Delft University of Technology
Master of Science Thesis in Embedded Systems

DynamicVLBC: Battery-less Visible Light Backscatter Communication in Dynamic Conditions

Lucan de Groot

Embedded Networked Systems

TU Delft Delft University of Technology

DynamicVLBC: Battery-less Visible Light Backscatter Communication in Dynamic Conditions

Master of Science Thesis in Embedded Systems

Embedded and Networked Systems Group
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Mekelweg 4, 2628 CD Delft, The Netherlands

Lucan de Groot

August 22nd, 2022

Author
Lucan de Groot

Title
DynamicVLBC: Battery-less Visible Light Backscatter Communication in Dynamic Conditions

MSc Presentation Date
August 30th, 2022

Graduation Committee

Assoc. Prof. dr. ir. Marco Zúñiga	Delft University of Technology
Asst. Prof. dr. ir. Jie Yang	Delft University of Technology
Ir. Talia Xu	Delft University of Technology

Abstract

Visible Light Communication (VLC) has been gaining interest in the industry and academia for the last decade. VLC enables a high-speed communication alternative to conventional radio such as Bluetooth or WiFi and presents a solution to the 'spectrum crunch'. More recently, the combination of energy harvesting and VLC has been explored to enable battery-less devices that can communicate bidirectionally using light. In parallel, drones are being used in the industry for tasks such as warehouse management. However, very little research has been done on the conjunction of VLC and drones, even though this offers interesting research opportunities and applications.

In this thesis, we perform the first evaluation of different types of modulation techniques between a drone and base station in the context of VLC. We present DynamicVLBC: a complete system consisting of a drone ('Reader') and base station ('Tag'). We optimize this system such that the Reader can fly and the Tag can operate battery-less at an ultra-low power level. We thoroughly evaluate the system in both indoor and outdoor conditions. Our evaluations show that when the Reader is static and the Tag is externally powered, that the system can communicate up to 200 cm with a BER <1%. Moreover, when the Tag is operating battery-less, the system can still effectively communicate up to 150 cm. Finally, when the Reader is airborne as well, we show that the system can still communicate up to 85 cm.

Preface

I have been interested in low-power embedded systems since an internship I did during my early studies. After finishing my bachelor I realized that I wanted to learn more, so I started a pre-master in preparation for a master in Embedded Systems. During my master I learned a ton about various topics, including smart and alternative communication methods. When searching for possible graduation projects, I found the Embedded and Networked Systems (ENS) group and contacted their representative: Marco Zúñiga. After Marco introduced me to the possible topics the ENS group researched, I quickly realized that VLC was something unique and fascinating and I wanted to do more with it.

This decision resulted in a wonderful graduation project that lasted nearly 10 months. I would like to thank my daily supervisors: Talia Xu for her valuable feedback, assistance and availability and Marco Zúñiga for his expert guidance, feedback and supervision. Also I would like to thank my partner, Amber, and my parents for their continued support and tolerance for the days where I locked myself in my room to work on this project day and night.

Lucan de Groot

Delft, The Netherlands
August 22nd, 2022

Contents

Chapter 1

Introduction

With the rise of connected devices and the exponential increase in mobile data traffic, the radio spectrum is becoming ever more crowded [1, 2, 3]. Visible Light Communication (VLC) offers a novel solution to this 'spectrum crunch'. VLC takes advantage of the visible light spectrum, which ranges from 380 nm to 750 nm and offers a bandwidth that is orders of magnitude wider than that of WiFi [4]. In VLC, wireless communication between two devices is achieved by modulating the intensity of an LED light source at high speed at the transmitter and detecting the change with an optical sensor at the receiver [2]. The high-speed modulation is invisible to the human eye and will not disturb the illumination functions of LEDs. In the meantime, the fast variations in the light intensity can be picked up by the optical receiver. Previous works have enabled communication with data rates up to 100 Gbps using off-the-shelf LEDs [2].

While VLC has the potential of addressing some of the challenges that RF communication is facing, it also has several limitations. For example, VLC requires access to the driver circuitry of LED lights, which might not be possible in infrastructure that has already been deployed. Additionally, LED lights can still consume significant power to communicate, making them non-ideal for pervasive deployment, especially in IoT settings. Visible Light Backscattering Communication (VLBC) addresses some of these issues by employing ambient light for communication. In VLBC, external optical surfaces, which consume µW of power, are used as transmitters for modulating ambient light. One example of such optical surfaces is the liquid crystal (LC) shutters. These shutters are transparent by default and can be made opaque by applying a small DC voltage to its terminals. By quickly changing between the two DC voltage levels, LC shutters can change between clear and opaque states, thus allowing or blocking light from going through its surface. LC shutters consume a small fraction of the power of LED lights, enabling ultra-low power applications, where devices can operate without a battery and rely on only power harvested using small solar panels [5, 6].

Active VLC and VLBC both have a wide range of applications. Active VLC can be used in high-speed Li-Fi communication, such as indoor communication and positioning; vehicle-to-infrastructure or vehicle-to-vehicle communication (Figure 1.1a); and underwater communication. VLBC on the other hand, allows applications such as interactive toys, passive smart tags (Figure 1.1b) or low power IoT communication.

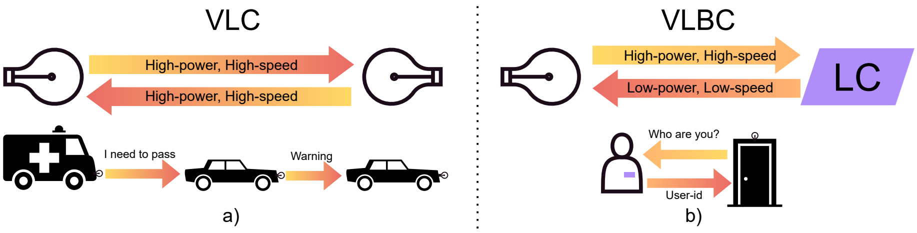

Figure 1.1: **Example illustration of conventional VLC and VLBC.**

Application of interest. While there are many applications that are taking advantage of Active and Passive VLC, there is an area that has not been investigated much: the use of VLC technologies for drones, and this thesis aim at exploring that direction. In recent years, the use of drones in sensing and delivery are in the rise. For example, utilizing drones as an alternative to perform visual inspections in areas that are difficult to reach or possibly unsafe for workers e.g. manufacturing environments and nuclear power plants. An additional use case of interest is for autonomous inventory management, where information can be collected by drones reading tags or barcodes across a warehouse [7]. Some companies like Eyesee, DroneScan, Doks Innovation and Infinium Robotics already provide warehousing services using drones [8, 9, 10, 11].

1.1 Problem statement

Generally, all drones use GPS, radio and cameras to position themselves and interact with the environment. These systems can fail, or be misused for unconsented recordings. Because of these reasons, recent regulations allows the use of commercial drones only at night [12].

In this thesis, we explore the use of both VLC and VLBC in drone communication as a secure and privacy-preserved alternative. Drones carry LED lights that can modulate to discover and send commands to tags, and tags use ultra-low power optical surfaces respond:

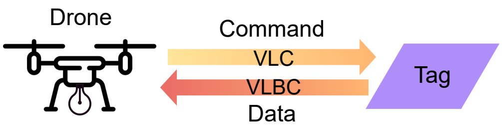

Figure 1.2: **Visualization of the proposed communication setup.**

This allows us to take advantage of the benefits of both VLC and VLBC: VLC for reliable high speed commands and VLBC for efficient low-power responses. In addition, the LEDs carried by the drone can also be used to power the tags to provide sufficient power for communication, even at night when no light is present.

An example application could be smart farming or warehouse management. In such an application, we envision a drone carrying a light that obtains information from pervasively deployed battery-less tags. The tags collect information about humidity, temperature, or contain inventory information, which is then collected by the drone to be processed in a central unit. Only the tags covered by the light beam of the drone will be able to communicate with it. There are two main components:

- Tags: They are stationary and battery-less and are able to collect local sensing data.
- Drones: They fly around to collect the sensed data from tags.

Using VLC and VLBC to communicate with drones present unique challenges that have not been explored previously. In the majority of state-of-the-art VLC studies, both the transmitters and the receivers are static, and careful alignment is required to establish point-to-point communication [2]. On top of that, decoding of the received signal is often done offline, providing a generous power budget [13, 14]. However, having the drones to stay completely still when communicating is unrealistic. The movements from a flying drone exacerbate the noise on the communication link between the drone and the tags. In addition, significant additional weight or power consumption can greatly restrain the flight time of a drone. This requires us to explore and design a system that is both resilient to noise and light-weight, which is able to establish reliable communication links on resource-constrained hardware. On the side of the tags, the battery-less nature of the tags can increase the response time in establishing a working communication. This requires us to optimize our system for low-power operation.

1.2 Contributions

In this thesis, we design and evaluate a new type of communication link: A bidirectional link between an airborne drone and a battery-less tag using VLC. Specifically, the contributions of this thesis are summarized as follows:

- Contribution 1: *Evaluation of different modulation schemes and an efficient implementation on resource-constraint hardware [chapter 4].* We evaluate how well different modulation schemes perform in establishing a reliable VLC link in a challenging environment with interference from drone movements and ambient light. Our evaluation shows that despite its higher complexities, frequency modulation provides the most reliable communication. Given this insight, we design and implement an efficient modulation and demodulation scheme and further optimize it to work on resource-constrained hardware.
- Contribution 2: *Design of a complete end-to-end system [chapter 5 and chapter 6].* We implement a complete system to allow a custom-designed reader mounted on a drone to discover and communicate with a battery-less tag. We identify and tackle the unique challenges on both sides: On the drone, we optimize our system to be lightweight and power conscious to function reliably when flying. On the tag, we conduct a thorough evaluation of the power profile of the system and optimize it to operate with only the harvested power.

- Contribution 3: *A thorough evaluation of the system [chapter 7].* We present a thorough evaluation of the system in both indoor and outdoor conditions. Our results show that the system can communicate up to 200 cm with a BER <1 % when the drone is not flying and the Tag is externally powered. The range decreases to 150 cm when the Tag is operating battery-less and to 85 cm when the drone is flying as well.

1.3 Organization

First, we briefly discuss the background of VLC, introduce some fundamental topics and present the state of the art. Then, we will present a high-level overview of the system, and underline the design challenges. After this, we will evaluate the commonly used modulation schemes in VLC, focusing on their trade-offs in the context of our implementation and detailing how we have implemented the most suitable technique. Next, we will detail the prototype of the drone-mounted and battery-less devices. Finally, we will evaluate the performance, give a conclusion and discuss any further work.

Chapter 2

Background

In this chapter, we present the fundamental concepts necessary in understanding this thesis. First, we briefly iterate the history of VLC to show its evolution over time into its modern form. After that, we will describe the system overview of VLC and VLBC to show how VLBC enables bidirectional communication using just a single light and is able to operate without a battery.

2.1 A brief history of VLC

The first documented use of visible light for communication in history starts with Polybius in ancient Greece, who developed a communication system using 5 torches [15]. In this system, messengers raise different numbers of torches to signal different messages to each other according to an agreed code. Later, in the 19th century, heliographs where used for positioning and communication during geodetic surveys [16]. These devices operate by moving a mirror back and forth and reflecting sunlight. The last major pre-electronic VLC invention is the Photophone developed by Alexander Bell, also the inventor of the regular telephone. This device worked with a vibrating mirror and photoresistor, and transmits messages using the vibration induced by sound [17].

Although the concept of using light as a medium to transmit information occurred early on, the modern form of VLC started recent in history with the rising popularity of LED. A major milestone in VLC is the 2011 TED talk by Harald Haas, who coined the term Li-Fi (Light Fidelity) [18]. This talk gained millions of views and significantly boosted the amount of research done in VLC [2]. Recent achievements include discussions to include VLC in future wireless technologies (5G, 6G); Pli-Fi, which combines Wi-Fi, power line communication and VLC; and the NASA Lunar Laser Communication Demonstration [2].

2.2 Architecture of VLC

Figure 2.1 provides an overview of a typical one-way VLC system. The system consists of three fundamental components: the transmitter, channel and receiver. This allows for one-way communication between two devices. For duplex communication using VLC, both sides need to be equipped with both a transmitter and a receiver. The following sections will give a brief overview of the aforementioned components.

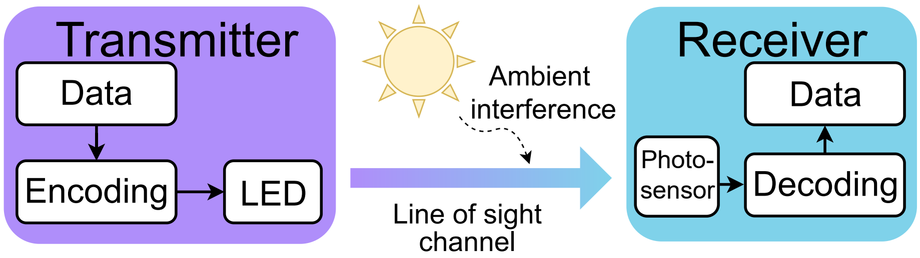

Figure 2.1: **A typical VLC system visualizing the three fundamental components: a transmitter, channel and receiver.**

2.2.1 Transmitter

The transmitter transmits data by repeatedly switching a light on and off, modulating `1`'s and `0`'s. Typically, a Light Emitting Diode (LED) is used due to their low cost, high popularity, high energy efficiency and low switching time compared to other lights. Moreover, even though white LEDs are most commonly used, RGB lighting can be used as well [2, 19]. The exact specifications of LEDs used for VLC transmitters vary widely and depend on the specific design parameters of the system. Key properties of a LED are the amount of luminous flux (lumens) produced and the Field of View (FoV) or beamwidth. Generally, as the transmitter transmits more energy (lumens), the receiver will receiver more energy, increasing the received signal strength and the signal to noise ratio (SNR). On the other hand, the beamwidth dictates how focused the light beam is, and a narrow beam makes it possible to focus more light on the receiver.

2.2.2 Channel

Once a signal is transmitted, it enters the communication channel between the transmitter and receiver. For most applications, VLC requires Line of Sight (LoS) to communicate effectively, since even a sheet of paper blocking the LoS could severely attenuate the signal. In the channel, the signal is mainly attenuated due to path loss and interference from ambient light [2].

2.2.3 Receiver

The modulated light is received using a photosensor, typically a photodiode (PD). However, other light-sensitive components such at phototransistors, LEDs or cameras can be used as well. PDs output a current proportional to the amount of light illuminating it, which is converted to a voltage using a transimpedance amplifier. To minimize interference, incoming light can be filtered using an optical bandpass filter. Additionally, lenses or other materials can be placed around the photosensor, such that they can be used to focus the PD to reduce the FoV and thus reducing the impact of ambient light.

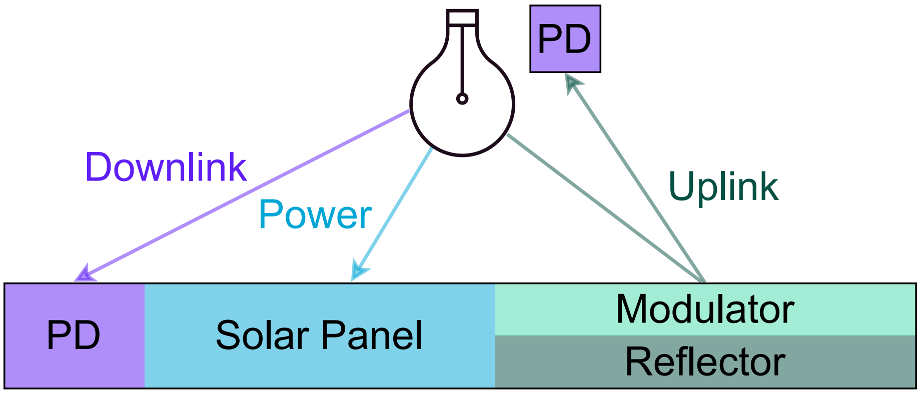

Figure 2.2: **An illustrative example of how VLBC works.**

2.3 Visible Light Backscattering Communication

To enable bi-directional VLC communication, both sides need to be equipped with an LED light and a photosensor. However, in many IoT applications, the two sides communicating with each other do not always have equal resources. For example, in the drone-tag communication scenario, the tag is a battery-less device operating only on harvested power to allow them to be pervasively deployed. A small LED consumes 75 mW of power [20], which is significant to a battery-less device, making it impractical to employ an LED for every tag. Therefore, low-power applications use *backscattering* in order to enable VLC while remaining on a low power budget. In the context of VLC, backscattering is called Visible Light Backscattering Communication (VLBC).

Figure 2.2 illustrates the general layout of a two-directional communication link with VLC and VLBC, which we will detail in the following sections. When a device is backscattering a signal, instead of producing its own light to modulate its data (a LED to turn on and off), it uses external ambient light. The ambient light can come from a device attempting to establish a communication link, or any ambient light in the environment. With VLC, backscattering is achieved by combining two separate components: a reflector and modulator. The reflector reflects the light back towards the source, while the modulator changes the intensity of the reflected light, in order to modulate data.

In our application we will have two devices: a *Reader* and *Tag*. The Reader will be mounted on a drone and operates as a regular VLC-capable device using a LED. The Tag will operate battery-less and use backscattering to achieve the ultra-low-power required for battery-less operation. In the following sections we detail the general components of a battery-less Tag, as illustrated in Figure 2.2.

2.3.1 Modulator

The modulator modulates the carrier, blocking or passing light to modulate 1s and 0s, creating the uplink. This is typically done with Liquid Crystal (LC) Shutters [5, 6, 21, 22, 23, 24, 25] mounted against the reflector.

Figure 2.3 visualizes how a LC shutter works. LCs consist of three layers: two polarizing layers with a liquid crystal layer in between. The first polarizer only

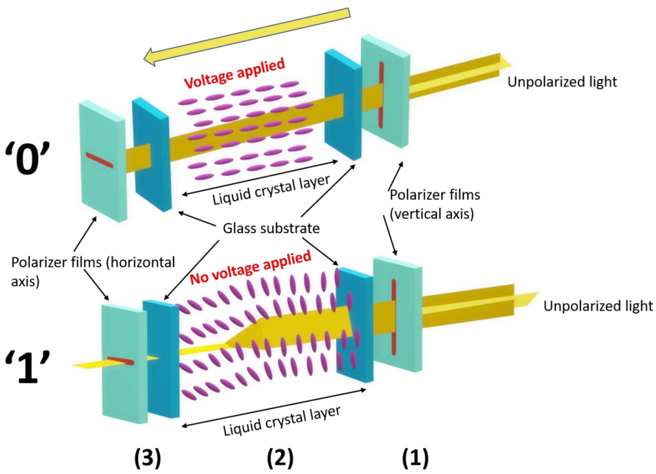

Figure 2.3: **Visualization of how a LC shutter becomes opaque and transparent. Figure sourced from [22].**

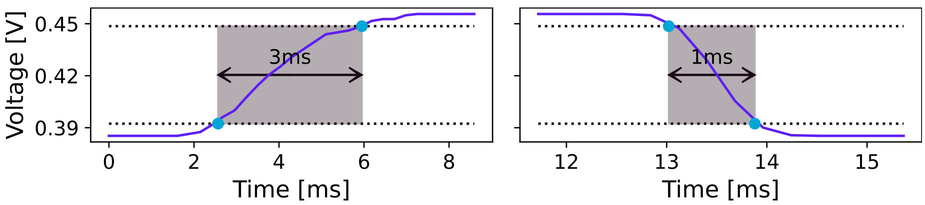

Figure 2.4: **Image visualizing the slow and asymmetrical rise times of a LC.**

allows vertically aligned light though. The second layer (liquid crystal) either lets the light through normally or rotates the light by 90 degrees. The final layer is a horizontal polarizer, which only allows horizontally aligned light through. If the light is rotated by the liquid crystal layer, then it will be blocked by the final polarizer layer, which requires light to be of the correct orientation. The liquid crystal layer can be switched between states by applying a voltage to its pins; creating an electric field over the liquid crystal, twisting them and making the LC opaque or transparent.

The mechanical process of physically twisting the liquid crystals is slow and takes somewhere in the order of milliseconds, depending on the specific LC and driving voltage [22]. Moreover, twisting the crystals into place is generally much slower than resetting their position back. This results in highly asymmetrical rising and falling edges, as shown in Figure 2.4.

2.3.2 Reflector

The purpose of the reflector is to reflect the modulated light back to the Reader, there are different ways to achieve this, which are illustrated in Figure 2.5. An obvious choice for a reflector would be a mirror. However, as illustrated, this requires ideal alignment between the Reader and the Tag, since a mirror reflects

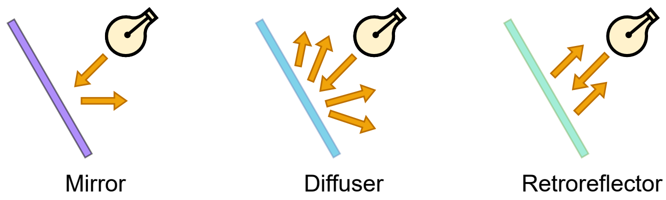

Figure 2.5: **Illustration of different types of reflectors for VLBC.**

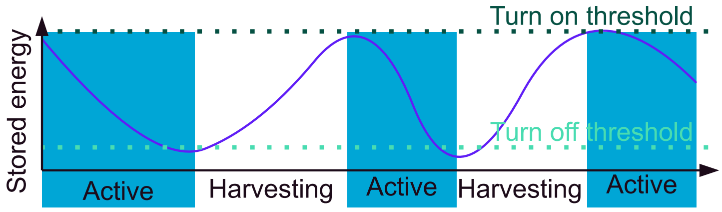

Figure 2.6: **Illustration visualizing how energy harvesting works.**

light outwards based on the incidence angle. Without near-perfect alignment between the Reader and the Tag the reflected light would miss the Reader.

To overcome this issue, one could use a diffuse material such as a sheet of white paper. However, this would significantly reduce the Signal-to-Noise Ratio (SNR) since it diffuses the light in all directions, most of which would miss the Reader. For this reason, it is typical to use retro-reflectors in VLBC [5, 6, 21, 23, 24]. Since they reflect light back to the source with little dispersion. A retro-reflector can be a large, three dimensional, device such as a corner cube, a flat tape or a fabric.

2.3.3 Energy harvesting

In order to allow the Tag to operate battery-less, we need to harvest energy from the surroundings to power the device. Since the LED used for communication illuminates the Tag, it is convenient to harvest energy from the downlink using solar panels. However, other methods such as piezoelectric, thermal or electromagnetic can also be used to harvest energy [26].

Figure 2.6 illustrates how energy harvesting works. Namely, an energy harvesting system consists of a producer, storage element and consumer. In the case of VLBC, the producer is a solar panel, which harvests energy from the LED mounted on the Reader. The storage element is typically a supercapacitor, but could be a small Li-Po battery. The consumer is the Tag and its components that receive the uplink, process the data and produce the downlink.

The power budget for a battery-less Tag is small, consuming between 90 µW and 350 µW [5, 6, 21]. Other works consume slightly more power, between 0.8 mW and 50 mW [22, 24].

Chapter 3

System overview

In our system, we establish a communication link between two devices. One of these devices is mounted on top of a drone and is dubbed the *Reader*, the other device will be placed in the environment and is called the *Tag*, as visualized in Figure 3.1.

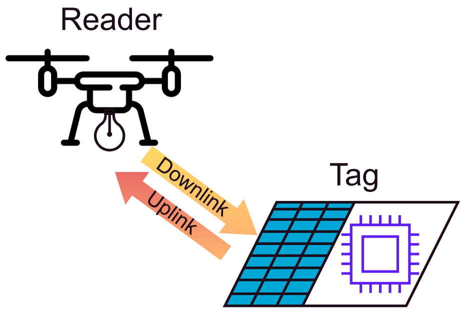

Figure 3.1: **Visualization of how the system is supposed to operate.**

Figure 3.2 gives a high level overview of the system. The Reader on the drone carries an LED light as a transmitter and a photodiode as the receiver. The Reader is powered by the battery of the drone. The Tag uses a retroreflector and a LC shutter as the transmitter, and a photodiode as the receiver. The Tag also has a small $46\,\text{cm}^2$ solar panel and is powered exclusively from the harvested power. The LED light carried by the Reader generates and transmits the downlink signals to the Tag. The retroreflector and LC shutter on the Tag then modulate and reflect this signal back to the Reader.

In our system, the communication is done through transactions, which are defined as follows:

1. The Reader initiates a transaction by sending a message to the Tag.
2. Once the Tag has harvested sufficient energy from ambient light, it will recognize and decode the message.
3. Once the message has been successfully decoded, the Tag will backscatter its data to the Reader by modulating its LC shutter.

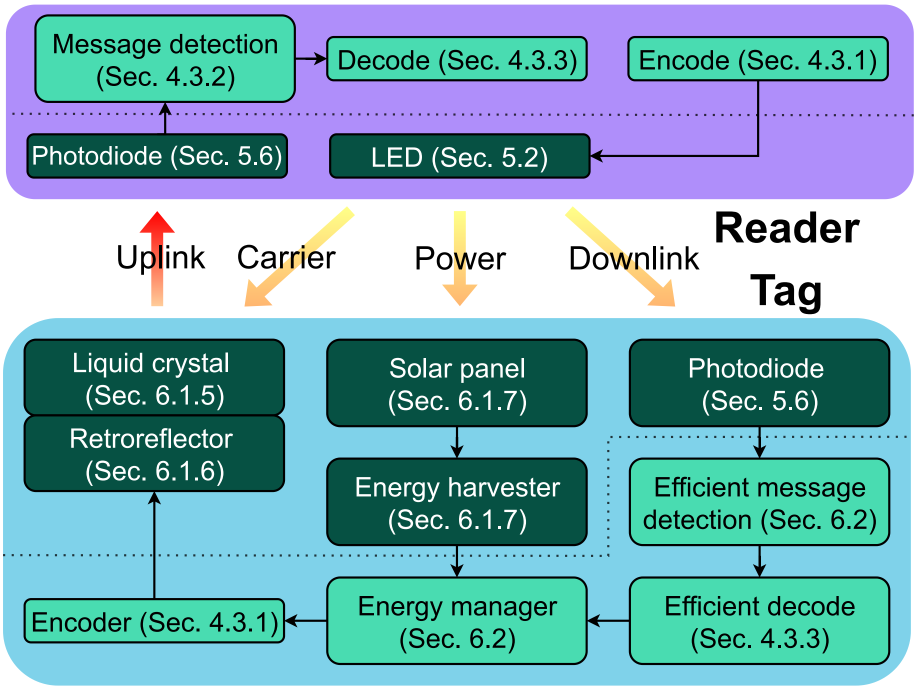

Figure 3.2: **High level overview of the different system components of the Reader and Tag. Dark green indicates that the component is part of the hardware, while light green means that it is software.**

4. The Reader decodes this message to finalize the transaction.

3.1 Design challenges

In the following sections, we discuss the constraints and challenges to be addressed in our system. These are divided into three categories: Reader, Tag and both.

3.1.1 Reader

Design challenge 1: The drone must be able to fly and maneuver while the Reader is mounted and operating. (Chapter 5)

Since the Reader is mounted on a drone, the main challenge is to maintain the flying capability of the drone. For example, the weight and size of the Reader must not be excessive compared to the drone, such that the drone is be able to fly with the Reader mounted on it. Moreover, the power consumption of the Reader (most notably the LED) must not be disproportionate compared to the power consumption and battery capacity of the drone.

3.1.2 Tag

Design challenge 2: To overcome the issues caused by the slow and asymmetrical switching times of the LC. (Chapter 4)

Due to the mechanical nature of LCs, the time it takes to switch from transparent to opaque and back again is in the order of milliseconds. Moreover, the rising and falling edges are highly asymmetrical. We need to make sure that the modulation scheme can tolerate these properties.

Design challenge 3: The Tag is able to operate without batteries, powered by the ambient light and communication carrier. (Chapter 6)

To reduce maintenance and increase ease of deployment, the Tag must be able to operate without batteries and harvest energy from the downlink. This means care must be taken into the power consumption of the hardware, software and communication infrastructure of the Tag.

3.1.3 Reader and Tag

Design challenge 4: The system is able to establish and maintain a communication link at a suitable range. (Chapter 4 and 6)

Due to the polarizing layers of the LC, the intensity of light passing through the LC is halved [20]. Moreover, since we are backscattering the uplink, the Tag transmission power is a fraction of the downlink power, since not all the light will hit the retroreflector and degrades quadratically [6, 23]. The aforementioned factors result in a weak uplink signal, which will be difficult to decode.

Design challenge 5: The link must be reliable when the drone is hovering. (Chapter 4)

The goal of the thesis is to create a communication link between a static object and dynamic drone. The drone will never hover perfectly still. Any slight variations will amplify interference from ambient sources. It is imperative that the link is reliable when the drone is hovering in front of the Tag, since this is its intended use case.

Chapter 4

Modulation

In order to realize communication between the Reader and Tag with light, the light needs to be modulated. Typically, simple and cost-effective VLC modulation is realized through intensity modulation with direct detection (IM/DD). Which means that the transmitted signal is modulated using the instantaneous power (amplitude or intensity) of the light [27, 28, 29]. In this section, we compare common modulation schemes and their suitability for our application. After selecting a suitable technique, we detail how the data is encoded, recognized on the receiver and decoded.

4.1 Common schemes

There are three fundamental modulation schemes for VLC: amplitude modulation, Pulse modulation and frequency modulation. Figure 4.1 gives a brief overview of various different modulation techniques, which we will describe in detail in the following sections.

4.1.1 Amplitude modulation

Amplitude-based modulation or Amplitude Shift Keying (ASK), modulates symbols by changing the amplitude of the signal. For VLC this means changing the brightness of the light. The simplest implementation of ASK is On-Off-Keying (OOK), which maps the transmitted bits to the presence or absence of a source signal. Generally, the bit ‘0‘ maps to the off state (low light intensity), while the ‘1‘ maps to the on state (high light intensity). The advantage of this scheme is its simplicity and high spectral efficiency.

However, OOK suffers from two disadvantages: it is sensitive to noise and operates asynchronous, i.e. there is no clock signal to synchronize the receiver and transmitter. If the transmitter transmits a long sequence of the same symbol, the receiver and transmitter might lose synchronization, which results in bit errors. To overcome this issue, the bitstream is typically coded using a self-clocking line code e.g. Manchester [5] or Miller [6] code to maintain synchronization when transmitting the same symbol consecutively.

Miller coding is implemented as follows. A ‘0‘ has keep the signal as-is, unless it precedes another ‘0‘, it then inverts the signal at beginning of its symbol

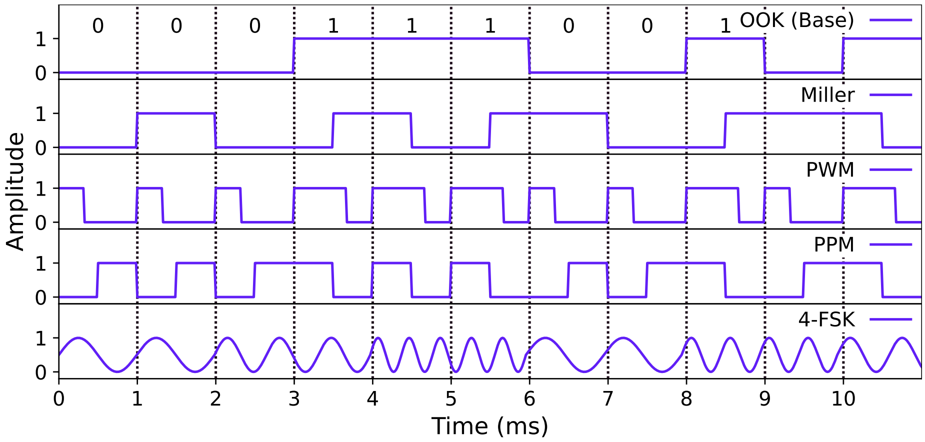

Figure 4.1: **Examples of various modulation schemes and variations. Amplitude-based: OOK, Miller. Pulse-based: PWM, PPM. Frequency-based: 4-FSK.**

period. Moreover, a '1' always inverts the signal in the middle of the symbol period. With miller coding the pulse length are either 1x, 1.5x or 2x the symbol period, these gaps assist in synchronizing the clock, but any variation will cause bit errors.

4.1.2 Pulse modulation

As a side step from amplitude modulation lies pulse modulation. Whereas amplitude modulation encodes data in the amplitude of the signal, pulse modulation encodes the data in where the pulse (change in amplitude) occurs and how long it is. There are two fundamental pulse modulation methods: Pulse Width Modulation (PWM) and Pulse Position Modulation (PPM). PWM transmits a pulse at a fixed frequency and encodes the symbol in the width of the pulse (duty cycle %). On the other hand, PPM transmits a pulse at a fixed frequency and encodes the symbol in the position of the pulse (i.e. start or end of the period). There are many variations and combinations of these methods that allow for e.g. light dimming control or higher spectral efficiency [1, 2].

Figure 4.1 visualizes PWM and PPM modulation. The PWM maps '0' and '1' to 33 % and 66 % duty cycle, respectively. Moreover, visually comparing Amplitude- and Pulse-based modulation techniques shows that they are very comparable. The main differentiating factor between these two methods is that pulse-based modulation allows for better and more flexible dimming [1, 2]. However, dimming support is not a focus of this work.

4.1.3 Frequency modulation

Frequency modulation or Frequency Shift Keying (FSK), modulates bits by transmitting a unique frequency per symbol. Generally, with M-ary FSK (M-FSK) one can employ M frequencies uniformly spaced with f_Δ, such that $f_i = f_0 + i \cdot f_\Delta$ for $i = 0, 1, ..., M - 1$. In this scenario, each frequency f_i maps to a distinct symbol s_i, representing $\log_2 M$ bits. For example, 4-FSK

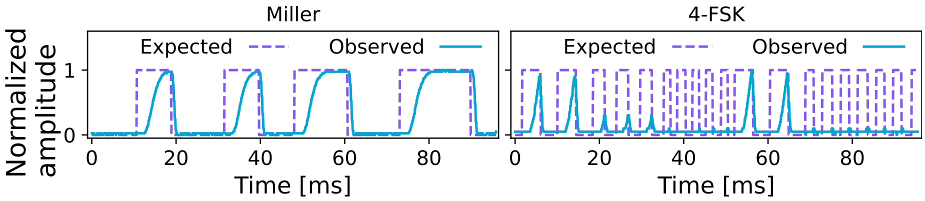

Figure 4.2: **Default Miller and 4-FSK encoded messages with the logic signal and driven LC signal.**

utilizes four frequencies: f_0, f_1, f_2, f_3. These frequencies map to four symbols: s_0, s_1, s_2, s_3 and resemble the following bit pairs: 00, 01, 10 and 11, respectively. By increasing the number of bits transmitted per symbol (symbol size) while keeping the symbol time fixed, one can increase the bitrate of the link, while keeping the symbol rate the same.

There are no restrictions on what frequencies can be used for FSK. However, there is a distinction between coherent and non-coherent FSK. Namely, coherent FSK guarantees a continuous phase without any abrupt transitions between symbols, while non-coherent makes no such promises. This is achieved by selecting f_0 and f_Δ as a multiple of the symbol rate f_s, i.e., $f_0 = n \cdot f_s$ and $f_\Delta = m \cdot f_s$. The advantage of coherent FSK is a higher spectral efficiency since abrupt transitions generate frequencies outside of the used frequency band. Moreover, as we will discuss later, the LC has a slow response time, which does not handle abrupt transitions well.

4.2 Selecting a suitable modulation technique considering the system constraints

Now, we will analyze what modulation technique is most suitable for our application. We will focus on just two: ASK with Miller coding [6] and M-FSK [20, 22]. Since pulse modulation suffers from the same drawbacks as amplitude modulation, and Miller coding is an improved version of Manchester coding.

4.2.1 LC response time

An inherent issue with LC shutters is their slow and asymmetrical response time (sum of the rise and fall time). We can analyse the response time by mounting a photodiode behind the LC and recording the received light while switching the LC. Figure 2.4 zooms in on one pulse and shows the response time of a LC shutter driven at 5 V and 60 Hz. It is clear that with a response time of approximately 4 ms that the uplink is severely bandwidth limited. Moreover, the rising edge is three times as long as the falling edge, making the pulses highly asymmetrical.

Figure 4.2 shows Miller and 4-FSK modulated versions of the signal from Figure 4.1 by a LC. These figures show that the LC has trouble modulating Miller encoded signals and completely fails to properly modulate an 4-FSK signal. This is caused by the slow rise times, requiring more than a full symbol period to plateau. To resolve this issue, one should lower the symbol rate, such

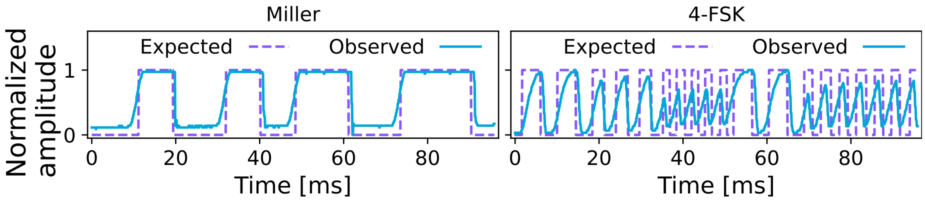

Figure 4.3: **Asymmetrical Miller and 4-FSK encoded messages with the logic signal and driven LC signal.**

that the rising edge has sufficient time to plateau. However, this is undesirable, since this would drastically reduce the symbol rate of the system.

Instead, one could increase the allocated time for the rising edge and drive the LC with an asymmetrical signal. This is trivially applied to FSK by simply changing the duty cycle % of the logic signal, which is originally 50 %. However, for a Miller coded signal it is vital that the pulse periods (1x, 1.5x or 2x) remain correct, this would not be the case if the duty cycle is applied to all '0' and '1' signals (i.e. make all '0' shorter and all '1' longer). To solve this issue, we can instead only focus on the edges. This way, the rising edge gets more time, which is compensated by reducing the time spend on the falling edge. Any symbols between the edges are left untouched.

Figure 4.3 shows the same Miller and 4-FSK encoded messages, but with a 16.67 % duty cycle instead. It is clear that driving the LC with a asymmetrical signal significantly improves the FSK signal and allows the miller pulses to plateau.

However, the signal is far from ideal. For example, the amplitude of the 4-FSK symbols decreases as the frequency increases. This is caused by the fact that the rise time is still not fast enough. As the symbol frequencies increase, the amplitude would decrease, until eventually, the oscillations approach a flat line (approx. 1000 Hz). Moreover, while this does lower the SNR, it does not drastically reduce performance, as it would with ASK, since the signal is decoded based on the frequencies, not the shape or amplitude of the signal. Moreover, the while the ASK signal clearly improved, the rising edges are still slow and can cause distortion when decoding.

4.2.2 Transmission during flight and ambient interference

Ambient light constantly illuminates the photodiode on the Reader, regardless of whether a Tag is present, modulating or not. This creates a DC-offset on the received signal. If this offset is constant, then it is trivial to remove or ignore it. For example, one could subtract the mean value of the signal from itself.

The amplitude of this offset increases as the amount of light illuminating the PD increases. For example, in a dimly lit room (75 lx) this offset might be 50 mV, while outside (10 klx) this offset is more than the PD can represent, saturating it. To prevent the PD from saturating, we can decrease the gain.

Moreover, ambient light rarely remains exactly constant, especially since the Reader is mounted on a drone. Even when hovering still, the drone will move slightly. Figure 4.4 gives a profile of the ambient light as perceived by a photodiode mounted the drone is hovering in-place approximately 1 m above the ground

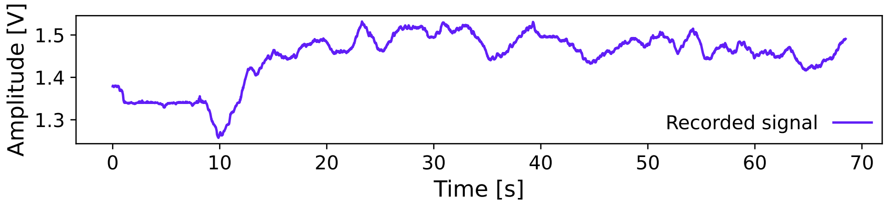

Figure 4.4: **Recording of the channel at a Reader mounted on a drone that is hovering in the air without a Tag nearby.**

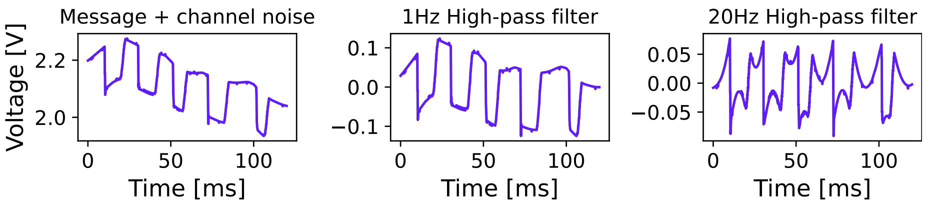

Figure 4.5: **The effect of high pass filters on a Miller encoded message with ambient channel noise.**

with a typical illumination 300 lx. This signal shows that the ambient interference consists mostly of low-frequency (<20 Hz) noise. The peak-to-peak change in this signal is 275 mV, which is 12 % of the dynamic range of the photodiode.

This low-frequency noise has little effect on the M-FSK signals, since decoding of these signals is focused on the specific frequencies $f_{0,\ldots,M-1}$. Therefore, as long as the photodiode does not saturate due to ambient light and the frequency spectrum of the ambient light is outside of the communication frequencies, then M-FSK will not degrade significantly due to interference from ambient light.

However, the same can not be said for miller encoded messages. Since these are decoded based on the distance between the peaks (1x, 1.5x or 2x). Figure 4.5 shows a Miller encoded message imposed on a randomly selected sample of ambient noise from Figure 4.4. In attempt to remove the ambient noise from the signal, the message is filtered with a high pass filter. While a second order high pass filter with a cutoff frequency at 1 Hz is sufficient to remove the DC-offset, there is still low-frequency noise present. If we increase the cutoff frequency to 20 Hz we also successfully remove this, but the signal gets severely distorted in the process. While the symbol pattern is still visible and it would be possible to successfully decode the signal, these results are undesirable and make decoding complicated. The distortion is caused by the fact that the Miller encoded signal is flat between edges. These flat sections will be distorted by the filter, since they have the same frequency characteristics as the noise.

This is shown in Figure 4.6, which plots the Power Density Spectrum (PSD) of both a Miller and M-FSK encoded signal. It is clear that the Miller encoded signal has no distinct peaks and the power is spread out over a wide bandwidth, especially compared to the 4-FSK signal, which focuses its power in the specific modulation frequencies ($f_i \ldots f_{M-1}$). In fact, for the M-FSK signal, 45.6 % of all power is concentrated within the four modulation frequencies ($f_i \pm 35Hz$ for $i = 0, 1, 2, 3$). On the other hand, 72.5 % of the miller coded signal power is

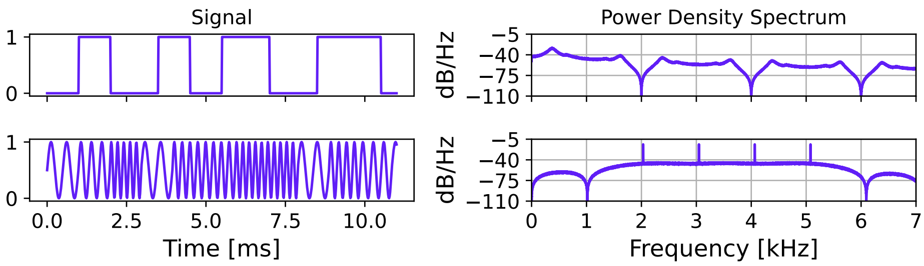

Figure 4.6: **The PSD of a Miller and 4-FSK encoded signal, both at 1000bps. Generated from a 100000 bit long random array.**

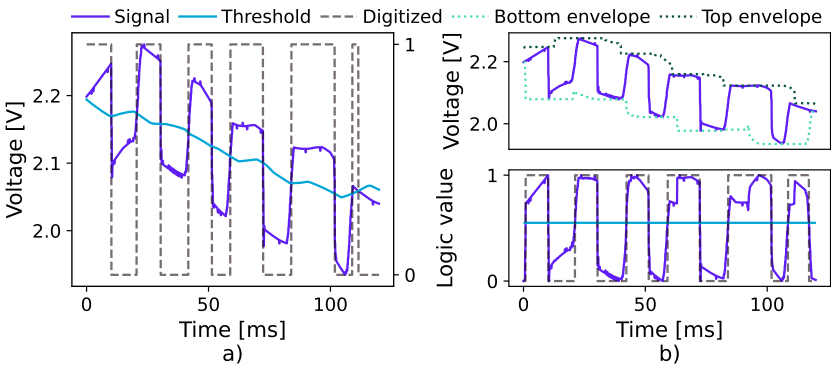

Figure 4.7: **Two alternative ways to threshold and digitize a Miller coded signal. a) Tracking the moving average and b) tracking the top and bottom envelope and normalizing the result before taking the mean threshold.**

concentrated before 500 Hz — even though the critical frequencies are 1000, 750 and 500Hz.

It is possible to decode the miller coded ASK signal, even in a dynamic channel. For example, one could locate the pulse edges by taking the differential [30]. However, this technique works best when the modulated signal has quick rise and fall times, since this would create a higher differential. This is not the case when we modulate with a LC, which has slow response times. Moreover, the response times are asymmetrical, which makes the differential lob-sided.

Alternatively, we could threshold the signal with a moving average, or normalizing it by tracking the top and bottom envelope (Figure 4.7). However, both these solutions work around the fundamental issue: amplitude-based modulation schemes inherently suffer more from noise compared to frequency-based schemes [20].

Other works try to work around this [5, 24] or even embrace it into the decoding scheme [6]. However, we realize that M-FSK can overcome the aforementioned issues, which is why we have selected M-FSK as the selected modulation scheme for this system.

Now, we can say that design challenge 2 and 5 have been resolved, since the selected modulation scheme can overcome the slow and asymmetrical LC response time and noise induced by the ambient light and mobility of the drone.

4.3 Implementation

Now that we have determined that FSK is the most suitable modulation technique for our application, we will detail how we encode, detect and decode M-FSK modulated messages. These techniques are split in parts: one implementation for the Reader and one for the Tag. As we will detail in the following sections, the implementation on the Reader does not work well on the Tag. However, due to time constraints the updated scheme was never backported to the Reader.

4.3.1 Encoding

Now that we have established that M-FSK is a better suited modulation technique compared to Miller coding (ASK) for our system, we will detail how we encode data into a M-FSK signal.

The frames used by the Reader and Tag are different. The Reader has a preamble that contains a 50 ms long alternating bit pattern of 1s and 0s, which helps synchronize the decoding scheme on the Tag. Right after the preamble, the Reader transmits the 56 bit ASCII code `Reader!`. Moreover, the Tag has a 16 bit preamble consisting of the SYN and STX ASCII symbols (`0001011000000010`), after which a 32 bit message `Tag!` is transmitted.

Given a bitstream $b_0, ..., b_{N-1}$, of length N, we want to encode this into symbols $s_0, ..., s_{\hat{N}-1}$, such that each symbol contains $S = \log_2 M$ bits. Where M is the number of distinct frequencies in M-FSK. This means that the number of symbols to be transmitted is $\hat{N} = \frac{N}{S}$.

To ensure proper encoding, the number of bits must be a clean multiple of S, i.e. $N \bmod S = 0$. If this is not the case, zeros are appended such that the last symbol is encoded correctly. These zeros are ignored on the receiver side.

Encoding the bits to symbols is done as follows:

Step 1 Take the first S bits from the bitstream:
$\hat{b}_0, ..., \hat{b}_{S-1}$.

Step 2 Left-shift and OR these bits together to get the symbol index:
$i = \hat{b}_0 \ll (S-1) \mid \hat{b}_1 \ll (S-2) \mid ... \mid \hat{b}_{S-1}$

Step 3 Append the found symbol index to the symbol stream, remove the first S bits from the bitstream and return to **Step 1** until all bits have been encoded.

Section 5.5 and 6.1.4 detail how the Reader and Tag convert these symbols to a signal in order to drive the LED and LC, since this is hardware-dependant and optimized for these devices.

4.3.2 Message detection

The Reader and Tag will not be continuously interacting. In fact, the channel will be idle most of the time. This means that the devices must detect when a message is transmitted, such that they can record and decode it. This detection is done with the help of a preamble, which is a predefined bit pattern that is transmitted right before the message.

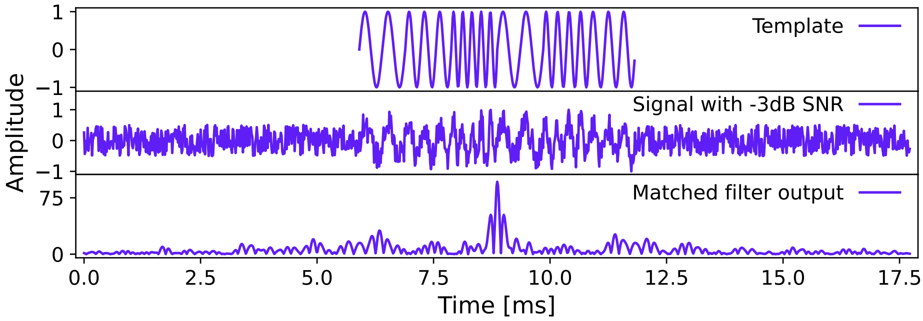

Figure 4.8: **Example detailing how a matched filter works. The template is convolved with a signal with a -3dB SNR. The matched filter clearly shows a peak in the middle of the template.**

Reader

Preamble detection on the Reader is done through a matched filter. A matched filter works on the basis of convolving a known time-reversed template signal with a random signal to see if the template is present in it. The correlation output will generate a significant peak in the middle of where the template is detected in the signal. Since the matched filter relies on a template signal, it requires that the transmission frequencies are predetermined.

In the case of the Reader, the template is the preamble transmitted by the Tag. Figure 4.8 shows how a matched filter works with an example 4-FSK message. The matched filter generates a significant peak in the middle of the template. By normalizing the input and continuously tracking the moving average of the peak we can determine when the preamble has been detected.

The advantage of this technique is that matched filters are known to be optimal for detecting a predetermined signal with additive Gaussian noise [31]. However, the disadvantage of this technique are that convolution scales $\mathcal{O}(n^2)$ and requires three buffers: One for the template, one for the random signal and one for the output, which increases the memory footprint and computation time. For these reasons, preamble detection is done differently on the Tag.

Tag

Matched-filter based detection is has good performance, but requires intense buffering. The selected MCU for the Tag only has 32KiB of RAM, a factor 10 less than what the Reader has available. For this reason, preamble detection on the Tag uses a less memory intensive method. Specifically, we continuously sample the photodiode and periodically apply a Fast Fourier Transform (FFT) to analyse the frequency domain and determine whether the M-FSK frequencies are present.

Figure 4.9 shows how the Fourier transform works. The top plot shows a 4-FSK modulated signal, the bottom plot shows the energy of the associated frequencies (f_0, f_1, f_2, f_3) over time. There are clear peaks in the middle of each symbol according to the associated frequency.

However, measuring the energy of different frequencies is only one part. Namely, we need to detect when a message is present, so when there is an in-

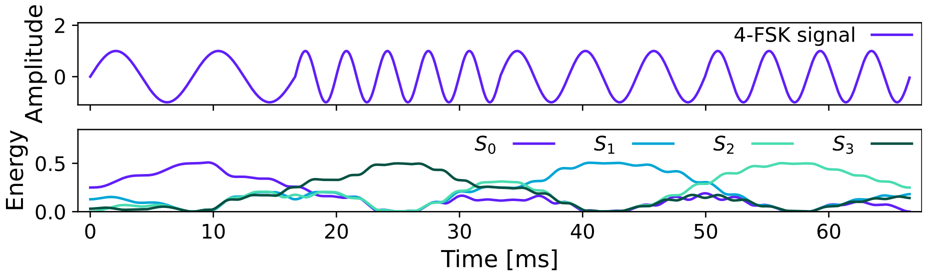

Figure 4.9: **Example showing how a real-valued Fourier Transform works. The top signal shows every symbol of 4-FSK once. The bottom plot shows the energy per symbol over time.**

crease in energy of the modulation frequencies. To generalize the thresholding, we create one metric, which is the sum of the energy representing the M-FSK frequencies $f_0, \ldots, f_{M-1}$. Then, we track the moving average of this metric, and detect that a message is present when the current energy is significantly greater than the moving average.

The advantage of this technique is that a FFT scales $\mathcal{O}(n \log n)$ and that we only require two buffers consisting of 32 samples, instead of three larger ones. These efficiency increases enable preamble scanning on the low-power microcontroller. However, as the results show (Chapter 7), this decoding scheme struggles to achieve consistent performance.

4.3.3 Decoding

After a message has been detected, it needs to be received and decoded. This is done differently for the Reader and the Tag, for the same reasons as mentioned previously. Both use the same technique for decoding as they do for message detection: The Reader decodes using a matched-filter based technique, while the Tag decodes based on the Fourier Transform.

Reader

After a preamble is detected, the Reader will immediately save the last $\hat{N}$ samples, where $\hat{N}$ is the length of the preamble in samples. Then, it will buffer samples and wait for the message to be fully received. In total, the buffered data contains N samples. Figure 4.10 shows the demodulator for the Reader, which is inspired by [32]. Moreover, Figure 4.11 visualizes the intermediate decoding steps.

First, the buffered signal is normalized between [-1, 1]. Then, the signal is processed through M matched filters, each tuned to $f_0, f_1, \ldots, f_{M-1}$. However, the output of the matched filter still contains the symbol frequency f_i, which makes determining symbol transitions challenging. For this reason, each output of the matched filter is passed through an envelope detector. This detector is implemented as a first order Butterworth low-pass filter with the cutoff frequency set to f_s.

Finally, we will align with the correct phase, such that we determine the symbols at the peak of the output signal, instead of, for example, between two

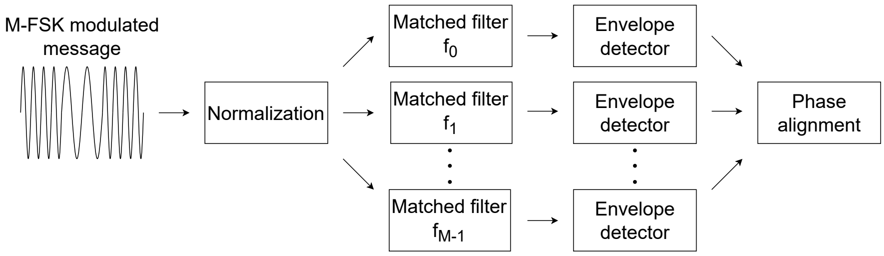

Figure 4.10: **M-FSK decoder for the Reader.**

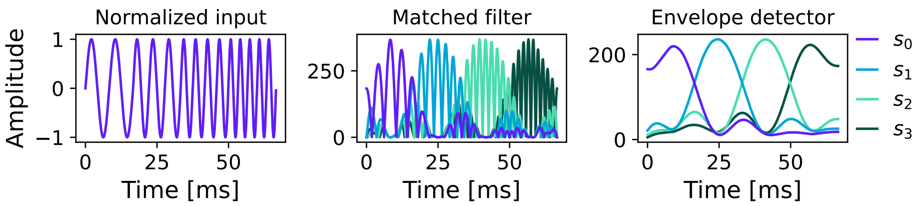

Figure 4.11: **Visualization of the Filter bank for four symbols of a 4-FSK encoded signal.**

symbols. The phase alignment is implemented as follows: Say that the output of the envelope detector is $y_0, y_1, \ldots, y_{M-1}$, then the standard deviation at sample i is:

$$STD(i) = STD(y_0[i],\ y_1[i],\ \ldots,\ y_{M-1}[i]) \tag{4.1}$$

Then, to simulate sampling the symbols at phase offset j:

$$s(j) = \sum_{i=o}^{\frac{N}{N_s}} STD(iN_s + j) \tag{4.2}$$

Where N_s is the number of samples per symbol. We can then find the phase offset k with the highest simulated sampled sum $s(k)$, which will align best with the peaks of the envelope detector:

$$k = \arg\max_j(s(j)) \tag{4.3}$$

The signal is converted to a list of symbols $\{s_0, \ldots, s_{\frac{N}{N_s}}\}$ by sampling the output buffers $y_0, y_1, ..., y_{M-1}$ with offset k and noting the symbol with the highest amplitude:

$$\{s_0, \ldots, s_{\frac{N}{N_s}}\} = \arg\max(y_0[iN_s + k], \ldots,\ y_{M-1}[iN_s + k]) \quad \text{for } i = 0, \ldots, \frac{N}{N_s} \tag{4.4}$$

Finally, the symbols are converted to bits through mapping, for example with 4-FSK:

$$b_i = \begin{cases} 0,0 & \text{if } s_i = 0 \\ 0,1 & \text{if } s_i = 1 \\ 1,0 & \text{if } s_i = 2 \\ 1,1 & \text{if } s_i = 3 \end{cases}$$

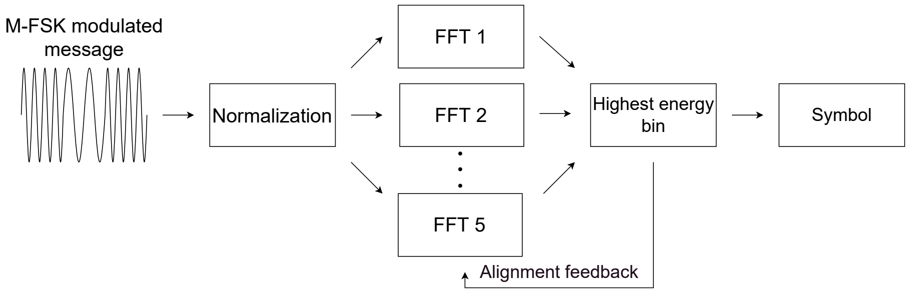

Figure 4.12: **M-FSK decoder for the Tag.**

This decoding scheme works quite well. However, it is resource intensive, requiring $M + 1$ processing buffers and M template buffers. For example, 4-FSK at a symbol rate of 60 baud, a sample rate of 4000 Hz and a message size of 48 bits, would require nearly 64 kilobytes of RAM, just for the buffers alone. It is clear that this decoding scheme cannot be directly ported to the low-powered microcontroller on the Tag, since it only has 32 kilobytes of RAM.

Tag

The original decoding scheme will not work well on the Tag, due to the significant amount of RAM required. Even though several improvements can be made to try and reduce the buffer size, such as: reducing the sample rate (number of samples per symbol); switching from 4 byte floating point representation to a 2 byte fixed-point representation; reducing M; calculating the templates at compile-time, such that they can be stored in flash; reducing the message length; and reducing the number of samples buffered. These improvements come at a cost of reduced accuracy and cannot guarantee that the RAM will not fill up.

Therefore, to decode messages in the resource-constrained Tags, we need a decoding scheme that requires minimal buffer space while being able to decode messages in real time. Our decoding scheme is inspired by [20, 22], which are also VLBC works that employ FSK.

The decoding scheme is visualized in Figure 4.12 implemented as follows:

1. A sliding windows with a (step) size of one symbol applies a FFT to the sampled data.

2. The M bins representing $f_0, \ldots, f_{M-1}$ are compared and the highest energy bin is selected as the associated symbol.

3. To align with the correct phase of the transmitted signal, we also apply four additional FFTs with a offset of ± 2 samples relative to the current window position.

4. If the resulting energy in one of these FFTs is higher than the previously found energy, we adjust the window to this position.

5. Finally, the found symbol is appended to an array. Once sufficient symbols are decoded, then the Tag converts the symbols to bits using the same mapping as the Reader.

This decoding scheme is much less resource intensive compared to the one employed on the Reader (a $\sim$ 128x improvement). This scheme allows us to decode symbols in real time, with a delay of just one symbol. Moreover, this improvement is the first step to solving design challenge 3 and operate the Tag battery-less.

Chapter 5

Reader

In the following sections, we detail the various hardware components of the Reader. Figure 5.1 shows the prototype and where the components are located.

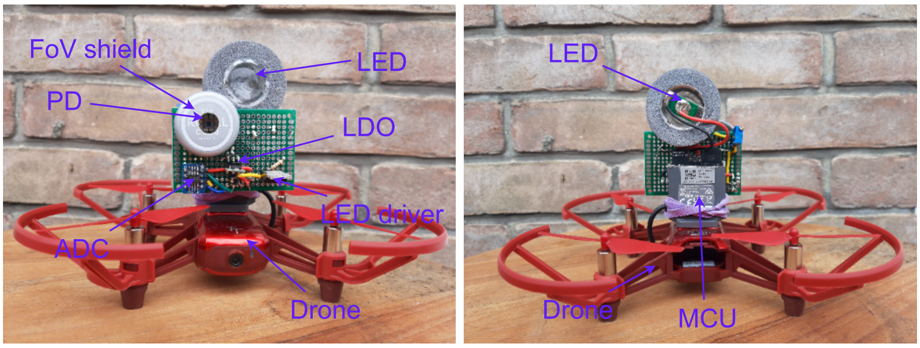

Figure 5.1: **The Reader prototype mounted on a drone.**

5.1 Drone

The Reader is mounted on a drone, which allows it to move around. For the drone, we selected the Robomaster TT [33], specifically because it allows for an extension board to be installed to conveniently prototype with the drone. This extension board exposes ten IO pins and can output up to 800 mA between its 3.3 V and 5 V rails. While the drone does not have a formally defined maximum recommended payload, it weighs 86 g in its default configuration. Moreover, since the drone formally supports an extension board, we expect that it is able to carry more weight that just itself.

To get some insight into the performance and flight time of the drone, we note the flight time while adding components to the drone:

- A flight-ready bare drone without any modules weights 86 g and can fly for 10 minutes.
- When the extension module (without reader attached) is mounted to it, the weight increases to 96 g and the drone can fly for 6 minutes and 45 seconds.
- When the Reader is mounted to the extension module, the drone weights 122 g and can fly for 4 minutes with the Reader turned off and 3 minutes when the Reader is turned on.

Importantly, the drone only behaves correctly when nothing is mounted to it. Even with just the extension module (not the Reader), the drone occasionally behaves erratically and will yaw and correct itself or even spin in circles eventually. This behaviour worsens as more weight is added to the drone, when the LED is turned on and over time (as the battery empties). With the Reader mounted and active, the drone can fly for 1 minute and 30 seconds before the first erratic behaviour occurs. Then, it can fly until 3 minutes before the drone is forced to land since it cannot consistently measure anymore. We gather that this is caused by the increased power consumption caused the added weight and LED, which in turn causes brown outs in the system.

With these optimizations, we are able to operate and fly the drone while the Reader it active, completing design challenge 1.

5.2 Light

The light powering the downlink is extracted from a generic off-the-shelf flashlight [34] and rated for 1 W. However, the drone behaves erratically and the flight time decreases significantly with the LED configured at its default power level, even though the extension board formally supports this power draw. To compensate, the LED is current limited to 0.5 W. We speculate that the increased power draw from the weight and LED causes voltage drops, which might affect sensor readings or cause brown-outs. The LED is significantly lower power compared to other VLBC works, which operate from 3 W to 30 W range [20].

Moreover, the integrated lens narrows the FoV to approximately 35.5°. The original housing incorporating the LED and lens was made of aluminium and was too heavy (21 g) for the drone to fly with. By placing the LED on the back of the PCB and replacing the aluminium housing with a foam tube balanced the drone sufficiently such that the drone can fly with the Reader mounted on it. With this modification, the total weight of the Reader is just 26 g.

5.3 Microcontroller

The MCU operating the Reader is contained in the extension board (the drone has a separate controller) and is able to send commands to the drone via a UART connection. Moreover, this extension board exposes some pins in order to prototype custom modules for the drone. The MCU located in the extension board is the ESP32-D2WD, which is a dual-core microcontroller with each core operating at 160 MHz with 520 KiB of RAM and is developed by Espressif. The ESP32 offers a wide array of peripherals, of which we will be using the Inter-IC Sound (I2S) module to sample the external ADC and the Remote Control Transceiver (RMT) to drive the FSK signal.

5.4 Analog front-end

The extension board only exposes IO pins linked to ADC2, which is shared with the WiFi module. This means that the WiFi module will periodically reserve and reset ADC2 for internal use. To overcome this issue, we are sampling the photodiode with an external ADC over I2S. The selected ADC is the MCP3201,

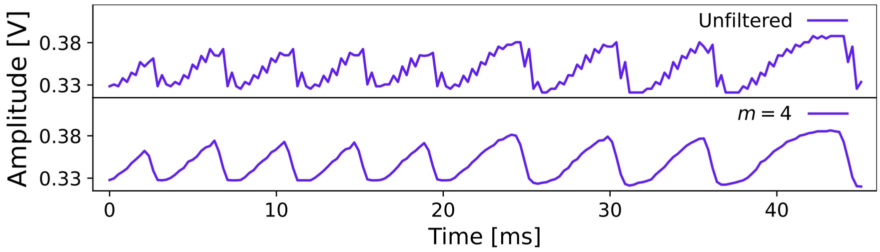

Figure 5.2: **Visualization of a decimation filter.**

which supports sample rates up to 50 ksps at 2.7 V and 100 ksps at 5 V, which is plenty for our application. In the next steps we iterate how we improved the noise floor on the analog signal. We define the noise floor as the standard deviation of the signal.

Step 0. First, we need to quantify the default noise floor of the analog input. We continuously sample the ADC and every 128 samples we calculate the STD. The default STD is 1.19.

Step 1. The first step to decrease the noise floor is to employ a decimation filter. A decimation filter increases the sample rate by an integer factor m, such that the new sample rate is $\hat{f}_s = mf_s$. Then, it averages m samples into a single sample — downsampling and decimating the signal. The resulting signal is sampled at the target sample rate of f_s, with a low-pass filter applied and an increased resolution of $\log_4(m)$ bits [35, 36]. This is visualized in Figure 5.2. Moreover, since the external ADC supports at most 50 ksps at 3 V, we have selected the following configuration: $\hat{f}_s = 50$ ksps, $m = 16$ and $f_s = \frac{\hat{f}_s}{m} =$ 3125 sps. Applying these settings improved the noise floor to 0.43 — a 177 % improvement.

Step 2. The MCP3201 has two voltage rails: V_{REF} and V_{DD}, which are used as analog and digital reference, respectively. To minimize parasitic noise in the analog voltage domain from the digital communication, these two voltage rails are separated. The digital input V_{DD} is connected directly to the 3.3 V rail of the ESP32, while the analog input V_{REF} is referenced to 3 V, which is regulated down from the 5 V rail of the ESP32 using an Low Dropout Regulator (LDO). Adopting this change decreased the noise floor from 0.43 to 0.32 — a 34 % improvement.

5.5 LED driver

The LED driver converts the symbol stream generated in Section 4.3.1 into a M-FSK modulated signal and drives the LED. For this we use the RMT peripheral, which is able to output arbitrary pulse trains. These pulses are defined by the following scheme: Every RMT period consists of two sub-periods, each with a 15 bit duration and a 1 bit level section. The duration of a sub-period is defined in the number of timer ticks, while the level field defines the logic output for this sub-period (0 or 1). These RMT periods can be chained and queued in memory to create a continuous signal.

Defining a RMT period with equal duration fields but different level values allows us to generate an arbitrary frequency with a 50 % duty cycle. We can

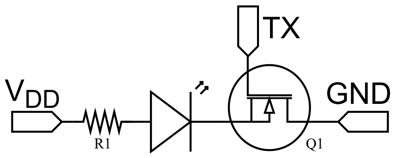

Figure 5.3: **Schematic overview of the low-side switch of the LED.**

generate the M-FSK signal by setting the duration fields each to a half period of the desired output frequency. Then, we can repeat this RMT period for the duration of that symbol and repeat this process until all symbols are transmitted.

However, the ESP32 cannot source or sink sufficient current to directly drive a 0.5 W LED (170 mA). So, we require a physical driver in order to toggle the LED with the previously generated signal. This is done with a low-side MOSFET switch, which connects or disconnects the ground based on whether the signal on the gate is high or low, respectively. If the ground path of the LED is disconnected, then no current can flow, turning the LED off. For the driver we have selected the N-channel IRLZ34N MOSFET, which has the right voltage thresholds such that we can toggle it with logic-level voltages (3.3 V). With this driver, the driving current is sourced from the voltage rail instead of the MCU. The LED driver is schematically shown in Figure 5.3.

5.6 Photodiode

The photodiode used for the prototype of both the Reader and Tag is the OP101, which is an active photodiode, since it contains a transimpedance amplifier that converts the output current to a voltage on the package. This PD operates from 2.7 V to 36 V and will be powered at 3.3 V to maintain analog compatibility with the MCUs. Moreover, the OPT101 has a wide FoV (nearly 180°), which means that it will pick up a considerable amount of ambient light that is not reflected from the Tag. This will cause three things: Add a DC-offset to the signal, add noise to the signal and cause the PD to saturate faster (requiring a lower gain).

To reduce these effects a shield or cover can be designed to block light from the sides, effectively reducing the FoV. Some works create a small enclosure surrounding the PD [24], while other works incorporate a lens [20, 22] to reduce the FoV. Incorporating a lens to reduce the FoV requires it to be placed at a specific focal length to operate optimally. This is not feasible, since the Reader must be lightweight and compact to be placed on the Drone. For this reason, a simple enclosure will be placed around the PD. This shield blocks most of the light coming from the sides of the PD, preventing PD saturation.

This modification allowed us to more than double the gain, effectively increasing SNR. Finally, the PD gain is set to 1.3 for the Reader and 0.5 for the Tag. These values are empirically chosen such that the PD does not saturate indoors. The Tag has a much lower gain, since the downlink is strong and the LED would otherwise saturate the PD.

Chapter 6

Tag

In the following sections, we will first detail the various hardware components of the Tag, then the optimizations required to operate battery-less are detailed. Figure 6.1 shows the prototype and where the components are located.

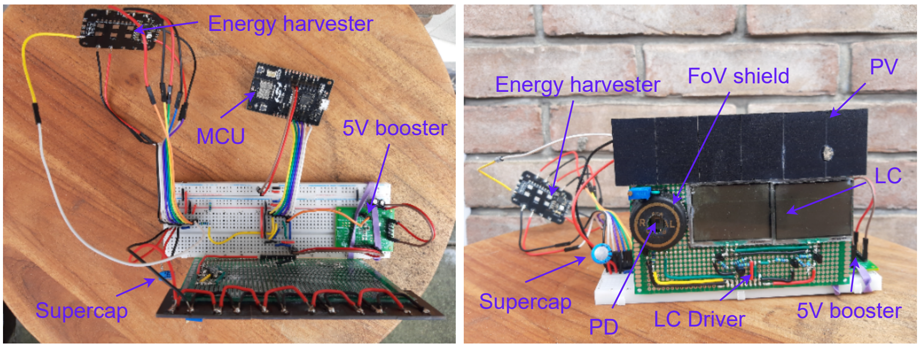

Figure 6.1: **Top and front view of the Tag prototype.**

6.1 Hardware

6.1.1 Microcontroller

The MCU operating the tag needs to be energy efficient, since the Tag will be operating on a tight energy budget. Because of this, the ESP32 selected for the Reader is not a good fit and a more efficient MCU is required. We have selected the BGM220P from Silicon Labs, which is a low-power oriented MCU with 32 KiB of RAM and an ARM Cortex-M33 operating at 76.8 MHz. Moreover, the included development board can accurately profile the power consumption of the system, which will be of use when optimizing the power consumption. We will be using the internal ADC and DMA controller for sampling the photodiode and the Low-Energy Timer (LETIMER) peripheral for generating the M-FSK signal that drives the LC. In section 6.2 we will show the power consumption per component.

6.1.2 ADC

The BGM220P has an internal incremental ADC that can achieve 1 Msps at a 12-bit resolution. However, it can use oversampling to achieve up to a 20-bit resolution, with a reduced sample rate. Moreover, it supports digital post-averaging to reduce noise. For our configuration, we set the oversampling factor to 16x, for an effective resolution of 14 bits.

The ADC will be used in deep sleep mode, where only the required peripherals and clocks are enabled and the samples are copied to a ping-pong buffer using DMA. Once a buffer is filled, the DMA controller automatically starts filling the other buffer and an interrupt triggers to process the samples. The ADC samples at 80 ksps and the DMA buffer contains 128 samples before triggering the interrupt.

6.1.3 DCDC booster

As mentioned previously, the LC operates at 5 V, while the MCU and photodiode require just 3.3 V. Moreover, the MCU is not 5 V tolerant. This means that the LC requires a dedicated voltage rail at 5 V. Originally, a generic hobbyist adjustable power supply was used. However, it operated at a low efficiency and high quiescent current (300 µA), so a more specialized low-power DC booster was selected: The TPS61099. This chip operates at approximately 75-90% efficiency and has a quiescent current of just a few µA. The exact impact of this improvement is detailed in section 6.2.4.

6.1.4 LC driver

The LC driver converts the symbol stream generated in Section 4.3.1 into an M-FSK modulated signal and drives the LC. We use the LETIMER in buffered PWM mode to convert the symbols into a physical signal. This is done as follows. The counter is initialized at *TOP*, counts to 0 and outputs a *HIGH* signal when *TOP* is lower than *COMPARE*, otherwise it will output a *LOW* signal. By adjusting the *TOP*, *COMPARE* and *COUNTER*, we can adjust the frequency and duty-cycle of the output signal. When writing the first bit, the repeat counter for both the current and next bit is set. The repeat counter and buffer ensure that the output contains the correct number of periods per symbol and is glitch-free. Once the repeat counter for the current symbol reaches zero, the buffered repeat count for the next symbol is automatically loaded. At the same time, an interrupt is triggered, which adjusts the *TOP*, *COMPARE* and counter values for the newly loaded symbol. Moreover, the buffered repeat counter for the next symbol is set. This processes is repeated until the final symbol is transmitted.

The LC operates at 5 V, while the MCU operates at 3.3 V. This means that we need to shift the M-FSK signal to the 5 V domain before driving the LC. This can be done with a simple level shifter (Figure 6.2a). However, as visualized, this design constantly sinks current when driving a low signal I_{low}. To manage this current loss, high resistor values for R_1 and R_2 should be used. However, since the LC is effectively a capacitor (approximately 9 nF [5]), this increases the charging times and reduces LC switching speeds. To overcome this issue, we base our design on the energy reuse circuit from [5]. Which is visualized in

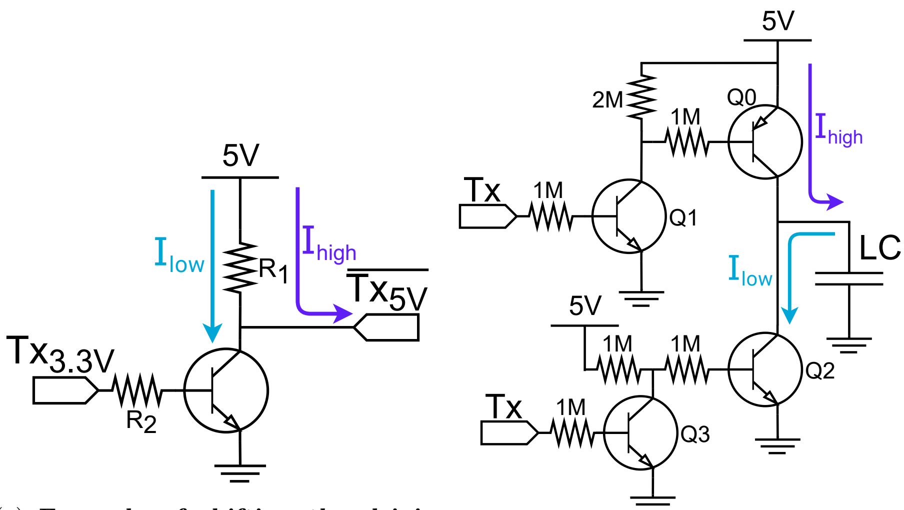

(a) **Example of shifting the driving signal from 3.3V to 5V.**

(b) **LC driver for the Tag.**

Figure 6.2: **Comparison of two different ways to drive the LC.**

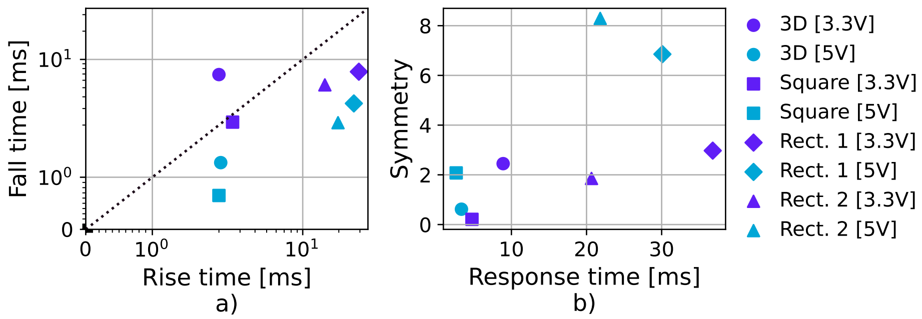

Figure 6.3: **Analysis of four different LC shutters at 3.3V and 5V. a) The rise and fall times. b) The response time compared to the symmetry of the edges.**

Figure 6.2b. The design ensures that the high current path (no resistors) either goes from 5 V to the LC (I_{high}), or from the LC to ground (I_{low}). While there is still some leakage current through the transistors, this is minimized since it is now possible to use high-valued resistors without significantly effecting the LC charging time.

6.1.5 Liquid crystal shutter

There is a wide selection of LCs available, and not all are made the same. For example, as LuxLink shows, the supported modulation frequency of an LC shutter can vary between 20 Hz and 147 Hz [22], depending model and driving voltage. To select a LC suited for our application, we will analyze the response time ($t_{rise} + t_{fall}$) and symmetry ($\frac{\max(t_{rise}, t_{fall})}{\min(t_{rise}, t_{fall})} - 1$) of the LCs. The driving voltage will be tested with 3.3 V and 5 V.

LC	Driving current‡ [µA]	Area [cm^2]
Rect. 1 [37]	54	35.2
Rect. 2 [38]	85	42.8
3D [39]	43	27.2$^{\dagger *}$
Square [40]	45	24.6†

Table 6.1: **Overview of the selected shutters. ‡At 200Hz and 5V † Due to small size, two shutters are used. *The area is approximated using an ellipse.**

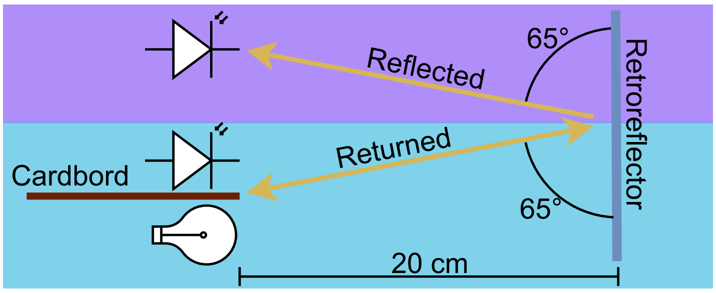

Figure 6.4: **The evaluation setup for various retroreflectors.**

Table 6.1 gives a brief overview of the different shutters tested. Moreover, Figure 6.3 gives insights regarding the performance of the different shutters. From these results it is clear that the two rectangular shutters [37, 38] do not perform well and significantly lag behind the other shutters. Interestingly, it appears that the driving voltage particularly affects the fall times. Moreover, both the 3D shutter [39] and square shutter [40] perform well: at 5 V, the 3D shutter is a bit slower than the square shutter. However, the rise and fall times of the 3D shutter are more symmetrical.

Since the decoding method of the uplink is FFT-based, and the shape of the signal (square, sine, sawtooth, etc.) does not affect the fundamental frequency found by the FFT, we don't need to focus on signal symmetry. Thus, because of the quicker response times, we have selected the square shutters for our prototype.

6.1.6 Retroreflector

There is no de-facto standard for retroreflectors in the literature. However, it is typical to use a retro-reflective fabric. For example, [5, 6, 23, 24] all use retro-reflective fabrics. Other works use larger, three dimensional, retroreflectors (Corner Cubes) [41, 42]. Moreover, mirrors can also be used [20, 21] for directional communication.

To determine what type of retroreflector is suitable for our use case, we have set up a small experiment: Two photodiodes are placed facing a retro-reflector with an incident angle of 65°. One photodiode is collocated with a LED and is dubbed 'returned', since it will receive retro-reflected light that is returned along the incidence angle. The other photodiode is dubbed 'reflected' and will receive

Reflector	3M 983-10 Tape	3M 8906 Fabric	Brand A Tape	Brand B Tape	Bike Reflector	Brand C Corner cube	Mirror	LC
Area [cm^2]	56.1	53.56	45.6	41.86	20.16	7.98	61.75	137.75

Table 6.2: **Summary of the tested reflectors.**

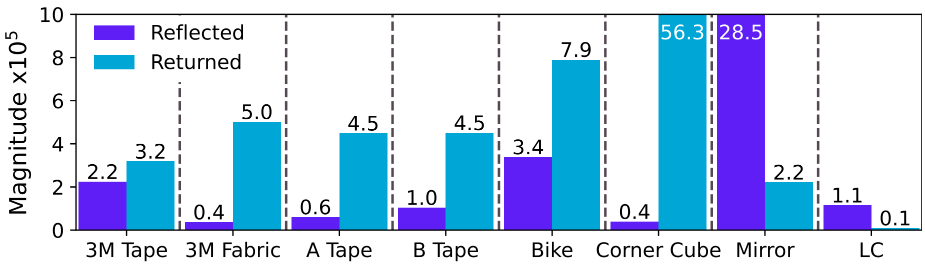

Figure 6.5: **Results of the reflectivity test.**

reflected light that is returned as if the retroreflector was operating as a mirror. This setup is visualized in Figure 6.4. The LED will modulate at 150 Hz and the received intensity is measured using the Fourier transform. First, a recording without any retro-reflector was made, to measure any environmental reflections. These results are subtracted from any subsequent recordings. Moreover, the received signal strength is divided by the total area of the reflector to normalize the results. Table 6.2 gives a brief overview of the different reflectors tested. The A/B Tapes and Corner cube are sourced from Aliexpress.

Figure 6.5 shows the results of the experiment. It is clear that indeed a mirror is an excellent reflector, reflecting an order of magnitude more light compared to any retroreflector, but comes up short regarding returning light. Moreover, we can see that the corner cube is clearly the best retroreflector, especially given its size. However, it is cumbersome to work with and has a fairly narrow FoV. The next best retroreflector is the bike, which clearly edges out the tape and fabric alternatives. However, this reflector is a fixed size and cannot be reshaped to fit the LC it will be mounted against. Finally, the 3M fabric narrowly outperforms the A/B tapes and greatly outperforms the 3M tape. Interestingly, the LC partially operates as a mirror, which will reduce the uplink SNR if the angle is not perpendicular.

From these results we can gather that there are a three suitable options: The corner cube, bike and 3M Fabric. Both the corner cube and bike reflectors suffer from their fixed size, reducing flexibility. Moreover, during testing it became clear that the corner cube is particular about its alignment and has a fairly narrow FoV. While the 3M tape might not have the best retroreflective properties, it does not suffer from any of these drawbacks, which is why we have selected the 3M tape for the prototype.

With this optimization, we try to maximize the signal strength of the uplink. Together with the optimized decoding scheme and FoV PD gain optimizations, we have overcome design challenge 4.

6.1.7 Energy harvester

The energy harvester selected for the battery-less Tag prototype is the 4Ever-Last3.0 from lightricity [43], which is build around the AEM10941 solar energy harvesting IC from e-peas. Our system is configured with three 200 µF capacitors in parallel for a total capacitance of 600 µF. Moreover, these capacitors are charged by six SM101K07L solar cells connected in parallel, for a total area of 46.2 cm^2. At full illumination under the sun, these solar cells would generate 924 mW of power. However, a typical indoor room is 100 times less bright and solar cells don't perform well at low illumination levels. So a power budget around 1 mW is a more reasonable estimate.

The energy harvester outputs two voltage rails: LV and HV, which are 1.8 V and 3.3 V, respectively. Also, a number of status lines are exposed: ST0, ST1 and ST2, these give insight into the status of the energy storage and energy harvesting. The BGM220P will be powered by the LV rail, since it can operate at just 1.8 V, and becomes more efficient at lower voltages. The PD and 5 V booster are powered by the HV rail.

6.2 Optimizing the Tag for battery-less operation

The following sections describe the challenges that arise when creating a battery-less system and how we overcame these. First, we will detail the base power profile and subsequently tackle the highest power component in the system.

The energy is profiled using the Advanced Energy Monitor (AEM) that is part of the development board for the BGM220P. This monitor can accurately measure currents down to 0.1 µA. Furthermore, the AEM works by measuring the voltage drop over a 2.35 Ω resistor, but only when the device is powered by the 3.3 V rail. This means that while these profiles are accurate, they are not fully representative of the current consumption when powered by the energy harvester, since the MCU will then be powered at 1.8 V.

Every profile is based on a transaction between the Reader and the Tag at 0.5 Hz. A transaction starts with the Tag scanning for a preamble, then the Tag will receive and decode the message and backscatter a response.

6.2.1 Base power profile

First, we will discuss the baseline power profile. The baseline is the initial functional prototype without any energy-oriented optimizations. Figure 6.6 shows the base power profile of the initial prototype.

The results show that the system consumes 10 mW on average, which is significantly over the estimated power budget of 1 mW. We can see that the preamble scan is the biggest energy consumer, requiring 11.46 mJ out of the 21.27 mJ total. Moreover, it is vital that the preamble scan is energy-efficient, since we cannot predict for how long the Tag will be scanning for a preamble before detecting a message. Ideally, it would be possible to scan for a preamble indefinitely.

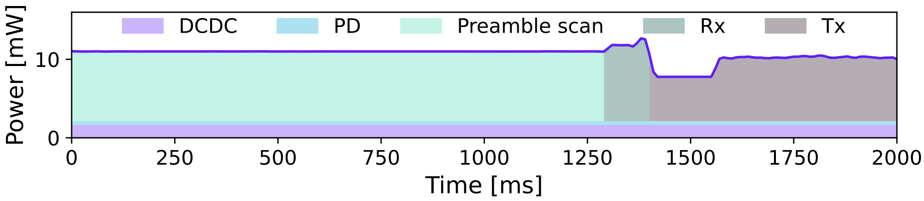

Figure 6.6: **The energy profile of the baseline prototype.**

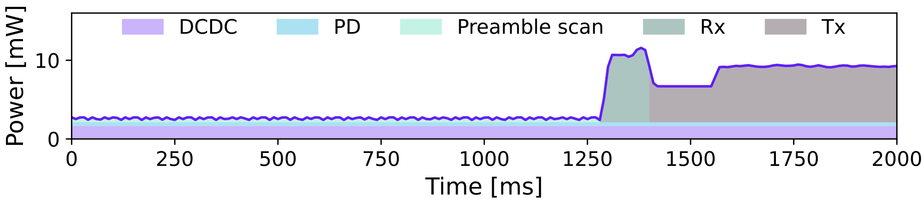

Figure 6.7: **Energy profile showing the efficient preamble scan optimization.**

6.2.2 Preamble scan

Now that a baseline is established, we can iteratively target the highest consuming components. First, we start with the preamble scan.

It is important to understand the function of the preamble scan: it needs to recognize the preamble and start the receiving process. There is no need to continuously scan for a preamble. Instead, if the preamble length is t_{pream}, then we only need to scan every $\frac{t_{pream}}{2}$ and still capture the preamble when it is transmitted. Moreover, we can save energy by putting the MCU to sleep between scans and when the ADC is accumulating.

For the following energy trace t_{pream} is set to 50 ms, which allows us to sleep for 25 ms. The process can be made more efficient by increasing t_{pream} further, however, this means that the responsiveness will decrease and the length (and total energy) of the Rx process will increase. To increase the likelihood of correctly detecting a message, we continuously scan for 2 ms after waking up and instead sleep for only 23 ms. Figure 6.7 shows the power profile with this optimization applied.

The results show that the energy consumed by the preamble scan is reduced by a factor 17.9, to just 0.64 mJ. This optimization reduced the total energy consumption from 21.27 mJ to 9.7 mJ, an improvement of 2.2. Following this improvement, the highest energy component is the backscatter process, which will be tackled next.

6.2.3 Backscattering

Now that we can scan efficiently, we need to backscatter efficiently as well. The uplink transmission signal is generated by the LETIMER, which is able to operate in deep sleep modes. Since the process is already interrupt driven, we can simply enter deep sleep while scanning for a preamble to conserve energy. Figure 6.8 shows the power profile of this optimization.

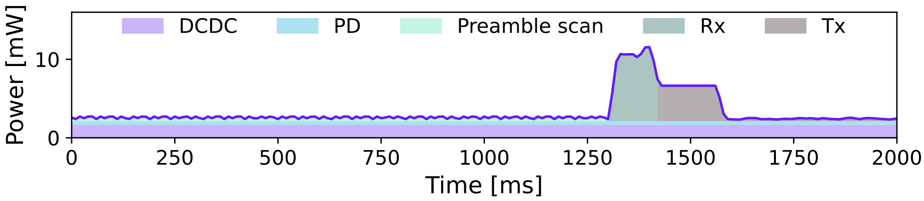

Figure 6.8: **Energy profile showing the efficient backscattering optimization.**

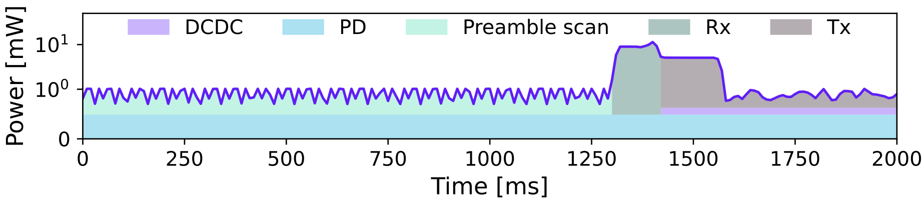

Figure 6.9: **Energy profile showing the efficient DCDC converter optimization. Note the adjusted log scale on the y-axis.**

The results show that the energy consumed by the uplink transmission is reduced from 3.89 mJ to just 0.84 mJ, a 4.6 times improvement. In turn, this reduces the total energy consumed by 32 %, to 6.58 mJ. Moreover, the timer initialization still consumes some energy, which we will tackle later. Now, DCDC converter is the highest energy component.

6.2.4 DCDC booster

Currently, the prototype uses an inefficient off-the-shelf DCDC booster to regulate the 5 V rail that powers the Tag. Moreover, this component is always enabled. By replacing it with a more efficient model and turning it off when the 5 V rail is not required, we can significantly reduce the energy consumed by this component. The DCDC booster is replaced by the TPS61099. This chip has a dedicated enable pin, which can be toggled from the MCU. Figure 6.9 shows the power profile of the this optimization.

The results show that the energy consumed by the DCDC converter is reduced significantly, from 3.26 mJ to just 0.08 mJ, a 40.75 times improvement. In turn, this reduces the total energy consumed by 49 % to 3.33 mJ. The next highest energy component is the PD.

6.2.5 Photodiode

As was the case with the DCDC converter, the PD is always enabled, even if we are not sampling the signal. Sampling the PD is only necessary when we are active during preamble scan and during the Rx state. So, disabling it when unused will preserve energy.

However, the OPT101 does not have a convenient disable pin like the DCDC converter has. Instead, we toggle the PD just like the transmission LED on the Reader: with a low-side switch. We use the same layout as Figure 5.3, except

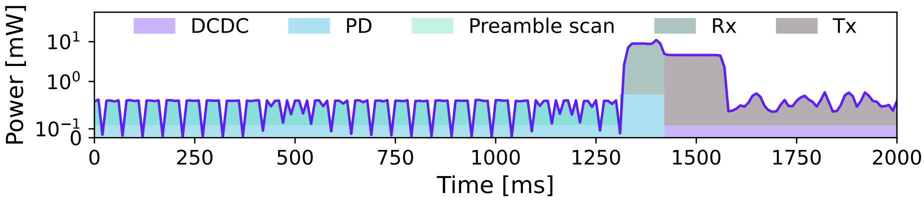

Figure 6.10: **Energy profile showing the efficient photodiode optimization.**

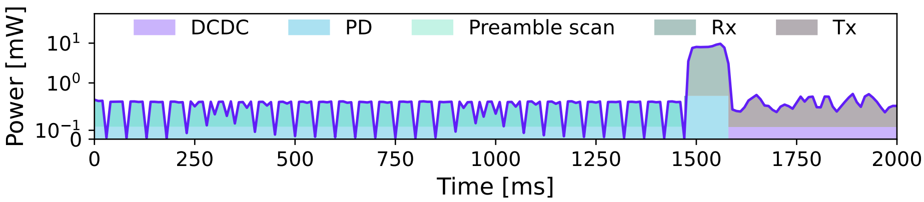

Figure 6.11: **Energy profile showing the optimized transmission initialization.**

Q1 is replaced by a N-Channel BJT transistor (2N3904) and the current limiting resistor R1 is placed between the MCU and Q1, instead of V_{DD} and the device.

Moreover, when the PD is disconnected from ground its output voltage is equal to V_{DD}. Once ground is connected again, the output voltage requires some time to settle. For this reason, the MCU waits for 250 µs before sampling the signal, to give the PD some time to settle. To increase the settling speed, a 2 MΩ pull-down resistor is connected to the output of the PD. Figure 6.10 shows the power profile of the this optimization.

The results show that the energy consumed by the photodiode is reduced moderately, by a factor of 2.88 to 0.24 mJ. In turn, this reduces the total energy consumed by 35 % to 2.16 mJ. The highest energy components are now the receiver and transmitter states. However, enabling deep sleep in the receiving state did not affect power consumption significantly.

6.2.6 Optimized transmission initialization

After the Tag has received and decoded a message, it backscatters a response. However, before the transmission begins, the transmitter needs to be initialized. This takes about 57 ms, where 54 ms is spend in the initialization of the LETIMER. This is caused by the default selected clock (LFXO), this is a crystal oscillator that requires approximately 55 ms to initialize. By changing the clock to one that has a much quicker startup time (LFRCO), we can reduce the total initialization time to just 3 ms. However, the precision of this clock is ±3%, compared to ±100ppm. Moreover, this can be further improved by hardcoding the preamble pattern and disabling excessive UART logging. The final initialization time is just 893 µs. Figure 6.11 shows the power profile of the this optimization.

The results show that the energy used for transmission is reduced significantly, to just 0.14 mJ, from 0.86 mJ, a 6.1 times improvement. This improvement brought the average power down to just 0.74 mW, a considerable improvement from 10 mW. These improvements make it feasible to run the Tag battery-less, powered by just the solar panels and completing design challenge 3.

Chapter 7

Evaluation

This chapter presents the evaluation of the final system. First, we analyze some fundamental performance metrics with everything powered externally, such as the working range, in *static* conditions. After that, we carry out additional experiments in batteryless and dynamic conditions when the drone is flying. Finally, we evaluate the power consumption of the Tag. There are a number of parameters that can be varied for these evaluations:

- Distance: The LoS distance as measured from the LED to the LC.
- Environmental lighting: The intensity of the environmental light as measured on the receiving photodiode, we discern three regions: Dark (<100 lx), Ambient (100 - 500 lx) and outside (>1000 lx).
- PD gain: The gain of the transimpedance amplifier of the PD.
- Scenario: Whether the drone is flying or not, discerned as Static and Dynamic. The Reader is powered from the drone battery in the Dynamic scenario, USB-powered otherwise.

Table 7.1 gives an overview of the system parameters, their possible and default values for both the Reader and Tag.

Parameter	**Possible values**		**Default values**	
	Tag	**Reader**	**Tag**	**Reader**
Distance [cm]	25 - 200		100	
Environmental Lightning	Dark/Ambient/Outside		Ambient	
PD Gain	0.5/0.15†	1.3/0.15†	0.5	1.3
Scenario	Powered/Battery-less	Fixed/Dynamic	Powered	Fixed
Data rate [bps]	120	2500	120	2500
Frame size [bits]	32	56	32	56

Table 7.1: **System parameters used for the evaluation. † For the outside environmental lightning, the gain is reduced to prevent saturation.**

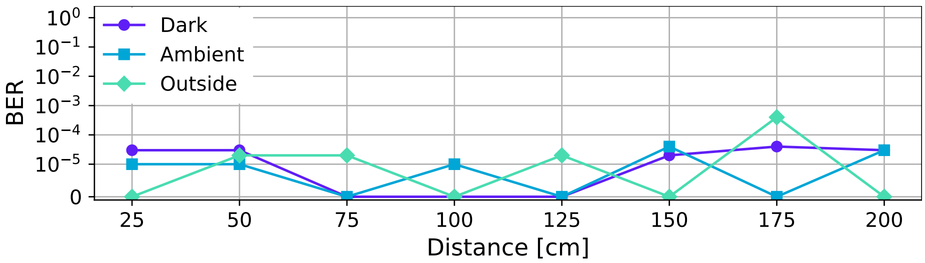

Figure 7.1: **The downlink bit error rate over distance in a static scenario with dark, ambient and outside light.**

7.1 System performance in static scenario

7.1.1 Communication distance

We first measure the impact of distance on the link by varying the LoS distance between the Reader and Tag from 25 cm to 200 cm in increments of 25 cm. In this evaluation, the Reader and Tag are perpendicular to each other (i.e. incidence angle = irradiance angle = 0°). These measurements are separated into two categories: Downlink (Reader to Tag) and Uplink (Tag to Reader). We determine the Bit Error Rate (BER) by transmitting known data and comparing the decoded message with the known value.

Downlink

First, we evaluate the downlink performance. Figure 7.1 shows the results for the downlink. From these results we can make three observations: 1) The downlink signal is strong and the Reader can effectively communicate up to at least 200 cm, even outside. 2) The BER hovers below 0.01 %, regardless of distance, with the exception of a single point in the outside scenario at 175 cm, which exceeds this with 0.05 %. The BER is not 0% because of an occasional bit-flip, setting the BER of that message to 1.8%. This phenomenon indicates that there is an issue in the decoding scheme, since the bit flip occurs regardless of SNR (distance). However, this error rate is consistently low, so Forward Error Correction (FEC) could trivially reduce this to 0 %. Finally, 3) the downlink range is likely much greater than 2 m, since even the reduced gain from the outside evaluation is able to perform well at this distance.

Uplink

We then measure the uplink performance. Figure 7.2 shows the results for the uplink. From these results we can gather that the uplink performs more consistent than the downlink, since the downlink was unable to obtain a BER of 0 % at any distance. This can be attributed to the different decoding schemes and data rates that the Reader and Tag use. Moreover, in a Dark or Ambient scenario, the Tag can effectively communicate up to 150 cm and maintains a BER<1% up to 200 cm. Outside, the effective range is significantly reduced. This is not unexpected, since the intensity of the light received from the Reader

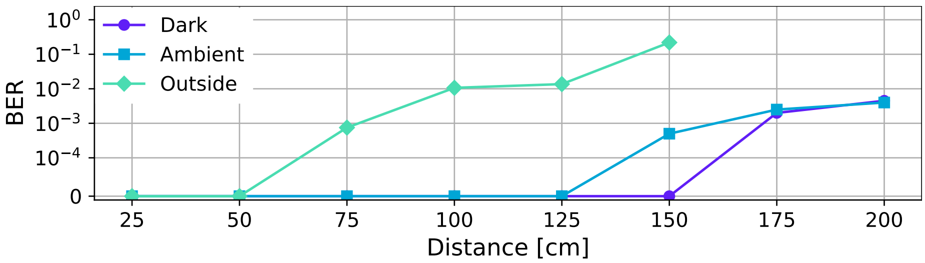

Figure 7.2: **The uplink bit error rate over distance in a static scenario with dark, ambient and outside light.**

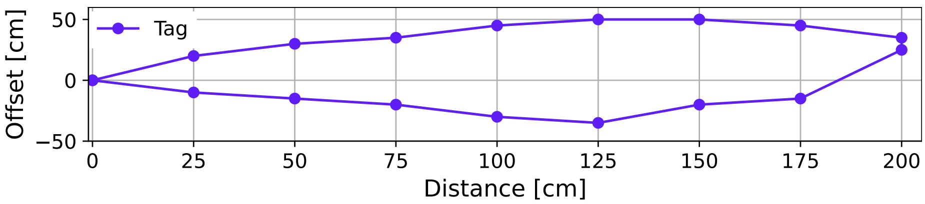

Figure 7.3: **The horizontal (top down view) working range of the Tag.**

on the PD is the same, but the PD gain is much lower. This results in a significantly reduced SNR.

7.1.2 Working range

Now that the effective communication range is known, it is interesting to discern the full working range of the system. For this experiment, the Tag is placed in a fixed position and the placement of the Reader is varied to find the edge of the working range. This edge is defined as the last point where the link BER is <1 %. Moreover, this test is done in the dark to reduce any variance induced by Ambient interference.

Figure 7.3 shows the horizontal working range of the Tag. The figure clearly shows that the range is slightly lob-sided, unable to achieve satisfying performance at 200 cm when perpendicular, but able to do so when slightly offset to the left. This could be caused by two things: One, both devices are asymmetrical, so not every side is treated equally by the channel; two, the LED mounted on the reader is attached at a slight angle, which slightly focuses the light in one direction. Regardless, the working range of the Tag can be considered good and comparable with other works [5, 6], although at a slower data rate. Achieving a working range of 200 cm with a maximum field of view of 35.5°.

7.1.3 Response time

To discern the maximum transaction rate, we need to determine the response time of the system. The response time is defined as the time it takes to complete a single transaction between the Reader and the Tag. A transaction starts once the Reader initiates the transmission and stops once it has decoded the response from the Tag.

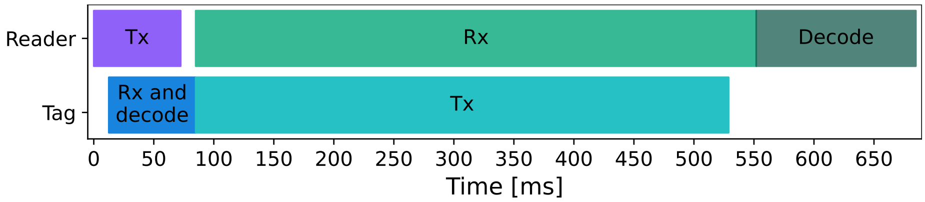

Figure 7.4: **The typical transaction response time.**

Figure 7.4 shows the typical response time of the system. Because the Tag does not continuously scan for a preamble and sleeps at a 25 ms interval, it offsets the Reader Tx and Tag Rx by 12.5 ms on average. Moreover, the Tag does not detect the end of the message while it is receiving data (there is no end of message marker). Instead, this is done afterwards. Consequently, the Tag will receive some irrelevant data after the full message has already been received, since it will receive the full length of the preamble and the data, regardless of when it detected the preamble. This is represented as the overshoot in the 'Rx and decode' state in the Figure.

Moreover, since the Reader continuously scans for a preamble, it immediately detects the uplink transmission from the Reader, but decodes it afterwards. The total measured response time of the system is 685 ms, just shy of 1.5 Hz. Moreover, 65 % of the transaction time is spend in the uplink transmission, the transaction rate could be increased significantly by increasing the uplink data rate or reducing the packet size.

7.2 Link reliability

Now that the fundamental performance is clear, we will focus on the performance when the system is operating in more challenging scenarios, such as battery-less or when the drone is flying. In the case of flying (dynamic) evaluations, the drone is positioned in front of the Tag using a smartphone as remote control.

For the following sections, the Transaction Success Rate (TSR) is used as a metric to determine link quality for this evaluation. A transaction is successful if the Reader can successfully decode a message from the Tag within 15 seconds of the initial transaction. Moreover, the Tag will only backscatter a response if it was able to successfully decode the downlink message. This duration was chosen empirically in order to obtain good performance in all scenarios and environments, we refer to the CDF for a detailed analysis of the response time. Moreover, this means that with the measured flight time of 1.5 to 3 minutes, that the system is able to scan approximately 6 to 12 tags before recharging.

7.2.1 Distance

Static battery-less

First, we evaluate the system performance when varying the LoS distance. In the first scenario, the Tag is operating battery-less, but the Reader remains static. Figure 7.5 presents the TSR and Figure 7.6 shows the CDF of select data points based.

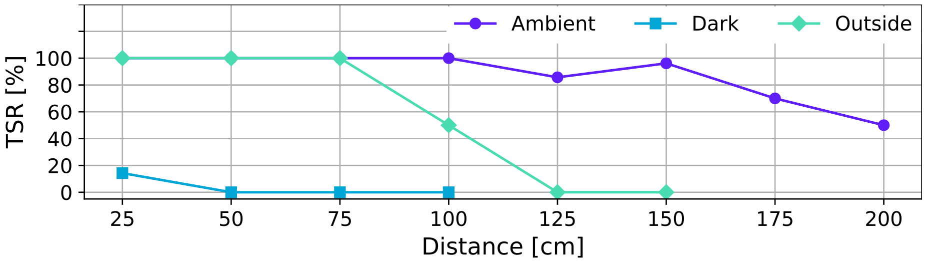

Figure 7.5: **TSR when the Tag is operating battery-less and the Reader is static.**

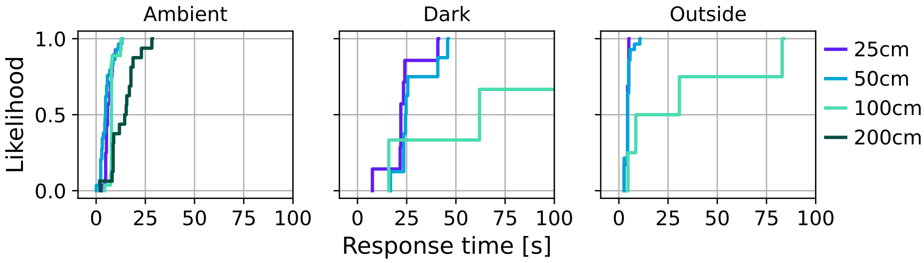

Figure 7.6: **CDF of a successful transaction when the Reader is static.**

These Figures clearly show that battery-less operation is not feasible in the dark, where the system was unable to achieve a reasonable TSR even at close range. As Figure 7.6 shows, the time it takes for the Tag to wake up increases significantly in the dark compared to outside or with ambient light. The response time hovers around the 15 to 25 seconds at close range and increases significantly at further distances.

Furthermore, Figure 7.5 also shows that the TSR degrades faster outside compared to ambient light. These results follow the same pattern as the fundamental results (Figure 7.2) and show that battery-less operation is possible with ambient lighting with a TSR above 84 % up to 150 cm. If we ignore the outlier at 125 cm, the minimum TSR increases to 96 %. Moreover, the system performs well outside up to 100 cm, where the performance takes a noise dive. This is caused by the fact that most messages are missed by the Reader due to the low SNR caused by the low PD gain.

In conclusion, the performance in a static and battery-less scenario is limited by the increased charging time in the dark. Moreover, in the ambient and outside lightning scenarios, the performance is limited by the PD gain. These factors decrease the likelihood of receiving a message correctly, which increases the response time and decreases the TSR.

Dynamic battery-less

In the second scenario, we operate the Tag battery-less and fly the drone. Figure 7.5 presents the TSR and Figure 7.6 shows the CDF of select data points. Additionally, we have evaluated flying the drone with the Tag externally powered, to compare the regression caused by just the dynamic channel.

The Figures show that the performance degrades significantly in a dynamic

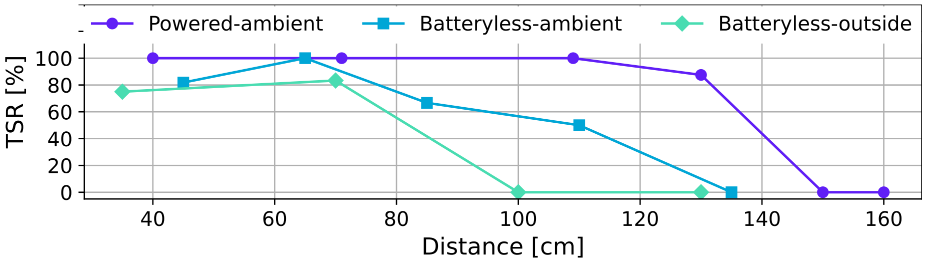

Figure 7.7: **TSR when the Reader is dynamic.**

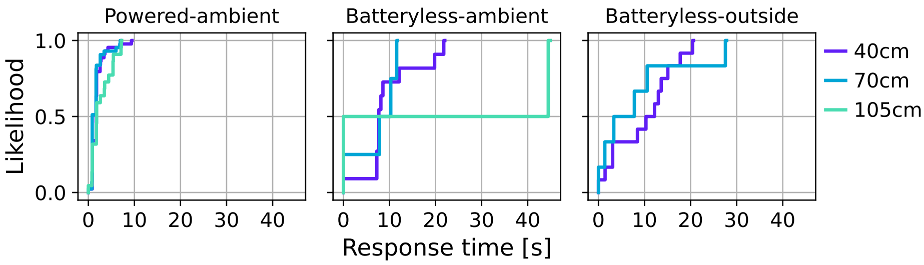

Figure 7.8: **CDF of a successful transaction when the Reader is dynamic. The distances are ±5cm, because the drone cannot be positioned in an exact position.**

and battery-less scenario, both outside and with ambient light. Operating battery-less in ambient light, the system struggles to achieve a TSR above 60 % at 85 cm, even though this is not a problem when the Tag is externally powered. This is caused by the fact that the Tag is able to transmit much more messages when externally powered in the 15 s time frame compared to when operating battery-less, since it does not need to harvest energy. This increases the likelihood of receiving a message correctly.

Moreover, the performance when operating outside degrades further, achieving an effective range of 75 cm. This is caused by the reduced PD gain, which decreases the received signal strength on the Reader.

Interestingly, in the battery-less ambient results, the response time at 70 cm is much lower compared to 40 cm. This could be caused by the slow moving and drifting of the drone when hovering (approximately) in-place. At short-range, these shifts can cause the illuminated area from the Reader to drift away from the Tag, severely attenuating the signal. Moreover, this is not the case for the outside or powered scenario, since the Tag is able to harvest more power, which allows it to transmit more messages. This, in turn increases the likelihood of correctly receiving a message in the 15 second deadline.

In conclusion, the performance in a dynamic battery-less scenario degrades significantly compared to static battery-less or dynamic powered scenarios. This is because at short range, the Reader occasionally fails to illuminate the Tag properly — due to small movements when hovering in place. This is exacerbated by the battery-less operation, which limits the number of messages transmitted, since the Tag needs to harvest its power. These factors increase the response time and, in turn, reduce the TSR.

Scenario	USB-powered		Battery-powered	
	LED off	LED on	LED on	LED on & flying
STD [x1000]	0.23	0.23	0.5	1.96

Table 7.2: **Overview of the Reader noise floor (standard deviation) in different scenarios**

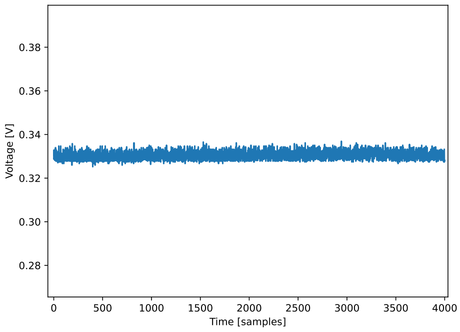

Figure 7.9: **A snapshot of the light measured on the Reader during the noise floor test. The time is measured in samples, and totals to 1.2 seconds.**

Response time

If we look at Figure 7.6 and 7.8, then we can see that the ambient response time generally remains within the 15 second deadline, especially when the Tag is externally powered. This indicates that performance degrades over distance not because the Tag fails to respond, but because the Reader fails to decode the message correctly. Moreover, doing the same comparison for the outside data shows that the majority of the transactions are within the 15 second deadline and the performance degrades because the Reader cannot correctly detect and decode messages anymore. Moreover, the system can effectively communicate in the dark, but requires significant time before the Tag responds. This shows that the Tag is heavily reliant on ambient light for power in order to obtain consistent performance.

7.2.2 Noise floor

The link reliability results show that the performance degrades significantly when operating the Reader in a dynamic channel. To determine why, we set up a comparable experiment as in Section 5.1. We evaluate the noise floor (standard deviation) of the PD input of the Reader prototype in four different scenarios: USB-powered, LED off; USB-powered, LED on; battery-powered, LED on; and battery-powered, LED on, flying. In all cases, the Reader is facing a plain white wall with little ambient light. Table 7.2 gives an overview of the noise in these four scenarios.

This table clearly shows that the noise floor is higher when the Reader is powered by the battery, and especially high when the drone is flying. Moreover, since the drone is facing a plain white wall, we know that there is little ambient

State	Scan		Rx		Tx		Average		Total
Part	**MCU**	**PD**	**MCU**	**PD**	**MCU**	**LC**	**MCU**	**Components**	
Power [uW]	403	184	6880	710	18.9	378	613	268	**881**
Energy [mJ]	528	241	688	71	11	223	1227	535	**1762**

Table 7.3: **An overview of the system power consumption at 3.3 V based on a single transaction in a two second time period.**

interference. A snapshot of the signal when the drone is flying is shown in Figure 7.9, which confirms that there is little-low frequency noise caused by ambient interference. However, the Figure does show significant high frequency noise of approximately 10 mV peak-to-peak. By applying a second order Butterworth low-pass filter we are able to improve the noise floor when flying from 1.96 to 1.20. This shows that improving the analog front-end of the Reader could yield significant improvements. For example, by separating the voltage rails from the drone, adding more decoupling capacitors or fine-tuned filtering.

7.3 Power consumption

The results show that the Tag can operate battery-less. However, it substantially relies on ambient light to properly operate and does not function properly in the dark. We evaluate the power consumption in order to gain an insight in where further optimization could be made. Table 7.3 gives an overview of the energy used per component.

As the results show, the average power consumption of the system is 881 µW at a two second transaction duty cycle. Moreover, the MCU is the largest power consumer, with the preamble scan and Rx at the forefront. Any future power optimization should be focused on these two states. Interestingly, the power consumption is slightly higher than measured at the end of Chapter 6.1. Specifically, the preamble scan state draws 111 µW more power. This could be caused by a higher illumination at the PD, which increases the amount of current used.

The state of the art are able to achieve an average power consumption in the range of 90 µW and 350 µW [5, 6, 21]. However, these systems all use amplitude- or pulse-based modulation technique, which can be decoded more efficiently compared to frequency-based modulation techniques. Moreover, [21] reports that the MCU uses 200 µW on average at 3.3 V. If we take this value as the average MCU power, then the total system power drops to just 468 µW, which is much more competitive.

However, the increase in power consumption is inherent with our frequency-based modulation scheme and the selected decoding method. If we look another work that uses FSK in a low-power context [22], we can see that their receiver consumes 36.1 mW, where the MCU contributes between 23.8 - 26.6 mW at 3.8 V. Compared to this work, our MCU outperforms it by a factor of 41 in terms of energy efficiency — while increasing the data rate by 50 %. Moreover, in Section 8.2.1 we discuss a potentially more efficient decoding method.

Chapter 8

Conclusions and future work

8.1 Conclusions

The aim of this thesis is to research the effect that flying has on the communication performance between a drone and battery-less Tag. First, we analyzed common modulation techniques used for VLC in the context of dynamic channels and LCs. From this analysis we selected M-FSK as the most suitable modulation technique. Then, we detailed how this technique is implemented on the Reader and Tag. After which, we iterate the design of two devices: a drone-mounted Reader and battery-less Tag; and discussed various compromises to make the system operate within the design constraints. Finally, we evaluate the system and show that while the fundamental performance is good, the link degrades significantly when operated outside or when the drone is flying. However, we determine that the performance degrades when flying not because the channel dynamics, but electrical noise imposed on the system because the Reader shares power with the rotors. Moreover, we show that the Tag is able to operate battery-less in all environments except in total darkness. Our results show that the system can communicate up to 200 cm with a BER <1 % when the drone is not flying and the Tag is externally powered. The range decreases to 150 cm when the Tag is operating battery-less and to 85 cm when the drone is flying as well.

8.2 Future work

This thesis presents only the first step in many towards a drone fully that is fully operable using just VLC. The following sections briefly illustrate potential ideas for future work.

8.2.1 Improved decoding scheme

The current decoding scheme relies on applying many FFTs. This requires significant processing power and time, which prevent the MCU from sleeping and saving energy. By developing an improved decoding scheme that offloads

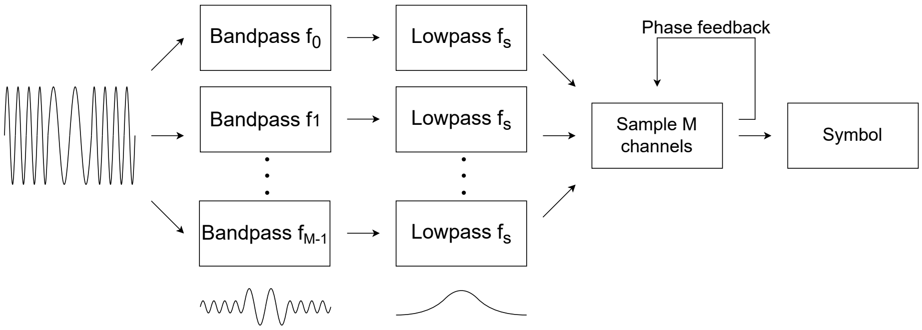

Figure 8.1: **Proposed energy efficient FSK decoder.**

most processing to passive circuitry, we could save an considerable amount of power.

Figure 8.1 shows the proposed decoder. This decoder first band-pass filters the signal for every M-FSK frequency and then low-pass filters these signals at the symbol rate to act as an envelope detector to create smooth peaks (comparable with Figure 4.11).

This decoder can be constructed fully in hardware with passive components (resistors, capacitors and inductors) and only requires the MCU to periodically scan M ADC channels and adjust the phase offset. Depending on the effectiveness of the filters, it is possible to use GPIO interrupt triggers instead. This decoder would allow the MCU to sleep for the majority of the time, which would reduce the effective power consumption considerably during reception.

8.2.2 Improved battery-less performance

While the Tag is able to operate battery-less, the performance is not competitive with other works [5, 6], which advertise a Packet Loss Rate (PLR) below 50 % and are able to operate in the dark. This is affected by two aspects: the Reader LED power and Tag energy efficiency.

For example, our LED is 0.5 W, which is 6 and 24 times less intense than the other works [20]. In turn, this will reduce the amount of energy harvested at the Tag, making it more reliant on ambient light. Moreover, the Tag in these works operates between 97 µW and 184 µW, which is 9 to 4.8 times more efficient than our Tag. These differences enable their system to operate at a lower PLR and in the dark.

Moreover, we have set our maximum response time to 15 seconds. However, comparable drone-based warehouse management solutions claim a response time of 2 to 5 seconds [9, 11]. In order to become competitive with these solutions, the response time of the Tag should be reduced significantly.

8.2.3 Backscattering from the drone

Currently, the LED power is 0.5 W, which is much less compared than other works [20]. Using a more powerful LED requires a bigger or separate battery and possibly a bigger drone. Instead, we could reverse the system and backscatter from the drone.

Given an application where the Tags are placed in a semi-permanent position with easy access to power, for example, a street or traffic light. We could mount and LED on the Tag, since there are no power constraints. Then, we could backscatter from the drone and save considerable power. However, LCs are not lightweight. Namely, the LC currently mounted on the Tag weights 18 g, without any PCB, retroreflector or other components. Given that the Reader currently weights 26 g, it could very well be possible that this method would become too heavy for the drone. However, this remains an interesting research direction that should be explored further.

8.2.4 Increasing drone flight time

Currently, the drone can fly for approximately 1.5 to 3 minutes before it needs new batteries, which is undesirable. This is due to the increased power draw from the Reader LED and weight. The flight time could be improved by:

- Drawing less power, for example, by backscattering from the drone, or using a lower power LED;
- Adding multiple or bigger batteries, such that the drone can better handle the increased power draw;
- Using a more powerful drone, which can supply more power and carry more weight;
- Externally powering the Drone. Some drone-based warehouse management solutions already tether the drone to a base station to allow for unlimited flight time [10, 11].

8.2.5 Increased data rates

While the uplink speed is faster than comparable FSK works [22], both the uplink and downlink are not competitive with the state of the art VLBC works. For example, [6] achieves a 10 kbps downlink and 1 kbps uplink, which is 4 and 8.3 times faster than our work, respectively.

Currently, the downlink is limited by the high processing power required by the decoding scheme on the Tag. Therefore, the downlink data rate could be increased by selecting a faster MCU or with an improved decoding scheme. Moreover, the uplink speed is limited by the response time of the LC and the SNR required to decode the signal. Selecting a faster LC; increasing the SNR by employing a stronger LED; or decreasing the SNR required to successfully decode a message would enable faster uplink speeds. Any future work focused on faster data rates should target these aspects.

8.2.6 Increased range

Currently the fundamental working range is about 1.5 m to 2 m, which is comparable with other works [5, 6]. However, the effective range is reduced to 85 cm when operating in a battery-less and in a dynamic environment. This is caused by electrical noise from the drone and slow response times of the battery-less Tag. Future research should focus on increasing the range of the system to improve general performance and applicability.

8.2.7 MAC layer

Currently, the system operates with a single Reader and Tag. However, some applications cannot guarantee that communications will never overlap, even though VLC is extremely directional. For example, in a warehouse management system, boxes or Tags could be placed near eachother. This scenario would require a basic MAC layer to prevent collisions. For example, sensing the carrier before transmitting would severely reduce the chance of a collision on the channel. Any future work exploring multiple devices should implement a MAC layer in order to communicate effectively in all scenarios.

8.2.8 Flicker-free operation

Currently, no attention has been paid to flicker-free operation. However, if the system is to operate in a human environment, this is something than needs to be taken into consideration. The VLC IEEE 802.15.7 standard recommends modulation speeds above 200 Hz [19] in order to avoid any human health issues. Moreover, research has already been done on flicker free VLC backscattering using FSK [22].

8.2.9 Localization and positioning

This thesis focused on the communication between a drone and a static object. Other works have already looked to positioning with visible light [12]. These works could be combined in order to enable communication, localization and positioning with VLC.

8.2.10 Tag recognition and search

Currently, the Reader assumes that there is always a Tag present in front of it and is constantly attempting to initialize a transaction. For a real-world application, this is not ideal. Instead, camera vision or even VLC could be used to intelligently search for Tags and intelligently start a transaction.

Bibliography

[1] Parth H Pathak, Xiaotao Feng, Pengfei Hu and Prasant Mohapatra. 'Visible light communication, networking, and sensing: A survey, potential and challenges'. In: *IEEE Commun. Surveys Tuts.* 17.4 (2015), pp. 2047–2077.

[2] Luiz Eduardo Mendes Matheus, Alex Borges Vieira, Luiz FM Vieira, Marcos AM Vieira and Omprakash Gnawali. 'Visible light communication: concepts, applications and challenges'. In: *IEEE Commun. Surveys Tuts.* 21.4 (2019), pp. 3204–3237.

[3] David Goldman. 'Sorry, America: your wireless airwaves are full'. In: *CNN Money* 3 (2012).

[4] Harald Haas. 'High-speed wireless networking using visible light'. In: *Spie Newsroom* 1.1 (2013), pp. 1–3.

[5] Jiangtao Li, Angli Liu, Guobin Shen, Liqun Li, Chao Sun and Feng Zhao. 'Retro-VLC: Enabling battery-free duplex visible light communication for mobile and IoT applications'. In: *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications.* ACM, 2015, pp. 21–26.

[6] Xieyang Xu, Yang Shen, Junrui Yang, Chenren Xu, Guobin Shen, Guojun Chen and Yunzhe Ni. 'Passivevlc: Enabling practical visible light backscatter communication for battery-free iot applications'. In: *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking.* ACM, 2017, pp. 180–192.

[7] Yousef Almadani, David Plets, Sander Bastiaens, Wout Joseph, Muhammad Ijaz, Zabih Ghassemlooy and Sujan Rajbhandari. 'Visible light communications for industrial applications—Challenges and potentials'. In: *Electronics* 9.12 (2020), p. 2157.

[8] *Warehouse drone.* May 2022. URL: `www.eyesee-drone.com`.

[9] *DroneScan as a solution.* May 2022. URL: `www.dronescan.co`.

[10] *The autonomous solution to your intralogistics.* Aug. 2022. URL: `doks-innovation.com/`.

[11] *Infinium scan.* Aug. 2022. URL: `infiniumrobotics.com/infinium-scan/`.

[12] Ricardo Ampudia Hernández, Talia Xu, Yanqiu Huang and Marco A Zúñiga Zamalloa. 'Firefly: Supporting Drone Localization With Visible Light Communication'. In: *arXiv preprint arXiv:2112.06677* (2021).

[13] Ahmad Helmi Azhar, Tuan-Anh Tran and Dominic O'Brien. 'A gigabit/s indoor wireless transmission using MIMO-OFDM visible-light communications'. In: *IEEE Photonics Technol. Lett.* 25.2 (2012), pp. 171–174.

[14] Ariel Gomez, Kai Shi, Crisanto Quintana, Masaki Sato, Grahame Faulkner, Benn C Thomsen and Dominic O'Brien. 'Beyond 100-Gb/s indoor wide field-of-view optical wireless communications'. In: *IEEE Photonics Technol. Lett.* 27.4 (2014), pp. 367–370.

[15] Gerard J Holzmann. 'Data Communications: The First 2500 Years.' In: *IFIP Congress (2)*. 1994, pp. 271–278.

[16] Gerard J Holzmann and Björn Pehrson. *The early history of data networks*. IEEE Computer Society Press, 1995.

[17] Alexander Graham Bell. 'The photophone'. In: *Science* 11 (1880), pp. 130–134.

[18] Dobroslav Tsonev, Stefan Videv and Harald Haas. 'Light fidelity (Li-Fi): towards all-optical networking'. In: *Broadband Access Communication Technologies VIII*. Ed. by Benjamin B. Dingel and Katsutoshi Tsukamoto. Vol. 9007. International Society for Optics and Photonics. SPIE, 2014, pp. 1–10. DOI: 10.1117/12.2044649. URL: https://doi.org/10.1117/12.2044649.

[19] Sridhar Rajagopal, Richard D Roberts and Sang-Kyu Lim. 'IEEE 802.15. 7 visible light communication: modulation schemes and dimming support'. In: *IEEE Commun. Mag.* 50.3 (2012), pp. 72–82.

[20] Talia Xu, Miguel Chávez Tapia and Marco Zúñiga. 'Exploiting Digital Micro-Mirror Devices for Ambient Light Communication'. In: *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. 2022, pp. 387–400.

[21] Sihua Shao, Abdallah Khreishah and Hany Elgala. 'Pixelated VLC-backscattering for self-charging indoor IoT devices'. In: *IEEE Photon. J.* 29.2 (2016), pp. 177–180.

[22] Rens Bloom, Marco Zúñiga Zamalloa and Chaitra Pai. 'LuxLink: creating a wireless link from ambient light'. In: *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. ACM, 2019, pp. 166–178.

[23] Purui Wang, Lilei Feng, Guojun Chen, Chenren Xu, Yue Wu, Kenuo Xu, Guobin Shen, Kuntai Du, Gang Huang and Xuanzhe Liu. 'Renovating road signs for infrastructure-to-vehicle networking: a visible light backscatter communication and networking approach'. In: *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. ACM, 2020, pp. 1–13.

[24] Yue Wu, Purui Wang, Kenuo Xu, Lilei Feng and Chenren Xu. 'Turboboosting visible light backscatter communication'. In: *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*. July 2020, pp. 186–197.

[25] Seyed Keyarash Ghiasi, Marco A Zúñiga Zamalloa and Koen Langendoen. 'A principled design for passive light communication'. In: *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. ACM, 2021, pp. 121–133.

[26] J Paulo and Pedro Dinis Gaspar. 'Review and future trend of energy harvesting methods for portable medical devices'. In: *Proceedings of the world congress on engineering.* Vol. 2. WCE, 2010, pp. 168–196.

[27] Taner Cevik and Serdar Yilmaz. 'An overview of visible light communication systems'. In: *arXiv preprint arXiv:1512.03568* (2015).

[28] Hany Elgala, Raed Mesleh and Harald Haas. 'Indoor optical wireless communication: potential and state-of-the-art'. In: *IEEE Commun. Mag.* 49.9 (2011), pp. 56–62.

[29] Suseela Vappangi and VV Mani. 'Performance analysis of dst-based intensity modulated/direct detection (im/dd) systems for vlc'. In: *IEEE Sens. J.* 19.4 (2018), pp. 1320–1337.

[30] Alin-Mihai Căilean, Barthélemy Cagneau, Luc Chassagne, Mihai Dimian and Valentin Popa. 'Novel receiver sensor for visible light communications in automotive applications'. In: *IEEE Sens. J.* 15.8 (2015), pp. 4632–4639.

[31] Ned J Corron, Jonathan N Blakely and Mark T Stahl. 'A matched filter for chaos'. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 20.2 (2010), p. 023123.

[32] Bob Watson. 'FSK: signals and demodulation'. In: *Watkins–Johnson Company Tech–notes* 7.5 (1980).

[33] DJI. *Robomaster TT.* URL: `www.dji.com/nl/robomaster-tt` (visited on 02/07/2022).

[34] Action. *Zaklamp.* URL: `www.action.com/nl-nl/p/zaklamp/` (visited on 02/07/2022). Product number: 2559483.

[35] *IMPROVING ADC RESOLUTION BY OVERSAMPLING AND AVERAGING.* AN118. Rev. 1.3. Silicon Labs. July 2022.

[36] Ronald E Crochiere and Lawrence R Rabiner. 'Interpolation and decimation of digital signals—A tutorial review'. In: *Proc. IEEE* 69.3 (1981), pp. 300–331.

[37] Pimoroni. *LCD Shutter.* URL: `shop.pimoroni.com/products/lcd-shutter?variant=47109536138` (visited on 06/07/2022).

[38] Adafruit. *Large Liquid Crystal Light Valve - Controllable Shutter Glass.* URL: `https://www.adafruit.com/product/3330` (visited on 06/07/2022).

[39] BoutiqueHouse. *2Pcs Active Shutter Eyewear.* URL: `aliexpress.com/item/1005001469733527.html` (visited on 29/07/2022).

[40] Liquid Crystal Technologies. *LCD Shutters.* URL: `www.liquidcrystaltechnologies.com/products/lcdshutters.htm` (visited on 07/07/2022).

[41] Sihua Shao, Abdallah Khreishah and Juan Paez. 'Passiveretro: Enabling completely passive visible light localization for IoT applications'. In: *IEEE INFOCOM 2019-IEEE Conference on Computer Communications.* IEEE, 2019, pp. 1540–1548.

[42] Sihua Shao, Abdallah Khreishah and Issa Khalil. 'RETRO: Retroreflector based visible light indoor localization for real-time tracking of IoT devices'. In: *IEEE Int. Conf. Intell. Comput.* IEEE. 2018, pp. 1025–1033.

[43] lightricity. *Power Sub-Systems.* URL: `lightricity.co.uk/power-sub-systems` (visited on 08/07/2022).