2
4	0	0	0	0	1	1
Sum	0	0	0	0	4	4

Table C-9. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 1 from the second staircase location

The three tables below show the results of the measurements described in Section 5.2, for the second corridor location.

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	0	0
2	0	0	0	0	0	0
3	0	1	0	0	0	1
4	0	0	0	1	1	2
Sum	0	1	0	1	1	3

Table C-10. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 5 from the second corridor location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	0	0
2	0	0	0	0	0	0
3	0	0	0	0	0	0
4	0	0	0	0	0	0
Sum	0	0	0	0	0	0

Table C-11. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 3 from the second corridor location

TX power / Coding Rate	1 (14 dBm)	2 (11 dBm)	3 (8 dBm)	4 (5 dBm)	5 (2 dBm)	Sum
1	0	0	0	0	0	0
2	0	0	1	0	0	1
3	0	0	0	0	0	0
4	1	1	0	0	0	2
Sum	1	1	1	0	0	3

Table C-12. Packet loss for different combinations of transmission power and coding Rate for packets send on Data Rate 1 from the second corridor location

Appendix D. Outdoor mobility measurement results

The images below show the RSSI values measured during the outdoor experiments per speed.

Appendix D.1 Images of measured RSSI values at walking speed

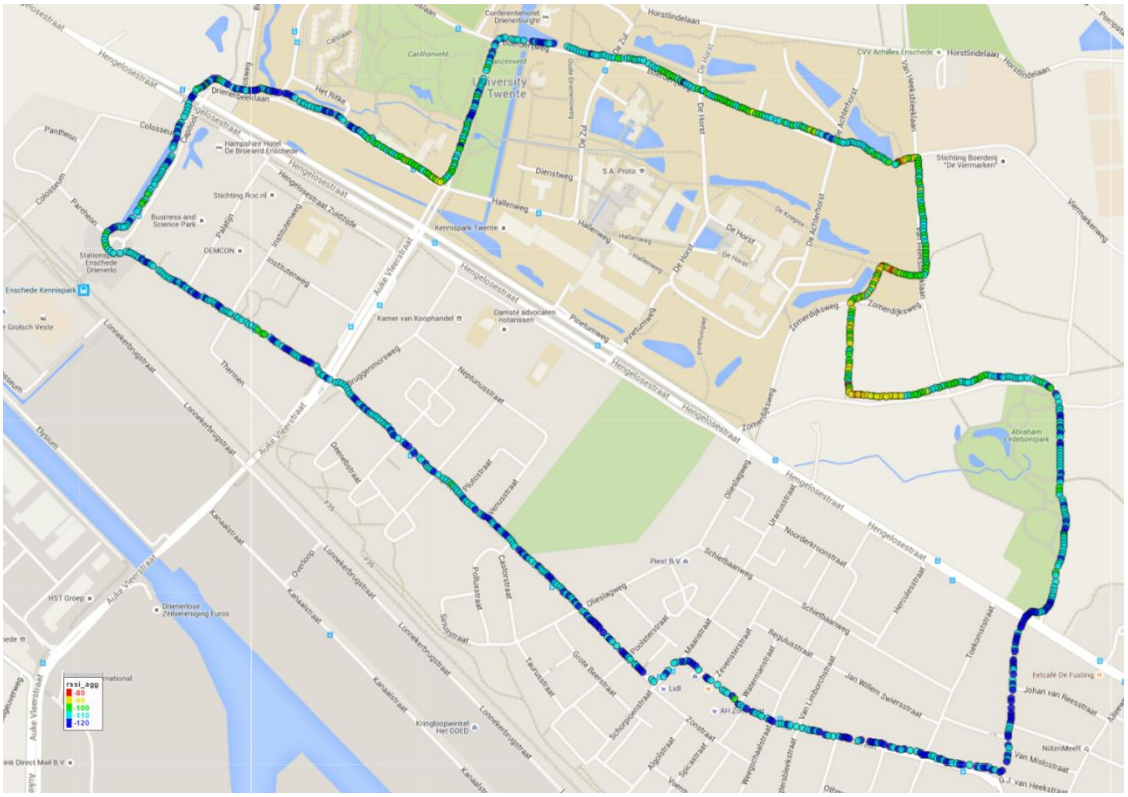

Figure D-1. RSSI values measured while transmitting packets at DR 5 and TX power 3 (8 dBm), at walking speed

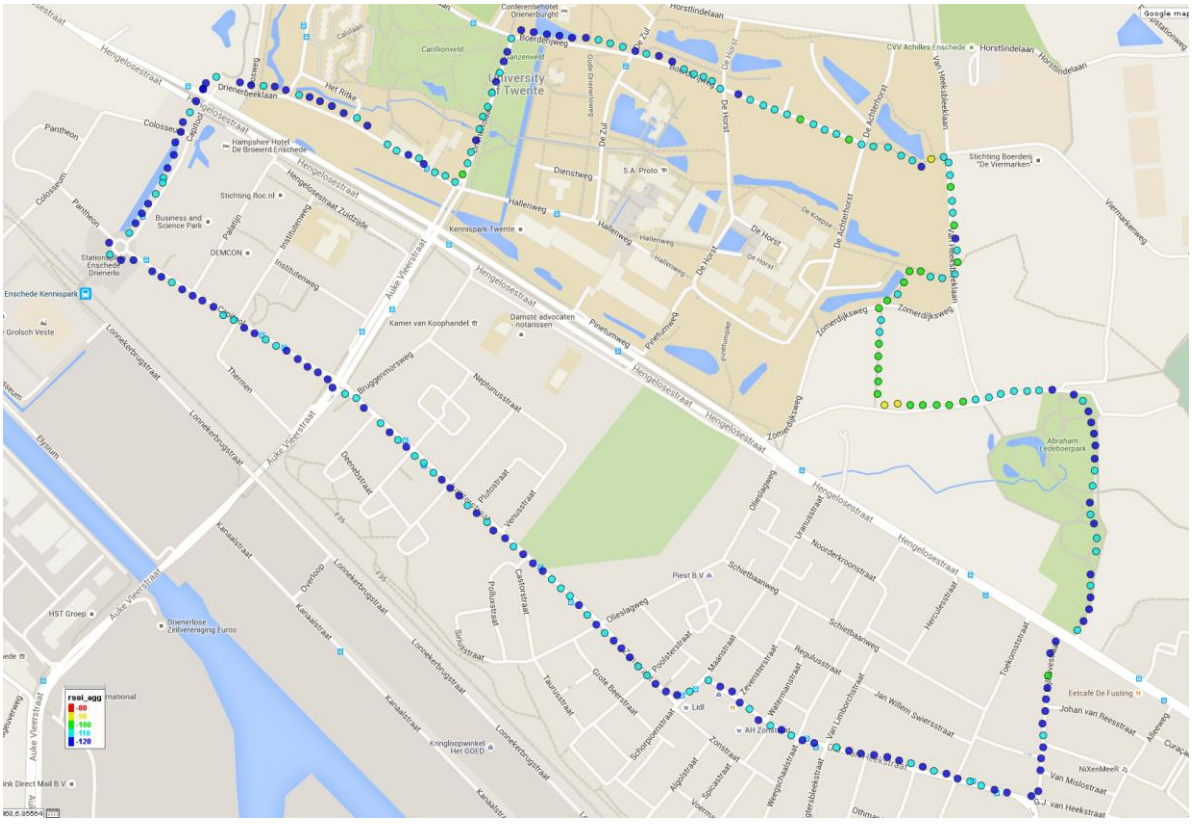

Figure D-2. RSSI values measured while transmitting packets at DR 3 and TX power 3 (8 dBm), at walking speed

Appendix D.2 Images of measured RSSI values at cycling speed

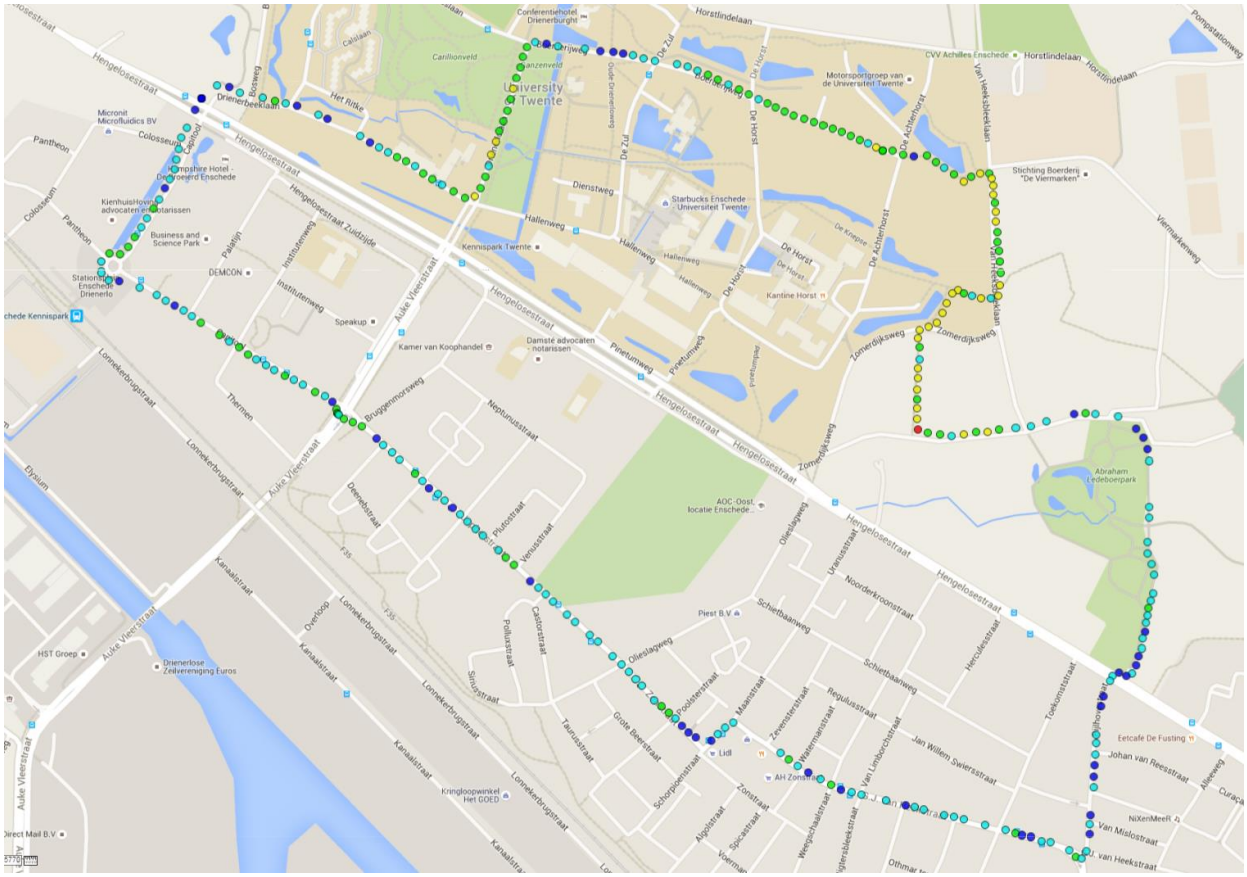

Figure D-3. RSSI values measured while transmitting packets at DR 5 and TX power 3 (8 dBm), at cycling speed

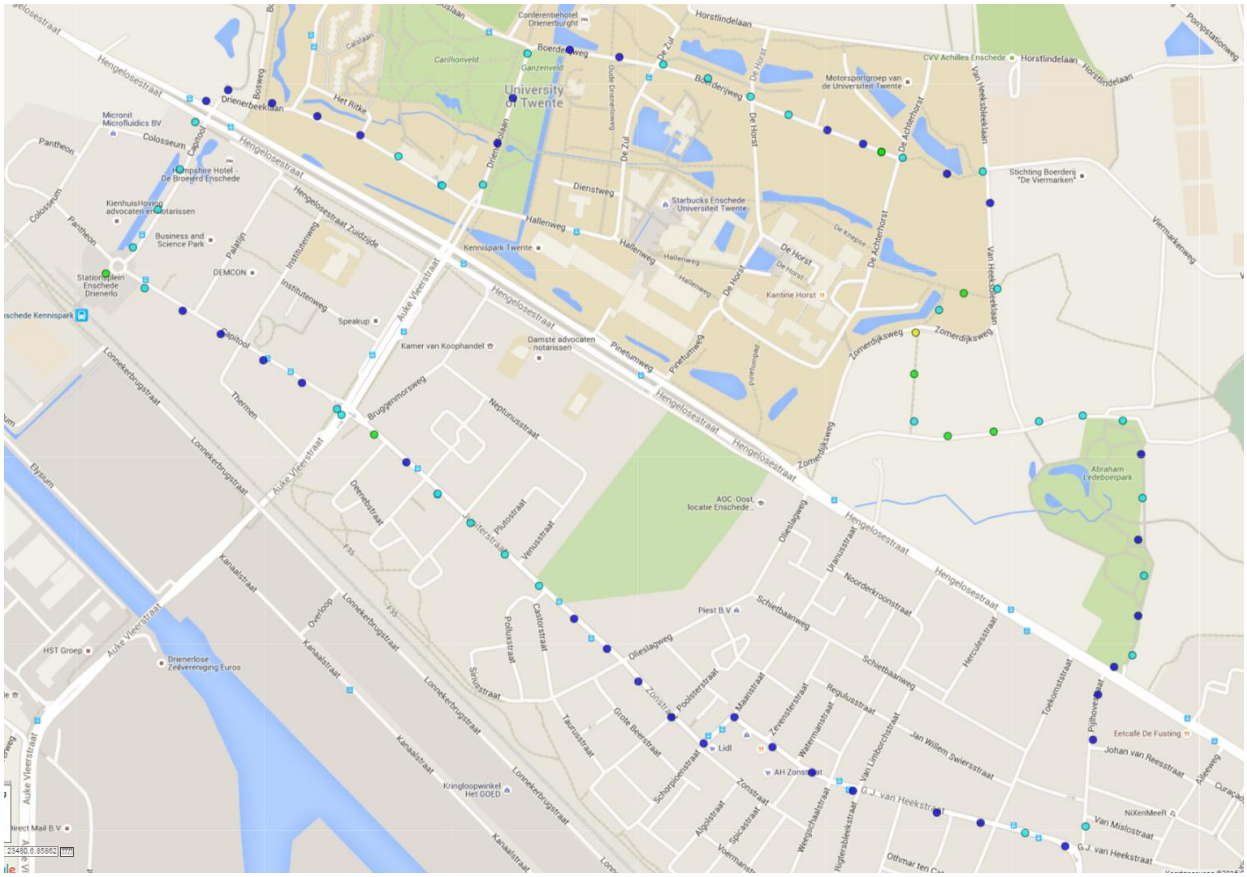

Figure D-4. RSSI values measured while transmitting packets at DR 3 and TX power 3 (8 dBm), at cycling speed

10 BIBLIOGRAPHY

[1] Dreyer, K. (2015, April 13). Mobile Internet Usage Skyrockets in Past 4 Years to Overtake Desktop as Most Used Digital Platform. Retrieved from https://www.comscore.com/dut/Insights/Blog/Mobile-Internet-Usage-Skyrockets-in-Past-4-Years-to-Overtake-Desktop-as-Most-Used-Digital-Platform

[2] Mahmoud, M. S., & Mohamad, A. A. (2016). A Study of Efficient Power Consumption Wireless Communication Techniques/Modules for Internet of Things (IoT) Applications. *Advances in Internet of Things*, 6(02), 19

[3] Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems, 29*(7), 1645-1660

[4] Palattella, M., Dohler, M., Grieco, A., Rizzo, G., Torsner, J., Engel, T., & Ladid, L. (2016). Internet of Things in the 5G Era: Enablers, Architecture and Business Models.

[5] Machina Research (2015) *LPWA Technologies—Unlock New IoT Market Potential* [White paper]. Retrieved from https://www.lora-alliance.org/portals/0/documents/whitepapers/ LoRa-Alliance-Whitepaper-LPWA-Technologies.pdf

[6] Andreadou, N., Guardiola, M. O., & Fulli, G. (2016). Telecommunication Technologies for Smart Grid Projects with Focus on Smart Metering Applications. In *Energies, 9*(5), 375.

[7] Andreev, S., Galinina, O., Pyattaev, A., Gerasimenko, M., Tirronen, T., Torsner, J., Sachs, J., Dohler, M., & Koucheryavy, Y. (2015). Understanding the IoT connectivity landscape: a contemporary M2M radio technology roadmap. *Communications Magazine, IEEE, 53*(9), 32-40.

[8] Vangelista, L., Zanella, A., & Zorzi, M. (2015). Long-Range IoT Technologies: The Dawn of LoRa™. *In Future Access Enablers for Ubiquitous and Intelligent Infrastructures* (pp. 51-58). Springer International Publishing

[9] Lawson, S. (2015, January 09). These IoT networks are 'unapologetically slow' Retrieved, from http://www.computerworld.com/article/2867331/these-iot-networks-are-unapologetically-slow.html

[10] How LPWAN drives long-range IoT applications. (2016, June 1). Retrieved from http://www.asmag.com/showpost/20578.aspx?name=news

[11] Xiong, X., Zheng, K., Xu, R., Xiang, W., & Chatzimisios, P. (2015). Low power wide area machine-to-machine networks: key techniques and prototype. *Communications Magazine, IEEE, 53*(9), 64-71

[12] Centenaro, M., Vangelista, L., Zanella, A., & Zorzi, M. (2015). Long-Range Communications in Unlicensed Bands: the Rising Stars in the IoT and Smart City Scenarios. *arXiv preprint arXiv:1510.00620.*

[13] Raza, U., Kulkarni, P., & Sooriyabandara, M. (2016). Low Power Wide Area Networks: A Survey. *arXiv preprint arXiv:1606.07360.*

[14] Hewlett Packard enterprise (May 2016). *Low Power Wide Area (LPWA) networks play an important role in connecting a range of devices* [White paper]. Retrieved from http://www8.hp.com/h20195/v2/GetPDF.aspx/4AA6-5354ENW.pdf

[15] Kranz, M. (2015, July 6). Number of Access Technologies and IoT Deployments Is Skyrocketing. Retrieved from http://blogs.cisco.com/digital/number-of-access-technologies-and-iot-deployments-is-skyrocketing

[16] The Past, Present, & Future Of LPWAN. (2015, December 9). Retrieved from http://www.link-labs.com/past-present-future-lpwan/

[17] Link Labs. (2016). A comprehensive look at Low Power, Wide Area Networks [Brochure]. Retrieved from http://cdn2.hubspot.net/hubfs/427771/LPWAN-Brochure-Interactive.pdf

[18] About SIGFOX. (n.d.). Retrieved from http://makers.sigfox.com/

[19] CNXSoft – Embedded Systems News. (2015, September 21). *Comparison Table of Low Power WAN Standards for Industrial Applications*. Retrieved from http://www.cnx-software.com/2015/09/21/comparison-table-of-low-power-wan-standards-for-industrial-applications

[20] Margelis, G., Piechocki, R., Kaleshi, D., & Thomas, P. (2015, December). Low throughput networks for the IoT: Lessons learned from industrial implementations. In *Internet of Things (WF-IoT), 2015 IEEE 2nd World Forum on* (pp. 181-186). IEEE.

[21] O'Connor, M. C. (2015, September 28). On-Ramp Becomes Ingenu, Announces Public Build-Out. Retrieved from http://www.iotjournal.com/articles/view?13553/

[22] Quinnell, R. (2015, September 15). Low power wide-area networking alternatives for the IoT. Retrieved from http://www.edn.com/design/systems-design/4440343/Low-power-wide-area-networking-alternatives-for-the-IoT

[23] LoRa Family | Wireless & RF ICs for ISM Band Applications | Semtech. (n.d.). [Online]. Retrieved June 18, 2016 from: https://www.semtech.com/wireless-rf/lora.html

[24] Abu-Rgheff, M. A. (2007). Fundamentals of spread-spectrum techniques. In *Introduction to CDMA wireless communications.* (pp. 153-194). Retrieved from http://booksite.elsevier.com/samplechapters/9780750652520/9780750652520.pdf.

[25] Nanotron Technologies GmbH. (n.d.). nanotron's Technology. Retrieved from Nanotron Technologies GmbH: http://www.nanotron.com/EN/CO_techn-css.php

[26] Semtech. (2015, May). *AN1200.22 LoRa™ Modulation Basics* [Data sheet]. Retrieved from http://www.semtech.com/images/datasheet/an1200.22.pdf

[27] MultiTech Developer Resources » Introduction to LoRa. (2016). Retrieved from http://www.multitech.net/developer/software/lora/introduction-to-lora/

[28] ETSI EN 300 220–1 V2.4.1 (2012 January), *Electromagnetic compatibility and Radio spectrum Matters (ERM); Short Range Devices (SRD); Radio equipment to be used in the 25 MHz to 1 000 MHz frequency range with power levels ranging up to 500 mW*, ETSI, 2012.

[29] ETSI TR 103 055 V1.1.1 (2011, September), *Electromagnetic compatibility and Radio spectrum Matters (ERM); System Reference document (SRdoc): Spectrum Requirements for Short Range Device, Metropolitan Mesh Machine Networks (M3N) and Smart Metering (SM) applications*, ETSI, 2012.

[30] Agentschap Telecom. (2014, June). Vergunningsvrije radiotoepassingen [Brochure]. Retrieved from https://www.agentschaptelecom.nl/sites/default/files/brochure-vergunningsvrije-radiotoepassingen.pdf

[31] European Radiocommunications Commission (ERC). (2015, September 30). Recommendation 70-03. Relating to the Use of Short Range Devices. Retrieved from http://www.erodocdb.dk/docs/doc98/official/pdf/rec7003e.pdf.

[32] LoRa Alliance (2016) “LoRa Alliance Wide Area Networks for IoT”, LoRa Alliance [Online]. Retrieved from: https://www.lora-alliance.org

[33] LoRa Alliance (2015, November) *LoRaWAN™ What is it? A technical overview of LoRa® and LoRaWAN™* [White paper]. Retrieved from https://www.lora-alliance.org/portals/0/documents/ whitepapers/LoRaWAN101.pdf

[34] LoRa Technology. (n.d.). Retrieved from https://www.lora-alliance.org/What-Is-LoRa/Technology

[35] Sornin, N., Luis, M., Eirich, T., Kramp, T., & Hersent, O., LoRa Alliance. (2015, January). LoRaWAN specification. Retrieved from https://www.lora-alliance.org/portals/0/specs/LoRaWAN%20Specification%201R0.pdf

[36] Semtech. (n.d.). LoRa Frequently Asked Questions. Retrieved from http://www.semtech.com/wireless-rf/lora/LoRa-FAQs.pdf

[37] Ritter, T. (1986). The great CRC mystery. *DR. DOBB'S J. SOFTWARE TOOLS., 11(2)*, 26-34.

[38] Semtech. (March 2015). *SX1276/77/78/79 - 137 MHz to 1020 MHz Low Power Long Range Transceiver* Data sheet]. Retrieved from http://www.semtech.com/images/datasheet/sx1276_77_78_79.pdf

[39] SX1276 137 MHz to 1020 MHz Low Power Long Range Transceiver. (n.d.). Retrieved from http://www.semtech.com/wireless-rf/rf-transceivers/sx1276/

[40] Microchip. (2015). *Low-Power Long Range LoRa® Technology Transceiver Module* [Data sheet]. Retrieved from http://ww1.microchip.com/downloads/en/DeviceDoc/ 50002346B.pdf

[41] Microchip. (2015). RN2483 LoRa™ Technology Module Command Reference User’s Guide. [Data sheet]. Retrieved from http://ww1.microchip.com/downloads/en/DeviceDoc/ 40001784B.pdf

[42] MultiConnect® Conduit™. (2016, April 19). Retrieved from http://www.multitech.com/models/94557202LF

[43] Multitech. (2016, April). *MultiConnect® Conduit™ Programmable Gateway for the Internet of Things* [Data sheet] Retrieved from http://www.multitech.com/datasheets/ 86002170.pdf

[44] Realtek RTL2832U DVB-T COFDM Demodulator USB 2.0. (2016). Retrieved from http://www.realtek.com.tw/products/productsView.aspx?Langid=1

[45] GNU Radio Companion. (n.d.). Retrieved May 31, 2016, from http://gnuradio.org/redmine/projects/gnuradio/wiki/GNURadioCompanion

[46] Semtech. (2015, January). *AN1200.24 Recommended SX1276 Settings for EU868 LoRaWAN Network Operation* [Data sheet]. Retrieved from http://www.semtech.com/images/datasheet/an1200.24.pdf

[47] ETSI TS 136 124 V13.0.0. (2016). *LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Electromagnetic compatibility (EMC) requirements for mobile terminals and ancillary equipment (3GPP TS 36.124 version 13.0.0 Release 13).* Retrieved from http://www.etsi.org/deliver/etsi_ts/136100_136199/136124/13.00.00_60/ts_136124v130000p.pdf

[48] Rodríguez Pallares, J. E. (2015). Wireless sensor network implementation with Arduino and Xbee.

[49] Schwartz, M. (2013, August 07). How to Run an Arduino for Years on a Battery - Open Home Automation. Retrieved, from https://www.openhomeautomation.net/arduino-battery

[50] Berg, O. (1998). Spread Spectrum in Mobile Communication (No. 40). IET.

[51] Patras, P., Qi, H., & Malone, D. (2012, June). Exploiting the capture effect to improve WLAN throughput. In *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2012 IEEE International Symposium on* a (pp. 1-9). IEEE.

[52] Whitehouse, K., Woo, A., Jiang, F., Polastre, J., & Culler, D. (2005, May). Exploiting the capture effect for collision detection and recovery. In *Proceedings of the 2nd IEEE workshop on Embedded Networked Sensors* (pp. 45-52).

[53] Oppermann, F. J., Boano, C. A., & Römer, K. (2014). A decade of wireless sensing applications: Survey and taxonomy. In *The Art of Wireless Sensor Networks* (pp. 11-50). Springer Berlin Heidelberg.

[54] Frenzel, L. (2012, October 11). The Fundamentals Of Short-Range Wireless Technology. Retrieved from http://electronicdesign.com/communications/fundamentals-short-range-wireless-technology

[55] Adafruit. (2016, May 17). Adafruit Ultimate GPS Logger Shield [Data sheet]. Retrieved from http://cdn-learn.adafruit.com/downloads/pdf/adafruit-ultimate-gps-logger-shield.pdf

[56] Goldsmith, A. (2005). *Wireless communications*. Cambridge university press.

[57] Abhayawardhana, V. S., Wassell, I. J., Crosby, D., Sellars, M. P., & Brown, M. G. (2005, May). Comparison of empirical propagation path loss models for fixed wireless access systems. In 2005 *IEEE 61st Vehicular Technology Conference* (Vol. 1, pp. 73-77). IEEE.

[58] Jalali, R., El-Khatib, K., & McGregor, C. (2015, February). Smart city architecture for community level services through the internet of things. In *Intelligence in Next Generation Networks (ICIN), 2015 18th International Conference on* (pp. 108-113). IEEE.

[59] Zanella, A., Bui, N., Castellani, A., Vangelista, L., & Zorzi, M. (2014). Internet of things for smart cities. *IEEE Internet of Things Journal, 1*(1), 22-32.

[60] Hancke, G. P., & Hancke Jr, G. P. (2012). The role of advanced sensing in smart cities. *Sensors, 13*(1), 393-425.