Visible Light Communication Networks for IoT and its Applications

by

Ander Galisteo Zabalo

in partial fulfillment of the requirements for the degree of Doctor in

Telematic Engineering

Universidad Carlos III de Madrid

Tutor & Advisor: Domenico Giustiniano

May 2020

Visible Light Communication Networks for IoT and its Applications

Prepared by:
Ander Galisteo Zabalo, IMDEA Networks Institute, Universidad Carlos III de Madrid
contact: ander.galisteo@imdea.org

Under the advice of:
Domenico Giustiniano, IMDEA Networks Institute
Telematic Engineering Department, Universidad Carlos III de Madrid

This work has been supported by:

To Irati, who gave me the support and emotional framework to navigate this journey.

To my family, who has always believed in me.

Thank you!

Acknowledgements

It is well known that doing a Ph.D. is a lonely journey, but it is not something you can do alone. Along the way, I have had the opportunity to work and interact with a significant number of people, amazing professionals, and beautiful human beings. This thesis would be two or three times longer if I had to mention everyone the way they deserve it. This is why, I will only mention some of them, but although others will not have a place in these pages, for sure will have one in my heart.

The first thank goes to my advisor, Dr. Domenico Giustiniano, who supervised and guided me during my Ph.D. His support during these years can not be described with words. I have learned so much from working with him that I leave believing that I am a much better scientist, engineer, and person than the day I first came to IMDEA Networks.

Secondly, I need to thank all the collaborators that I had during all these years. People like Dr. Marco Zuniga, Dr. Sofie Pollin, Dr. Lorenzo Mucchi and Dr. Thiemo Voigt have been a massive help for me. Nevertheless, I have work more closely with people like Patrizio Marcocci and Jona Baysens with whom I have shared long days in the lab and longer nights... Unluckily in the lab too. You are a clear example of professional behavior, and I do not doubt that you will find success wherever you go. Special thanks to Diego Juara too, who was an incredible help when designing and testing the OpenVLC platform and Héctor Cordobés, who help with the video demo.

Nevertheless, there are two people that I would like to thank above all as they have been my mentors along this way. The first one is Dr. Qing Wang. Without his help, especially in the first half of the thesis, my progress would have been much, MUCH slower. He has been and still is a guidance figure for me. I hope that my work, which is a clear inheritance from his work, lives up to the expectations of the incredible platform he envisioned. The second one is Dr. Ambuj Varshney. He inspired me to do what I consider one of the most fantastic works of these years, your support has been immeasurable, and your desire to do excellent quality work will remain with me for years to come.

Obviously, I can not forget about two groups of people that, although contributed more to my mental state than to my research, are equally important to me.

The first one is all the people from IMDEA Networks. As an institution that is growing (incredibly) fast, IMDEA will face the challenges that institutions that go from small to

large have to face. Still, when I joined, it felt like a small group of people that were excited to research and push the borders of knowledge, always laughing, joking, going to Piratas Rock Wey for lunch (the best restaurant in Madrid) and attending table tennis tournaments. I even won one of those! (I did technically not, but I got a T-shirt, which is nice). The coffee breaks, the memes, the discussions... Will be with me forever. Special thanks to all the people from IMDEA and Comunidad de Madrid that helped me with my communication activities, a not so valued part of the research life that I believe is one of the paths to bring science to society.

The second group is the people from La Caixa scholarship. Getting the La Caixa grant was not only good for the monetary appraisal itself (although that was not bad) but also because of the people I met. The annual retreats that we did were a breath of fresh air where I could interact with non-engineer scientists, which is surprisingly tricky when you are an engineer. My admiration and love for them are immeasurable. And of course, to the people from La Caixa organization itself, they treated us way better than what I would have expected in the best of my dreams.

As an honorable mention, I should thank all the people that I interacted with in University of Uppsala, Sweden. Coming from a mediterranean country, their warm hearts compensated the cold climatology.

Of course, thank you to my girlfriend, my family, and friends to support me when I was not in my best moments and for celebrating with me the sweetest of my victories. This thesis is yours too.

Finally, to anyone reading this starting its Ph.D.: the thesis is longer than it looks. Once you start, there are no breaks until you defend it or drop it, so try to find your pace, do not rush for an ending that is too far. Deadlines come and go, but your mental health is yours only. Take care of it; take care of yourself.

Published Content

This thesis is based on the following papers:

[1] **Ander Galisteo**, Diego Juara, Qing Wang, Domenico Giustiniano. OpenVLC1.2: Achieving higher throughput in low-end visible light communication networks. Published in *14th Annual Conference on Wireless On-Demand Network Systems and Services, WONS 2018* , 6 - 8 February 2018, Isola 2000, France. 10.23919/WONS.2018.8311672

- This work is partially included and its content is reported in Chapter 2.
- The role of the author in this publication is co-designing, implementing and testing the proposed platform with Diego Juara as well as leading the main writing.

[2] **Ander Galisteo**, Diego Juara, Domenico Giustiniano. Research in Visible Light Communication Systems with OpenVLC1.3. Published in *IEEE 5th World Forum on Internet of Things, WF-IoT 2019* , 15-18 April 2019, Limerick, Ireland. 10.1109/WF-IoT.2019.8767252

- This work is fully included and its content is reported in Chapter 2.
- The role of the author in this publication is designing, implementing and testing the proposed platform as well as leading the main writing.

[3] Jona Baysens, **Ander Galisteo**, Qing Wang, Diego Jara, Domenico Giustiniano, Sofie Pollin. DenseVLC: A Cell-Free Massive MIMO System with Distributed LEDs. Published in *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, December 4-7 2018, Heraklion/Crete, Greece. 10.1145/3281411.3281423

- This work is fully included and its content is reported in Chapter 3.
- The role of the author in this publication is co-designing the MAC protocol, leading the design of the synchronization system and the Programmable Real-time Unit (PRU) software programming, collaborated on setting up the experimental setup, testing and writing.

[4] Jona Baysens, Qing Wang, **Ander Galisteo**, Domenico Giustiniano, Sofie Pollin. A Cell-free Networking System with Visible Light. *IEEE/ACM Transactions on Networking*, on February 20, 2020, 2019. 10.1109/TNET.2020.2966322

- This work is partially included and its content is reported in Chapter 3.
- The role of the author in this publication is co-designing the MAC protocol, leading the design of the synchronization system and the PRU software programming, collaborated on setting up the experimental setup, testing and writing.

[Accepted] **Ander Galisteo**, Ambuj Varshney, Domenico Giustiniano. Two to Tango: Hybrid Light and Backscatter Networks for Next Billion Devices. Published in *The 18th ACM International Conference on Mobile Systems, Applications, and Services (Mobisys)*, June 15-19, Toronto, Canada.

- This work is fully included and its content is reported in Chapter 4.
- The role of the author in this publication is to program the processor, designing the protocol and the communication structure, testing, support on the design on the Visible Light Communication (VLC) receiver and writing.

[5] **Ander Galisteo**, Qing Wang, Aniruda Deshpande, Marco Zuniga, Domenico Giustiniano. Follow that Light: Leveraging LEDs for Relative Two-Dimensional Localization. Published in *Proceedings of the 2017 13th International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, December 12-15, 2017, Seoul/Incheon, South Korea. 10.1145/3143361.3143371

- This work is fully included and its content is reported in Chapter 3.
- The role of the author in this publication is working on the mathematical design, simulating it, implementing and testing the proposed platform as well as leading the writing.

[Accepted] **Ander Galisteo**, Patrizio Marcocci, Marco Zuniga, Lorenzo Mucchi, Borja Genovés Guzmán, Domenico Giustiniano. Filtering Visible Light Reflections with a Single-Pixel Photodetector . Published in *IEEE International Conference on Sensing, Communication and Networking (SECON)*, June 22-25 2020, Como, Italy.

- This work is fully included and its content is reported in Chapter 6.
- The role of the author in this publication is working on the design of the experiments, automatization of the testbed, implementation of the receiver side and processing and correcting the samples as well as leading the writing.

Additional Published Content

[6] Ander Galisteo, Diego Juara, Héctor Cordobés, Domenico Giustiniano. Demo: Video transmission using low-cost visible light communication. Published in *ACM Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)* , July 2-5, 2019, Catania, Italy. 10.1145/3323679.3326617

- This work is fully included and its content is reported in Appendix B.
- The role of the author in this publication is co-designing, implementing and testing the proposed platform as well as leading the writing.

[7] Ander Galisteo, Hongjia Wu, Qing Wang, Diego Juara, Marco Zuniga, Domenico Giustiniano. OpenVLC1.2 for Increased Data Rate with Embedded Systems. Published in *Proceedings of the 4th ACM Workshop on Visible Light Communication Systems (VLCS)*, October 2017, Snowbird, Utah, USA . 10.1145/3129881.3129895

- This work is partially included and its content is reported in Chapter 2.
- The role of the author in this publication is co-designing, implementing and testing the proposed platform as well as leading the writing.

Abstract

Visible Light Communication(VLC) has emerged in the last years as a new way to communicate. Using the existing lighting infrastructure, it has great potential to provide high bandwidth and communication security, making it a strong alternative against conventional Radio Frequency (RF) communications. Although VLC addresses several of the problems that RF communications have for specific scenarios, its potential for Internet-of-Things (IoT) applications must still be unleashed.

IoT deployments are, by nature, limited in some way. The limitation could be given by the hardware used, as the cost may need to be minimal to have dense realistic deployments; the energy available, which depends on the battery size of the device; and computing power available, which is given by the available energy and the processing power of the device, among others. Therefore, there is an interest in studying the advantages, drawbacks, and limitations of integrating VLC in IoT scenarios with the constraints mentioned above. This is especially necessary in the case of real-life deployments, mainly if the devices used are multi-purpose, and they need to perform other tasks, such as sensing, in addition to communicating.

First of all, VLC deployments for IoT use the available dense lighting infrastructure to achieve communication on top of illumination in indoor scenarios. The advantage of such an approach is that it allows reusing the existing infrastructure, improving the coverage, and the energy consumed. Although VLC is energy efficient, it consumes more than just illuminating. If the luminaries are not correctly controlled, energy could be wasted by transmitting from a luminary with little or no effect into the receiver. In DenseVLC, we explore the energy consumption of luminaries in dense deployments, and we propose an approach to optimize the Signal-to-Interference-plus-Noise Ratio (SINR) given an energy budget. In order to do so, we propose to coordinate the transmission done by several independent devices concurrently. We introduce a novel synchronization method that uses the Non Line-Of-Sight (NLOS) component of the signal to tackle this problem. Our approach can improve the average system throughput by 45%, or improve the average power efficiency by 2.3 times, compared to existing solutions.

Secondly, even if the required infrastructure follows the design presented above, IoT deployments still need to face one fundamental problem: power management at deployed

mobile devices. Having batteries increases the price of the device, its size, the maintenance required, and the ecological impact that the product has. Removing the battery while still being able to operate under realistic circumstances would be desired. In this work, we study what limitations such a system has. We then propose a new communication scheme, combining VLC and RF backscattering, that allows having continuous end-to-end communication with a custom-designed battery-free device. We design the hardware, software, and protocol that optimizes each aspect of the system to decrease the power requirement of each component. Finally, we evaluate our system and show that it can run with consumptions as low as 95µW, transmit continuously at 500 bits/second, and achieve more than 20 meters on backscattering distance, even with blockage elements as glass and walls covering the Line-Of-Sight (LOS).

Thirdly, we explore one of the multiple applications that the designed VLC systems for IoT allow to implement; device positioning. The majority of the literature requires to have multiple transmitters and/or receivers to achieve localization. The objective of this work is to localize with the minimum amount of necessary hardware, which is critical for IoT applications. We investigate how VLC systems could be used for positioning in dynamic scenarios. Exploiting the fact that, in our scenario, the transmitter and receiver are relatively moving, we propose a mathematical solution that, just using one VLC transmitter and one receiver both equipped with a compass, computes the correct relative position. We then implement our solution in a modified version of OpenVLC and achieve accuracies with less than 5 cm of error. Nevertheless, in this work we assume that the NLOS is non-existant, which may not always be the case.

Finally, we try to overcome the problem mentioned above of NLOS reflections for device positioning with a low resource consumption NLOS component detector. Similar work try to solve this problem computing the Channel Impulse Response (CIR), but for systems with limited resources this is impractical because *1)* The VLC front-end may not be fast enough for acquiring required data for the CIR calculation or *2)* IoT boards are not able to run computationally expensive algorithms in real time. In this thesis, we propose a solution that in complex environments, reduces the localization error using LEDs up to 93%.

In order to perform experimental research in VLC for IoT, a research platform is needed. In this thesis, we also present the latest version of an open-source, software-based, VLC platform, OpenVLC. OpenVLC was first introduced as part of the thesis of Dr. Qing Wang. During this thesis, the platform has been re-designed on both hardware and software. The throughput improved more than 23 times and the transmission distance increased by a factor of 4. In this thesis, OpenVLC, parts of it, or modified versions have been used as a framework to create new IoT systems and explore the practical side of VLC.

As a summary, in this work, we explore how VLC can be leveraged for IoT

deployments. We study the features of such real-world deployments from different perspectives in a variety of scenarios, and we show that realistic implementations of VLC systems are not only possible but doable, enabling new features that IoT developer can exploit.

Table of Contents

List of Tables

List of Figures

List of Acronyms

ADC Analog-to-Digital Converter

AWGN Additive White Gaussian Noise

EM Electromagnetic

FSK Frequency-Shift Keying

GPIO General Purpose Input/Output

GPSDO GPS Disciplined Oscillator

IoT Internet-of-Things

IP Internet Protocol

I/Q In-phase & Quadrature

JCR Journal Citation Reports

LoRa Long-Range

NTP Network Time Protocol

OFDM Orthogonal Frequency Division Multiplexing

RF Radio Frequency

RMSE Root Mean Square Error

RSS Received Signal Strength

RSSI Received Signal Strength Indicator

SNR Signal-to-Noise Ratio

UDP User Datagram Protocol

USRP Universal Software Radio Peripheral

AC Alternating Current

ACK Acknowledgement

BBB BeagleBone Black

BER Bit Error Rate

CDF Cumulative Distribution Function

CFM-Multiple-Input Multiple-Output (MIMO) Cell-Free Massive MIMO

CIR Channel Impulse Response

CTI Cross-Technology Interference

CPU Central Processing Unit

DC Direct Current

DT Decision Tree

D-MISO Distributed Multiple-Input-Single-Output

FoV Field-of-View

FPGA Field-Programmable Gate Array

FW Firmware

HD High Definition

HOS Higher-Order Statistics

HW Hardware

LCD Liquid Crystal Display

LED Light Emitting Diode

LiFi Light Fidelity

LoRa Long Range

LOS Line-Of-Sight

MAC Medium Access Control

MCU Microcontroller

MIMO Multiple-Input Multiple-Output

MISO Multiple-Input-Single-Output

NLOS Non Line-Of-Sight

NTP Network Time Protocol

OOK On-Off-Keying

PCB Printed Circuit Board

PD Photodiode

PER Packet Error Rate

PHY Physical Layer

PRU Programmable Real-time Unit

PTP Precission Time Protocol

RFID Radio-Frequency IDentification

RoI Region of Interest

RSS Received Signal Strength

RX Receiver

SFD Start-Frame-Delimiter

SINR Signal-to-Interference-plus-Noise Ratio

SISO Single-Input Single-Output

SJR Signal-to-Jamming Ratio

SPI Serial Peripheral Interface

SoA State-of-the-Art

SW Software

TCP Transmission Control Protocol

TIA Trans-Impedance-Amplifier

TX Transmitter

V2V Vehicle-to-Vehicle

V2I Vehicle-to-Infrastructure

VLC Visible Light Communication

LPS Light-based Positioning Systems

VLPS Visible Light Positioning System

Wi-Fi Wireless Fidelity

Part I

Introduction

“ And God said, "Let there be light," and there was light. ”

Genesis 1:3, Bible

1 Introduction

Light is one of the fundamental elements that supports life [8] [9]. The energy that it provides has allowed life to flourish. Every single living organism on Earth depends on light to survive directly or indirectly. Thus, for millions of years, life has evolved, optimizing the use of light. While plants use it as an energy source, the evolution has provided animals with mechanisms to see part of the Electromagnetic (EM) spectrum, usually in the form of eyesight. We use it to navigate through the environment, interact with each other, and to get an accurate representation of the world.

Ancient greeks envisioned light as a tool for communicating [10]. With a system of torches located on top of tall towers, they could communicate at vast distances relatively fast for the time. The main drawback of those systems was the maintenance and low communication rate, as the time to turn on and off the torches and assure that the message was received at the other end was high.

The invention of incandescent light was one of the key elements that pushed citizens to include electrical infrastructure in their buildings. They worked by having a thin filament of a difficult to melt metal in semi-vacuum and making an electric current run through it, heat it, and make it glow [11].

In the early XX century, this changed with the invention of the Light Emitting Diode (LED) light. Unlike its incandescent counterpart, LEDs do not require heat to produce light. When an electric current flows through a p-n junction, the recombination of electrons with holes in a lower energy state generates a photon. The frequency of the photon, and therefore the frequency of the light, depends on the band-gap between electrons and holes. In order to obtain different gaps, the use of different materials is required.

Nevertheless, manipulating materials to obtain LEDs with such behavior requires high amounts of precision and technical improvement. This is why it took nearly 40 years since the invention of the first red LED to achieve a blue LED. The invention of the blue LED was so important that the Nobel award of physics was given to its inventors. The blue LED is essential because it was the final piece of the puzzle to achieve white LEDs.

Combining the technology of red, green and blue LEDs, the first white LED was created, and with it, a new era for efficient illumination [12]. Creating light this way is much more energy efficient than heating a thin wire [13] and with that, another advantage was found. In LED illumination, the time that it takes to change the light's properties is orders of magnitude smaller. This allows controlling the light fast enough to embed information on it at rates comparable to Radio Frequency (RF) communications [14].

Visible Light Communication (VLC) and Light Fidelity (LiFi) appears as an alternative to indoor, and vehicular communications. Having great lighting infrastructure deployments, especially in indoor environments, presents itself as an opportunity to implement such communication. This is useful as the number of devices using free bands for communicating has increased dramatically in the last years [15], which increases the probability of finding concurrent communications happening at the same frequency.

Protocols like Wireless Fidelity (Wi-Fi), SigFox, and Zigbee are more popular every day, and thus, the number of devices that uses them increases, saturating the RF bands. Due to the higher directionality of LEDs and the fact that light can not travel through opaque materials, VLC has higher space reusability than other RF alternatives, which is especially important in dense deployments. This also makes it more secure as it is more difficult to intercept the signal without being noticed.

The research community was soon interested in this new type of technology [16]. As the bandwidth of VLC signals is high, some researchers pushed the links to the limit to achieve extremely high throughputs. Nevertheless, although doing practical VLC research is necessary, it has a higher entry barrier compared to doing it with other RF technologies [6]. There are hundreds of out-of-the-box RF solutions available for researchers and designers using different protocols (such as Wi-Fi, Sigfox, Zigbee, GSM, LTE, etc.) and different bandwidths, and even customizable frameworks able be configured for different hardware such as GNU-Radio [17], where users can program their intended behaviour to try their approach without having to do all the hardware and software development by themselves, unlike in VLC.

Component manufacturers do not tend to show the critical features required for VLC in a component's datasheet, which forces designers to do quite a lot of trial-and-error to characterize specific properties of the device they want to use. Examples of this could be the rising and falling time of LEDs designed for illumination, the LEDs bandwidth, the communication bandwidth of the optical receivers, their degradation over time, on top of the expertise required to design, manufacture, solder and test the components in customized devices. All this forces researchers to chose between only doing theoretical work or having to spend a considerable amount of resources and time into developing their VLC platform from zero.

The practical work done around VLC is divided into two groups. The first one is trying to push the limits of what is possible to do with VLC, reaching high data rates

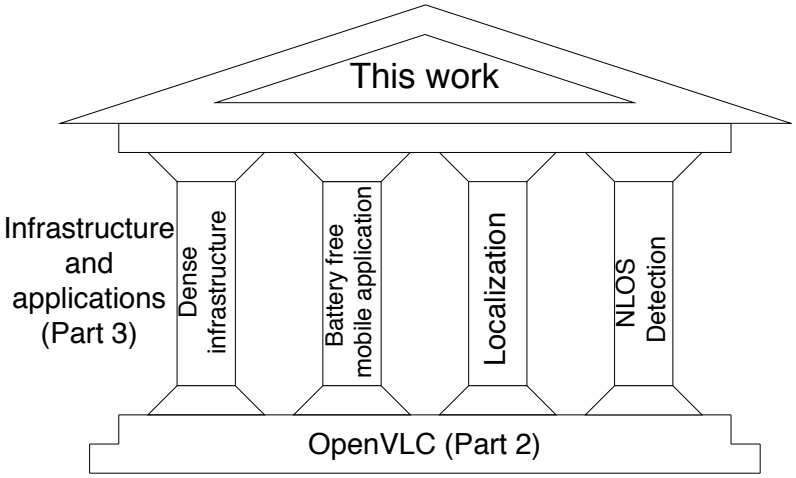

Figure 1.1: Thesis content diagram.

with customized hardware and/or very computation extensive signal processing done in both ends of the communication [18] [19] [20]. Although this type of work is essential to push the limits of what is possible, in the majority of cases is far from being practical in real-life scenarios, as the number of resources needed both in hardware and in processing power makes it impractical or very costly [21]. There are some spin-offs and companies such as PureLifi and Lvx Systems that already do deployments. However, they offer final products, with limited capabilities to optimize them for research-oriented work.

The second group, where this work is located, acknowledges that not all the application scenarios require very high throughputs [22]. For some applications, such as Internet-of-Things (IoT), having lower throughputs is acceptable if the cost of deployment per link is small enough to make it affordable in realistic scenarios. In recent years, several platforms have emerged as the need to have a lower barrier for VLC research is evident [23] [24] [25]. These platforms try to get the cost of the system down by reducing the amount of processing required, thus using simple modulation schemes and off-the-shelf components for their designs.

A diagram showing the structure and content of this thesis is shown in Fig. 1.1.

In this thesis, we present the newest version of OpenVLC, an open-source, embedded, software-based VLC platform. It runs on top of a BeagleBone Black (BBB) and takes advantage of its features to boost the communication capabilities of the system. The custom-designed Hardware (HW), Firmware (FW) and Software (SW) are able, to the best of our knowledge, to achieve the highest communication rates of any embedded system used for VLC, while still having a simple interface that makes it easy to use for a wide range of applications, adapting to the users' needs. This platform is not envisioned as a final product, but more as a tool for researchers and developers to see if VLC fits their needs. During this thesis, OpenVLC has been used as a framework for implementing, prototyping and testing.

As mentioned before, this platform has been designed with the IoT world in mind. Every day, new IoT devices are deployed in a wide variety of scenarios. It is expected that 500 billion devices will be connected to the internet by 2030 [15]. These devices will be deployed both indoor and outdoor, with sensors and/or actuators, wearables, or integrated into the infrastructure, with multiple purposes and design complexities, but all of them will have one thing in common: they will need to be connected to exchange information.

This thesis is presented as a work that tries to explore how VLC can improve IoT designs. We try to explore synergies between both technologies and see in which scenarios VLC could improve the performance of IoT systems. The following designs are the main contributions of this thesis, as shown in Fig.1.1:

- The majority of indoor VLC deployments rely on the current lighting infrastructure. In Chapter 3, we explore how communication affects consumption on the infrastructure's side, and we propose an algorithm to choose transmitters given a power budget. Then, we propose a cell-free approach for VLC communication using a novel synchronization method that allows unconnected transmitters to transmit a frame at the same time.

- One of the main problems of IoT deployments is the energy consumption of the devices, which limits their life-cycle and increases the maintenance costs. In Chapter 4, we present our novel battery-free design, both in hardware and in the software running on it, that uses a hybrid VLC-backscattering communication protocol to achieve ultra-low energy consumption under realistic scenarios.

- In Chapter 4, we explore one of the applications that exploits the properties of VLC; localization. We propose a localization method for dynamic scenarios where the minimum amount of transmitters and receivers needs to be used. Adding a shared reference direction, in this case, a compass, we show mathematically that localization is possible. We further implement the system in real hardware to prove its feasibility.

- Nevertheless, Received Signal Strength (RSS) based localization methods are vulnerable to Non Line-Of-Sight (NLOS) components, common in both indoor and outdoor environments. Usually, Channel Impulse Response (CIR) calculations are required to characterize it, which requires high amounts of computation, not acceptable in system with limited computing power, or multiple VLC frontends with different orientations, which require extra hardware. In Chapter 6, we present a method that uses easily computable statistical values to detect and correct the NLOS component that requires little processing per sample.

1.1. Challenges

During this thesis, uncountable challenges were encountered. Some of them very small, some of them incredibly complex. As VLC is usually not used for IoT, the majority of the challenges came from trying to adapt the VLC technology to IoT environments. The lack of computational power, available energy and device's size forced us to think out-of-the-box to achieve VLC communication with limited resources. Moreover, creating solutions for applications with different requirements is even more challenging, as in this type of systems there is no universal solution. More details about the research challenges of this work are described below.

- *The design of a custom VLC frontend:* Designing the HW for the VLC frontend of the IoT devices supposes a challenge as it is required to have communication on top of illumination. As IoT devices usually require dense deployments, the design needs to be as low-cost as possible, maintaining the capabilities of the device. Balancing cost and functionality was one of the most significant challenges faced in this work.
- *Making it run in real-time:* Our system wants to be affordable, which means that the available processing power is going to be limited. Nevertheless, the system should be deployable in real scenarios. This means that protocols and devices designed and implemented need to work in real-time with limited resources.
- *Realistic scenarios:* The main objective of our research is to be practical. Practicality can not be achieved if the results are only implemented in laboratories under strictly controlled conditions.
- *Having the final user in mind:* The main goal of any practical scientific study like this should be to generate a great impact on the research community and the society. The best way to achieve that is to make it easy for others to use your work. Making the tools and devices created easy to use for others adds another layer of complexity to the design and implementation of any system.

1.2. Contributions

This thesis is supported by 9 publications. 1 in *Transaction on Networking* (indexed in Journal Citation Reports (JCR)). 2 in the tier-1 conference *ACM CoNEXT*[1], 1 in tier-1 conference *ACM MOBISYS*, 1 *IEEE SECON*, 1 in *IEEE WONS* and another 1 in *IEEE WF-IOT*. One workshop paper and a demo are also presented. The contributions of the thesis are,

[1]CORE2014: http://portal.core.edu.au/conf-ranks/ or ERA2010: http://www.conferenceranks.com/

Contribution 1. *New version of OpenVLC.*
Although the platform is not new [26], it has undergone major changes since the beginning of this thesis. The hardware has suffered dramatic improvements, and the communication pipeline has been completely modified, boosting the throughput and increasing the reliability of the system, still having an intuitive interface to use in a wide range of applications. This allows users and developers from all around the globe to use the platform for their projects. For the first time, the throughput of the platform is big enough to stream video.

- **Ander Galisteo**, Diego Juara, Qing Wang, Domenico Giustiniano.OpenVLC1.2: Achieving Higher Throughput in Low-End Visible Light Communication Networks Published in *IEEE 14th Annual Conference on Wireless On-demand Network Systems and Services (WONS).*
- **Ander Galisteo**, Diego Juara, Domenico Giustiniano. Research in visible light communication systems with OpenVLC1.3. Published in *IEEE 5th World Forum on Internet of Things (WF-IoT).*
- **Ander Galisteo**, Diego Juara, Héctor Cordobés, Domenico Giustiniano. Video Transmission Using Low-Cost Visible Light Communication. Published in *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC).*

Contribution 2. *Energy budget optimization and synchronization with NLOS algorithms.*
One of the reasons to use VLC is its power efficiency. As it is also used for illumination, it is expected that it will be densely distributed with multiple users concurrently. As VLC consumes more than pure illumination, selecting which transmitters are used for a particular user is essential to achieve power-efficient communication in Multiple-Input Multiple-Output (MIMO) system. This also requires to synchronize several transmitters, which is achieved by doing it through the channel using the NLOS signal.

- Jona Beysens, **Ander Galisteo**, Qing Wang, Diego Juara, Domenico Giustiniano, and Sofie Pollin. DenseVLC: A Cell-Free Massive MIMO System with Distributed LEDs. Published in *The 14th International Conference on emerging Networking EXperiments and Technologies (ACM CoNEXT).*
- Jona Beysens ,Qing Wang, **Ander Galisteo**, Domenico Giustiniano, and Sofie Pollin. A Cell-free Networking System with Visible Light. Accepted in *IEEE/ACM Transactions on Networking (TON).*

Contribution 3. *A battery-free VLC backscattering hybrid IoT device.*
One of the greatest problems of IoT devices is the requirement for a battery or a connection to the electric grid. In this work, a novel design for combining VLC and backscattering for a battery-free IoT device is proposed, implemented and tested, showing that such a device is not only possible but also useful in real-world applications.

- **Ander Galisteo**, Ambuj Varshney, Domenico Giustiniano. Two to Tango: Hybrid Light and Backscatter Networks for Next Billion Devices. Accepted in *The 18th ACM International Conference on Mobile Systems, Applications, and Services (ACM Mobisys).*

Contribution 4. *A relative localization algorithm for dynamic scenarios and NLOS correction*
Light is being used not only for illumination and communication but also for localization. Nevertheless, the majority of work done assumes not only that at least part of the system is static, but also that multiple light references are available. We propose and implement a new localization technique that only using the minimum amount of transmitters and receivers and a reference direction, can locate the devices relative to each other. We also confront the typical assumption that NLOS components do not affect such systems and implement a solution to detect them using a single pixel photodetector.

- **Ander Galisteo**, Qing Wang, Aniruddha Deshpande, Marco Zuniga, Domenico Giustiniano. Follow that Light: Leveraging LEDs for Relative Two-Dimensional Localization. Published in *the 13th International Conference on emerging Networking EXperiments and Technologies (ACM CoNEXT).*
- **Ander Galisteo**, Patrizio Marcocci, Marco Zuniga, Lorenzo Mucchi, Borja Genovés Guzmán, Domenico Giustiniano. Filtering Visible Light Reflections with a Single-Pixel Photodetector . Accepted in *IEEE International Conference on Sensing, Communication and Networking (SECON).*

1.2.1. Open Source available projects

During the development of the thesis, several pieces of code have been released as open-source. The majority of them are related to OpenVLC, but there are also parts of some deployments and demos.

1.3. Outline of the thesis

This thesis contains several chapters, each one of them presenting the contributions above mentioned.

Table 1.1: Software released as Open Source.

Name	License	Link	Paper
OpenVLC1.3	GPLv3	OpenVLC1.3_revA[1]	[2]
DenseVLC	GPLv3	CoNEXT18_DenseVLC[2]	[3]
Video Demo	GPLv3	Video_demo[3]	[6]

[1] github.com/openvlc/OpenVLC/tree/master/OpenVLC1.3_revA
[2] github.com/jonabeysens/CoNEXT18_DenseVL
[3] github.com/openvlc/OpenVLC/tree/master/Video_demo

We present the OpenVLC platform in Chapter 2, its purpose, the structure of HW, SW, and FW, the reasoning behind the decisions taken and a little processing frame detection technique that allows communicating in real-time with commodity components.

Chapter 3 describes how to optimize the use of resources in a lighting infrastructure given a power budget, and we present a practical solution to synchronize different independent luminaries. In Chapter 4, we present, implement, and test a new system architecture that combines VLC and backscattering to achieve ultra-low energy consumption in a battery-free device. Chapter 5 presents a novel localization system for dynamic scenarios that requires only one transmitter and one receiver equipped with a compass to be able to localize, assuming no NLOS component in the signal. In Chapter 6, we present a solution to detect these components using a single pixel photodetector and Higher-Order Statistics (HOS).

Finally, Chapter 7 summarizes the most important conclusions obtained during this thesis.

Each chapter of this thesis has been written with the goal of being self-contained. This is why there will be a small introduction to get the context and a conclusion at the end.

Part II

OpenVLC: Open, Software based, Embedded VLC Platform

The number of Internet-of-Things (IoT) devices is increasing every day, and thus, the connectivity requirements. The majority of deployed devices use Radio Frequency (RF) technologies to communicate, which leads to a saturation of the available frequencies that, in some cases, may affect the performance of the deployment. Therefore, researchers and developers are continually trying to find new ways to communicate for specific use cases. In the last years, Visible Light Communication (VLC) is starting to emerge as an alternative to conventional RF technologies. Under specific scenarios, the lighting infrastructure can be reused to provide VLC to the devices. Nevertheless, doing dense deployments is costly with the available alternatives, and prototyping is demanding as access to the technology is still not extended. Consequently, a platform that could be used as a framework would not only be exciting but also desired.

In this part, we present the latest version of OpenVLC, OpenVLC 1.3, an open, software-based, embedded VLC platform aiming to fill the gap between high-end VLC systems and IoT deployments. In Sec.2, we show the difference between OpenVLC versions and the structure on Hardware (HW), Software (SW), and Firmware (FW). We present our approach to making it easy to use and, thus, very versatile in order to accommodate the userâs requirements. We show how each element of the system works interconnected and the low resource consumption techniques that we use to assure reliable communication using the limited available processing power.

"Not everything of value in life comes from books- experience the world."
Thomas Alva Edison (1847 – 1931)

2 Affordable VLC communication for IoT devices

Visible Light Communication (VLC) is gaining significant interest as a medium to connect to the Internet [27–29]. In the last few years, a range of applications have been developed with low-end VLC platforms: human sensing [30], communication with toys [31], mobile interaction [32], indoor localization [33, 34] and passive VLC [35, 36]. Industry interest is also resulting in the establishment of the IEEE 802.11bb task group, where the objective is to amend the Medium Access Control (MAC) and Physical Layer (PHY) of IEEE 802.11 with Light Communications [37].

To solve the lack of an open-source and flexible platform for low-end VLC research, we introduced OpenVLC at the VLCS'14 workshop [26], which allowed for quick and flexible testing of new VLC protocols and applications. More recently, we introduced OpenVLC1.2 [38] in an attempt to increase the data rate. However, the board was still working only at relatively short range and was largely affected by light interference.

In this Chapter, we introduce the latest version of our platform, OpenVLC1.3, that increases the data rate and the communication range without adding any hardware cost to the platform. Our contributions are as follows:

- We make a design that occupies a smaller physical space, improve both the Light Emitting Diode (LED) and Photodiode (PD) of OpenVLC1.3 and add high-pass and low-pass filters to minimize the effect of noise in the system, including the Direct Current (DC) from other illumination sources and high-frequency components from the circuitry such as the overshooting generated by the amplification stages.

- We improve software stability and make a new design in software to modulate the LED light. This allows us to increase our sampling rate over 2 MHz and achieve a User Datagram Protocol (UDP) throughput of about 400 kb/s. This throughput could fulfill the needs of a range of applications with only off-the-shelf low-end hardware.

- We design a new technique for computation- and memory-limited fast frame detection to solve the problem of constrained memory in the microprocessor.

Figure 2.1: OpenVLC1.3 cape on top of an embedded board.

- We design and implement a new reception mechanism to avoid the synchronization problems present in previous versions of OpenVLC.

The rest of this Chapter is organized as follows. Section 2.1 introduces the background on OpenVLC and a high-level view of the system architecture. Details on the design of Hardware (HW), Firmware (FW) and Software (SW) for the Transmitter (TX) and Receiver (RX) are presented in Section 2.2 and Section 2.3, respectively. The evaluation results are reported in Section 2.4 and the limits of the OpenVLC1.3 in Section 2.5. Finally, discussions and conclusions are drawn in Section 2.6.

2.1. New System Architecture

The new version of OpenVLC consists of four parts: the BeagleBone Black (BBB) embedded board [39], the OpenVLC cape, the OpenVLC firmware and the OpenVLC driver. The OpenVLC cape is the front-end transceiver that is attached directly to the BBB. The OpenVLC firmware uses real-time processing in BBB's Programmable Real-time Unit (PRU), that work as microprocessors. The OpenVLC driver is a module in the Linux kernel. Both the firmware and the kernel module implement the VLC MAC and Physical Layer (PHY) layers and implement primitives such as sampling, symbol detection, coding/decoding and Internet protocol interoperability. Also, OpenVLC1.3 retains the best characteristics of previous versions, being flexible and open-source and communicating with a low-cost front-end.

With respect to its predecessor, the architecture of OpenVLC1.3 has been re-designed to increase the network performance. The new hardware (HW), called OpenVLC1.3 cape, is shown in Fig. 2.1. The OpenVLC1.3 cape has been modified, reducing its surface by

Table 2.1: Comparison between OpenVLC versions.

Version	TX HW	TX SW	RX HW	RX SW	Data rate
OpenVLC1.0	High Power LED	Kernel software	Basic components	Running in Kernel	18kb/s
OpenVLC1.2	Support for higher power LED and faster modulation	Firmware with user space connection	Faster PD and external amplifier	New frame detection and faster reception in firmware	100kb/s
OpenVLC1.3	More powerful LED and external power	Faster firmware and direct connection to Kernel	Filters to remove interferences and reduced cape size	New frame and symbol detection	400kb/s

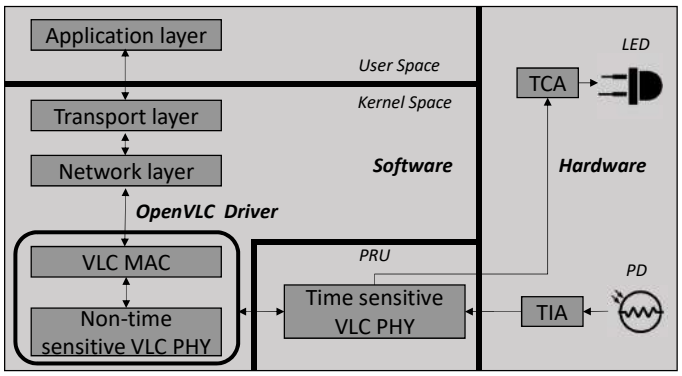

Figure 2.2: The diagram of OpenVLC1.3.

more than 50%. This also allows to use the remaining pins to connect sensors, for instance, for Internet-of-Things (IoT) applications.

The system architecture of OpenVLC1.3 is shown in Fig. 2.2. The hardware runs on external power to allow higher power consumption and harnesses the new LED and PD, together with ancillary circuits, to transmit and receive visible light signals, respectively. The software is responsible for modulating the LED light in order to transmit and sample the incoming signals to receive, both implemented in the OpenVLC1.3 firmware. The software also implements the MAC layer and part of PHY layer in the OpenVLC1.3 driver.

There are three main differences in the design comparing OpenVLC1.3 to its predecessors:

- A new design of the OpenVLC cape (HW).
- A new system architecture (both in SW and HW).
- A FW implemented in the PRUs for data transmission as well as frame and symbol detection.

To boost the date rate in OpenVLC1.3, we exploit the PRUs of the BBB. Time-sensitive operations are implemented in the 2 PRUs that control the General Purpose

Table 2.2: Frame format and sizes (in bytes).

Preamble	SFD	Frame Length	Dst. Address	Src. address	Payload	Reed-Solomon
3	1	2	2	2	0-MAX	16

Input/Output (GPIO) to modulate LED light and perform sampling of incoming signals. This separation was also proposed in OpenVLC1.2, but resulted in overall lower performance and required some module in user space. Communication between the driver and the firmware is now performed using a shared memory. A new technique for computation- and memory-limited frame detection also resides in the firmware (the details are presented in Sec. 2.3). The OpenVLC driver implements the MAC protocol and non-time sensitive PHY operations. This maintains the advantages of software-based flexibility and programmability while increasing its performance.

OpenVLC1.3 is already available to the research community[1]. A summary of the improvements of each version can be found in Tab. 2.1.

2.1.1. Data exchange

The data stream is received in the driver from upper layers. The VLC frame is prepared and then the symbol stream is sent to the shared memory from where it is read by the firmware in the PRU, as seen in Fig. 2.2. The PRU then controls the GPIOs to modulate the LED light for data transmission.

At the receiver, light signals are detected by the PD and sampled by the firmware in the PRUs. Once a valid preamble and Start-Frame-Delimiter (SFD) are detected, received data is sent to the shared memory, and then received and processed by the OpenVLC driver. Finally, the received data is sent to the network layer, where it is handled using the Transmission Control Protocol (TCP)/Internet Protocol (IP) Linux kernel.

2.1.2. Firmware

The firmware of OpenVLC runs in the PRUs of the BBB, which operates at 200 MHz, meaning that each instruction takes 5 ns. Each BBB has its own memory and a shared one between the two. The size of each memory of the PRUs is 8KB and the shared memory is 12KB. The reason behind adopting the PRUs in OpenVLC is to increase the data rate and handle a higher sampling frequency of the Analog-to-Digital Converter (ADC). Nevertheless, this effort also requires a tight timing precision in both the modulation and sampling processes. For this reason, *assembly* is used to program the PRUs. In this way, the code of the PRUs has been designed to know the exact number of instructions executed and, subsequently, the time required to execute them. In addition, the memory

[1] www.openvlc.org

space in each PRU is minimal and this requires careful optimization of all instructions. Finally, there is no enough memory to implement queues, and as such, the communication between the PRUs and upper layers must be handled carefully.

2.1.3. Kernel Driver

The main objective of OpenVLC is to have a flexible, low-cost and reconfigurable system for communication using visible light. In order to do so, OpenVLC1.3 has been designed to be as versatile and easy to use as possible. For this reason, we have taken two design decisions:

- OpenVLC is mostly code-based and the use of VLC hardware is as small as possible. This makes easy to modify the behavior of the platform just by modifying the software code, such as introducing new MAC protocols.
- OpenVLC should be easy to use and adaptable to most use case scenarios. Taking this into account, OpenVLC's interface has been designed as a Linux kernel module.

The OpenVLC kernel module allows us to create a network interface. This means that any user will see the OpenVLC module as just any other network device such as Wireless Fidelity (Wi-Fi) or Ethernet and any application that we would like to run would be connected through the VLC network interface. As the kernel runs in the processor of the BBB, its processing power is much higher than the one of the PRU microcontroller. For this reason, the most computationally demanding tasks are performed in this module.

2.2. Transmitter

In this section we present all the different parts that allow OpenVLC to perform VLC transmission.

2.2.1. Kernel module for transmission

When a user-space application transmits data, first the packet is received from the IP layer of the kernel. After unwrapping the frame, the driver prepares the header for the VLC MAC layer. The frame structure is presented in Table 2.2. Each frame starts with a frame header that contains the following fields: preamble, SFD, frame length, destination address and source address. Each field in the MAC header (starting from the frame length) can be freely modified, for instance, to adapt it to the IEEE 802.15.7-2011 standard for VLC [40] and upcoming 802.11bb standard for integration with Wi-Fi [37].

The preamble consists of 24 alternating HIGH and LOW symbols. After that, the SFD is appended to avoid false positives. The next field denotes the length of frame in bytes, followed by the destination and source addresses.

We use Reed-Solomon code to correct errors in the data field during the transmission. The bits for Reed-Solomon are appended to the frame. We use Reed-Solomon (216,200) error correction code in our default configuration. Subsequently, we use Manchester line encoding, which encodes one bit into two symbols with On-Off-Keying (OOK) modulation (a symbol is either a HIGH or a LOW) and it ensures that the average signal power remains constant. In particular, Manchester line coding converts a 1 bit into a LOW-HIGH symbol pair and a 0 into a HIGH-LOW. This is done to avoid flickering, no matter what data is transmitted. Both Reed-Solomon code and Manchester encoding are also used in the 802.15.7-2011 standard [41]. Finally, the driver places the VLC frame in a shared memory, so that the OpenVLC FW in the PRU can access it.

2.2.2. Shared memory

In Fig. 2.3a we show how the shared memory is used in OpenVLC. The kernel driver transmit data to the PRU using a shared memory. The first 32-bits (referred as the first 'register' in the rest of this Chapter) of the shared memory are the only space where the PRU can write data. This register is used to exchange status flags between the kernel and the PRU.

The communication works as follows: the PRU continuously reads the value of the first register. If it is zero, it rereads it in a loop, waiting for its value to change. If the kernel receives data from upper layers, it will modulate it and put it the shared memory. When it finishes, it writes in the first memory address the number of registers that the PRU should read. The kernel will not be able to write into the shared memory again until the PRU finishes the transmission. The PRU will then transmit the data and once it finishes, it changes the flag in the first memory register to zero so that the kernel knows that the memory is available again.

2.2.3. Firmware for signal transmission

For the transmission, the PRU is used for the sole purpose of emitting the visible light signal according to the pattern of HIGH and LOW symbols stored in the shared memory. Two PRUs are available in the BBB, and only one PRU is used for transmitting the signal. We implement a counter to track the time between symbols. When it reaches zero, a new symbol is transmitted. In order to transmit each symbol, a HIGH or LOW signal is sent through one of the pins of the BBB. In the current implementation, HIGH corresponds to emitting the visible light signal and LOW to not emitting any signal.

2.2.4. Hardware improvement

The main purpose of the TX circuit is to take the signal given by the BBB and amplify it to turn the LED on/off. As mentioned above, the HW is controlled using the firmware

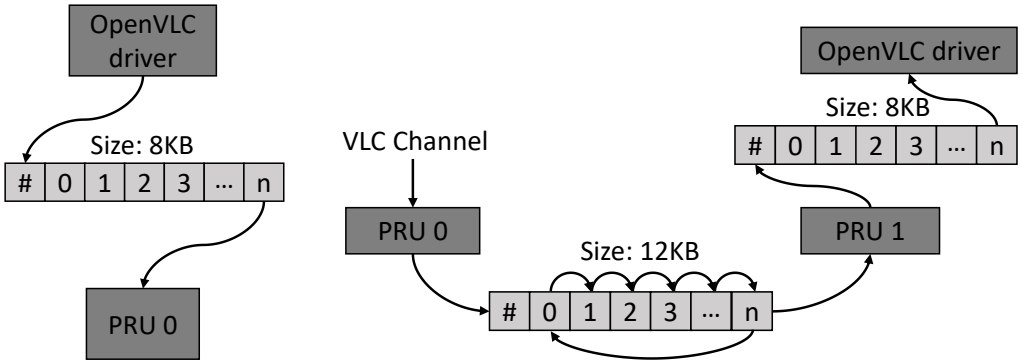

(a) Memory sharing between the OpenVLC driver and the PRU_0 at the transmitter.

(b) Memory sharing between the PRUs and between PRU_1 and the OpenVLC driver at the receiver.

Figure 2.3: Memory interface of OpenVLC1.3 ('#' stores the physical address of the latest updated data).

implemented in the BBB's PRUs. The HW used for the OpenVLC transmitter can be seen in Tab. 2.3. We have improved the HW circuit compared with previous versions mainly to support a higher transmission rate and a larger communication range.

- *Increase the transmission rate:* we use a PRU at the TX to modulate the LED light at higher speeds. Moreover, we use a MOSFET gate driver transistor to control the current flowing to the LED and support a faster switch.
- *To expand the communication range:* a LED that maximizes the power supported by the transmission circuit, working at 2.8 W with a luminous flux of 400 lm. A lens has been attached on top of the LED to better concentrate the optical power and thus, reach further distance. A heat-sink is attached to dissipate better the heat generated by the high-power LED.

2.3. Receiver

In this section we explain the design and behavior of the RX.

2.3.1. Hardware for reception

In the previous versions of OpenVLC, the bottleneck for the throughput was the RX's sampling rate. In OpenVLC1.3, this is solved partly by introducing a new faster PD. This PD does not have its own amplifying circuit. Thus, we add an external amplifier to the RX. The PD's position in the cape is also adjusted for better detection of visible light. The most important components are shown in Tab. 2.3.

Table 2.3: Main components of OpenVLC1.3.

Component	Name
ADC	ADS7883
OP-AMP	LTC6269
MOSFET	FQPF30N06L
LED	XHP35A-01-0000-0D0HC40E7CT
Lens	TINA FA10645
PD	SFH206K
DC-DC	LM2585SX-ADJ

The bottleneck of the system for the transmission distance on the reception circuit was the high sensitivity to noise on the receiver circuit. For this reason, a reception chain has been added between the PD and the ADC. In previous versions, there was only an amplification stage between the PD and the ADC. In this version, as seen in Fig. 2.4, the first amplification stage is a low-noise Trans-Impedance-Amplifier (TIA) that converts the current of the PD into voltage.

Subsequently, a high-pass filter is used in order to remove the low-frequency components (specially the DC component from other illumination sources). The cut-off frequency of this filter is 10 KHz. This filter allows us to remove the DC light component and other sources of interference. Although non-visible for the human eye, light flickering at this frequency would distort the VLC signals. After this, a DC component of 2.5 V is added to the signal so that the signal is centered at half the span of the ADC. Then, the second amplification stage prepares the signal for the dynamic range of the ADC. Finally, before the ADC, a low-pass filter with a cut-off frequency of 1.1 MHz removes the higher frequency noise components mainly due to the overshooting of the amplifiers.

2.3.2. Firmware for signal reception

The configuration of the RX is more complex than the one of the TX. It requires two PRUs. One for handling the HW in a very precise manner and another for processing the received signal.

One of them, PRU_0, performs signal sampling from the ADC and obtains the Received Signal Strength (RSS) and sends it to the other PRU, PRU_1, that handles signal detection and the process of converting the raw signal into bits. PRU_0 reads the RSSs from the ADC at a frequency higher than twice the symbol rate.

Then, the raw value from the channel is shared with PRU_1. PRU_1 interprets the RSSs into symbols for frame detection. PRU_1 continuously checks if a new RSS has been read

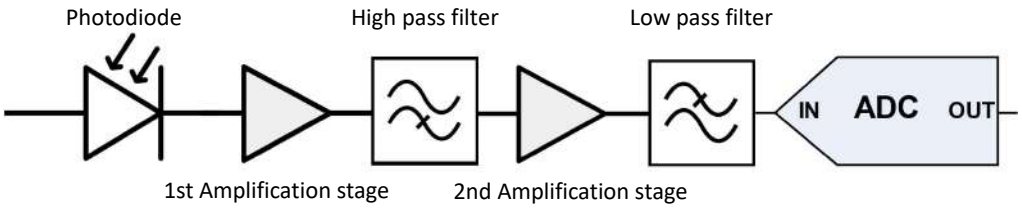

Figure 2.4: Reception chain.

by PRU_0. If yes, PRU_1 processes it immediately.

2.3.2.1. The bit slip problem

One of the most sensitive stages in a communication system is the correct reception of transmitted symbols. One of the main problems with low-cost systems is that TX and RX get easily desynchronized over time. Their clocks are not exactly the same and the frequencies at which they run are slightly different. This could make the system sample a symbol twice or miss a symbol. This problem is known as "bit slip" [42]. In the older OpenVLC1.2, the TX and RX frequency were adjusted to the *instruction level*. This meant that there are exactly the same number of instructions between two symbols transmission and between two symbols reading. However, the clocks in the TX and RX always run at slightly different frequency, and thus, part of the synchronization problems were still present.

In order to solve this problem, we need to make sure that:

- All the symbols are sampled at least twice.
- The system should detect if a symbol has been sampled more than twice.

OpenVLC1.3 over-samples the signal to assure that all the symbols are sampled at least twice. The higher the oversampling rate, the more information the system is going to have to detect the symbol correctly. Nevertheless, a high oversampling rate requires fast processing. In our case, in order to fulfill the requirements mentioned above, the sampling frequency $f_{sampling}$ should be:

$$2f_{symbol} < f_{sampling} < 3f_{symbol} \tag{2.1}$$

With this configuration, OpenVLC1.3 makes sure that we always receive at least 2 sample per symbol and a maximum of 3. It is not possible to receive 4 samples per symbol, which is necessary to assure the second condition. In the implementation, we modulate at 1 MHz and sample at the receiver at a rate of 2.1 MHz.

The symbol detection in OpenVLC1.2 was just a thresholding algorithm with one sample per symbol with the consequent bit sleep problem. In OpenVLC1.3, we avoid this

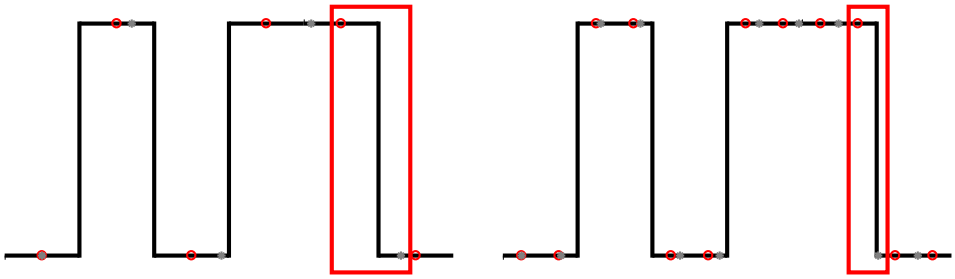

(a) 1 sample per symbol (symbol lost). (b) 2 sample per symbol (no symbol lost).

Figure 2.5: Bit sleep problem. The red circles represent the ideal sampling time. The grey stars represent the real sampling time given by the drift.

by using a pseudo-edge detection algorithm. Manchester modulation converts a 1 bit into a LOW-HIGH symbol pair and a 0 into a HIGH-LOW. This means that the maximum number of symbols with the same value is 2. There cannot be more than 2 consecutive HIGHs or LOWs. With two samples per symbol, the system makes sure to read at least each symbol twice, as it can be seen in Fig. 2.5.

The FW of OpenVLC1.3 counts the number of samples with the same value. If the number of samples is 2 or 3, only 1 symbol has been received. If the number of samples is 4 or 5, 2 symbols have been received. This method can be thought as a very simple and rudimentary edge detection system, as it looks for changes in the signal to see when a new symbol (or pair of symbols with the same value, depending on how long ago was the last change in the signal) has been received.

2.3.2.2. Frame detection

The system assumes that a new frame is being detected when the last samples received correspond to the preamble + SFD. The frame detection technique is the one used in OpenVLC1.2 [1] as it showed to be both low-complexity in processing (required by our application) and very effective. This technique works as follows: First, the RSS values read are compared with their previous values. Because at the beginning of a frame in the preamble (0xAAAAAA) every HIGH is between two LOWs and every LOWs between two HIGHs, every symbol is different from the previous one. If the value of the last 24 bits received is not the same as the preamble, the system continues to collect samples. If it does, it continues receiving data.

Once the preamble is detected, it continues receiving the rest of the frame. Once received and demodulated, the data is sent to the OpenVLC driver for further processing.

2.3.2.3. Communication between PRUs

The signal reception starts with the PRU_0 reading values from the ADC and putting them in a memory shared by both PRUs. This memory is used as a circular memory. When it reaches the end, it continues filling the beginning of the shared memory. In the first register, the address of the latest memory where data has been written is stored as illustrated in Fig. 2.3b. In this way, PRU_1 is able to keep track of the RSS obtained by PRU_0 in real-time. Then, the PRU_1 processes the samples taking two symbols at a time, to determine if the encoded Manchester bit is a 1 or a 0. If they contain valid data, it is decoded and shared with the kernel using the same process as for the transmitter.

2.3.3. Kernel for reception

The frame is received by the kernel after being converted from symbols to bits in the PRU. The Reed-Solomon code is checked with three possible outputs:

- Reed-Solomon reports no errors, so the packet is forwarded.
- Reed-Solomon shows some errors that is able to correct, so corrects them and forwards the packet.
- Reed-Solomon code shows that there are too many errors, so discards the packet.

If the packet is forwarded, the kernel encapsulates the packet so that upper layers can manage it.

Both in the transmission and in the reception of packets, OpenVLC considers that, although more powerful, the kernel cannot run in real-time. For this reason, two driver queues are implemented, one for transmission and one for reception. Every incoming packet that arrives to the kernel is queued and transmitted to the PRU or to upper layers as soon as the resources are available. In this way, we minimize the likelihood of losing a frame because the Central Processing Unit (CPU) is occupied.

2.4. Evaluation

In this section we evaluate the performance of OpenVLC1.3.

2.4.1. Reception chain

In order to understand the behavior of the reception chain we have measured the raw signal with an oscilloscope. As it can be seen in Fig. 2.6, the signal is noisy and small after the first amplifier. Then, the signal is filtered and centered around the center of the ADC's span. Finally, the signal is amplified and cleaned, to improve the reception.

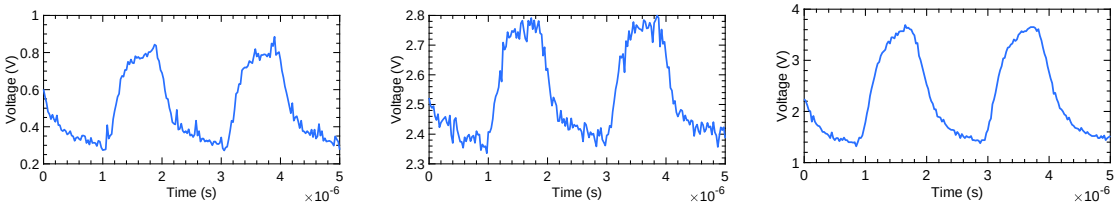

Figure 2.6: Received signal after the first amplifier (left), after the high-pass filter (center) and at the entrance of the ADC (right).

2.4.2. Throughput vs. payload

Setup. We use two OpenVLC1.3 nodes, one as TX and the other as a RX. Since OpenVLC1.3 provides a new network interface that can be easily accessed by upper-layer applications, we use the tool *iperf* to evaluate the UDP performance of OpenVLC1.3.

2.4.2.1. Results

The first test performed is to see how the system works depending on its payload. In previous versions of OpenVLC, the payload had a huge impact on the system. If the payload was too short, the overhead due to the physical layer headers was too big, decreasing the throughput. Nevertheless, if the payload was too big, the reception was desynchronized and the frame lost. As we modify the symbol detection technique, now no frames are lost due to the size of the payload. For this reason, the bigger the payload, the better. All the following tests are done with payloads of 800 bytes.

Throughput vs. distance. This test contains the two most important parameters regarding VLC. The first one is the distance at which the VLC communication takes place. The second one is the maximum throughput achievable by the system. OpenVLC has been tested over distance under 3 different scenarios. In the first one, the system has been deployed in a realistic scenario with no artificial lights on (here OpenVLC is seen as the primary illumination source), but with the windows open during the day ("W. open"). In the second one, OpenVLC is tested without any external light interference and the window shutter closed ("W. closed"). In the last one, we open the windows, and we add an artificial fluorescent light source with frequency components that enter in the frequency band of the OpenVLC receiver ("Interference").

As it can be seen in Fig. 2.7, the maximum throughput that OpenVLC achieves is 400 kb/s. Until 3.5 meters, the difference between being in a completely dark environment ("W. closed") and with external light ("W. open") is negligible, which did not happen in previous versions. This is due to the filters added to the reception chain. Then, at 3.5 meters the intensity of the external light level becomes similar to transmitted light, making it more difficult for the ADC to differentiate between HIGH and LOW symbols, so the throughput starts to drop. When the windows are closed, the communication is

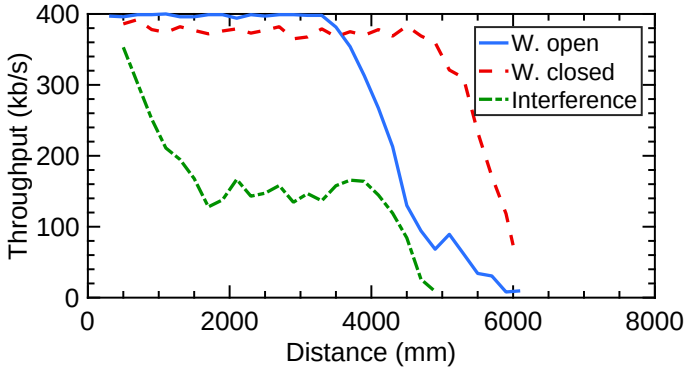

Figure 2.7: UDP throughput as a function of the distance.

possible at longer ranges, reaching 6 meters.

In summary, the maximum achievable distance is more than three times better than the previous version and 6 times better than the original one. Also, the throughput is 4 times better than the previous version and more than 22 times better than the first version [26] [38].

2.4.3. Practical uses

With the achieved throughput and distance, our system is able, for the first time, to support more applications outside IoT. This is why we implemented a video transmission solution using OpenVLC1.3, to show the capabilities of our system. The implementation can be seen in Appendix B.

2.5. Limitations of the system

The OpenVLC1.3 platform has several limitations, as any real system. The first one is that the throughput can not be improved without major changes in both software and hardware. Changing the ADC would increase the cost of the board and the processing power required to perform the reception. Also, although technically possible to use OpenVLC as a transceiver, in the current version, the throughput is maximized using one board as transmitter and one as receiver.

In addition, the communication is UDP, as there is no return VLC link in the system. This design decision has been taken after realizing that:

- The processing power of the BBB's PRUs is limited, and having a bidirectional system requires at least twice as processing power as a one-way link.
- The current trend in networking is that VLC will operate in hybrid systems, where the downlink is VLC and the uplink is. By using the USB interface in the BBB, users may, for instance, use Wi-Fi to send uplink data as well as Acknowledgement (ACK) for downlink VLC.

2.6. Conclusion

In this Chapter, we have presented our latest OpenVLC version and we have evaluated its performance improvements. To the best of our knowledge, OpenVLC1.3 is the first low-cost research platform that achieves a UDP throughput of 400 kb/s using only low-end off-the-shelf hardware. Apart from being used for research and teaching as its predecessors, OpenVLC1.3 can enable real-world applications such as real-time video transmission.

Part III

Real-world VLC infrastructure and applications for IoT

Every year, the deployment of Visible Light Communication (VLC) systems increases in a great variety of scenarios. For Internet-of-Things (IoT) scenarios having a VLC platform like the one shown in Part II helps to accelerate the creation of new systems and applications.

There are several structures that a VLC deployment could have, but in general, there are 3 main elements in the VLC link:

- The VLC Transmitter (TX).
- The VLC Receiver (RX).
- The desired functionality.

The TX and RX are necessary by default for any communication. In some cases, the ultimate goal will be the communication itself. In others, the characteristics of the system will allow to have other features on top of the communication. Each scenario will have unique requirements and properties. It is desired to try to exploit those properties in new imaginative ways to improve the capabilities of the system. In our case, the desired behavior will allow us to localize the devices relative to each other.

In this part of the thesis, we analyze each one of the elements of VLC for IoT. In Chapter 3 we focus on the TX side. We propose a method to synchronize independent dense deployed luminaries and an algorithm to optimize their behavior given a power budget. In Chapter 4 we primarily focus on the mobile device, on a hybrid VLC-Radio Frequency (RF) bidirectional communication method to reduce the consumption of the IoT device and thus, make it battery-free. In Chapter 5, we study the use of light for localization. We designed a model for Visible Light Positioning System (VLPS) with minimum component requirements, use simulations to verify it and implement it using real components. In Chapter 6, the effects of Non Line-Of-Sight (NLOS) components on Received Signal Strength (RSS) based localization systems are studied and a solution is present to detect them using easily computable statistics with a single-pixel photodetector.

"Coming together is a beginning. Keeping together is progress. Working together is success"

Henry Ford (1863 - 1947)

3 Energy-adaptive cell-free massive MIMO for dense VLC deployments

Artificial illumination consumes about 20% of the world's electricity and produces carbon emissions that are comparable to the global automobile fleet [43]. To reduce this high energy consumption, traditional fluorescent and incandescent bulbs are being replaced by energy-efficient Light Emitting Diode (LED)s [44]. Deploying LEDs is not only beneficial for illumination, but also enables Visible Light Communication (VLC) to piggyback on LED's illumination [45, 46]. Yet, the primary function of an LED remains illumination. To increase user comfort, uniform illumination has become an essential requirement in advanced lighting systems. However, due to the well-known Lambertian propagation property of LEDs, it is challenging to illuminate a large area uniformly and efficiently with a single LED [47]. To tackle this problem, deploying LEDs in arrays was proposed in [48] and has been investigated extensively to obtain uniform illumination [49–51].

With this dense-luminaries infrastructure, new techniques for communication have been explored to exploit the LEDs' spatial diversity by enabling the Distributed Multiple-Input-Single-Output (D-MISO) in VLC networks [52–54]. In these previous works, all the LEDs are synchronized to increase the received signal strength at the Receiver (RX). However, these works are highly energy inefficient. As saving energy is the key reason for deploying LEDs for illumination, it is essential that (1) VLC incurs into limited extra power, and (2) no power is wasted. The above works do not satisfy these requirements, as all LEDs acting as a Transmitter (TX) send the same data to a RX and consume the same amount of power, even if they do not contribute equally to the performance improvement at the RX. To improve the power efficiency in D-MISO VLC, recent works have proposed new precoding schemes to serve multiple users [55, 56]. However, these past works are purely based on simulations, only consider a limited number of TXs, static RXs and make simplistic assumptions for the synchronization among the LEDs.

Recently, researchers in Bell Labs proposed the concept of Cell-Free Massive Multiple-Input Multiple-Output (MIMO) (CFM-MIMO) in [57]. In CFM-MIMO, a large number

of distributed, low-cost, and low power access point antennas are connected to a controller to serve a much smaller number of users. The proposed system is not partitioned into cells anymore and all the users are simultaneously served by all LEDs within the RX's field of view. This facilitates mobility and improves dynamic performance, compared to the conventional small cell-based design [58]. CFM-MIMO has now attracted increasing attention from the massive MIMO society, and has been proved to improve the system performance greatly in terms of outage rate, system throughput, and so on [59–62].

Inspired by this new concept, in this Chapter we propose, design and implement *DenseVLC*, a practical CFM-MIMO system enabled by densely distributed LED luminaries. With the intrinsic characteristics of being low-cost, low power consumption, and densely distributed, DenseVLC can enable adaptive and power-efficient CFM-MIMO beamspots to serve multiple RXs simultaneously. Given the measured link qualities between the LED TX and RX, DenseVLC allocates dynamically the power budget for communication among the LEDs such that the total system throughput is maximized. DenseVLC also maintains a constant and uniform illuminance level. Our contributions are summarized as follows:

- We propose DenseVLC and design its system architecture. It enables power-efficient and adaptive CFM-MIMO with VLC to serve multiple RXs. We formulate the policy on adapting the CFM-MIMO beamspots to achieve the best system throughput as an optimization problem and derive the solution (Sec. 3.2).

- We show that given a constraint on power budget, using all the power budget does not necessarily mean the system will operate in the most power-efficient state. We also show that only two modes of operations for the LEDs are enough to efficiently allocate the power budget, either only illumination at the bias current or full swing in the LED communication region. The latter allows us to simplify the hardware and system design of DenseVLC (Sec. 3.3).

- Based on the insights obtained in Sec. 3.3, we further design a heuristic algorithm that can reduce the complexity of finding the optimal adaptation policy by 99.96%, at the cost of sacrificing the system throughput only by 1.8%. This allows DenseVLC to quickly react to users' dynamics (movement, etc.), which is important to maintain the optimum of the system performance in mobile scenarios (Sec. 3.4).

- We propose a novel and practical synchronization method by exploiting non-line-of-sight VLC to synchronize all the TXs that serve the same RX. This synchronization method is more fine-grained, scalable and flexible compared to traditional methods (Sec. 3.5).

- We implement DenseVLC with off-the-shelf devices. We design and develop 36 RXs and 4 RXs from scratch to build a real CFM-MIMOCFM-MIMO system with VLC.

We solve practical challenges in the system implementation, such as those related to the proposed synchronization (Sec. 3.6).

- We evaluate DenseVLC in different scenarios, considering limitations caused by real systems, and compare it with existing solutions. Our results show that DenseVLC improves the average system throughput by 45%, or the average power efficiency by 2.3 times, compared to existing solutions. These improvements are obtained without affecting the uniformity of illumination (Sec. 3.7).

3.1. Design Space and Background

3.1.1. Design space

Our goal in DenseVLC is to design and develop a practical CFM-MIMO system that addresses the following challenges:

Interference. In DenseVLC, multiple RXs can be served simultaneously. Due to the existence of densely distributed LEDs in the system and mobility of RXs, severe interference can arise and degrade the performance of the system greatly. Thus, strong interference among different beamspots must be avoided.

Power efficient. DenseVLC decides which LEDs of the array form beamspots. This adaptive decision should be aware of the power consumption for communication and the Signal-to-Interference-plus-Noise Ratio (SINR) at RXs, and should always lead to the best system throughput within the allowed power budget for communication, without affecting the illumination uniformity.

Fast adaptation. VLC links exhibit high dynamics when the TXs and RXs are not static [28]. We are interested in mobile RXs which are more practical in reality. Therefore, algorithms are required that can quickly determine how to form the beamspots for the RXs.

Scalable and flexible synchronization. Synchronization plays an important role in practical CFM-MIMO systems because TXs need to send the same signal simultaneously. However, none of the existing related VLC systems have proposed a practical solution. As we will show in Sec. 3.5.1, synchronization based on traditional methods such as Network Time Protocol are neither flexible nor scalable for dense deployments. Other methods such as GPS Disciplined Oscillator (GPSDO) [63] cannot be applied, as DenseVLC will be typically deployed indoors.

Minimal circuitry to drive LEDs. The LEDs that form CFM-MIMO beamspots are spatially distributed within the area-of-interest and should have simple electronics and logics to transmit data. This allows their integration into today's infrastructure, retrofitting off-the-shelf LEDs, and further reduce their power consumption for communication.

3.1.2. Background

Channel propagation. In VLC, the intensity of light beam from an LED is captured by the Lambertian propagation model. Denoting P_t as the optical power of the transmitter, the received optical power P_r at the receiver is given as:

$$P_r = H \cdot P_t + n \ , \tag{3.1}$$

where n is the noise, which is modeled as Additive White Gaussian Noise (AWGN), H is the Line-Of-Sight (LOS) path loss, and

$$H = \begin{cases} \frac{(m+1)A_{pd}}{2\pi d^2} \cos^m(\phi) g(\psi) \cos(\psi), & 0 \leq \psi \leq \Psi_c \\ 0, & \text{otherwise} \end{cases} \tag{3.2}$$

where $m = \frac{-\ln(2)}{\ln(\cos \phi_{1/2})}$ is the LED's Lambertian order in which $\phi_{1/2}$ denotes the half power semi-angle, A_{pd} the collection area of the photodiode, d the distance between the transmitter and the receiver, ϕ and ψ the irradiation angle and incidence angle, respectively, $g(\psi)$ the concentrator and filter gain of the photodiode, and Ψ_c the field-of-view of the receiver.

Operation mode of LEDs in VLC. In VLC systems, LEDs can have two operating modes: illumination mode and illumination+communication mode:

- *Illumination mode:* The LED is solely used for illumination in this mode, which can be achieved by applying a constant current to the LED.
- *Illumination+communication mode:* In this mode, the light intensity emitted from LEDs is modulated to transmit data. To avoid flickering, the average brightness of LEDs must remain the same as in *illumination mode*.

3.2. DenseVLC Design

We consider an indoor area with a grid of LEDs distributed on the ceiling. The LEDs illuminate the area and act as TXs to communicate with multiple RXs. The TXs are synchronized dynamically to enable CFM-MIMO beamspots.

3.2.1. System architecture

The system architecture we propose for DenseVLC consists of a controller, a grid of densely distributed TXs, and multiple RXs, as shown in Fig. 3.1.

Controller. It orchestrates the CFM-MIMO beamspots, hosts the decision logic and the Medium Access Control (MAC) protocol, and guides the wireless synchronization among the TXs.

TXs. Densely distributed TXs integrate simple electronics to seamlessly switch between illumination mode and illumination+com-munication mode. In the latter mode, relevant TXs send synchronized signals to the RXs. This simple design of a TX facilitates its implementation in today's indoor environments. *Synchronization among the TXs is achieved through VLC*, where synchronization signal piggybacking on visible light is sent by a leading TX, reflected by the ground, then received and decoded by other TXs. This design for the synchronization makes DenseVLC scalable and flexible.

Mobile RXs. They can move within the area-of-the-interest and receive data from the TXs. A RX is able to measure the channel qualities between the TXs and itself, and send this information back to the controller for the orchestration of adaptive CFM-MIMO beamspots. Note that DenseVLC only requires channel measurements, and does not need to know the location of the RXs.

3.2.2. MAC protocol

The MAC protocol that can exploit the dynamic link qualities between the TXs and the RXs, due to the RXs' mobility or environment changes, works as follows:

- The controller sends pilot signals in a time-division scheme to each LED of the array. Each RX measures the downlink channel qualities between each TX and itself, and reports these measurements to the controller.
- Based on the measured downlink channel qualities and the allowed power budget for communication, the controller implements a decision logic to allocate the communication power among the TXs, with the goal to optimize the system performance. Based on this allocation, TXs form CFM-MIMO beamspots to serve multiple RXs.
- For each CFM-MIMO beamspot, a TX is appointed by the controller as the leading TX that sends synchronization signals through VLC to synchronize all TXs that form the same CFM-MIMO beamspot. The visible light carrying the synchronization signal is reflected by the ground and then detected by the TXs to perform synchronization.
- Synchronized TXs switch their LEDs to operate in illumination+communication mode (cf. Sec. 3.1.2), send synchronized signals to the desired RX. If a TX is not assigned any communication power, it's LED operates in asynchronous illumination mode to only provide illumination.

Next, we introduce our system model that will be used to derive the controller's policy on allocating the communication power among the TXs that form CFM-MIMO beamspots.

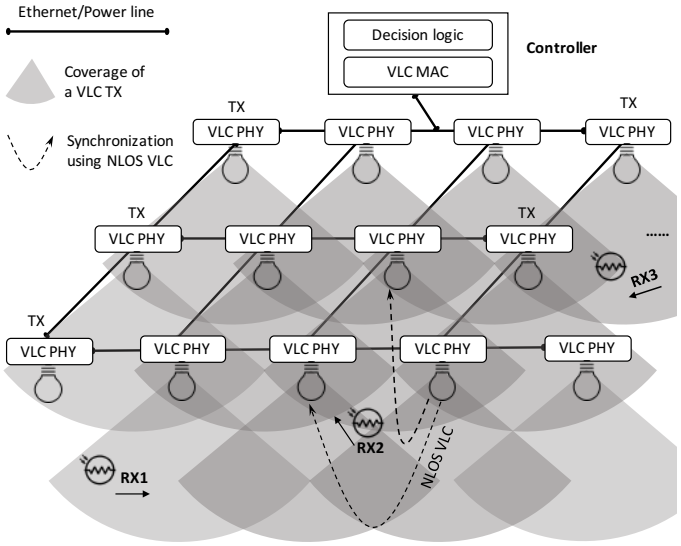

Figure 3.1: DenseVLC's system architecture.

3.2.3. System model

We consider an area with N LED transmitters deployed in a grid on the ceiling. For simplicity, we assume that each TX uses a single LED[1]. Each LED has two operating modes, as introduced in Sec. 5.2, and can switch seamlessly among them. When an LED is in the illumination mode, a constant current I_b flows through the LED to achieve the desired illumination. When it operates in the illumination+communication mode, the light intensity emitted from the LED is modulated to transmit data using a modified On-Off-Keying (OOK) modulation, where *variable swing levels* I_{sw} around the bias are adopted to represent the symbol HIGH and LOW. A large I_{sw} can increase the received signal strength at the RX, but also consumes more power and may generate more interference to other RXs. A current $I_h = I_b + I_{sw}/2$ is used to represent a HIGH symbol, while a current $I_l = I_b - I_{sw}/2$ represents a LOW symbol. This ensures that illumination mode and illumination+communication mode generate the same brightness. The current I_b depends on the desired illumination level. A graphical representation is given in Fig. 3.2.

To avoid flickering and ensure equal probability of the HIGH and LOW levels, Manchester encoding is employed where a transition $I_l \rightarrow I_h$ denotes a binary 0 and a transition $I_h \rightarrow I_l$ denotes a binary 1. As a result, the same brightness of the LED is realized in the two operation modes.

Let $\mathbf{P_t} = (P_t^1, P_t^2, ..., P_t^N)^T$ denote the optical power of the TXs for communication,

[1]In a more general case, a total of M LEDs can be used at each TX to satisfy the illumination level where the power consumed by each TX increases linearly with M.

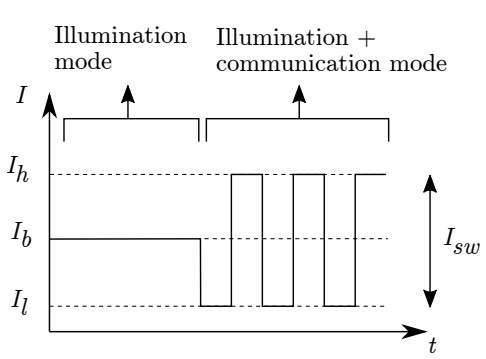

Figure 3.2: Operating modes.

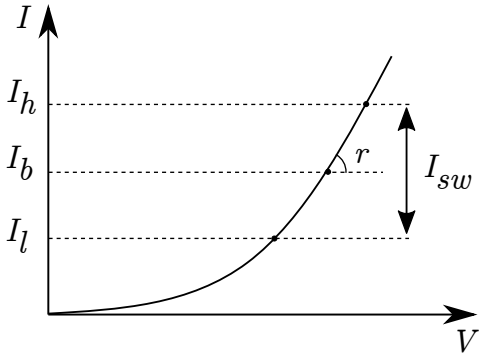

Figure 3.3: LED I-V curve.

and P_r the received optical power at a RX, then from Eq. (3.1) we have

$$P_r = \mathbf{H} \cdot \mathbf{P_t} + n \ , \tag{3.3}$$

where $\mathbf{H} = (H_1, H_2, ..., H_N)$ is the LOS path loss that can be derived from Eq. (3.2).
Objective: Let SINR_i be the received signal-to-interference-plus-noise ratio at RX i, and $P_{C,tot}$ the total power consumption for communication of all the TXs averaged over time. $P_{C,tot}$ can be expressed as:

$$P_{C,tot} = \sum_{j=1}^{N} \overline{P}_C^i, \tag{3.4}$$

where $\overline{P}_C^j$ is the power consumed by TX j for communication averaged over time. Let $\mathbf{I_{sw}} = (I_{sw}^1, I_{sw}^2, ..., I_{sw}^N)^T$ denote the swing level of the TXs. Our objective in this work is to select the swings $\mathbf{I_{sw}}$ that optimize the sum log throughput, ensuring user proportional fairness [64]:

$$\max_{\mathbf{I_{sw}}} \sum_{i=1}^{M} \log\left(B \log_2\left(\text{SINR}_i + 1\right)\right) \tag{3.5}$$

subject to

$$0 \leq \sum_{k=1}^{M} I_{sw}^{j,k} \leq I_{sw,max}, \ j = 1, \ldots, N. \tag{3.6}$$

$$P_{C,tot} = \sum_{j=1}^{N} r\left(\frac{\sum_{k=1}^{M} I_{sw}^{j,k}}{2}\right)^2 \leq \widetilde{P}_{C,tot} \tag{3.7}$$

where M denotes the number of RXs, $I_{sw,max}$ the maximal swing of the LED and $\widetilde{P}_{C,tot}$ the total power budget for communication. The inequality imposes both a minimum bound on the swing to ensure non-negativity and a maximum bound to ensure to operate in the communication region (derived further in Eq. (3.9)).

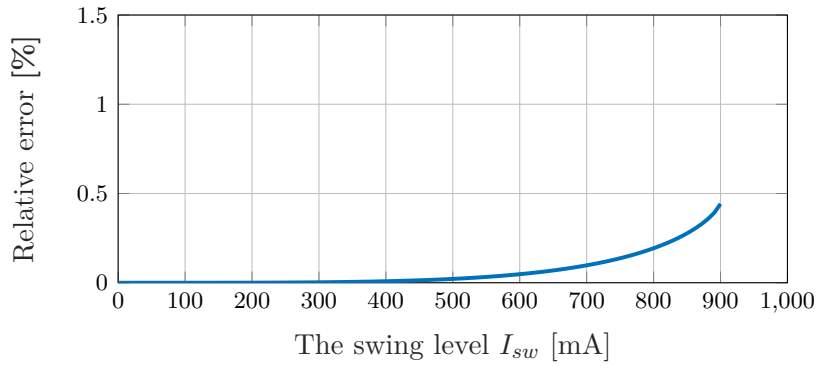

Figure 3.4: Approximation error on power consumption vs. swing level, with $I_b = 450$mA (CREE XT-E LED).

3.2.4. Analysis

To solve the optimization function (3.5) presented above, we need to derive the expressions for $P_{C,tot}$ and SINR_i.

3.2.4.1. Extra power consumption for communication

VLC consumes extra power beyond that for pure illumination [65]. To quantify the extra power $\overline{P}_C^i$ consumed by each LED transmitter i for communication, we use the following equation that models the power consumption of an LED as a function of the current I that flows from the anode to the cathode:

$$P_{led}(I) = k\, V_t\, \ln\left(\frac{I}{I_s} + 1\right)\, I + R_s I^2\ , \tag{3.8}$$

where k is the diode ideality factor, V_t the thermal voltage, I_s the reverse bias saturation current and R_s the series resistance of the LED [66]. As the modulation is performed around the bias current I_b, the Taylor expansion[2] of Eq. (3.8) around I_b gives the following expression:

$$P_{led}(I) \approx \underbrace{P_{led}(I_b)}_{P_I} + \underbrace{\frac{P'_{led}(I_b)}{1!}\,(I - I_b) + \frac{P''_{led}(I_b)}{2!}\,(I - I_b)^2}_{P_C}\ , \tag{3.9}$$

where the first term $P_{led}(I_b)$ represents the power consumption for illumination P_I, and the second and third terms refer to the power consumption P_C for communication. Since we use Manchester coding, the probabilities of symbol HIGH and symbol LOW in the modulated data are equal. The average power consumed on communication over time $\overline{P}_C$

[2]Note that we consider the Taylor expansion up to the second-order term for the approximation in Eq. (3.9), sufficient for small variations around I_b.

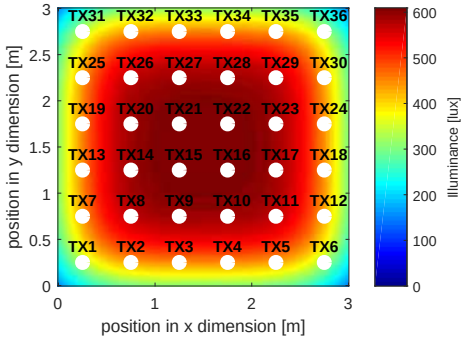

Figure 3.5: Illumination distribution.

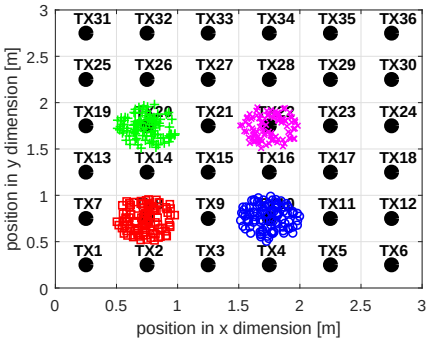

Figure 3.6: 100 random instances.

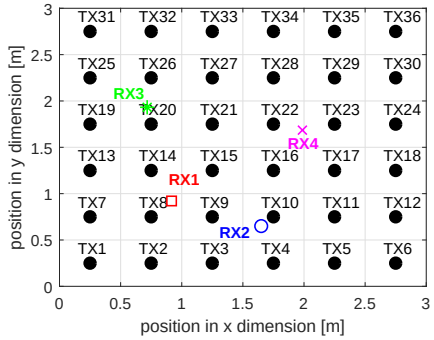

Figure 3.7: An illustrated instance.

is:

$$\overline{P}_C = \mathbb{E}[P_C(t)] = \frac{\frac{kV_t}{I_b} + 2R_s}{2!} E[(I(t) - I_b)^2] = r\left(\frac{I_{sw}}{2}\right)^2, \tag{3.10}$$

with $r = \frac{kV_t}{2I_b} + R_s$ the LED's dynamic resistance at the bias working point I_b. A graphical representation is shown in Fig. 3.3. Note that when the LED is used purely for illumination, the swing $I_{sw} = 0$ and $\overline{P}_C = 0$, confirming our derivation. From Eq. (3.10), we can obtain the total extra power consumption for communication in the channel:

$$P_{C,tot} = \sum_{i=1}^{N} \overline{P}_C^i = \sum_{i=1}^{N} r\left(\frac{I_{sw}^i}{2}\right)^2. \tag{3.11}$$

Notice that the Taylor expansion is valid for small excitations around I_b. However, since the LED I-V curve can be approximated by a second order function, the approximation is also valid for larger swing levels, as shown in Fig. 3.4. Using the maximum swing level of $I_{sw} = 900$mA results in an error of only 0.45%, which is acceptable.

3.2.4.2. Signal-to-interference-plus-noise ratio

The SINR at RX i is computed as follows:

$$\mathrm{SINR}_i = \frac{\left(R\eta r \sum_{j=1}^{N} H_{j,i}(\frac{I_{sw}^{j,i}}{2})^2\right)^2}{N_0\, B + \left(R\eta r \sum_{\substack{k=1 \\ k\neq i}}^{M} \sum_{j=1}^{N} H_{j,i}(\frac{I_{sw}^{j,k}}{2})^2\right)^2}\,, \tag{3.12}$$

with R the responsitivity of the photodiode, η the wall-plug efficiency of the LED (i.e., the energy conversion efficiency from electrical to optical power), N_0 the single-sided spectral power density and B the communication bandwidth.

Note that the bias current I_b is not included to calculate the SINR, since it does not contribute to data transmission.

Since we have derived the expressions for $P_{C,tot}$ and SINR_i, then we can solve the optimization function (3.5) with nonlinear programming tools, such as `fmincon` in Matlab. Note that inequality (3.7) depends on the bias current I_b necessary to guarantee the required illuminance of the environment. Setting the bias I_b at the center of the linear region of the LED allows us to use a larger $I_{sw,max}$ (cf. Fig. 3.3). The opposite holds for a smaller or larger value of I_b, as the Taylor approximation will only be valid in a smaller region.

3.3. Insights into the System Design

In this section, we provide an in-depth study of the system model introduced in Sec. 3.2, and present the main insights that will drive us in the design of a practical D-MISO system.

Simulation setup. We consider an indoor area of 3 m × 3 m × 2.8 m consisting of a grid of $N = 36$ TXs and $M = 4$ RXs. The TXs are aligned in a 6 × 6 array with an inter-node distance of 0.5 m. They are attached at a height of 2.8 m from the ground, facing downwards. The 4 RXs are located at a table of height 0.8m, facing upwards to the ceiling. The LED data is based on the off-the-shelf CREE XT-E LED [67], which is also used in our experimental validation in Sec. 3.7. Table 3.1 lists additional relevant system parameters.

Illuminance distribution. Fig. 3.5 shows the spatial illuminance distribution. According to the ISO 8995-1 illumination standard for indoor office premises [68], the average illuminance should not fall below 500 lux and the illuminance uniformity (the ratio between the minimum illuminance and the average illuminance) shall not be less than 70%. We define an area of interest of 2.2 m × 2.2 m (i.e. excluding the boundary), positioned in the center of the room. In this area, the average illuminance is 564 lux and the uniformity equals 74%, satisfying the illumination requirements.

Table 3.1: System parameters for simulations.

Parameter	Notation	Value
General		
Single-sided spectral power density	N_0	7.02×10^{-23} A^2/Hz
Communication bandwidth	B	1 MHz
LED		
Half power semi-angle	$\phi_{1/2}$	15°
Reverse bias saturation current	I_s	1.44×10^{-18} A
Ideality factor, series resistance	k, R_s	2.68, 0.19 Ω
Bias current, wall-plug efficiency	I_b, η	450 mA, 0.4
Maximum swing current	$I_{sw,max}$	900 mA
Wall-plug efficiency	η	0.40
Receiver		
Field of view, collection area	Ψ_c, A_{pd}	90°, 1.1 mm^2
Responsitivity	R	0.40 A/W

3.3.1. Throughput vs. power consumption

Without loss of generality, we assume that the 4 RXs are located around certain TXs in the XY-plane, as shown in Fig. 3.6. For each RX, we generate 100 random positions around these TXs. We gradually increase the power budget $\widetilde{P}_{C,tot}$ for communication, and derive the optimal swings that lead to the maximal sum log throughput, by solving the optimization problem in Eq. (3.5). The results of the system throughput and each RX's throughput, both with 95% confidence interval, are depicted in Fig. 3.8.

We can observe that either the system or the individual's throughput increases in accordance with $\widetilde{P}_{C,tot}$. Since we optimize the sum log throughput, the throughput of the RXs is balanced independent of $\widetilde{P}_{C,tot}$, ensuring user fairness. When $\widetilde{P}_{C,tot}$ is large, RX3 and RX4 achieve better performance than RX1 and RX2, due to the availability of more non-interfering TXs. However, we should note that the power efficiency on communication decreases when $\widetilde{P}_{C,tot}$ exceeds a certain threshold. For instance, we can notice in Fig. 3.8 that the system throughput increases more slowly with the same extra power consumption when $\widetilde{P}_{C,tot}$ exceeds 1.2 W.

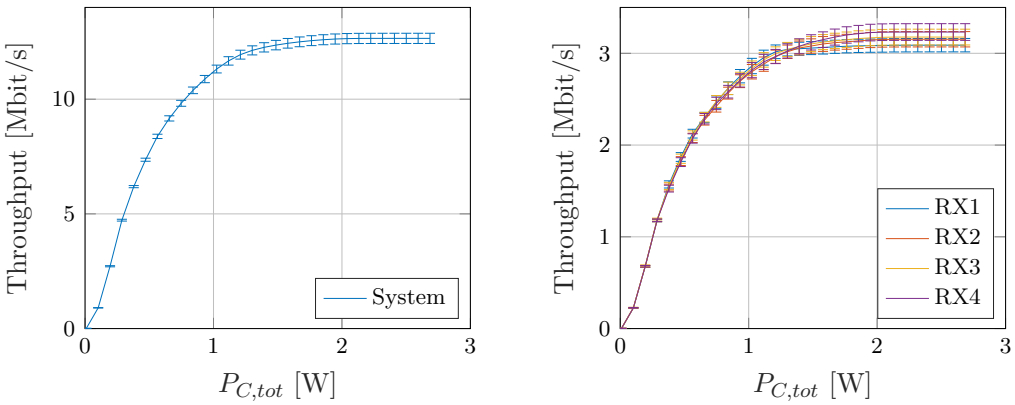

Figure 3.8: Average throughput over positions with 95% confidence interval versus communication power.

3.3.2. Optimal swing levels

To show the results clearly, we consider a random instance of the RXs' positions used in the previous section. For this instance, the positions of the four RXs are shown in Fig. 3.7. The optimal swing levels leading to the maximized system throughput under different power budget $\widetilde{P}_{C,tot}$ are shown in Fig. 3.9. Due to the space limitation, we only show the results related to RX1 and RX2. We can observe that with the increase of $\widetilde{P}_{C,tot}$, TX8 and TX10 are assigned first to RX1 and RX2, respectively, i.e. the swing level of a TX that poses the best channel increases first. When TX8 and TX10 reach the maximal allowed swing, the consumed communication power equals $P_{C,tx,max} = r\left(\frac{I_{sw,max}}{2}\right)^2 = 74.42\,\text{mW}$, where r denotes the dynamic resistance. Increasing $\widetilde{P}_{C,tot}$ further, more TXs start sequentially from zero-swing to operate at full-swing. For example to serve RX1, a group of TXs, forming the same CFM-MIMO beamspot, operate at full swing in the order of TX8→TX14→TX7→TX2→TX1→TX13.

Notice that the transition of a TX's optimal swing level from zero-swing to full-swing is fast. As shown in Fig. 3.9, with the increase of $\widetilde{P}_{C,tot}$, the occurrence of the gray areas (representing that the TXs are operating at neither zero-swing nor full-swing) is negligible compared to the occurrence of the black (representing zero-swing) and white (representing full-swing) areas. This phenomenon applies to the 100 instances shown in Fig. 3.6.

For four representing TXs (TX3, TX5, TX10, and TX15), we depict the empirical Cumulative Distribution Function (CDF) in Fig. 3.10 for five randomly selected instances[3]. TX10 has a very steep edge in the CDF at the maximum swing level. This means that TX10 mostly operates at full-swing because it possesses the best channel to RX2 among all the TXs. A similar trend is observed for TX5, although it is assigned later because it is farther from RX2 compared to TX10. Therefore, we see an offset in

[3]For better visualization, we only show the results for five instances.

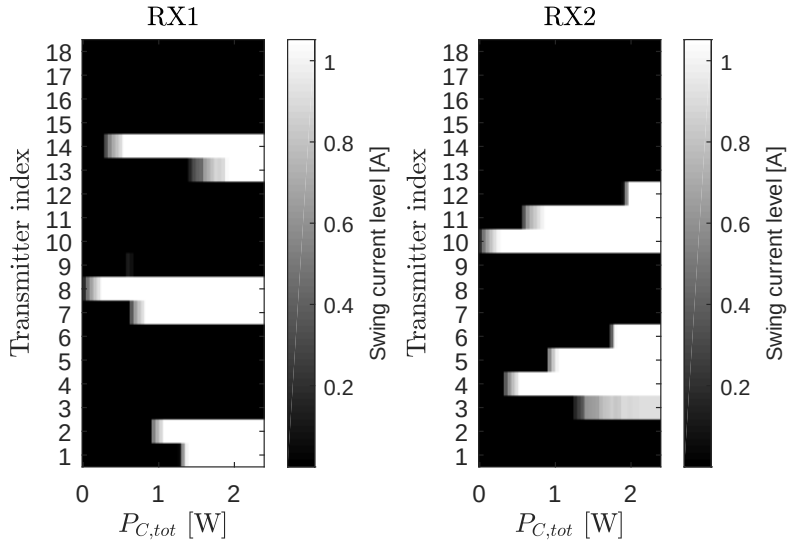

Figure 3.9: Swing levels vs. communication power.

the CDF with respect to TX10.

An interesting observation in Fig. 3.10 is the empirical CDF for TX3. Instead of being a step function, the CDF of the optimal swing level of TX3 increases more smoothly and does not always reach the maximal swing, meaning that optimally it operates more often than other TXs at neither zero-swing nor full-swing. However, discretizing TX3's swing level to either zero or maximum results in a little loss with respect to the optimal swing level. The gain contributed by TX3 to the total system throughput is rather limited: if it is not used, the average loss in the total system throughput is only 0.48%. Finally, TX15 is not used at all because it would generate too much interference to the RXs.

Based on the above-presented characteristics of the optimal swing levels, we have the following three insights:

Insight 1: Given a power budget for communication, for each RX, the optimal solution always assigns the available power sequentially to its preferred TXs. This means that the optimal solution starts to assign the remaining communication power to the next preferred TX, only after the swing level of the RX's first preferred TX reaches the full-swing.

Insight 2: With the increase of the power budget, the transition of a TX's optimal swing level from zero-swing to full-swing is fast. That is, the probability of a TX operating at at neither zero-swing nor full-swing are negligible compared to the probability of operating at either zero-swing or full-swing. Therefore, to achieve a near-optimal throughput in practice, it is sufficient for each TX to either use (i) I_{sw}=0 if operating in illumination mode, or (ii) the maximum swing $I_{sw,max}$ if operating in illumination+communication mode.

Insight 3: If the assignment of a TX to serve an intended RX generates much

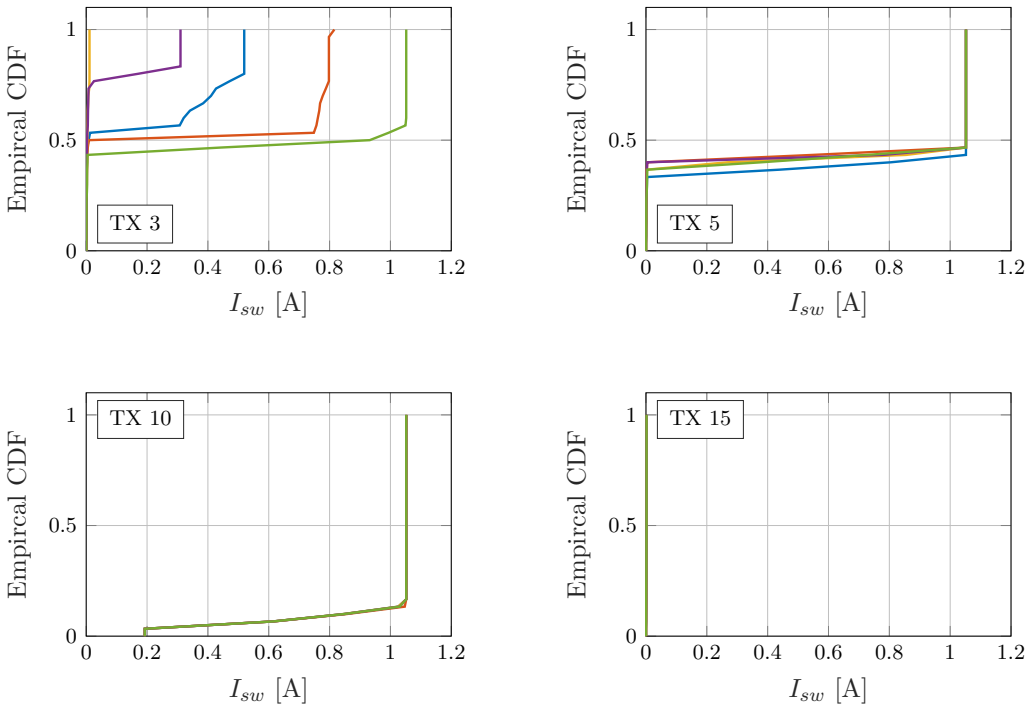

Figure 3.10: Empirical CDF vs. swing levels of RX2.

interference to other RXs, then that TX should be given lower priority in the sequence for power allocation. This means the TX will be assigned later, or not at all.

3.4. Heuristic algorithm

Solving the optimization problem in Eq. (3.5) takes 165 seconds in Matlab, under the setup of 36 TXs and 4 RXs as presented in Sec. 3.3. This is not acceptable for a practical communication system. To reduce this complexity and facilitate a practical D-MISO system, in this section, we propose a ranking-based heuristic inspired by the insights gained in Sec. 3.3.

The key in our proposed heuristic is the custom-defined *Signal-to-Jamming Ratio (SJR)*. Its definition is inspired by Insight 3. Let the $N \times M$ matrix $\mathbf{SJR}$ denote the SJRs of all combinations of TXs and RXs:

$$\mathbf{SJR} = \begin{bmatrix} \mathrm{SJR}_{1,1} & \cdots & \mathrm{SJR}_{1,M} \\ \cdots & \mathrm{SJR}_{n,m} & \cdots \\ \mathrm{SJR}_{N,1} & \cdots & \mathrm{SJR}_{N,M} \end{bmatrix}, \tag{3.13}$$

where $\mathrm{SJR}_{i,j}$ of the TX i to the RX j is defined as:

$$\mathrm{SJR}_{i,j} = \frac{H_{i,j}^{\kappa}}{\sum_{j'=1}^{M} H_{i,j'}} \ , \tag{3.14}$$

with κ a parameter that can be used to tune the weight of the desired channel to RX j with respect to the interference generated at other RXs. The higher the κ, the more weight is given to channel of TX i to RX j, and the less attention is paid to the interference generated by TX i to other RXs.

The proposed heuristic is outlined in Algorithm 1. It works as follows: after calculating all the SJRs, i.e. obtaining the **SJR** (line 1 to line 3), the heuristic searches sequentially for the TX having the maximum SJR (line 5). The tuple $< i^*, j^* >$ that correspond to the maximum SJR is recorded and added to the ranking list $RankedTX$ (line 6), which is a $1 \times N$ vector where its kth element $RankedTX_k$ is a tuple $< i, j >$, and $i \in \mathcal{N}, j \in \mathcal{M}$. Afterwards, the entire ith row of **SJR** is removed from the search space (line 7), since the ith TX is already assigned to a RX. This process is repeated until all TXs are added to ranking list $RankedTX$, i.e. when the search space of **SJR** is empty.

Algorithm 1: Proposed ranking-based heuristic

Input : κ; $H_{i,j}$
Output: $RankedTX$: a $1 \times N$ vector
for $i \leftarrow 1$ **to** N **do**
  for $j \leftarrow 1$ **to** M **do**
    $\mathrm{SJR}_{i,j} \leftarrow \frac{H_{i,j}^{\kappa}}{\sum_{j'=1}^{M} H_{i,j'}}$
for $k \leftarrow 1$ **to** N **do**
  $(i^*, j^*) \leftarrow \arg\max_{i,j} \mathbf{SJR}$
  $RankedTX[k] \leftarrow (i^*, j^*)$
  $\mathbf{SJR} \leftarrow \mathbf{SJR} \setminus \mathbf{SJR}_{i^*,:}$

The proper value of the parameter κ depends on the system setup, such as the density of the TXs and the positions of the RXs. To verify the performance of our proposed heuristic, we set κ to different values and present the result of the ranking algorithm in Fig. 3.11. For a lower κ, i.e. $\kappa = 1.0$, the TXs give higher priority to avoid generating interference to other unintended RXs, therefore, resulting a low system throughput when the power budget is low. For a larger value of κ, e.g. $\kappa = 1.5$, the TXs prefer to serve the intended RXs, which generates more interference when the budget is large, also resulting in a lower throughput. $\kappa = 1.2$ and $\kappa = 1.3$ show a good balance between the channel quality and generated interference. Compared to the optimal system throughput gained by solving the optimizing problem, the average throughputs with $\kappa = 1.0$, $\kappa = 1.2$, $\kappa = 1.3$, $\kappa = 1.5$ are decreased by 40.3%, 2.4%, 1.8%, 2.6%, respectively. The histogram

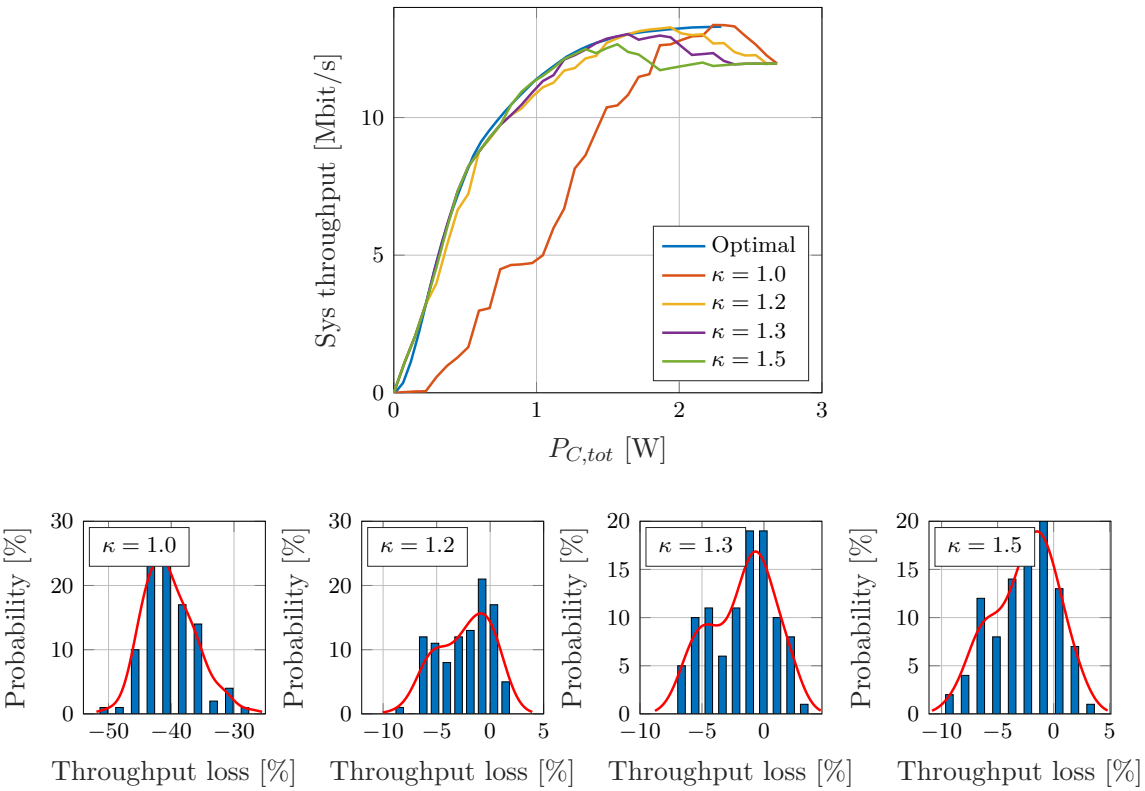

Figure 3.11: Heuristic verification: top) system throughput obtained with heuristic, for the instance shown in Fig. 3.7; bottom) average loss in system throughput, for the 100 random instances shown in Fig. 3.6.

of average loss in system throughput with different κ for the 100 random instances (cf. Fig. 3.6) is further shown in Fig. 3.11. So in this setup, $\kappa = 1.3$ is the best choice.

Note that our proposed ranking-based heuristic only takes 0.07 seconds, i.e. reducing the complexity by 99.96% compared to the optimal solution presented before, at the cost of scarifying only 1.8% of the system throughput. Besides, our heuristic can be easily implemented in a microcontroller.

3.5. Synchronization

For each CFM-MIMO beamspot in DenseVLC, there is a varying number of TXs that are dynamically synchronized to send the same information signal to an RX. Thus, fine-grained, flexible and scalable synchronization among the TXs is essential. In this section, we first state the problems of using existing protocols in DenseVLC, then propose a novel and practical synchronization method based on Non Line-Of-Sight (NLOS) VLC.

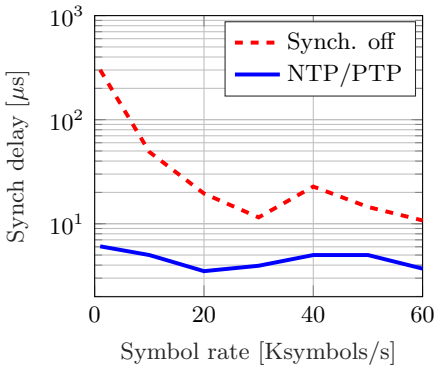

Figure 3.12: Synch. delay.

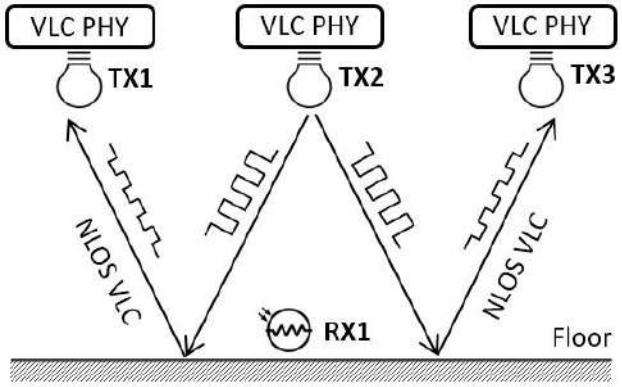

Figure 3.13: Synch. with NLOS VLC.

3.5.1. Synchronization with NTP/PTP

We first test two synchronization protocols together in DenseVLC: Network Time Protocol (NTP) [69] and Precission Time Protocol (PTP) [70]. NTP is used to coarsely synchronize the controller's time to the time of an external server, while PTP is in charge for a finer time synchronization among VLC TXs. In addition, in this test, the frame structure of DenseVLC (cf. Table 3.3) is modified such that each TX will start transmission at the absolute time specified in μs in the Sync field, appended to the frame. Knowing this *absolute time* to transmit, each TX waits until the specified absolute time after which they transmit the data simultaneously.

We implement the above approach and measure the synchronization delay between two received signals transmitted simultaneously by two TXs. The time difference of each two 'synchronized' symbols from the two TXs is computed, and the median delay over an entire frame is derived. This process is repeated 10 times (for 10 different frames) and we average out the 10 median delays as the final synchronization delay. We measure the synchronization delay at several different symbol rates, for which the results are shown in Fig. 3.12. We observe that the NTP/PTP-based synchronization method significantly improves the delay, by at least a factor of two between two VLC TXs. If the maximum acceptable overlap between 'synchronized' symbols is 10% of the symbol width, then the maximum symbol rate at which we can transmit with two synchronized TXs is *14.28 Ksymbols/s.*

Therefore, *this approach has a fundamental limitation in symbol rate.* The TXs cannot be synchronized with a higher accuracy than the one showed above, because it relies on external libraries running on top of an operating system. This approach could work for low transmission rates, but is unreliable for faster communication. Moreover, it requires Central Processing Unit (CPU) time and synchronization messages through Ethernet, which can further degrade the synchronization accuracy when the network workload is higher.

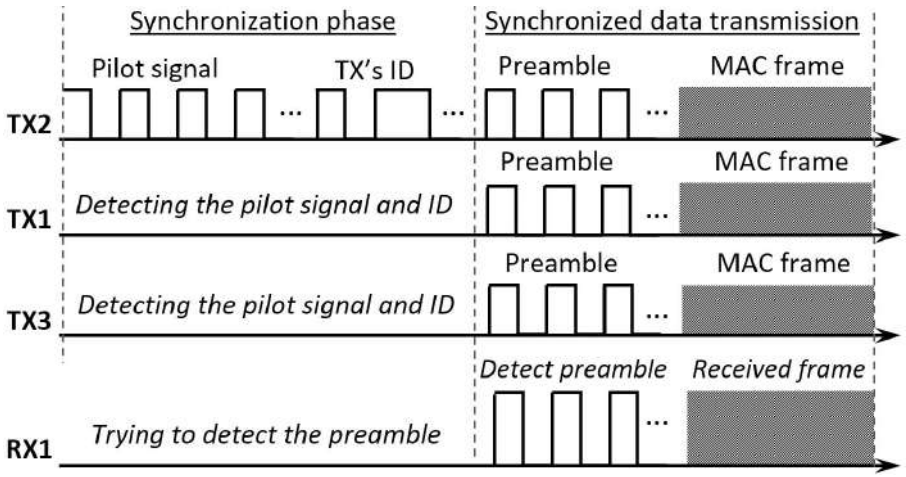

Figure 3.14: The proposed method that exploits NLOS VLC for wireless synchronization.

3.5.2. Synchronization with NLOS VLC

In DenseVLC, we propose a novel synchronization method by exploiting NLOS VLC. For each RX, the controller appoints a leading TX from those that will jointly serve the RX. The leading TX transmits the pilot signal to synchronize the rest of the TXs. Then they concurrently send a synchronized signal to the RX. This method can reduce synchronization delay and does not rely on external libraries running on top of an operating system. It uses relative time instead of absolute time, and therefore, does not need to connect to an external time server. The proposed method works as follows:

- For each RX, the controller sends data to the desired TXs (resulting from our heuristic) and appoints one of them as the leading TX. The leading TX, e.g. TX2 as illustrated in Fig. 3.14, first transmits the pilot signal for synchronization and its own ID, then transmits the real frame (Preamble + MAC frame). The symbol rate at TX2 is denoted as f_{tx}.
- The non-leading TXs, e.g. TX1 and TX3 in Fig. 3.14, listen for the synchronization signal with a sampling rate of f_{rx} that is much higher than f_{tx}. The higher the f_{rx} compared to f_{tx}, the better synchronization we can achieve.
- After detecting the pilot signal and the leading TX's ID, RXs compare the ID to their desired one. If matched, they start sending the preamble and MAC frame synchronously after a short pre-defined guard period.

The proposed synchronization method is scalable in terms of the number of TXs and RXs because the synchronization is performed only by listening to the pilot signal from nearby leading TXs. The synchronization method is also flexible since we relay on wireless visible light communications and do not need to wire each two TXs physically to achieve proper synchronization.

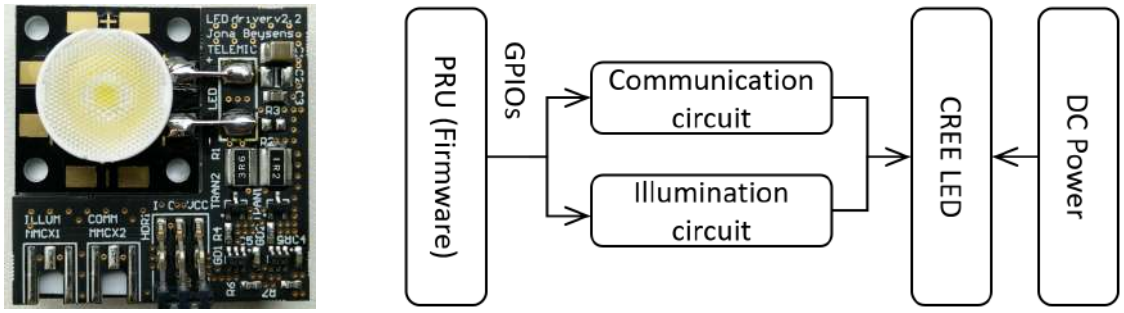

(a) PCB (b) Block diagram

Figure 3.15: Design of the TX front-end.

3.6. System Implementation

In this section we present the implementation of DenseVLC, including both the hardware and software implementations.

3.6.1. Hardware

We use the BeagleBone Black (BBB) as an embedded computer in our design. The controller has a dedicated BBB to orchestrate the CFM-MIMO beamspots and run the VLC MAC. Each TX has a front-end for transmission, and a receiving front-end for synchronization. The VLC Physical Layer (PHY) of four TXs is managed by 1 BBB, so 9 BBBs are used in total for the 36 TXs. For each RX, its front-end is mounted on a dedicated BBB Wireless. Next, we present the two front-end designs.

TX front-end: We use the high performance LED CREE XT-E, covered by a lens to limit the field of view of the LED. We design a circuit to emit light with three different intensity levels (symbol LOW, illumination, symbol HIGH), instead of two levels (light and no light) as done in typical low-end VLC systems [26]. In our design, the LED emits no light when transmitting the symbol LOW. For illumination and symbol HIGH, we use two parallel branches, containing a power transistor and a resistor in series, to drive the LED. The two resistor values are tuned such that the average luminous flux from the LED does not change when going from illumination mode to 50% duty-cycled communication mode and vice versa. The average measured electrical power consumption is 2.51W for illumination and 3.04W for 50% duty-cycled communication. The TX front-end supports a transmission rate of up to 2MHz. The developed Printed Circuit Board (PCB) board and block diagram of the TX front-end are presented in Fig. 3.15.

RX front-end: We design a three-stage analog RX front-end to amplify the signal with minimal noise. In the first stage, the Photodiode (PD) current is amplified to a voltage by a low noise Trans-Impedance-Amplifier (TIA). The second stage contains an Alternating Current (AC) coupled amplifier, which filters out low-frequency ambient light

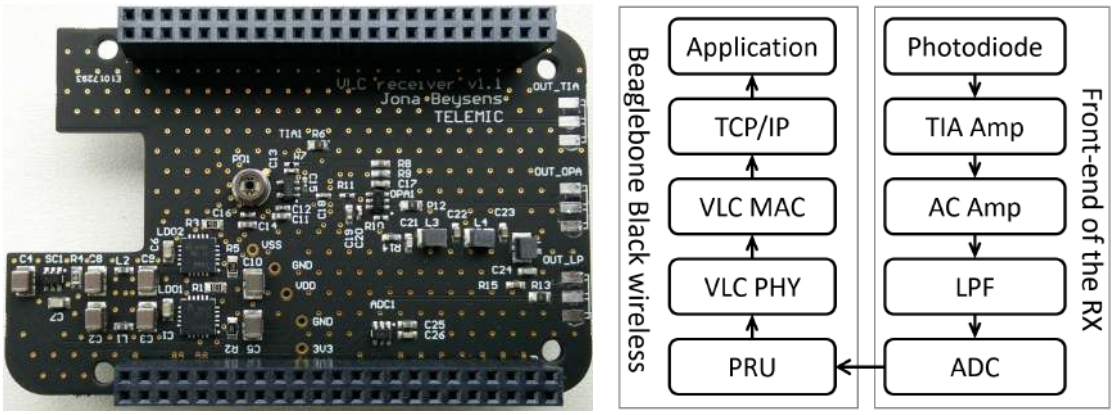

(a) PCB (BBB cape) (b) Block diagram

Figure 3.16: Design of the RX front-end.

Table 3.2: Hardware components.

TX		RX	
LED	CREE XT-E	Photodiode	S5971
Lens	TINA FA10645	TIA / AC Amp	OPA659 / OPA355
Power transistor	NTR4501	ADC	ADS7883

and amplifies the signal further, and thus enables the detection of a very weak signal, e.g., a reflected pilot signal. In the third stage, the high-frequency noise components are removed by a 7th order passive low pass Butterworth filter, to avoid aliasing in the subsequent Analog-to-Digital Converter (ADC). The ADC digitizes the signal at 1 Msamples/s and sends it via Serial Peripheral Interface (SPI) to the Programmable Real-time Unit (PRU) of the BBB Wireless. The developed BBB cape and the block diagram of the RX front-end are presented in Fig. 3.16. Note that this front-end is used for both the receiver on the ground and the transmitter on the ceiling, but with different software; the former to receive data packets, and the latter to receive synchronization pilot signals.

3.6.2. Software

The DenseVLC software is developed partly based on the OpenVLC platform [1]. The source codes of DenseVLC have been made public and can be found on https://github.com/jonabeysens/CoNEXT18_DenseVLC.[4]

Data communication from controller to TXs. After encapsulating the data with an Ethernet header, the frames are sent from the controller via Ethernet multicast to the TXs. A high-level illustration of the frame structure is shown in Table 3.3. This header

[4]This repository is also made available at https://github.com/openvlc/CoNEXT18_DenseVLC.

Table 3.3: Frame structure (Controller to VLC TXs; B: bytes).

ETH PHY + MAC header	TX ID	Pilot sig.	Preamble	SFD	Length	Dst	Src	Protocol	Payload	Reed-Solomon
	8B	32 symbols	32 symbols	1B	2B	2B	2B	2B	xB	$\lceil x/200 \rceil$ 16B

is decapsulated once it reaches the VLC TX. The start of the VLC PHY header contains an 8 bytes field specifying the IDs of the TXs that should transmit the data. Each TX checks this field and acts upon it accordingly. The RX receives the data and if the frame is decoded successfully, it sends a MAC acknowledgement frame back to the controller using WiFi [71]. We have taken advantage of the fact that BBB wireless has Wireless Fidelity (Wi-Fi) connectivity already integrated. As such, the system has full mobility. Note that uplink packets are usually smaller in quantity and size compared to downlink packets. Therefore, the Wi-Fi link is not easily congested.

Channel measurements. The controller selects the optimal operation mode of the VLC TXs based on periodic channel measurements. To quantify the channel quality, each TX sends a predefined bit stream sequentially to the RX. The RX measures received signal strength and sends this information to the controller. To minimize the resulting signaling overhead, this response is fit in a frame with minimal length and send when the channel is idle. To estimate the Signal-to-Noise Ratio (SNR) from the experimental data, the M2M4 estimator is used since (1) it shows good performance for AWGN channels and (2) it can use symbols of the frame after the ADC, without the need for estimating the channel before communication [72]. As the noise can differ during the reception (depending on the circuitry of the RX and the PD's negative bias), the latter property is useful to characterize the SNR better.

Decision logic. Once the controller has received the channel updates for all the links from the TXs to the RXs, it selects $I_{sw} = \{0, I_{sw,max}\}$ for each VLC TX based on our ranking algorithm presented in Sec. 3.4.

3.7. Performance evaluation

In this section we evaluate the performance of DenseVLC.

Experimental setup. We evaluate the performance of DenseVLC in a system of 36 VLC TXs and 4 RXs. The TXs are deployed 6×6 within an area of 3m $\times$ 3m, with 0.5 m inter-TX distance, and a height of 2 m from the floor. The 4 RXs are placed on the floor, controlled by 4 OpenBuilds ACRO System [73] and can be moved to any position within the 3m $\times$ 3m area. The experimental setup is depicted in Fig. 3.17.

Illuminance distribution. DenseVLC provides an average illumination of 530 lux and a uniformity of 81%. The measurements were performed with the HS1010 lux meter.

Figure 3.17: Experimental setup.

Table 3.4: Evaluation of the proposed synchronization.

	No Synchronization	NTP/PTP	NLOS VLC
Median error	10.040 μs	4.565 μs	0.575 μs

3.7.1. Synchronization evaluation

The synchronization between TXs is one of the key enablers in DenseVLC. To evaluate the proposed method that exploits NLOS VLC for synchronization, we first randomly choose two neighboring TXs, RX2 and RX3. RX2 is appointed as the leading TX to send the pilot signal for synchronization. The symbol rate f_{tx} at RX2 and the sampling rate f_{rx} at RX3 are set to 100 Ksymbols/s and 1 Msamples/s, respectively. We connect the anodes of the LEDs at RX2 and RX3 to an oscilloscope (RIGOL MSO1104) to capture the transmitted signals. The delay between the corresponding symbol edges of the two signals are measured. As in Sec. 3.5.1, we calculate the median of the synchronization delay and compare it to those from the method based on NTP/PTP and to the one without synchronization. The results are shown in Table 3.4. We can see that the median synchronization delay of our method with NLOS VLC is only 0.575 μs, improving the synchronization granularity by nearly *an order of magnitude* compared to the one using NTP/PTP. Note that with advanced devices supporting a higher sampling rate of f_{rx}, the synchronization granularity supported by our NLOS VLC based method can be further improved.

To test the synchronization performance, we perform iperf measurements for 100 seconds under three different scenarios. For all the scenarios, there is one RX, located

Table 3.5: Experimental result using iperf.

Scenario	Throughput [Kbit/s]	PER [%]
2 TXs	33.9	0.19
4 TXs (no sync)	0	100
4 TXs (with our sync)	33.8	0.55

in the center of RX2, RX3, RX8 and RX9. The results are shown in Table 3.5. In the first scenario, only RX2 and RX8 serve the RX. Since RX2 and RX8 are managed by the same BBB, no synchronization is required. The Packet Error Rate (PER) is low, due to the strong signal strength and low noise at the RX. The achieved throughput is lower than the used symbol rate of 100 Ksymbols/s, due to Manchester encoding, PHY and MAC layer overhead and Reed-Solomon error correcting. In the second scenario, RX3 and RX9, managed by another BBB, also serve the RX. However, no synchronization is enabled. No packets are received, due to improper alignment of the frames in time. In the last scenario, synchronization is added. Very low packet loss is observed again, showing that our NLOS VLC based synchronization works.

3.7.2. Heuristic evaluation

To evaluate our proposed ranking algorithm in DenseVLC, we carry out experiments under three representing scenarios, for which the RX positions are listed in Table 3.6:

- *Scenario 1:* interference-free; no dominating TX.
- *Scenario 2:* with interference; no dominating TX.
- *Scenario 3:* with interference; with dominating TX.

First, we perform experimental channel measurements from the 36 TXs to the 4 RXs. Afterward, the path loss is computed as the received swing level at the RX and reported to the controller. Using the path loss data, the controller runs the ranking-based heuristic as presented in Algorithm 1 for different values of κ. We assign the TXs from the ranked list one by one (i.e. increasing the allowed power budget on communication step by step) to the corresponding RXs, and calculate the SINR based on Eq. (3.12) with the experimental data. The system throughput is obtained based on Eq. 3.5.

Scenario 1: In this scenario, the inter-RX distance is 2 m. The result of the ranking algorithm is shown in Fig. 3.18. When we gradually increase the $P_{C,tot}$ by assigning more and more TXs, we observe that assigning a TX to one RX results in no throughput drop to the other RXs, which means the interference is limited. The system throughput for

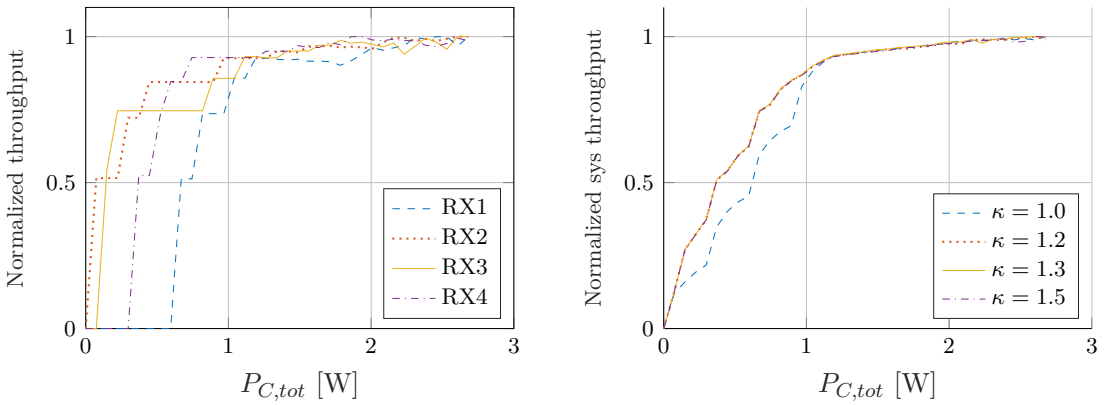

Figure 3.18: Experimental results for Scenario 1.

the different κ values is similar, although $\kappa = 1.0$ performs a little bit worse because with a smaller κ, a little bit of interference could prevent using more TXs for a RX.

Scenario 2: The RX positions are now identical to the simulation setup presented in Fig. 3.7. We observe that the experimental results show a similar trend for the system throughput as the simulations, presented in Fig. 3.11. As in the simulations, $\kappa = 1.3$ also shows good performance here. Further, we can observe from Fig. 3.19 that in the beginning the throughput of all the RXs increases at the same rate, but afterward RX1 achieves a lower throughput than the other RXs. This is because RX1 is positioned closer to the TXs that generate most interference compared to the other RXs. Next, assigning a TX to one RX has more influence on the throughput of the other RXs than in Scenario 1, especially at higher $P_{C,tot}$. Finally, $\kappa = 1.0$ pays too much attention to interference at low $P_{C,tot}$, resulting in a lower system throughput.

Scenario 3: The inter-RX distance in this scenario is 1 m, and every RX is located exactly under 1 TX. From Fig. 3.20, we observe similar results as in Scenario 2. However, RX1 achieves now comparable throughput to the other RXs. Further, due to slightly more interference as in Scenario 2, the system throughput drops when assigning many TXs.

Table 3.6: RX positions in the experiments (in meter).

Scenario	RX1	RX2	RX3	RX4
1	(0.50,0.50)	(2.50,0.50)	(0.50,2.50)	(2.50,2.50)
2	(0.92,0.92)	(1.65,0.65)	(0.72,1.93)	(1.99,1.69)
3	(0.75,0.75)	(1.75,0.75)	(0.75,1.75)	(1.75,1.75)

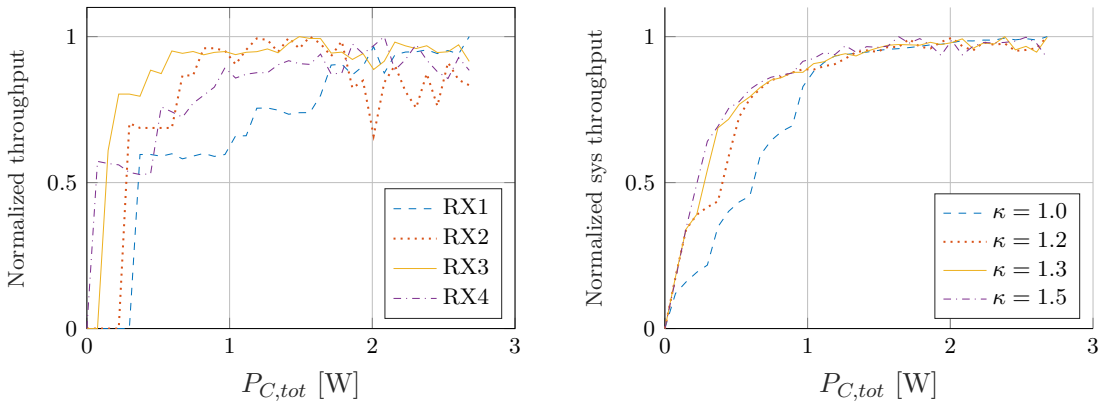

Figure 3.19: Experimental results for Scenario 2

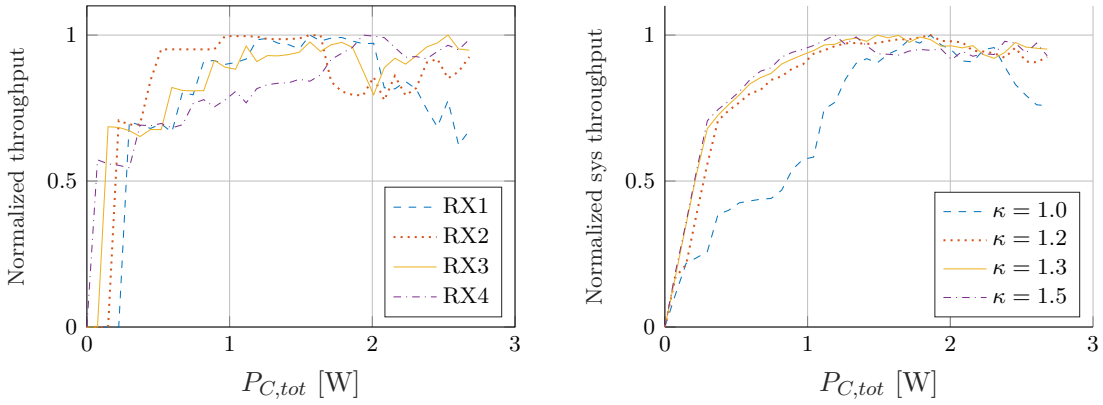

Figure 3.20: Experimental results for Scenario 3.

3.7.3. Power efficiency

In order to benchmark the performance of our method, we compare DenseVLC with following two techniques:

- *Nearest-TX communicating* (Single-Input Single-Output (SISO)): only the nearest TX to the RX is used for communication, resulting in 4 assigned TXs in total. The others are used solely for illumination.
- *All-TXs communicating* (D-MISO): all TXs are used for communication, independent of the position of the receivers. For this setup, this means that each RX is assigned 9 surrounding TXs.

Fig. 3.21 depicts the system throughput in Scenario 2 for DenseVLC with $\kappa = 1.3$, SISO and D-MISO. The markers denote the operating point of SISO and D-MISO, and a horizontal line is added to facilitate the comparison. We observe that the operating point of SISO crosses with DenseVLC, meaning that it achieves the same power efficiency, i.e. throughput versus power. The power consumption is 298mW, for a normalized throughput

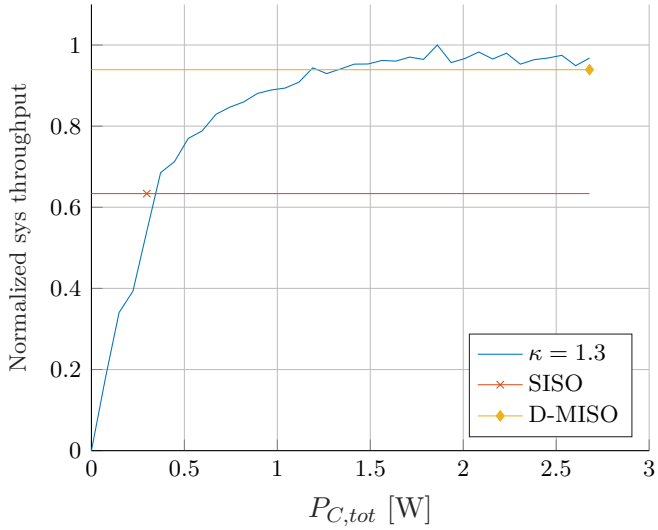

Figure 3.21: DenseVLC vs. SISO and D-MISO.

of 0.63. However, as opposed to SISO, DenseVLC can provide more throughput by adding extra TXs. By adding more TXs, the gap between DenseVLC and D-MISO reduces. At $P_{C,tot} = 1.19W$, DenseVLC achieves the same normalized throughput of 0.94 as D-MISO. However, D-MISO would require power consumption of 2.68W to achieve the same result. Therefore, DenseVLC can improve power efficiency by 2.3 times. The gain in throughput with respect to SISO for this operating point equals 45%. Due to space limitations, we only show the result for Scenario 3, but the conclusion is also valid for the other scenarios.

3.8. Limitations and Discussions

We discuss the current limitations of DenseVLC in this section.

Personalized and adaptive κ. In the heuristic presented in Sec. 3.4, we use the same κ for all the TXs when calculating the signal-to-jamming ratio. In a real cell-free massive MIMO system, TXs will cause different interference to unintended RXs, depending on RXs' positions and the surroundings. Therefore, accurately personalized and adaptive κs can boost the system performance towards the optimal result obtained from solving the optimizing problem in Eq. (3.7).

Advanced hardware. One of the limitations in DenseVLC is the low-cost PRU/BBB used for sampling. With advanced dedicated hardware such as Field-Programmable Gate Array (FPGA), we can perform the sampling much faster and as such achieve better synchronization granularity, as well as to exploit advanced modulation schemes such as Orthogonal Frequency Division Multiplexing (OFDM) in VLC.

Blockage. In traditional VLC, blockage can degrade the system performance greatly because it can break visible light communication links. In cell-free massive MIMO systems

with VLC, however, blockage could bring benefit to the system since it can reduce the interference from other TXs. We will investigate this in DenseVLC in the near future.

TX and RX density. Ideally, a user is served by a transmitter which is as close as possible, because then the received desired signal is the highest, and the interference the lowest. The lower the TX density, the less degrees of freedom we have to serve the users. This results in both a lower system throughput and user fairness. Next to this, also the impact of the number of users on the system performance. We will evaluate the impact of TX and RX density in the future work.

RX orientation. Although receiver orientation has an impact on the receiver's SINR, this is not studied here since it is not the focus of this work. Both the optimization problem and the heuristic are not limited to facing up receivers, and work for all receiver orientation.

NLOS synchronization. The proposed NLOS synchronization approach relies on the reflected pilot signals from the floor. From the preliminary study on the performance for different floor materials, we observe that the pilot signal can also be detected with less reflective floor materials. Further, the evaluation is done without human motion. However, we notice that even when a person is walking by, the pilot signals are still received and as such this does not harm the synchronization significantly. An extensive evaluation has been planned for future work.

3.9. Related work

We summarize the most relevant works in this section.

MIMO/(MU-)MISO in VLC. There is many research done on general MIMO schemes in VLC networks. Here, we only cover those that are most related to our work. In [74], the authors consider a MIMO system where each TX has multiple independent LEDs. They design an algorithm that can select the best subset of the LEDs of each TX based on the channel qualities, to optimize the system capacity, while satisfying the required uniform illumination. In [75], the authors design precoders to improve SNR performance of a non-distributed MIMO system with constraint on the LED-dependent swing. They show that the larger the swing, the lower the achieved Bit Error Rate (BER). The authors in [76] target at maximizing the SNR with the constraint on power consumption. The problem is formulated as a non-linear optimization and solved numerically. However, it cannot provide stable illumination. In [58], they introduce a cell-free VLC system that can achieve proportional fairness with three transmitters. Compared to the above work, we study a scenario with massive VLC TXs and multiple RXs, allowing us to build a real CFM-MIMO system. We design the system architecture, optimize throughput, design and build a testbed to evaluate the performance. We also propose a novel method to synchronize multiple distributed TXs with NLOS VLC. None

of the above work considers practical synchronization, and none of them has a complete system design and experimental evaluation.

Power-efficient VLC. VLC consumes extra power beyond illumination. The authors in [65] show that this extra power consumption depends on the LED, the illumination pattern, and the driver circuit. To have a power-efficient VLC system, these aspects can be exploited. For example, [77] exploits the illumination patterns to achieve an energy efficient system while the authors in [78] design an energy-efficient LED driver for VLC. In this work, our analysis on the extra power consumption for communication is inspired by [65] and our analytical results and findings guide us to design the system components of DenseVLC and build a real system.

3.10. Conclusion

Motivated by the dense-luminaries infrastructure and the concept of Cell-Free Massive MIMO, we proposed DenseVLC, a novel massive VLC system enabled by densely distributed LEDs and visible light communications. DenseVLC adapts the beamspots based on system dynamics to serve multiple users simultaneously. We designed a heuristic to reduce the complexity on adaptation and proposed to synchronize distributed TXs with NLOS VLC. We implemented DenseVLC with off-the-shelf devices and evaluated its performance in extensive experiments. Our results demonstrated the feasibility of building a practical CFM-MIMO system with VLC. Going forward, we envision that DenseVLC can inspire follow-up system research towards this promising direction.

"Battery life is currently one of the biggest hurdles for IoT"
Vanja Plicanic Samuelsson, founder and CEO of Qoitech

4 Hybrid VLC - Backscattering Battery-Free communication platform

Internet-of-Things (IoT) devices are entering into our life. However, all forecasts in terms of number of sold IoT devices are consistently below the expectations [79]. In 2012, IBM predicted a total of 1 trillion IoT devices sold by 2015. In 2017, Cisco predicted 50 billion devices by 2020. In 2018, GSMA predicted 25 billion devices by 2025. As a reference, in 2018 there have been 7 billion connected IoT devices [80]. Most likely, the *need for batteries* is playing a tremendous role in this deceleration. Scalability issues caused by the replacement of batteries, the limited sensor functionality caused by the physical dimension of batteries, and the environmental harm of battery disposal [81, 82] make IoT systems heavily dependent on batteries to operate.

While battery-free devices that operate on harvested energy have long been touted as a panacea [83–86] fundamental roadblocks impede this vision. In particular, the lack of sufficiently powerful energy harvesting sources forces the devices to work under extreme energy constraints. In fact, battery-free devices primarily rely on Radio Frequency (RF)-based energy sources, but the RF energy harvesting works at a much shorter range than communication. The limited amount of harvested energy means that battery-free devices operate under extreme energy constraints and need to sacrifice functionality and efficiency to achieve continuous operation. As shown in this work, this in turn leads to low wireless performance across several dimensions including low spectral efficiency, uni-directional communication, and low receiver sensitivity. For instance, existing work primarily uses backscatter for communication from the battery-free device, but they are inefficient when downlink communication is needed. We propose a different approach, where we leverage the complementary properties of modulated light and radio for passive communication. An example of our battery-free device is shown in Fig. 4.1.

Light fixtures densification. Dense lighting networks are emerging due to technology push driven by the need for providing uniform illumination for user comfort. In our view, this also offers an opportunity to design dense networks that harvest energy for battery-free devices with orders of magnitude better performance. There will be no

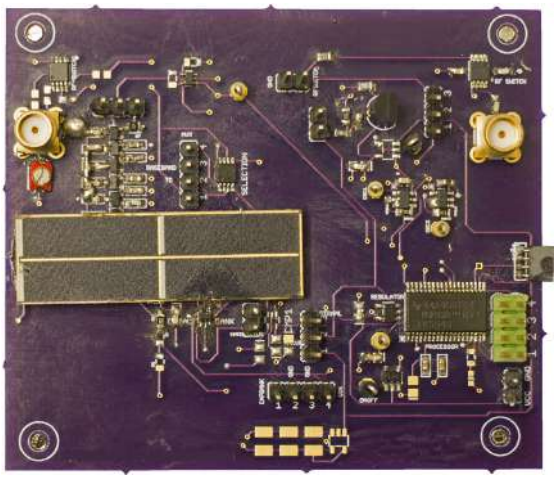

Figure 4.1: *Battery-free LiFi and RF-Backscatter tag.* The tag operates on the energy harvested from ambient light or RF carrier signal. It supports uplink using RF-backscatter, and downlink using a small solar cell based LiFi receiver.

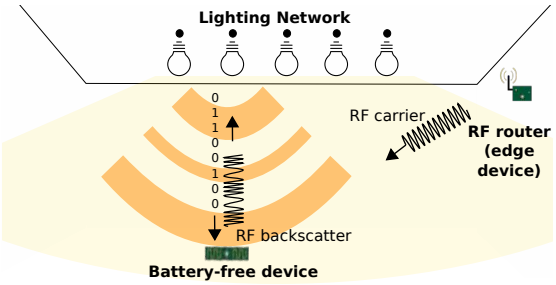

Figure 4.2: High level view of the system.

need anymore to stay close to the RF source to harvest enough energy. Light bulbs are becoming also modulated to transmit data through Visible Light Communication (VLC) and networking solutions are commonly referred as Light Fidelity (LiFi) systems. This opens the opportunity to communicate and harvest through light. But visible light is not always better than radios – light suffers in non-line-of-sight and uplink. We find that LiFi and RF backscatter are individually incomplete (see Section 4.1).

Vision for the deployment. Light Emitting Diode (LED) bulbs will be uniformly deployed for uniform illumination. In the same environment, we expect that the RF router that provides Internet connection will operate as an edge device for passive communication (cf. Fig. 4.2). LED bulbs and RF router will be retrofitted with realistic changes, yet vast potential to enable our passive communication system. Going beyond the transceiver circuitry required by LiFi technology, LED bulbs will require an RF receiver operating in the frequency of interest. Intelligent light bulbs from major vendors already provide similar functionalities, for instance to control the hue and brightness of the bulb itself. Our bulbs transmit LiFi data to the battery-free device and receive their RF data. RF carrier generator for backscatter is outsourced to the RF router. This modality of operation is often already possible in current Wireless Fidelity (Wi-Fi) chipsets.

Main contribution. We present the design and quantitative assessment of EDISON, a bi-directional and passive communication system that operates under overall consumption much smaller than of conventional RF designs. EDISON integrates LiFi and RF for passive communication. The key technical contributions of this Chapter are the following:

- This is the first work that demonstrates that LiFi downlink and RF backscatter uplink solve several bottlenecks of these technologies when taken in isolation;
- Our tag is the first one integrating capabilities to both harvest and receive with solar cells in passive communication systems;
- We propose a low-power processing mechanism of symbols with Microcontroller (MCU) triggered only by wake-up mechanisms and low-power timers;
- We make several other contributions, such as investigating the usage of the LiFi transmitter as an oscillator for RF backscatter.

Our experiments demonstrate that our system even allows continuous reception of LiFi data without intermittent behavior typical to battery-free systems up to 500 symbols/s in a limited illumination environment.

4.1. Challenges

We discuss in this section the challenges in using visible light and backscatter as a communication mechanism on battery-free sensor platforms.

4.1.1. LiFi Bottlenecks

The fundamental problem of using LiFi systems comes from the high power consumption for receptions and transmissions which limits the use on battery-free platforms. This is caused by the following bottlenecks:

- **B1**: LiFi front-ends employed in state-of-the-art systems highly depend on power-hungry amplifiers to amplify the weak signals received by Photodiode (PD)s. New low-power designs still keep some of the old problems (and add new ones).
- **B2**: LiFi uses intensity modulation for transmission. But sampling of light intensity for reception is energy expensive. Current low-power solutions remove sampling mechanisms but still consume a large amount of energy resources.
- **B3**: uplink transmission with LED or retroreflected light is power-hungry and/or cause disturbance to user comfort.

We discuss these issues in detail in what follows. In order to address **B1** and **B2**, we need a new design that is more energy efficient in the whole chain of LiFi data reception. For **B3**, we propose to explore the usage of RF backscatter as an alternative solution (cf. Section 4.2).

LiFi Front-end. LiFi front-ends employed in state-of-the-art systems are energy expensive consuming few-to-tens of milliwatts of power for their operation [87, 88]. A

high power consumption restricts the operation of such receivers on the energy harvested from the ambient environment, which is commonly several orders of magnitude lower. The reason for the high power consumption is the use of PDs for their operation. PDs have a good responsivity, but are fundamentally characterized by having a small reception area. As such, they always need energy-expensive Trans-Impedance-Amplifier (TIA)s to amplify the weak signals received.

Recent works such as RetroVLC [89] and PassiveVLC [36] design LiFi receivers that increase the circuitry complexity in analog domain to reduce the power hungry processing in digital domain. However, they still rely on PDs and energy expensive TIAs in the analog domain. In addition, it does not allow to have flexibility for future changes in the system.

Data acquisition. Most existing LiFi systems employ a direct detection method to receive and process LiFi data [90]. It consists of sampling the light intensity periodically using an Analog-to-Digital Converter (ADC) and interpreting the received bits to recover the data. Periodic sampling with the ADC requires continuous operation of energy-expensive timers and frequent waking up of the MCU. Current low-power solutions for communication [36, 89] and sensing [91, 92] remove sampling mechanisms and introduce a comparator. Visible light sensing mechanisms operate at a much lower frequency rate, with limited impact on the harvested energy. And designs for data communication are dependent on access to timing information which depletes the scarce harvested energy.

Uplink transmissions. Generating light for uplink transmissions is energy expensive. A VLC transmitter typically relies on LEDs for transmission. However, even ultra low-power LEDs consume energy in the order of 50-80 mW for transmission. Similarly, a conventional radio transceiver is extremely energy expensive to operate within the constraints of the harvested energy.

RetroVLC and PassiveVLC demonstrate uplink by reflecting light. Their tags backscatter the light they receive. As impinging light from LEDs in the ceiling is in the visible spectrum, reflected light covers the visible spectrum as well. This brings practical issues: the illumination of these tags can be distracting in a living environment, they require larger tag size due to the need of Liquid Crystal Display (LCD) screens for modulating the light, and increase their uplink range at the cost of high directionality and reduced communication reliability.

4.1.2. Backscatter Downlink Bottleneck

Ultra-low power tags use simple receivers such as passive envelope detectors as opposed to In-phase & Quadrature (I/Q) receivers [93]. This causes the following major bottlenecks:

- **B4**: passive RF receivers are affected by false detections, as any ambient traffic within the frequency band would trigger simple envelope detectors.

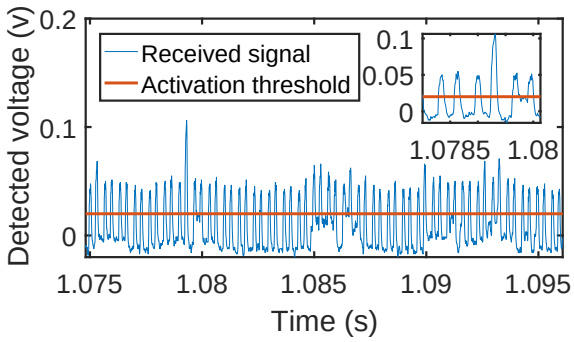

Figure 4.3: Example of the output of our envelope detector measuring at 868 MHz.

- **B5**: the sensitivity of a passive receiver is very low, in the order of -50 dBm [94, 95]. This also limits the amount of RF energy that can be harvested.

We discuss **B4** and **B5** in details in this section to motivate our choice for downlink LiFi (cf. Section 4.2).

Ambient traffic. In recent few years wireless sensors and other wireless applications have seen significant growth. This has increased the Cross-Technology Interference (CTI) significantly. While in traditional active communication systems, CTI mainly causes packet losses, CTI is particularly challenging in battery-free platforms. A constant CTI level is beneficial for battery-free devices as energy could be harvested for the operation of the tag. But any changes in the CTI level is problematic, as it triggers the envelope detectors commonly employed for receptions on these platforms. The precious harvested energy collected over a significant period is served to process false wakeup events.

We experimentally study this problem and observe that the trend towards large scale deployments of low-power wide-area network standards such as Long Range (LoRa) negatively affects our envelope detectors. To demonstrate this phenomenon, we perform an experiment where we collect energy measurements using a zero consumption envelope detector and a logic analyzer at 5 different locations of our campus and city. We calculate the wakeup events encountered in these experiments. Without powering the tag, we measure the output of the envelope detector at each location for 3 minutes.

An example of the collected results can be seen in Fig. 4.3. The figure at the top right corner is an amplified version of a section of the figure, to better appreciate the results obtained. As the comparator of our design may be triggered with voltages as low as 20 mV, low-power RF receivers can detect communication of other devices. In turn, this may trigger the MCU causing an undesired increase in the consumption of the tag as well as errors in the communication. We quantify this problem by measuring the percentage of channel occupancy time caused by CTI. In Fig. 4.4 we can see that although indoor occupancy may be still relatively low, in an outdoor scenario the occupancy can already reach over 40%.

Density of Carrier Emitter. Battery-free devices largely rely on RF-based energy

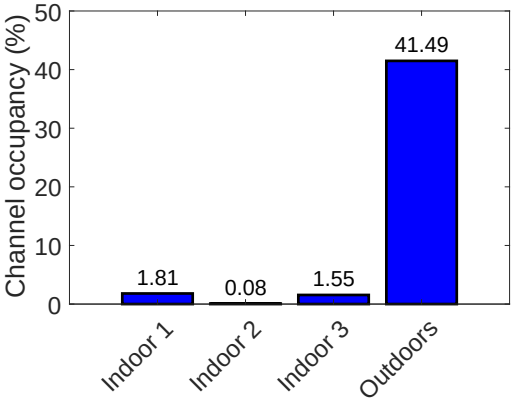

Figure 4.4: Channel occupancy at different locations.

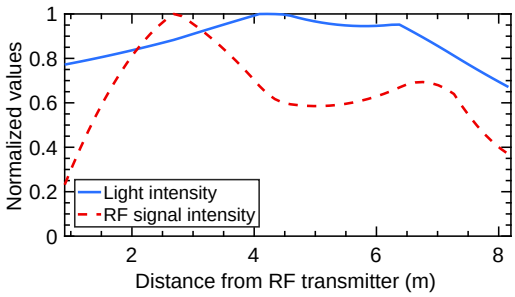

Figure 4.5: Normalized intensity values for RF and light. In the room, we measure light homogeneity factor twelve times higher than RF.

sources, but the RF energy harvesting works at a much shorter range than communication. In fact, tags need to harvest a minimum amount of energy to turn on (*harvesting threshold*). A tag typically needs a few tens of μW to turn on and begin operation [94], hence the range is usually limited by whether such energy can be obtained from the harvesting source. This presents a scalability challenge: we need dense deployments of high-power RF emitters for ubiquitous harvesting. While dense Radio-Frequency IDentification (RFID) deployments have been touted as a solution, technology trends do not suggest this as a viable solution, and facts show that there are only isolated deployments such as industrial IoT (e.g. supply chain warehouses).

We instead expect a typical indoor deployment to have a single device such as a Wi-Fi router that provides the necessary carrier signal to support downlink communication. However, such a scenario limits the operational range to proximity to the carrier emitter device due to the aforementioned RF effects. In order to quantify this problem, we place a carrier generator which is a Universal Software Radio Peripheral (USRP) Ettus B200 at the corner of the room, an ideal location for a Wi-Fi router in a home, and we program it to generate a signal with a maximum strength of 17 dBm. Next, we observe the signal strength of the carrier signal across the room using an RF spectrum analyzer as a receiver together with the light intensity measured with a luxmeter. We measure the uniformity of RF and light strength samples at different 2D positions inside the room. The values of RF and light signal's intensities have been normalized and plotted across the diagonal of the room.

The results are displayed in Fig. 4.5. We observe that close to the RF emitter, contrary to our expectation, the RF signal intensity is weaker, because we are close to the wall and a metallic stand which might have impacted the signal. However, overall, it can clearly be seen that the RF signal has a large peak to average ratios making it less reliable for powering a battery-free device.

4.2. EDISON

We advocate a design that uses LiFi downlink and RF backscatter uplink. Our solution is called EDISON. We design the system to be able to communicate using a hybrid RF-light medium while having a very low-power consumption to support the operation on the harvested energy.

In Section 4.2.1, we first show that the LiFi downlink offers the opportunity to solve the bottlenecks in backscatter downlink, and then explore our main design choices to address the LiFi bottlenecks. The main components of our system are then presented in Section 4.2.2, and the details of the main components are then presented in Section 4.2.3 (tag).

4.2.1. Design principles

Light infrastructure is getting denser at a higher pace than RF infrastructure, and RF alone is not deployed so pervasively. This has consequences on the intensity level measured with RF and light. To illustrate this point, in the same setting as for the RF experiments conducted in Fig. 4.5, we also measure the light levels using a sensor tag light meter. The room uses standard light fixtures for illumination. As for RF, we collect light strength at different locations in the room and plot the normalized values across the diagonal of the room. The results are displayed in Fig. 4.5. We measure a light homogeneity factor twelve times higher than RF (0.502 versus 0.040). Our experiments have been performed in a relatively old office environment. We expect that new rooms will have a factor of illuminance uniformity above 0.6 [96]. In conclusion, dense deployments of LED bulbs allow us to address bottleneck **B5** and provide denser infrastructural elements for energy harvesting.

Downlink LiFi addresses bottleneck **B4**. In indoor environments with uniform lighting, any changes in ambient light caused by human activity will cause infrequent changes in a small area avoiding to trigger the low-power receiver. Further, human activity such as walking, which can disturb the ambient light environment, occurs at a much slower rate when compared to traffic from the ambient signals. In addition, because of the directionality nature of LiFi downlink, LiFi receivers are also significantly less sensitive to changes caused by other LiFi communication in the area.

In order to solve bottleneck **B1**, we propose to leverage a solar cell both to harvest energy and receive LiFi data at very low power. RF backscatter provides a solution to **B3**, but it requires careful design of the elements in the system and their interaction. Our approach to addressing **B2** is finally presented in Section 4.3.

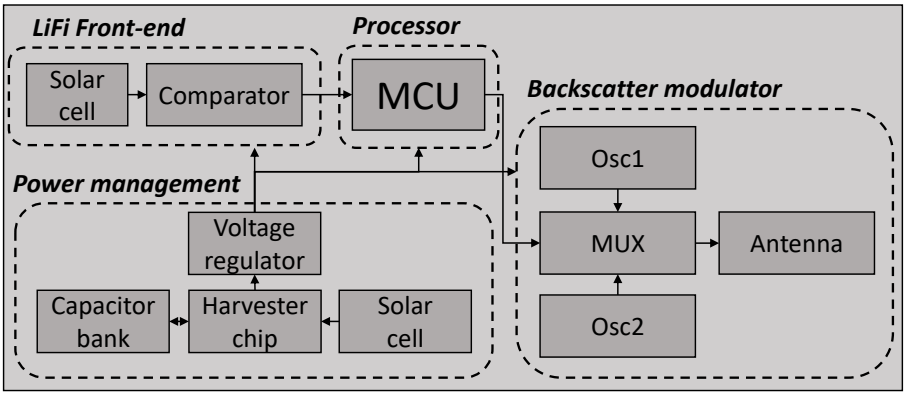

Figure 4.6: Schematic of our tag. The tag consists of an MCU, power management unit, LiFi receiver to support downlink communication, and RF backscatter modulator for uplink transmissions.

4.2.2. Overview

Our system can be broadly divided into three main components.

Tag. The tag is the key component of the system. It harvests energy from the ambient light using a solar cell and stores this energy on a small capacitor using a harvester circuit. The tag uses this energy to power a low-power MCU, the VLC receiver, the RF envelope detector and the backscatter transmitter. While operating on this harvested energy, the tag does baseband processing to receive transmissions using light, and can also transmit messages such as uplink transmissions using RF backscatter mechanism.

LED bulb. LiFi bulbs jointly provide the energy required by the tag to operate and downlink communication data. We also integrate an RF-transceiver in the bulb to support the reception of the backscatter message. We note that most smart LED bulbs already have radio transceivers integrated and we might require minimal changes in the hardware to support the reception of RF backscatter signals.

Edge Device. The final component of our system is an edge device that is responsible for coordinating the downlink transmissions from the LiFi Transmitter. Further, the edge device is also equipped with a carrier generating device to provide the necessary external carrier signal to enable uplink transmissions through the backscatter mechanism.

4.2.3. Tag

In this section, we discuss the tag. A high-level overview of the tag is shown in the Figure 4.6. It consists of four main components - LiFi front-end, backscatter modulator, power management unit, and processor. We present the first three components in this section, while the processor is presented in Section 4.3 together with the design of the protocol.

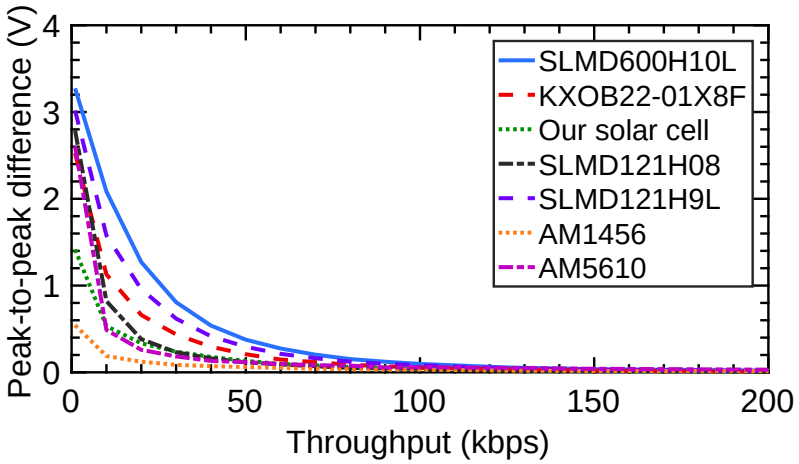

Figure 4.7: Performance of commercially-available solar cells. Distinguishing between the ON and OFF states of a LiFi transmitter becomes difficult at higher bitrates since the peak-to-peak voltage difference becomes too small.

4.2.3.1. LiFi Front-end

Instead of using power-hungry amplifiers and photodiodes (bottleneck **B1**), in our system we have a solar panel operating in photovoltaic mode. This allows having a zero power consumption for the LiFi front-end in our tags. Solar cells have larger reception areas than photodiodes. This is an advantage for harvesting, but reduces the response time compared to photodiodes. However, this does not impact us, as the baseband processing on low-power MCU limits our ability to support high bitrate transmissions.

Choosing solar cells. We evaluate the responsiveness to changes in light levels of several existing solar cells. As a transmitter device, we use a controllable LED connected to a pulse-wave generator that creates an alternating sequence of 1s (LED on) and 0s (LED off). We test seven solar cells with a form factor suitable for wearable applications. Five of these are monocrystalline with different dimensions and parameters, while two are amorphous silicon cells. We connect all solar cells to the ADC of a logic analyzer to find the peak-to-peak difference in the signal amplitude.

Figure 4.7 demonstrates that all seven solar cells have similar patterns. As the sending bitrate increases, the solar cells' ability to distinguish between the two LED states diminishes. We choose the SLMD121H04L [97] (< \$6) cell for our tag. This solar cell allows us to detect transmissions at a frequency as high as 90 kHz. It is also small and generates high short-circuit currents.

Digitalizing solar cell output. The output of the solar cell is an analog signal representing the light conditions. This needs to be digitalized, and on active LiFi receivers, this step is performed using an ADC converter.

We overcome this challenge by building on recent low-power visible light systems that have used a comparator coupled with a low pass filter as thresholding circuit for

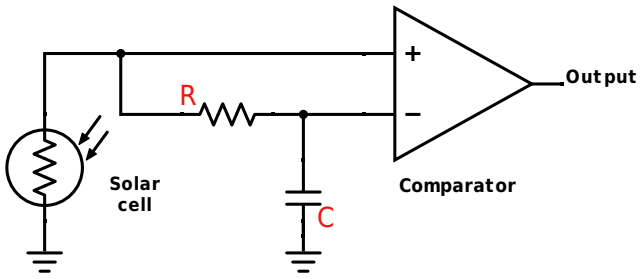

Figure 4.8: Circuit diagram of the low-power receiver designed using a solar cell.

communication [36, 89] and sensing [91, 92]. We build on these designs and demonstrate it can enable downlink LiFi communication even while operating on the harvested energy from a solar cell. The schematic is illustrated in Fig. 4.8.

Figure 4.9 compares the energy efficiency of our receiver to conventional radio transceivers and also to a LiFi receiver implemented using a transimpedance amplifier. We are orders of magnitude more energy-efficient when compared to these designs.

4.2.3.2. Backscatter Modulator

To support up-link transmissions, we use the backscatter mechanism. Backscatter enables transmissions at significantly lower energy costs compared to conventional transceivers by reflecting or absorbing ambient wireless signals [93, 98].

The ability to harvest a small amount of energy with dense LED bulbs deployment can allow us to avoid relying purely on energy from the RF router for harvesting. We exploit this energy budget for performing RF backscatter with frequency shifting for increased RF communication range [99]. Using frequency shifting, the range of such a transmitter can be quite high; for example, recent work has shown that the backscattered signal can be received across several walls in a home [98, 100]. Thus, we do not need dense edge devices deployed as RF carrier generators, and we demonstrate in Sec. 4.4 that only one device per home (e.g., integrated with a Wi-Fi router) can suffice for uplink.

At a high level, the modulator works using two low-power oscillators, that generate frequency representing the two symbols 0 and 1, generating Frequency-Shift Keying (FSK) modulation similarly to [98]. We select these symbols depending on the payload to be transmitted using a low-power multiplexer chip. This signal is used to control the backscatter frontend, which reflects or absorbs the incident carrier signal. We can support operation both in the 868 MHz and the 2.4 GHz band, while consuming a peak power of 70 µW and 650 µW, respectively.

LiFi Transmitter as an Oscillator. In the backscatter modulator, the oscillator is the most energy-expensive component. All the components excluding the oscillator consume

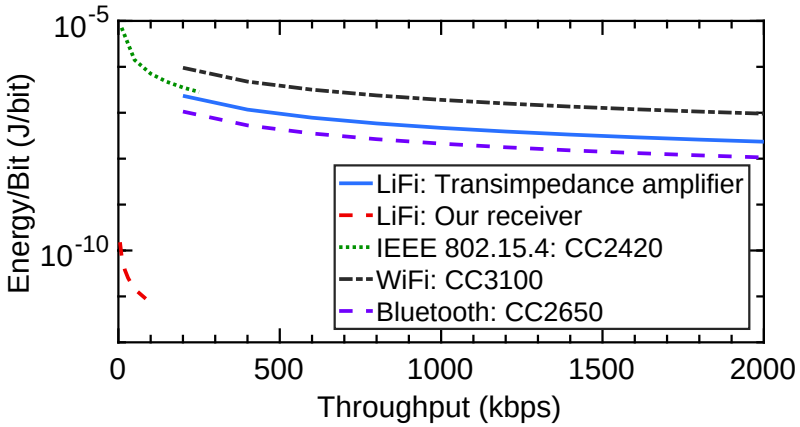

Figure 4.9: Energy consumption/throughput comparison among LiFi receivers and modern RF chips. The ultra-low-power receiver meets its performance goals of large energy-saving while providing a sufficient data rate for passive downlink communication.

sub-µW of power for their operation. If we eliminate the oscillator from the modulator and yet support frequency shift backscatter communication, we can reduce the power consumption very significantly.

4.2.3.3. Power Management

Power management is an essential operation of the tag due to the battery-free nature. Our tag can harvest energy from either ambient light or RF carrier signal, store the energy onto a small capacitor, and provide it to the components on the board for their operation. The fluctuations in voltage due to change in the voltage of the capacitor are challenging for battery-free systems. We overcome this challenge by using a low-power regulator and operate the entire board at a low-voltage of 2.2 V to minimize the power consumption.

4.3. Protocol

Design of protocol to leverage the hybrid medium while operating under the constraints of the harvested energy is one of the vital contributions that we make in this work. We design the protocol to achieve meager power consumption which, as an example, in our experiments, even allows continuous reception of LiFi data without intermittent behavior typical to battery-free systems. We discuss the design of the protocol next.

4.3.1. Energy-efficient LiFi Reception

Most existing LiFi systems use energy-expensive mechanisms to perform the necessary baseband processing to receive LiFi transmissions [38]. This involves sensing the light signal using a PD coupled with a TIA. The light signal is digitized and processed using

ADCs. However, this is a prohibitively expensive operation when operating on harvested energy. This is because it requires continuously polling and processing from the MCU.

Recent low-power LiFi systems overcome this limitation by using an event-based mechanism [36, 89]. These systems react to an external event such as interrupts instead of periodically sampling the channel to detect changes in the light intensity levels. These interrupts are generated using low-power comparators, which, as we had discussed in Section 4.2.3.1, help us to avoid the energy-expensive ADCs. Our protocol also uses event-based design. We keep the energy-expensive MCU to be in a low-power state for a vast majority of the time. The MCU awakens only when there is an external event caused by changes in the light intensity levels. However, we overcome the limitation of existing event-based LiFi receivers that perform significant processing to measure the symbol period. Instead of keeping track of time by involving the MCU to measure clock cycle count, our system simplifies this design by using a low-power timer. This enables us to place the energy-expensive MCU to low-power state in between the processing of the symbol.

At a high level, our LiFi receiver mechanism works as follows: We monitor the output of the comparator in the LiFi receiver. Next, the change in the light level due to LiFi transmission causes the MCU to be woken out of the sleep state. The MCU's low-power timer is compared against the symbol period and interprets the incoming symbol. Finally, once the symbols are received, the MCU constructs the frame and processes the data.

4.3.2. Design and Implementation

To prevent flickering and support uniform illuminance during the downlink communication, we use Manchester coding. This ensures that even if data is transmitted, the LED bulb has a constant Direct Current (DC) value of 50% of the high level of the signal. This encoding has another feature that we exploit: It ensures that irrespective of the content of the bitstream, the maximum consecutive symbols with the same value can be only two.

The reception logic operates in the following way: As the tag switches on, it places the MCU to the low-power state to preserve the scarce harvested energy. Next, it waits for interrupts from the LiFi receiver's comparator to receive ongoing LiFi transmission. Due to the inherent 1-bit nature of the comparator, we are limited to only the information about the transitions between the high and low states. We precisely encode LiFi data within the timing information when the transition happens. This, however, requires us to maintain a notion of time at the tag which we achieve using a 32 kHz timer. This helps us consume a much lower power consumption to maintain time as compared to state-of-the-art works [36, 89] enabling us to sleep between the reception of symbols. A question we may ask is: *What should be the period of the timer event to aid in symbol detection ?*

The duration of the timer has to be such that is between the symbol duration and

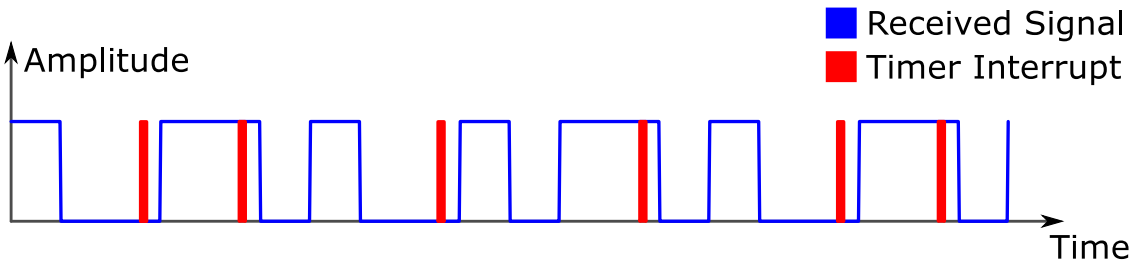

Figure 4.10: Behavior of the timer. The timer resets each time a rising or falling edge is found. If it triggers, it means that there is no edge between two symbols and thus, two consecutive symbols with the same value have been received.

Preamble	Start-Frame-Delimiter (SFD)	Src. address	Dst. Address	Frame Length	Payload
32	16	8	8	32	0-MAX

Table 4.1: Frame structure and size (symbols).

twice the symbol duration. In order to avoid problems such as frequency offset of the oscillator, we set a conservative value of 1.5 times the symbol duration for the timer event. We illustrate the behavior in Figure 4.10. From an implementation point of view, we detect each symbol in the following way. We set the timer once the interrupt due to a rising or falling edge is found. Next, we distinguish if we have received a single symbol or contiguous symbols in the following way. If we receive a timer event before receiving an interrupt from the comparator, we know contiguous symbols have been received, and vice-versa. We end transmission when we receive two continuous timer events without the comparator being triggered. This is because due to Manchester encoding, it is not possible to receive two symbols with similar value. An event with timer firing twice means that three contiguous symbols have been received which cannot be possible unless transmission has concluded or there is an error which results in the case that we start to assemble the LiFi frame.

We have two key advantage when using the aforementioned approach when compared to traditional sampling method: (i) There are hardly any synchronization related issues between transmitter and receiver which is common to sampling-based state-of-the-art LiFi systems, and (ii) Baseband processing is minimized due to timer-based events which helps to conserve scarce energy and maintain low power mode.

Frame structure. We show the frame structure in the Table 4.1. We store the received symbols into a circular memory structure. Our protocol detects a valid frame by comparing the preamble and SFD. To support deployments with a dense number of tags, we also maintain space for Transmitter (TX) and Receiver (RX) ID.

Component	Name
Solar Cell Data	SLMD121H04L [97]
Harvester	Texas Instruments BQ25570 [101]
Solar Cell Harvesting	MP3-37 [102]
Comparator	NCS2200 [103]
MCU	MSP430FR5969 [104]
Oscillator	Linear Technology LTC 6906 [105]
Multiplexer	ADG904 multiplexer [106]
Regulator	S-1313 [107]
VLC Processor	BeagleBone Black [108]
RF receiver	CC1310 [109]
Carrier Generator	B200 USRP [110]

Table 4.2: Summary of the main components used in the implementation.

4.3.3. Backscatter Uplink

Our protocol also takes care of the uplink needed to have a bi-directional system. We achieve this using the backscatter mechanism. We implement the uplink baseband processing using timers similar to LiFi downlink logic. The timer to generate interrupts at the symbol rate of uplink transmission. During these timer events, we select the oscillator that corresponds to the bit to be transmitted.

4.4. Evaluation

The tag has been designed with the components shown in Table 4.2. On the infrastructure's side we use a BeagleBone Black (BBB) and the software and firmware of OpenVLC [2]. We modulate the data to a General Purpose Input/Output (GPIO). Then, we use an external driver to power up the voltage and connect it to a LED bulb of 4.3 W. The BBB is also connected to a Texas Instruments CC1310 that has been configured as an FSK receiver. This allows the bulb to directly control the handshake of communication with the tag. Both are connected through serial.

For the carrier generator, we use a B200 USRP generating a carrier at 868 MHz. In the future, this element could also be designed and integrated with the infrastructure to minimize the cost of it.

Next, we evaluate our system in a range of conditions that applications might encounter.

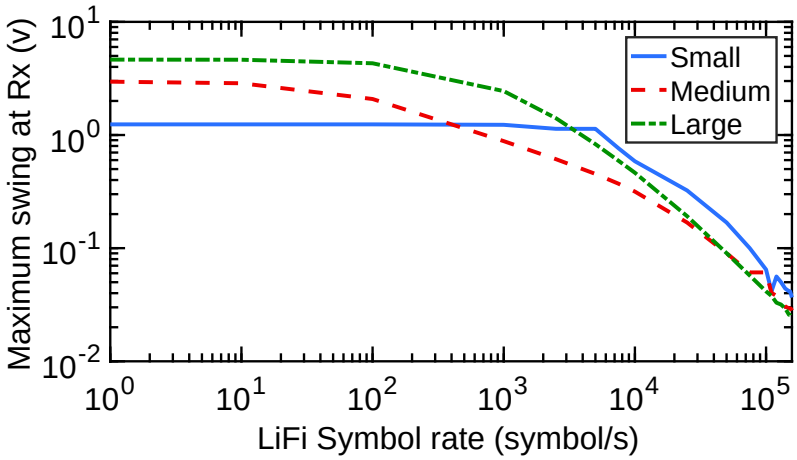

Figure 4.11: Maximum swing of the signal at the output of the solar cell. Capacitance effects reduce the swing at high symbol rates.

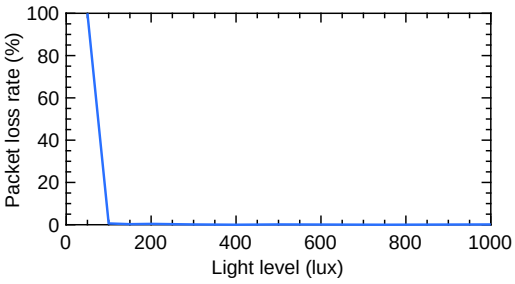

Figure 4.12: Packet loss rate as a function of luminous intensity at the solar cell. If the intensity is too small, the comparator would not trigger.

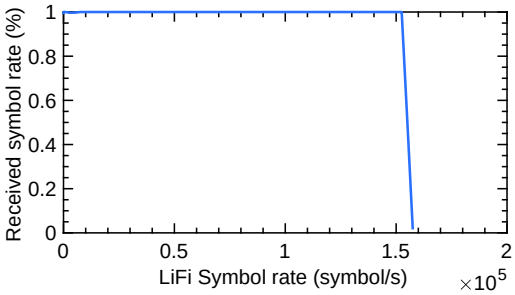

Figure 4.13: Received symbol rate at a function of the LiFi communication frequency. At 15 ksamples/s the swing is so small that it does not trigger the comparator anymore.

4.4.1. Wake up mechanism

Our LiFi front-end has two key components, the solar cell and the comparator. We measure the swing (difference between high and low symbols) obtained at different transmission rates with a constant illumination of 500 Lux (a typical illumination level for indoor environment). The swing is necessary for the comparator to be triggered.

The results are presented in Fig. 4.11. At low frequencies, no matter the size, the solar cell is able to fully transition. This means that bigger areas can collect more energy and thus, show bigger swings. Nevertheless, the solar cells have capacitance effects that increase with the reception area. The capacitance limits its frequency response. In Fig. 4.11, we observe that, at higher frequencies, smaller solar cells show better behavior for communication as their transitions happen faster. For the rest of the experiments with solar cells, we will use the smallest one, as its behavior is the best at high frequencies.

Fig. 4.12 shows the packet loss rate as a function of the light intensity impinging on

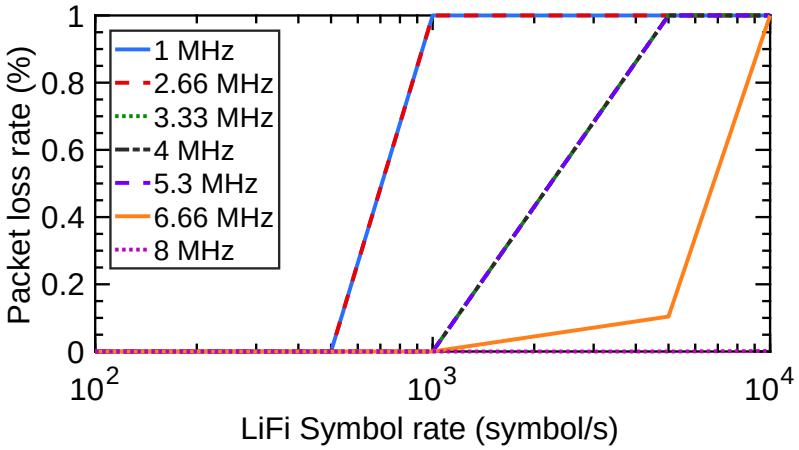

Figure 4.14: Packet loss rate as a function of MCU clock. At some point the processing time becomes greater than the symbol duration, making it impossible to transmit data.

the solar cell. For very low light intensity, there is no enough voltage change to trigger the wake-up mechanism and the communication is affected by a high packet loss. However, this occurs at very low light level (in the order of 100 Lux), much smaller than typical suggested indoor illumination (of about 500 Lux).

Another parameter affecting the wake-up mechanism is the relation between the received symbol rate and the communication frequency, with results presented in Fig. 4.13. At frequencies above 15 ksamples per second, the sensitivity of the comparator is not large enough to sense the difference between high and low symbols received with the solar cell.

4.4.2. LiFi Reception ability

The performance of the tag as a LiFi receiver will depend not only on the previously mentioned wake-up mechanism but also on the capacity of the tag to process the data. The latter depends on the clock frequency of the MCU. Nevertheless, as the MCU is the primary source of consumption on our system, it is desired to have a clock frequency as low as possible, in order to reduce the overall consumption. We evaluate the tag under different transmission rates and different MCU clock frequencies. The frequencies selected are the ones given by the manufacturer. The results are presented in Fig. 4.14. Below 3.33 MHz the maximum rate at which no packet is lost is 500 symbols/s, between 3.33 MHz and 6.66 MHz 1 ksymbols/s, 5ksymbol/s at 6.66 MHz and 10 ksymbols/s at 8 MHz. At each rate, from some point on, the tag is unable to receive anymore. This is the result of the MCU taking more time to process a symbol than the symbol duration.

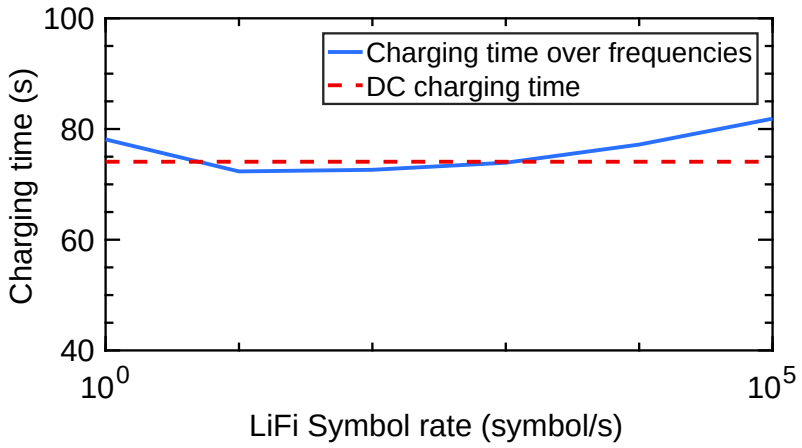

Figure 4.15: Charging time of a capacitor for different transmission frequencies. It shows that the harvester is able to harvest modulated light with an efficiency similar to harvesting a DC signal.

4.4.3. Energy Harvesting

The operation of our tag on energy harvested from the ambient environment is an essential functionality for our system. In this section, we present experiments that we perform to evaluate this ability of our system to operate under varied energy harvesting conditions.

LiFi Transmission Rate. We explore how the LiFi transmission rate impacts the energy harvesting ability. We measure the energy harvesting time as we vary the symbol rate, which impacts the incident light intensity. We maintain an average luminosity flux of 550 Lux. Figure 4.15 shows the result of the experiment. We observe varying the symbol transmission rate does not significantly impact the energy harvesting rate when compared to an unmodulated LiFi transmission. We believe this is because the amount of energy incident on the solar cell changes negligibly with the LiFi transmission rate. The high capacitance effects of the solar cell and the harvester chip further mitigate this problem.

Light Intensity. Next, we evaluate the ability of our system to harvest energy under different light intensity conditions. The incident light levels impact the energy available to harvest impacting the harvesting time. In this experiment, we explore the time taken to charge a capacitor of size 2.2 mF as we change the light intensity levels. Figure 4.16 shows the result of the experiment. We observe that the charging time decreases exponentially with the increase in the light intensity levels. At a light level of approx 3000 Lux we observe that increasing further light level does not change the harvesting time indicating a saturation point. We also note, such high levels are usually not present in the indoor environment and can damage the energy harvester chip.

Capacitor Size. In this experiment, we evaluate the impact of capacitor size on the energy harvesting time. A larger capacitor size can sustain the operation of the tag for a

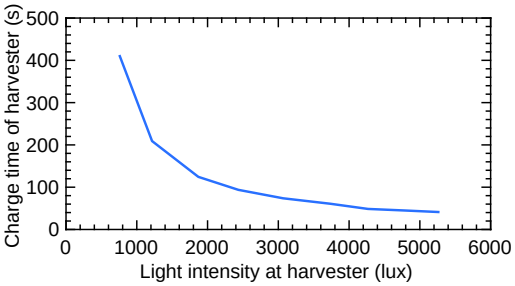

Figure 4.16: Charging time of a capacitor for different light intensities. At low intensities the harvester chip is less efficient.

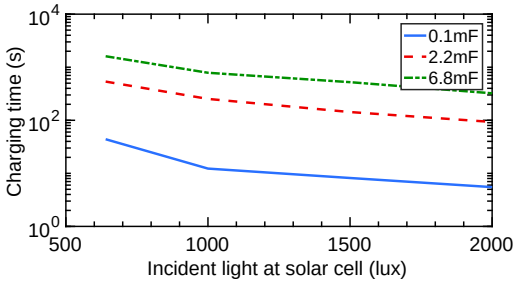

Figure 4.17: Charging time of the harvesting circuit for different capacitors.

longer time duration but also takes a much longer time to charge. We experiment with similar conditions as the last experiment. We evaluate three different capacitor sizes, 100 µF, 2.2 mF and 6.8 mF. Figure 4.17 shows the result of the experiment. We observe that as expected, a large capacitor or brighter light conditions result in faster charging time. We note that, even under normal indoor illumination, we can harvest sufficient energy to sustain the operation of our tag for a few minutes.

LiFi Reception. We evaluate the ability of our system to receive LiFi transmissions under the harvested energy. We experiment with light conditions of 540 Lux, which is similar to levels commonly found in the indoor environment. We configure the LiFi transmitter to send at different symbol rates. We set the MCU at the lowest clock frequency to reduce power consumption. Figure 4.19 shows the result of the experiment. We observe that under such conditions, our system can continuously receive data up to 500 symbols/s without intermittent behavior commonly shown by battery-free systems. This is because of the optimized design of our system. As we increase the LiFi transmission rate, the energy harvested and used becomes asymmetric impacting the reception ability, as the tag starts to lose energy during the reception.

4.4.4. Power Consumption

In this experiment, we measure the power consumption of the tag. Two main contributors influence the power consumption of the tag: the MCU; and the consumption of the other units of the circuit. Its operating state influences the power consumption of MCU itself. We measure the power consumption using a highly sensitive Keysight E36313A power supply [111].

We measure the power consumption of the tag with the MCU in the lowest power state, and find that it consumes a peak power of $224\mu W$, independent from the MCU's clock rate, as it is turned off in low power mode. The high power consumption is due to all the units, such as, backscatter modulator, LiFi receiver being active. In the future version of the

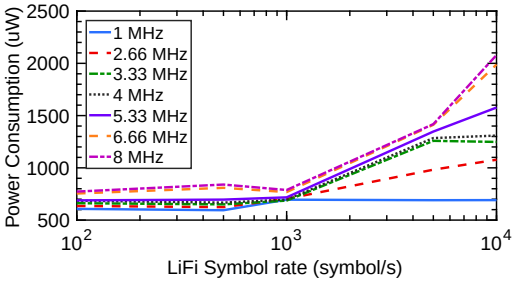

Figure 4.18: Power consumption for different symbol rates and processor's clocks.

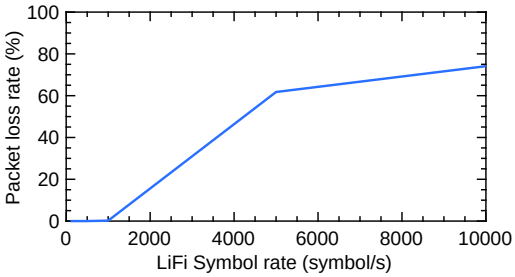

Figure 4.19: Packet loss due to intermittent behavior.

tag, we will aim to turn off these units individually. When receiving and processing LiFi data, the power consumption of the tag scales with the symbol rate and clock frequency. The amount of processing required is fixed per symbol. This means that the higher the transmission rate, the more processing MCU needs to perform, which pushes the power consumption. We illustrate these results in Figure 4.18. As we increase the transmission rate, we need higher MCU frequency to be able to process symbols fast enough before the next symbol is received. Further, a higher transmission rate also prevents MCU from transiting to a low-power state which also pushes the power consumption. We can achieve communication of 500 symbols/second at 620 μW and 10 ksymbols/second at 2 mW.

4.4.5. LiFi as Oscillator

In this experiment, we evaluate the ability of LiFi transmitter to eliminate onboard oscillators and support frequency shift backscatter. This capability can significantly help to lower the power consumption of the tag.

We program the LiFi transmitter to generate an unmodulated carrier signal up to the frequency of 90 kHz, which is the maximum frequency up to which we see a noticeable voltage swing across solar cell output to enable digitization operation (cf. Section 4.2.3.1). Next, we connect the output of the digitization mechanism to the backscatter modulator. Thus, the overall power consumption is dictated by the comparator, which results in approximately 20 µW. We generate a carrier signal with a strength of 17 dBm located several meters away. We locate an RF spectrum analyzer close to our battery-free tag. We keep track of the noise floor and the backscattered signal strength as we change the frequency of unmodulated carrier transmitted from LiFi emitter. Figure 4.20 shows the result of the experiment. We observe that, as we change the clock frequency of the LiFi transmitter, we also observe a high backscattered signal strength which implies we can support frequency shift backscatter without requiring onboard oscillators. The noise floor was measured by min. holding the signal analyzer at 868 MHz.

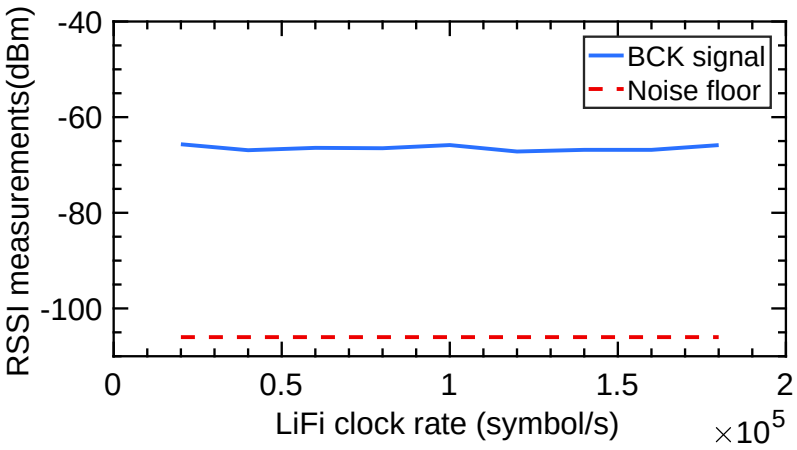

Figure 4.20: Modulating light with a signal of frequency required for frequency shift backscatter transmissions helps us to eliminate onboard oscillators. We couple LiFi receiver directly to backscatter modulator. This allows the ability to support dynamic changes in the frequency of the backscattered signal which can be dictated by the LiFi transmitter.

4.4.6. Communication Range

In this experiment, we evaluate the communication range of our system. Due to the hybrid design, the communication range includes light-based downlink communication, and uplink through backscatter mechanism.

LiFi Downlink. The range of the LiFi downlink system depends mainly on the intensity of the light and the sensitivity of the LiFi receiver. As our system can communicate with light intensity as small as 100 Lux and we use a 4.3 W LED as a transmitter, and our maximum communication distance is 3 meters. The use of a collimator would allow increasing the transmission distance at the cost of reducing the field-of-view. A typical envelope detector has a sensitivity of around -50 dBm, as we had discussed earlier, which means with a source transmitting at a maximum strength of 4 W, we can achieve a similar or larger range when compared to LiFi transmission. However, as we had discussed earlier, RF signals suffer from many unintended effects and also make dense deployment challenging.

Backscatter Uplink. In the case of the backscattering communication, the communication distance depends on the relative distance between the carrier generator and tag, tag and receiver in the LED bulb, and carrier generator's intensity. To measure the capabilities of our system as a backscatter under real deployment conditions, we place the tag in the same room as the receiver and place the carrier generator at different locations. The kitchen, corridor 1, room 1, room 2 and corridor 2 are located at a distance of 6, 4.5, 8.5, 15 and 20 meters, respectively and we collect the Received Signal Strength Indicator (RSSI) measurements of the received packets. The kitchen has direct Line-Of-

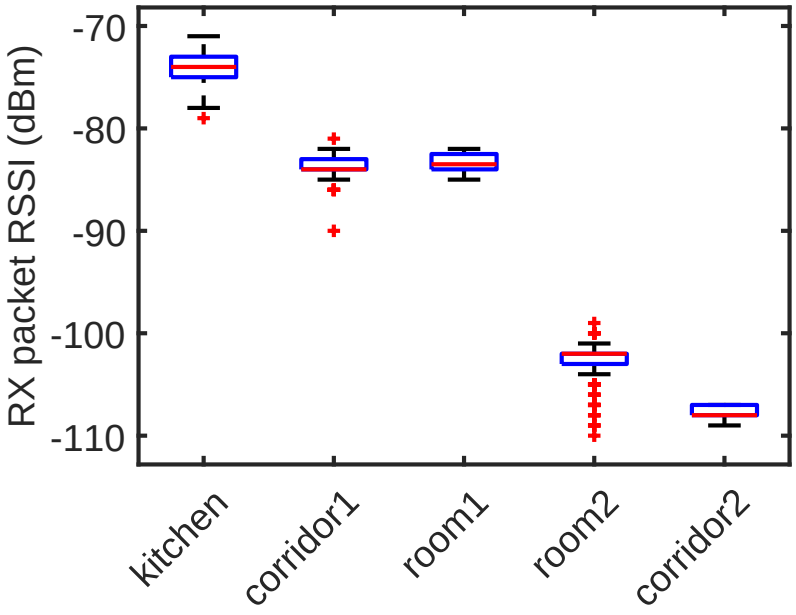

Figure 4.21: RSSI measurements after placing the carrier generator at different locations.

Sight (LOS) with the room where the tag is located. The rest of the measurements are taken in the corridor with no LOS with the tag or in offices in the same corridor.

At each measurement point, we take measurements for 1 minute. The results of this experiment can be seen in Fig. 4.21. In general, the further we are from the tag, the smaller the RSSI is. Nevertheless, here we show that our system can communicate at tens of meters even when objects such as walls are blocking the direct LOS of the system. The sensitivity of the RF receiver in the LED bulb is around -110 dBm.

4.5. Related Work

Passive communication. There is a growing number of applications that have been recently investigated for passive communication, including a battery-free eye-tracker for augmented reality [83], wristband for hand gesture recognition [84], the first phone to make calls without batteries [85], and the first attempt, partially in simulations, for battery-free High Definition (HD) video streaming [86]. However, all these works share the limitations that mobile devices mainly transmit data, while reception of data occurs at very low rate. Current approaches must sacrifice either the data rate or operate at a very close distance from the harvesting source.

Network densification. Networks are getting more densely deployed to increase capacity and reduce latency, rather than for communicating to battery-free devices. The underlying radio technologies, mm-wave and Massive Multiple-Input Multiple-Output (MIMO), are power-hungry and use high frequencies to optimize the spectrum

reuse [112, 113]. There has been work to make battery-free devices talking to current architectures and exploiting existing RF waves [93, 114–116], rather than relying on isolated, costly and low-performance deployments of RFID readers. However, all these works are affected by fundamental limitations, such as (i) they can not scavenge enough RF energy as legacy transmissions are typically wideband, while peak power of energy harvested occurs for narrowband transmissions; (ii) they all work only for uplink communication from the battery-free device.

Low-power visible light communication. Another body of work has explored passive uplink communication for VLC where the receiver replies using an LCD shutter to change the reflected signal in time and sends the signal in the direction of the transmitter (LED) using a retro-reflector [36, 89]. An alternative scheme proposed in [117] uses the mobility and the reflective properties of materials to send data to the receiver. Schemes using uplink VLC reach a low data rate in the order of $\approx$ 1 kb/s, perform poorly in terms of power consumption, and create disturbing illumination artifacts.

Intersection of different technologies. There has been limited investigation on the study of the interface between different technologies. A low-power transceiver called Morpho was presented in [94] to integrate active radio components like oscillators and active amplifiers with passive radio components like backscatter reflectors and envelope detectors. Similar interest has emerged to understand the interaction between radio and acoustic [118]. Yet, these works do not operate in battery-free regime. There has also been limited research to integrate visible light and radio-based communication. [71, 119] did not consider battery-free devices and purely aimed at optimizing the throughput. Related work has investigated how to sense coarse light variations in order to modulate the RF backscatter signal [91]. The proposed solution did not allow to send data to the IoT tag. [120] proposed to retrofit LEDs with an RFID reader to communicate to battery-free device. Their design is cost prohibitive ($\approx$ \$300) and subjected to low sensitivity and high transmit power beyond regulatory compliance. This work is inspired by BackVLC [121]. In their position paper, the authors proposed to integrate Visible Light Communication and RF backscatter. In their approach, they integrated an RF carrier generator in each LED bulb, which would increase the complexity in coordinating RF emitters. They present a feasibility study and show they could transmit one frame per second to the tag through VLC and receive its acknowledgment through RF without packet losses.

4.6. Conclusion

In this Chapter we have proposed a novel LiFi-RF hybrid communication system that addresses some of the fundamental roadblocks that impede the vision of battery-free communicaiton. The system allows to exploit the advantages of LiFi and RF to provide a battery-free communication and provide technical solutions to solve their bottlenecks. We

have presented several contributions to communicate at very low power, and have shown that our system can even allow continuous reception of LiFi data without intermittent behaviour typical to battery-free systems up to 500 symbols/s.

"I do not know where I am going, but I am on my way."
Voltaire (1694 - 1778)

5 Leveraging LEDs for Relative Two-Dimensional Localization

During the last five years, there has been an increasing number of studies exploiting visible light for localization [122]. There are two main reasons for this trend: the pervasive presence of artificial lighting in our environments and the rather deterministic propagation properties of visible light waves, which makes them easy to model and predict [123].

Most of these studies focus on indoor positioning [124, 125]. Ceiling luminaries in our buildings and homes are seen as anchor points [126], and users with photoreceptors estimate their position based on the information obtained from nearby luminaries [127]. The strong focus on indoor positioning has lead many of these methods to make two important assumptions: (i) light sources are *static*, a realistic assumption given the fact that the location and orientation of most light fixtures are fixed in our ceilings; and (ii) users can expect *line-of-sight with many luminaries*, a stronger assumption that may not always be satisfied since it depends on the density of luminaries in the area and the Field-of-View (FoV) of the receiver.

While indoor positioning is justifiably considered the most relevant application of visible light methods, it may be valuable to remove the assumptions mentioned above, and extend these methods to other scenarios. In that manner, we could pave the way to create a new type of general localization methods that could be applied to *any* object (static or mobile) as long as it contains at least one Light Emitting Diode (LED) light. This work is a step in such direction.

Scenario of interest. Consider two nodes, A and B, where node A would like to know the relative position of node B with respect to its position. These nodes could be any object containing a single light such as motorbikes or robots. Considering this scenario, our goal is the following: *as long as a node is within the illumination coverage of a neighbor, the node should be able to obtain its relative position without any prior knowledge of its surroundings.* These relative positions could be used for task coordination in the case of robots, for Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications in the case of smart vehicles, or can be leveraged by the beamforming

techniques to increase the communication data rate.

Research problem. Localization methods based on visible light can be divided into two macro groups depending on the type of photoreceptor being used: cameras [33,128] or Photodiode (PD)s [129]. Cameras are a popular option but they have severe limitations decoding information at variable distances. PD, on the other hand, are known to operate well at a wide range of distances. Thus, our work focuses on PDs. The challenge of using PDs is that the system needs to consider Lambertian radiation patterns. Several studies have looked into methods to exploit the properties of Lambertian patterns for localization [122, 124, 130]. Next we describe the key concepts we build upon from the State-of-the-Art (SoA) and the novelty of our work.

Building Block 1: Inertial sensors. The first idea we build upon from the SoA is the use of inertial sensors. Contrary to isotropic sources, where a receiver can determine its unique *circular* iso-countour (locus in which every point satisfies the received power and orientation requirements) based on the Received Signal Strength (RSS) observed from the transmitter, with LED sources the RSS information is not sufficient. To define a single *Lambertian* iso-contour the receiver needs information about its relative angle w.r.t. to the sender. To overcome this problem, several studies use accelerometers to determine the orientation of the receiver with respect to the ceiling [131], so a unique iso-countour can be defined per light. Then the iso-contours of several lights are used to localize the receiver (similar to standard trilateration methods). We also use inertial sensors to detect the direction of nodes, but we cannot use gravity as a frame of reference to derive orientations, because in our case lights are not fixed at ceilings.

Building Block 2: Rotations. The second idea we build upon from the SoA is the use of rotations to obtain more location information [132]. Without inertial sensors, receivers cannot pinpoint a unique iso-contour, but with inertial sensors receivers cannot only obtain their current iso-countour, but they can obtain more iso-countours by simply rotating on its same position. The reason for this phenomenon is that two different receiving angles create two different iso-countours. Note that with traditional isotropic sources rotations of the transmitter or receiver provide no extra information because the RSS remains the same. Rotation-based approaches have been used to add information (equations) to under-determined localization systems in order to obtain unique solutions. We also exploit rotations but considering more complex scenarios: rotations at the receiver, the transmitter or both. In traditional indoor settings only the simplest case is considered: rotations of the receiver.

Novelty: a localization framework for single mobile lights. Our key insight to map the traditional indoor localization problem to the mobile case is the following: *In the most general sense, mobility patterns involve changes in distance and direction*[1]. *Any change*

[1]This is the highest level of abstraction used by one of the most popular mobility models: the random

in direction –either by the transmitter, the receiver, or both– provides extra information to pinpoint the relative location of nodes, because for that short period of time (before and after the rotation), the distance between nodes could be assumed to be (roughly) the same.

To achieve this goal, we assume that nodes have (i) a LED source, a compass sensor, and a PD (PD); and (ii) visible light communication capabilities, that is, nodes can modulate their lights to transmit information, and use their PDs to decode information. With these basic capabilities, our relative positioning algorithm works in the following manner. *First*, nodes use their lights not only for illumination but also to broadcast continuously information about three parameters: their direction (provided by the compass sensor), the light's output power (which determines the length of the beam) and the light's Lambertian order (which determines the width of the beam). *Then*, neighboring nodes use their PDs to decode the sender's information, and use our mathematical model to obtain its relative location based on (i) the information received from the sender, (ii) its own direction and (iii) the detected signal strength.

Our contributions. Considering the above mentioned challenges the key contributions of our work are:

- Method [Sections 5.3 and 5.4]. We derive closed-form expressions to obtain unique localization solutions for cases where either the transmitter or the receiver rotate. For the case where both nodes rotate, we show that multiple solutions are possible, but we use simulations to provide some basic rules that help with identifying a single solution.

- Evaluation [Sections 5.5 and 5.6]. We implement our model with off-the-shelf components and perform a systematic evaluation of our approach. We analyze the effects of single and concurrent rotations and perform evaluations with mobile nodes. Our results show that our method can provide localization errors below 5 cm.

5.1. Main Concept

In this section, we present background information to understand the concept of Lambertian patterns and the basic localization principle behind this work.

5.1.1. Lambertian patterns

Consider an LED light source (Transmitter (TX)) and an optical Receiver (RX), as illustrated in Fig. 5.1. Given any output power at the TX (LED), the received signal strength at the RX (PD) depends on three key parameters: the distance between them (d), the irradiation angle (ψ) and the incidence angle (θ). The longer the distance, or

way-point model.

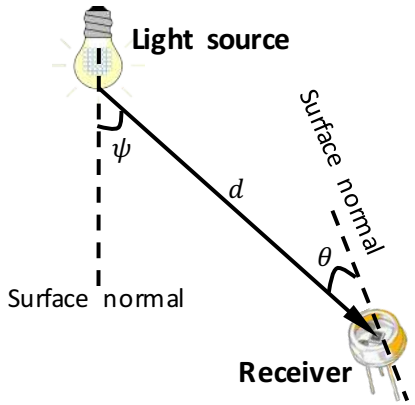

Figure 5.1: Propagation properties of LEDs.

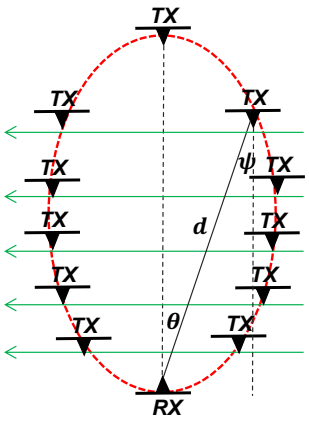

Figure 5.2: Iso-contours of received power.

the wider any of these angles, the lower the RSS. The interplay among these three parameters leads to Lambertian radiation patterns, where the maximum length of the coverage is mainly determined by the output power of the LED and the maximum width is determined by the Lambertian order (m). A small value of m leads to a broad coverage of the LED; a large m leads to a long but narrow coverage.

Formally, this pattern is captured by the well-known Lambert's cosine law [123]:

$$R_t(\psi) = \frac{m+1}{2\pi} \cos^m(\psi) \tag{5.1}$$

The channel loss $H(0)$ between the TX and RX is

$$H(0) = \begin{cases} A_{\mathrm{RX}} \cdot \frac{m+1}{2\pi d^2} \cos^m(\psi) \cos\theta & \text{for } \theta \in [0, \Theta_c], \\ 0 & \text{for } \theta > \Theta_c \end{cases} \tag{5.2}$$

where A_{RX} is the sensing area of the PD of the RX, and Θ_c is the PD's FoV.

Letting P_t and P_r denote the optical transmission power of the LED and the received power at the PD, respectively; and letting N refer to the sum of ambient noise and the PD's shot and thermal noise; then P_r can be written as:

$$P_r = P_t \cdot H(0) \cdot g_r(\theta) + N \tag{5.3}$$

where $g_r(\theta)$ is the optical gain of the PD. $g_r(\theta)$ is a non-zero *constant* when $\theta \in [0, \Theta_c]$, and is zero otherwise [133]. Therefore, in the rest of this Chapter, we simply use g_r to denote $g_r(\theta)$ in the calculation of P_r.

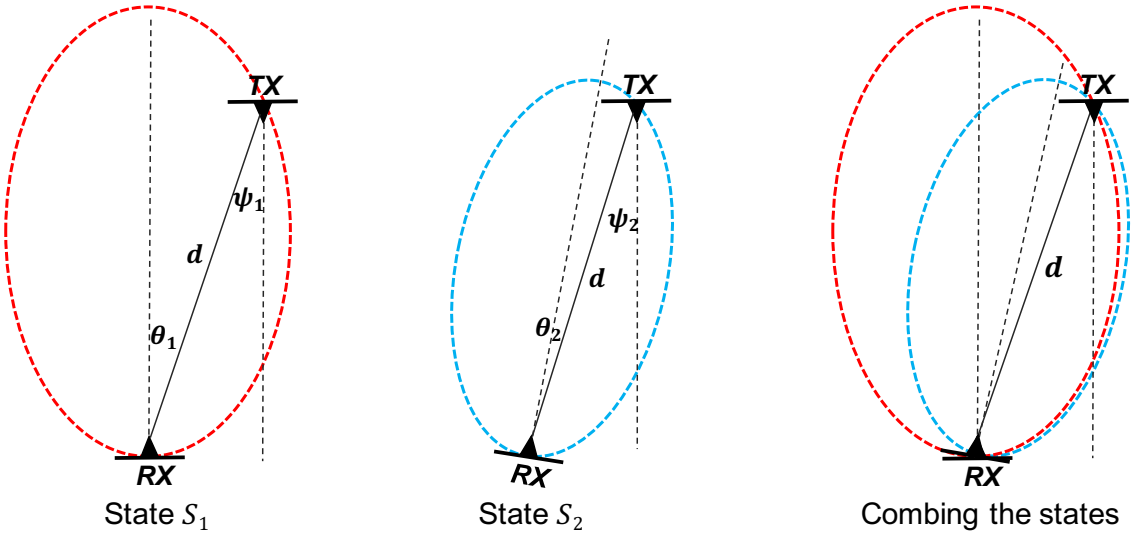

Figure 5.3: Localization using iso-contours intersection.

5.1.2. Basic localization principle

Given a received power P_r, the can be present in multiple positions with respect to the , cf. Eq. (5.2). For example, for a low P_r, the can be far away but aligned to the ($\psi = \theta = 0$), or it can be nearby but misaligned. If we fix the TX's orientation and move it in a horizontal 'scanning' motion at different perpendicular distances from the , as illustrated in Fig. 5.2, each scan provides two locations where the received power is measured as P_r. All these locations form an *iso-contour* where the received power is the same.

The principle behind this work is to exploit changes in the iso-contours due to nodes' movements, cf. Fig. 5.3. If the RX rotates, the iso-contour changes from its original shape (red) to a new shape (blue). In this particular case, the change in the iso-contour is caused by the change in the incidence angle θ. The intersection of these two iso-contours can be used to estimate the relative position of the TX. In this work, we exploit this property in *mobile scenarios*, where any type of node, or , can move freely.

5.2. System Model

Consider a system of two *mobile* nodes: one acts as a reference point (referred as receiver) that has a PD; and the other is a target to be localized (referred as transmitter) that has an LED light source. Both nodes can measure their orientations and movement through on-board sensors, compasses and accelerometers respectively. The transmitter has information of the optical properties of its LED lights, such as the transmission power and Lambertian order. The transmission power and Lambertian order together with the transmitter's real-time orientation are shared with the receiver via visible light

communication. The receiver can decode the transmitted information through its PD and it can also measure the received power.

Next we define the *state* of our system. Since we assume that both nodes are mobile, we define a state S as follows:

$$S = (\alpha_{\text{tx}}, \alpha_{\text{rx}}, d, P_r) \tag{5.4}$$

where α_{tx} and α_{rx} are the orientations of the transmitter and receiver with respect to North, respectively, d is the relative distance of the transmitter with respect to the receiver, and P_r is the received power at the receiver.

As discussed in Sec. 5.1.2, given fixed positions and orientations of the transmitter and receiver, there exists an iso-contour for the received power P_r. But we need more than one iso-countour to pinpoint a unique location. The problem we face in our scenario is two-fold. First, for the iso-contours to be useful, we need them to be rooted at the same location, but given that we are dealing with mobile nodes each iso-countour could be rooted at different locations, rendering the information useless for our localization process. Second, we face an under-determined system, where we have many variables but not enough equations.

Let us describe our second challenge in a more formal manner. Assume our system is currently in state $S_1(\alpha^1_{\text{tx}}, \alpha^1_{\text{rx}}, d_1, P^1_r)$. According to Eq. (5.3), the received power P^1_r for state S_1 can be written as follows:

$$P^1_r = P_t A_{\text{RX}} \frac{m+1}{2\pi (d_1)^2} \cos^m(\psi_1) \cos(\theta_1) \cdot g_r + N \tag{5.5}$$

A similar equation can be derived for a later state $S_2(\alpha^2_{\text{tx}}, \alpha^2_{\text{rx}}, d_2, P^2_r)$, where the relative position and the orientations of the transmitter and receiver change, and are denoted by d_2, α^2_{tx}, and α^2_{tx}, respectively; and the measured received power is P^2_r:

$$P^2_r = P_t A_{\text{RX}} \frac{m+1}{2\pi (d_2)^2} \cos^m(\psi_2) \cos(\theta_2) \cdot g_r + N \tag{5.6}$$

In the above two equations, we have six variables: the irradiation angles ψ_1, ψ_2, the incidence angles θ_1, θ_2, and the relative distance d_1 and d_2 between the transmitter and receiver. So we can not solve them directly.

To overcome the challenges above we propose two approaches. First, we force nodes to broadcast continuously and periodically their orientation and LED parameters via visible light communication. And, at the receiver, we only consider changes in orientation detected within a very short period of time, so we can assume that the locations of the and remain "constant". This approach has two important consequences, first it allows us to assume that the two iso-countours are rooted at the same point; and second, it reduces the number of variables in the above equations from six to five, because over a short period of time we can assume that $d_1 = d_2 = d$. Second, we exploit the information

coming from the compasses to derive dependencies among the five remaining variables, enabling us to identify a single solution in most cases.

5.3. Analysis

In this section we derive the dependencies that exist among the unknown variables described in Eqs. (5.5) and (5.6). Our analysis focuses on changes in orientations, but we also provide some results for scenarios where there are no changes in orientations but there are changes in distance, for example, two mobile nodes in a straight road moving at variable speeds.

5.3.1. Deriving dependencies

For these derivations, we consider a short period of time where the relative distance between the transmitter and receiver is 'fixed' while their relative orientations change.

First, we will identify dependencies for the incidence angles, which will remove two unknowns (θ_1 and θ_2). The incidence angles θ_i are a function of the orientations of the TX and RX as well as the irradiation angles ψ_i. Considering states S_1 and S_2, these dependencies can be described as follows.

Dependency 1: $\theta_1 = f(\alpha^1_{\text{tx}}, \alpha^1_{\text{rx}}, \psi_1)$

Dependency 2: $\theta_2 = f(\alpha^2_{\text{tx}}, \alpha^2_{\text{rx}}, \psi_2)$

This relationship is illustrated in Fig. 5.4. Let $\Delta\alpha^i$ be the relative angle between the orientations of the RX and TX at state S_i, then we have

$$\Delta\alpha^i = \alpha^i_{\text{rx}} - \alpha^i_{\text{tx}}, \quad i \in [1, 2] \tag{5.7}$$

Further, the incidence angle θ_i can be expressed as:

$$\theta_i = \pi + \psi_i + \Delta\alpha^i = -\pi + \psi_i + \alpha^i_{\text{rx}} - \alpha^i_{\text{tx}}, \quad i \in [1, 2] \tag{5.8}$$

Second, we will identify a dependency for the irradiation angle ψ_2, which will remove one extra unknown. The irradiation angle ψ_2 is a function of the irradiation angle ψ_1 and the orientations of the TX in both states:

Dependency 3: $\psi_2 = g(\alpha^1_{\text{tx}}, \alpha^2_{\text{tx}}, \psi_1)$

This dependency is illustrated in Figure 5.5. Let $\Delta\alpha_{\text{tx}}$ be the difference of the TX's orientations in states S_1 and S_2, then

$$\Delta\alpha_{\text{tx}} = \alpha^2_{\text{tx}} - \alpha^1_{\text{tx}} \tag{5.9}$$

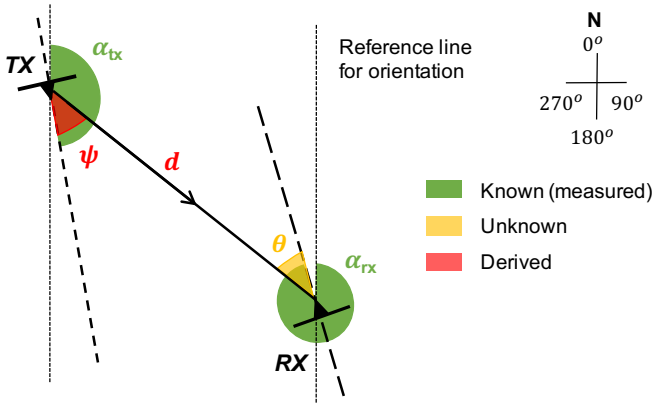

Figure 5.4: Relation between irradiation angle, incidence angle, and orientations.

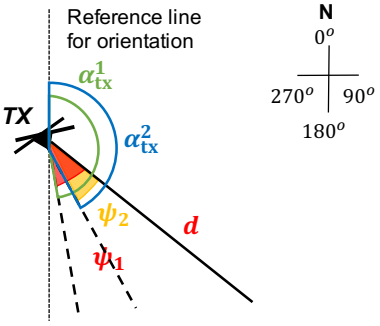

Figure 5.5: Relation between irradiation angles.

Based on this, we can derive the expression of ψ_2 as follows:

$$\psi_2 = \Delta\alpha_{\text{tx}} + \psi_1 = \alpha^2_{\text{tx}} - \alpha^1_{\text{tx}} + \psi_1 \tag{5.10}$$

Finally, substituting the above three dependencies into Eq. (5.5) and Eq. (5.6), we have

$$P^1_r = P_t A_{\text{RX}} \frac{m+1}{2\pi d^2} \cos^m(\psi_1) \cos(f(\alpha^1_{\text{tx}}, \alpha^1_{\text{rx}}, \psi_1)) \cdot g_r + N \tag{5.11}$$

$$P^2_r = P_t A_{\text{RX}} \frac{m+1}{2\pi d^2} \cos^m(g(\alpha^1_{\text{tx}}, \alpha^2_{\text{tx}}, \psi_1)) \cos(f(\alpha^1_{\text{tx}}, \alpha^2_{\text{tx}}, \alpha^1_{\text{rx}}, \alpha^2_{\text{rx}}, \psi_1)) \cdot g_r + N \tag{5.12}$$

In these two equations there are only two unknowns (ψ_1 and d), and thus, they can be solved numerically. Next, we will describe how to solve these functions for two cases: non-simultaneous and simultaneous rotations. We define *simultaneous* rotations as those events where both, the TX and RX, change their orientations within the same short period of time. If the changes of orientations occur in different periods, we consider the rotations to be non-simultaneous. Notice that our definition depends on how often we send up-to-date information over the visible light channel. For our implementation, simultaneous rotations are those occurring within 2 ms of each other (due to the sampling rate of our sensors).

5.3.2. Non-simultaneous rotations

We will first consider the scenario where the RX rotates, and then, we will consider the scenario where the TX rotates. The first scenario is simpler and is the one we find in indoor setups (because the lights, TX, are fixed).

1) RX rotates. For the scenario where the RX rotates, we have the following proposition. If the RX rotates and the TX is fixed, i.e., TX does not rotate between the two states S_1 and S_2, then we have the following closed-form expression of θ_2:

$$\theta_2 = \arctan\left(\frac{(P_r^1 - N)\cos(\alpha_{\text{rx}}^2) - (P_r^2 - N)\cos(\alpha_{\text{rx}}^1)}{(P_r^1 - N)\sin(\alpha_{\text{rx}}^2) - (P_r^2 - N)\sin(\alpha_{\text{rx}}^1)}\right) - \pi + \alpha_{rx}^2 \tag{5.13}$$

Proof: The proof is presented in Appendix A.1. ■

2) TX rotates. For the scenario where the TX rotates, the analysis is different from the scenario where the RX rotates. This occurs because the path loss caused by the TX is magnified by the Lambertian order m: $\cos^m(\psi)$, cf. Eq. (5.2), whereas the path loss caused by the RX is not ($\cos(\theta)$). Therefore, for $m > 1$ (which are common Lambertian orders for most LEDs), the same rotation angle at the TX and the RX will have a different impact. For the scenario where the TX rotates, we have the following proposition. If the TX rotates and the RX is fixed, i.e., RX does not rotate between the two states S_1 and S_2, then we have the following closed-form expression of θ_2:

$$\theta_2 = \arctan(j) - \pi + \alpha_{rx}^2 - \alpha_{tx}^1 \tag{5.14}$$

where

$$j = \frac{\cos(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1) - \sqrt[m]{\frac{(P_r^2 - N)}{(P_r^1 - N)}}}{\sin(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1)}$$

Proof: The proof is presented in Appendix A.2. ■

For both non-simultaneous rotations, we have closed-form expressions for θ_2. Thus, we can obtain *the closed-form expression of the distance* d based on Eq. (5.11) and (5.12):

$$d = \sqrt{\frac{(m+1)A_{\text{RX}}}{2\pi(P_r^1 - N)}\cos^m(\psi_1)\cos(\psi_1 - \pi - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1)\cdot g_r}$$

Based on the closed-form expressions of θ_2 and d, we obtain the relative position of the TX with respect to the RX, and thus, we can conclude that for non-simultaneous rotations, we have *unique* solutions to Eq. (5.11) and (5.12).

5.3.3. Simultaneous rotations

For these scenarios, we are not able to derive closed-form expressions for θ_2 and d. This is due to the fact that the transmission pattern is Lambertian with the form $\cos^m(\cdot)$ where m is positive and real. We can only solve the equations numerically. Based on Eq. (5.11) and (5.12), we first derive an equation where ψ_1 is the only unknown variable.

By transforming the problem from the ψ_1 domain to the k domain, we have the

following polynomial expression for k:

$$b \cdot \sin(h) \cdot k^{m+1} + b \cdot y \cdot k^m - \sin(c) \cdot k - x = 0 \tag{5.15}$$

where k is a function of ψ_1:

$$k = \cos e - \tan(\psi_1) \cdot \sin e \tag{5.16}$$

and

$$\begin{cases} b = \frac{P_r^1 - N}{P_r^2 - N} \\ c = \alpha_{\text{rx}}^1 - \alpha_{\text{tx}}^1 \\ e = \alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1 \\ h = \alpha_{\text{rx}}^2 - \alpha_{\text{tx}}^1 \\ x = \sin(\alpha_{\text{tx}}^2 - \alpha_{\text{rx}}^1) \\ y = \sin(\alpha_{\text{tx}}^2 - \alpha_{\text{rx}}^2) \end{cases}$$

Proof: The proof is presented in Appendix A.2.1. ■ Proposition 5.3.3 implies that by solving the polynomial in Eq. (5.15), we can get all possible solutions for k; and once k is known, we can use Eq. (5.16) to obtain:

$$\theta_2 = \arctan\left(\frac{\cos e - k}{\sin e}\right) - \pi + \alpha_{rx}^2 - \alpha_{tx}^1 \tag{5.17}$$

And d can be obtained by substituting Eq. (5.17) into Eq. (5.15).

From Eq. (5.15) and (5.17) we observe that we can have up to $m + 1$ solutions for our localization problem. However, by considering some physical constraints such as the FoVs of the TX and the RX, we can filter out some of the results. We will look further into this issue in Sec. 5.4.

5.3.4. Impact of measurement errors

Based on the dependencies derived earlier in this section, we now analyze the effects of measurements errors in our system. We focus on orientation errors (α_{rx} and α_{tx}) and errors in the received power (P_r).

From Eq. (5.12), we can derive the distances d as a function of the other parameters:

$$d^2 \propto \frac{\cos^m(\psi_1)\cos(f(\alpha_{\text{tx}}^1, \alpha_{\text{rx}}^1, \psi_1))}{P_r^1}$$
$$d^2 \propto \frac{\cos^m(g(\alpha_{\text{tx}}^1, \alpha_{\text{tx}}^2, \psi_1))\cos(f(\alpha_{\text{tx}}^1, \alpha_{\text{tx}}^2, \alpha_{\text{rx}}^1, \alpha_{\text{rx}}^2, \psi_1))}{P_r^2}$$

Based on these equations, we derive the following insights:

Insight 1: Lights with wider beams ameliorate the effect of orientation errors in the TX. Given an error in the orientation of the TX, a wider beam (small m) will lead to more accurate location estimations than a narrower beam (large m). This occurs because the cos() function takes values between 0 and 1. If the Lambertian order is high (narrow LED beam), the error will be exacerbated because it will be elevated to a power m. Thus, for mobile scenarios, Lambertian sources with $m = 1$ are the preferred choice. Note that in indoor scenarios this is a lesser problem because lights are fixed to the ceiling, and thus, the orientation is known and fixed.

Insight 2: Orientation errors at the RX are more detrimental when the RX lies at the edges of the illumination area. Errors on the RX's orientation will have a particularly negative impact as the RX moves away from the normal surface of the transmitter. This occurs because the cos() function decreases rapidly as the incidence angle moves away from 0^o. Thus, the more aligned the TX and RX are, the higher the accuracy of our estimations.

Insight 3: Measurement errors in the received power will have a larger effect at longer distances. This is a property that is common to all electromagnetic sources because the intensity decays exponentially with distance. Thus at close distances, errors in P_r will lead to small errors in localization, but at longer distances the same error in P_r can lead to exponentially larger errors in localization.

5.3.5. Variable distances

Until now, our analysis has focused solely on events where rotations are expected. But some mobile scenarios may not have any rotations. For example, nodes moving on a straight path at variable speeds will not change their orientations, but their relative distances will change. For these types of scenarios, the equations describing the system at states S_1 and S_2 are

$$P_r^i = P_t A_{\text{RX}} \frac{m+1}{2\pi d_i^2} \cos^m(\psi) \cos(\theta) + N, \quad i \in \{1, 2\} \tag{5.18}$$

If we assume that we can estimate the difference in distances between states S_1 and S_2 (by for example performing a double integral of an accelerometer):

$$\varepsilon = d_2 - d_1 \tag{5.19}$$

Then, based on Eq. (5.18) and Eq. (5.19), we can obtain a closed-form expression for the relative distance as follows:

$$d_2 = \frac{\varepsilon}{1 - \sqrt{\frac{P_r^2}{P_r^1}}} \tag{5.20}$$

From Eqs. (5.20), (5.8), and (5.18), we have the following expression

$$\frac{2\pi d_2^2 P_r^2}{(m+1)A_{\text{RX}}} - N = \cos^m(\theta + \pi + \alpha_{tx} - \alpha_{rx})\cos(\theta) \tag{5.21}$$

where θ can be calculated numerically.

An important point to consider for these no-rotations cases is that the relative location is no longer obtained by intersecting two Lambertian iso-contours, but by intersecting a single Lambertian iso-contour and a circle. This implies that unless the TX and RX are aligned ($\psi = \theta = 0$), we will always obtain two possible locations. Nevertheless, due to physical restrictions in the system in practical, one of the location could be discarded.

5.4. Model validation and insights for the design

In this section, we perform extensive simulations to validate the models in Sec. 5.3 and provide significant insights in the design of the systems. For each scenario, the values of ψ_1, ψ_2, θ_1 and θ_1 are randomly chosen following a uniform distribution between $\frac{-\pi}{2}$ and $\frac{\pi}{2}$. The maximum *coverage* distance between a TX and RX is then computed by taking into account that the maximum path loss can be 20 dB (beyond that threshold no communication can be reached). We perform 10,000 simulations per scenario.

Based on Eqs. (5.5) and (5.6), in what follows, we call *iso-contour* the shape that represents all the possible positions that the receiver could be, given some values of α_{tx}^i, α_{rx}^i and P_r^i. The iso-contour shape contains two main components:

- The component relative to the transmitter ($\cos(\psi_i)^m$).
- The component relative to the receiver ($\cos(\theta_i)$).

After a rotation or a distance change (see previous section), a new iso-contour is generated. Each point (2D coordinates) where both iso-contours cross is the result returned by the algorithm.

Based on these considerations, in what follows we present our study for the different setups, starting with the non-simultaneous rotations case.

5.4.1. Non-simultaneous rotations

To validate Proposition 5.3.2 and Proposition 5.3.2 (non-simultaneous cases), we perform a brute-force study in a noiseless scenario. We then compute the Root Mean Square Error (RMSE) of the position error varying m. The results are presented in Fig. 5.6a. The figure shows that the RMSE is less than 10^{-11}, regardless of m (this error is due to the computational limits of the simulator). This shows that the models for non-simultaneous scenarios are correct. In addition, we always find one unique solution

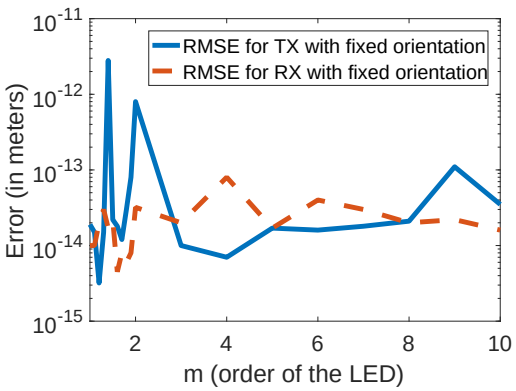

(a) Model validation (RMSE of the positioning error between model and simulation).

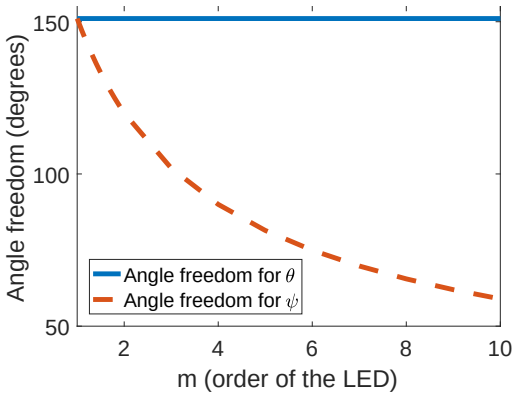

(b) Angle freedom over m.

Figure 5.6: (a) model validation when either the transmitter or the receiver is fixed (non-simultaneous case). (b) angle freedom for θ (blue) and ψ (orange) for different values of m.

of the system. This is expected, as two Lambertian iso-contours generated from the same source with a different rotation angle can only cross in one point in space.

Due to the directionality of LEDs, rotating the transmitter will have a much higher impact on the received power than rotating the receiver. The difference increases with m. In order to confirm this, we place the transmitter and the receiver at half the distance from the maximum coverage distance for the configured sensitivity. We then rotate the transmitter until the received signal is below the sensitivity threshold. Afterward, the same is done at the receiver. The range of angles where the received power is above the sensitivity threshold is called angle freedom. Fig. 5.6b shows that, while the angle freedom remains constant for the receiver, it decreases for the transmitter as m increases. This confirms that rotations in the transmitter are more critical than at the receiver.

5.4.2. Simultaneous rotations

We now study Proposition 5.3.3 for the simultaneous scenario. Unlike the non-simultaneous case, in this case both transmitter and receiver are free to rotate. Solving Eq. (5.15), we may get either a unique solution or more than one solution, depending on m. From a physical point of view, both components $\cos(\psi_i)^m$ of the transmitter and $\cos(\theta_i)$ of the receiver change, and thus the shape of the iso-contour changes, making *several* results possible. But for location purposes, it is desired to have only *one* solution.

Next we show that we can follow some basic rules to discard many potential solutions to reach (in most cases) a single solution. Some solutions can be discarded as follows: i) discard solutions that give complex numbers, ii) discard solutions with no physical meaning, iii) discard solutions unlikely to happen in real environments.

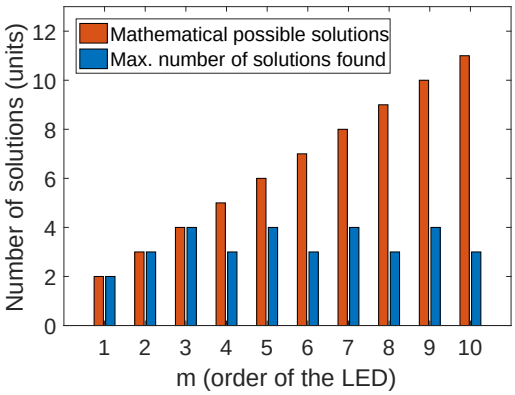

(a) Total vs real solutions varying m.

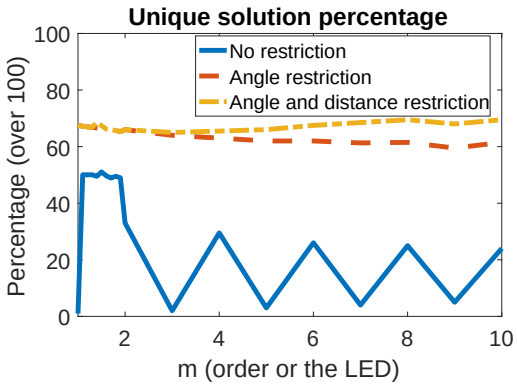

(b) Unique solutions.

Figure 5.7: (a) number of mathematically possible solutions and number of real solutions through numerical analysis. (b) percentage of unique solutions with and without physical constraints.

For natural m, as it can be seen in Fig. 5.7a and as expected from Eq. (5.15), the number of possible solutions that can be computed numerically increases with m. Nevertheless, the number of maximum real solutions, and thus with a physical meaning, is three or four depending on m. Intuitively, the fact that maximum number of real solutions is four, can be explained by the fact that two concave surfaces intersect in at most four points.

The next physical constraint that can be introduced into the system is that ψ_1, ψ_2, θ_1 and θ_1 need to be between $-\frac{\pi}{2}$ and $-\frac{\pi}{2}$. If any of them is not, the communication would not have been possible, because that will require the transmitter (in the case of ψ) or the receiver (in the case of θ) to have transmitted or received with the back of the LED or the PD, respectively, so the solution could be discarded.

Finally, the last restriction that has been introduced in the system is the distance restriction. It is assumed that the distance is not going to be very short (less than 1cm), so solutions that lay inside this region are discarded too.

An analysis on how the aforementioned constraints affect the number of possible solutions is performed, and the results are shown in Fig. 5.7b. If we do not apply our constraints, the number of scenarios with unique solutions is low. But when we take into consideration the constraints we derived, the number of scenarios with unique solutions go from around 30% to close to 70%.

5.4.2.1. Measurement error

Finally, we study how errors in the compass readings affect the system. For this study, we assume a Gaussian noise with $\sigma = 5°$ for both TX's and RX's compasses[2]. The RMSE of the distance is in the order of 10^{-1} meters, which means accuracy of decimeters. The accuracy of the TX's orientation is the key to minimize the error as the TX is more directional than the RX. Comparing the two non-simultaneous scenarios, it can be seen that when we fix the TX's orientation, the error is 50% smaller than when we fix the RX's orientation. The error in the simultaneous scenario is comparable to the worst error in the non-simultaneous scenario. These results highlight an important difference between our work and the SoA, which focus on indoor scenarios where the TXs are fixed, and thus are less challenging.

5.4.2.2. Insights

First, we observe that a compromise is needed when choosing the order m of the LED. A small m implies a wider covered area (good), smaller power distributed in the direction of maximum emission (not good) and a slightly smaller number of real solutions considering physical constraints (good). The opposite holds when increasing m. Second, it is preferable to have mobility patterns where the receiver is the one rotating. This is due to the fact that LEDs have Lambertian sources and rotation in transmitters may cause larger power received differences. Finally, in the simultaneous scenario, it may be desirable to have more than two measurements. In this way, several results may be discarded and the one common result should remain as the correct one.

5.4.3. Variable distances

We verify the correctness of the closed-form expression for fixed angle introduced in Eq. (5.20). For this study, we compute the position considering a brute force analysis varying m. The RMSE in distance between the simulated position and the one computed by the model is below 10^{-12} for 100% of the time (the difference is due, as above, to computational limitations), that confirms the correctness of the model.

5.5. Implementation

We implement our localization algorithm in the OpenVLC platform [135, 136] for low-cost Visible Light Communication (VLC). OpenVLC consists of three parts: a BeagleBone Black (BBB) board, an optical front-end transceiver and the software solution. We customize it to satisfy the requirements of our localization method. The

[2]This value of σ characterizes the effect of buildings on compass errors [134]. Outdoors the errors in compass are far lower.

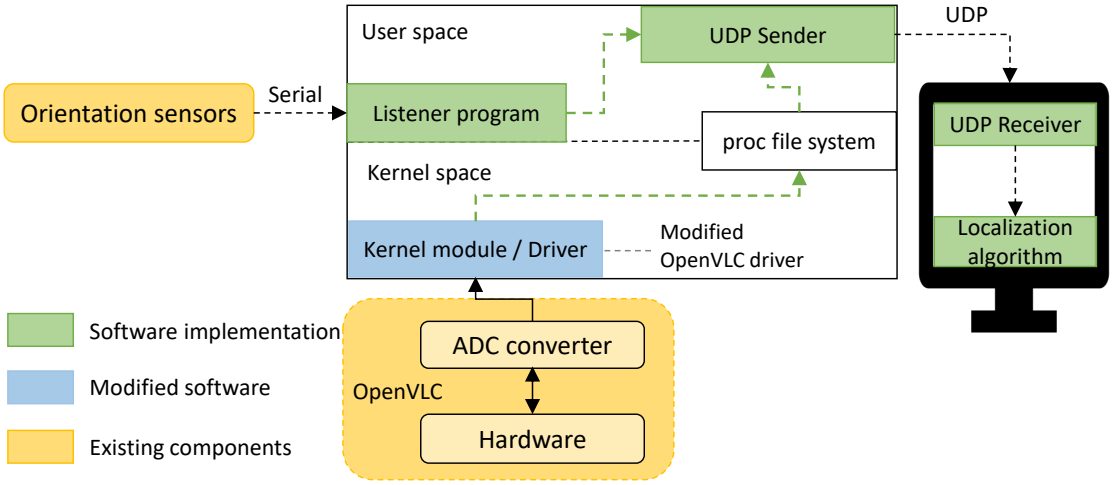

Figure 5.8: Functional blocks of our system.

Table 5.1: Modified frame format in our system.

Preamble	...	*Transmission power*	*Lambertian order*	*Orientation*	...
3 Bytes	...	*1B*	*1B*	*2B*	...

functional blocks of our system are shown in Fig. 5.8. As indicated in Eq. (5.12), to run our localization method, we need the following information from the TX: *transmitted power, Lambertian order m of its transmitting LED and its orientation.* And the following information from the RX: *received power and its orientation.*

Measuring the orientation of nodes. Orientation plays a crucial part in our system and it can be measured by inertial sensors. The sensor used is *Adafruit 9-DOF Absolute Orientation IMU Fusion Breakout - BNO055*, which does sensor fusion of the magnetometer and accelerometers for higher accuracy. Our OpenVLC node communicates to the sensor using the serial port, as shown in Fig. 5.9. The sampling rate of the sensors is equal to 500 Hz.

Sharing information between nodes. *The transmission power, Lambertian order of the LED at the TX, and the TX's orientation* must be shared with the RX. In our work, we use visible light communication to achieve this. We modify the original frame format of OpenVLC by adding three fields to convey this information, as shown in Table 5.1.

Calculating the received power. The preamble of each frame is used to determine the received power. The length of the preamble is 24 bits (cf. Table 5.1). These bits are composed of alternate HIGH and LOW symbols, primarily used for synchronization. Information pertaining to the received power can be extracted based on the preamble. We measure P_r for each HIGH symbol. The measurement comprises the total power received from the LED plus ambient noise and PD's shot and thermal noise (cf. Eq. (5.3)). In

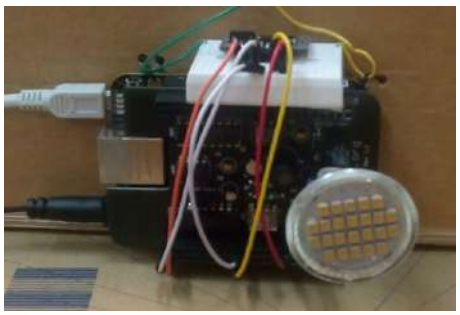

Figure 5.9: Sensors connected through serial to the BBB, and the BBB to the computer via USB.

Figure 5.10: A snapshot of the experiment setup

order to have controlled environment for our evaluation, measurements are done in a dark environment without external light sources.

Localization. The localization algorithm is implemented to operate in real time, with interactions among the blocks illustrated in Fig. 5.8. First, we modify the OpenVLC driver to collect raw power readings from the data frame preamble, and share these readings with the user space using the *proc* file system. Second, a listener program collects the orientation information from sensors by continuously polling the serial port in the BBB. The listener thread then sends the data to the computer every second using User Datagram Protocol (UDP). In the computer, a program reads the UDP messages and every two messages, runs the localization algorithm based on Eq. (5.12).

5.6. Experimental Evaluation

We evaluate the proposed relative localization method with experiments under various settings. When the LED emits the symbol HIGH, it is supplied at a voltage of 12 V. The measured current (in general, depending on the LED characteristics) is equal to 52 mA. This results in a total power emitted by the LED close to 0.6 W (some power in the order of 0.1-0.2 W is dissipated by the driving circuit). The PD used at the is the OPT101[3]. A snapshot of the experiment setup is shown in Fig. 5.10. The system is located in a 15.2 × 5.8 meter empty room, which is dark (<3 lux). The and are located at a sufficiently high distance from the ground to minimize reflections from the floor. In all the tests, the is located in the position $(0, 0)$ while the transmitter is placed in different locations.

Next, we first characterize the system and then we evaluate our localization method in several key scenarios.

[3]http://www.ti.com/lit/ds/symlink/opt101.pdf.

5.6.1. System characterization

The first step is to characterize the system for our particular hardware. From Eq. (5.5), we observe that the values we need to characterize are m, A_{RX}, and g_r.

5.6.1.1. m calculation

Dividing Eq. (5.5) and Eq. (5.6), configuring $\theta = 0$ and knowing ψ_1 and ψ_2 we get the following expression:

$$m = \frac{\log\left(\frac{Pr_1}{Pr_2}\right)}{\log\left(\frac{\cos(\psi_1)}{\cos(\psi_2)}\right)} \tag{5.22}$$

The value of m has been computed for different values of ψ_1 and ψ_2 and the average is taken as the final value, which leads to $m = 1$. Following the insights from the previous sections, the use of an LED with low m has several benefits and the only drawback of reduced sensitivity in the direction of maximum light emission. The maximum communication range in our setup is of approximately 1 meter, when the transmitter and receiver are perfectly aligned. The range could be increased with LEDs of higher nominal power and with more sensitive PDs.

5.6.1.2. A_{RX} and g_r calculation

For a single measurement, $A_{\mathrm{RX}} \cdot g_r$ can be calculated as:

$$A_{\mathrm{RX}} \cdot g_r = \frac{2\pi d^2}{(m+1) \cdot \cos^m(\psi_1) \cdot \cos(\theta_1) \cdot P_r^1} \tag{5.23}$$

The system is measured for several values of d so that ψ_1 and θ_1 are 0. The average value of $A_{\mathrm{RX}} \cdot g_r$ is 20902.

5.6.2. Relative angles variation

We consider three scenarios: TX rotation, RX rotation or both. The rotation angles both for transmitter and receiver are (0, 20, 30, 45, 60, -20, -30, -45, -60) degrees and the real locations of the transmitter in cm units are equal to (0, -10, -20) for x and (0, 10, 20, 30, 40, 50, 60) for y.

The results of all these experiments are shown in Fig. 5.11 and a summary of the average error is presented in Table 5.2. They show an average error in the x- and y-axis of less than 3.5 cm for all the cases, resulting in a position error of less than 5 cm. Position errors could be caused, in part, by small errors in the estimation of the Lambertian order m that has been performed manually (cf. Eq. (5.22)). More accurate values could be either available in datasheets of professional LEDs or measured through advanced photometers tools.

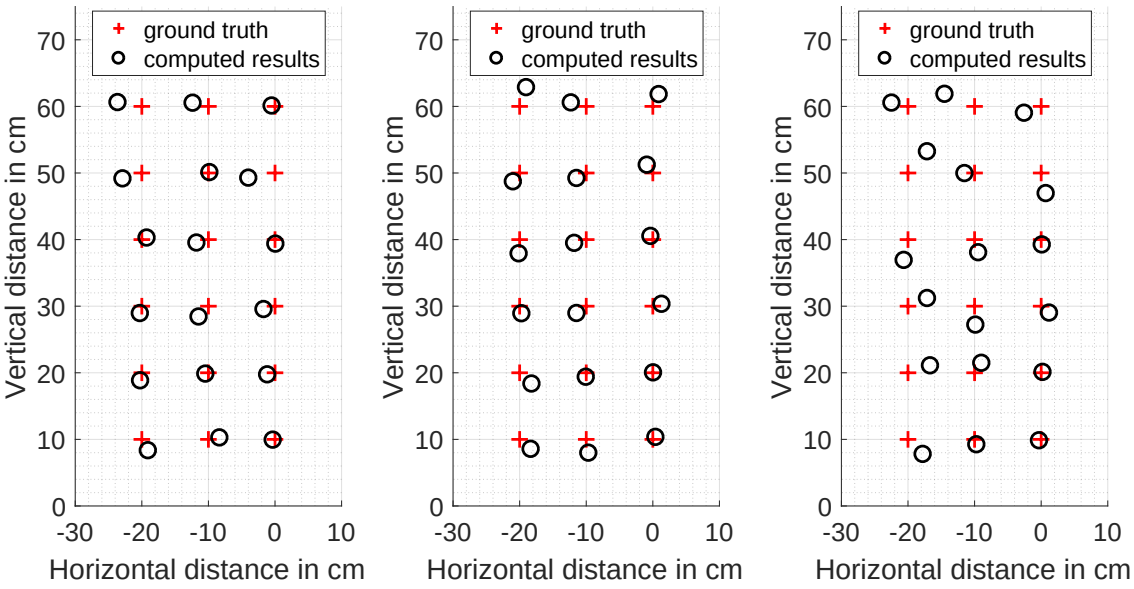

(a) Only RX rotates. (b) Only TX rotates. (c) Both rotate.

Figure 5.11: Accuracy results for rotations.

Table 5.2: Mean error in the x- and y-axis for rotations.

	Mean x-axis error	Mean y-axis error
Only RX rotates	2.61 cm	1.63 cm
Only TX rotates	3.20 cm	2.11 cm
Both rotate	3.83 cm	3.47 cm

It can be also seen than the error increases with the distance (Insight 3 in Section 4.4). This occurs because the sensitivity of the PD affects more the results when the received signal is closer to the receiver's noise. An illustration of this phenomenon is presented in Fig. 5.12. Another problem of PDs is that we cannot distinguish well scenarios with very short distance (less than 10 cm in our experiments) as PDs get saturated. However, a more advanced receiver can reduce the impact of this issue.

For the case of both devices rotating, it must be noted that the number of solutions we find is two for all the configurations. This is expected from Fig. 5.7a and the selection of Lambertian order close to $m = 1$. Nevertheless, one of the solutions could always be discarded because it belongs to a scenario where no transmission could be accomplished (cf. Section 5.4). The discarded solution requires the receiver to be in an invalid orientation with respect to the transmitter.

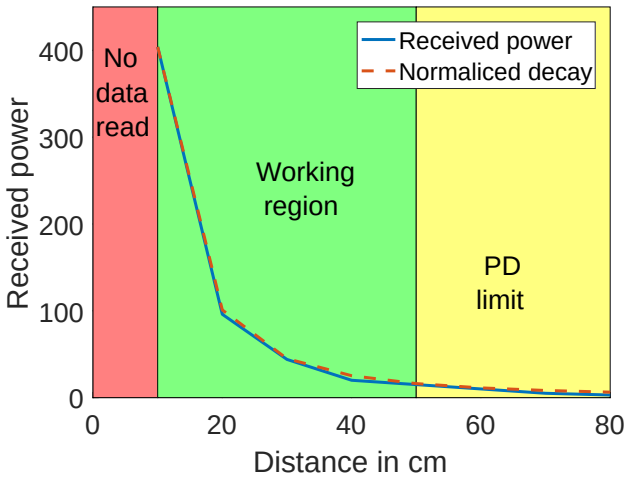

Figure 5.12: Real received power and theoretical distance decay over distance.

Table 5.3: Mean error in the x- and y-axis for distance changes.

	Mean x-axis error	Mean y-axis error
Distance change	1.65 cm	1.13 cm

5.6.3. Relative distance variation

In these experiments both the and are free to move assuming that they do not change their orientation with respect to each other. In order to compute the correct value of the system, the distance difference needs to be known as well as the received power. In these experiments the distance change is assumed to be known (in a real system it can be estimated with a double integration of the accelerometer input).

In the experiments, the distance between the TX and RX changes between the first and the second measurement with different steps (10, 20, 40, 50 cm). This is performed for different angles in degrees between the TX and RX (0, 20, 45, 60, -20, -45, -60). The results of the experiments are depicted in Fig. 5.13 and the average distance error in the x- and y-axis are shown in Table 5.2. The table shows a small position error (less than 2 cm) which validates the accuracy of our system.

It is important to mention that the system returns a second different solution around 25% of the time. This is due to the specific configuration of the system that allows, mathematically, more than one solution. Nevertheless, once the angle restrictions derived in Section 5.4 are applied, the solution remained unique for all the experiments.

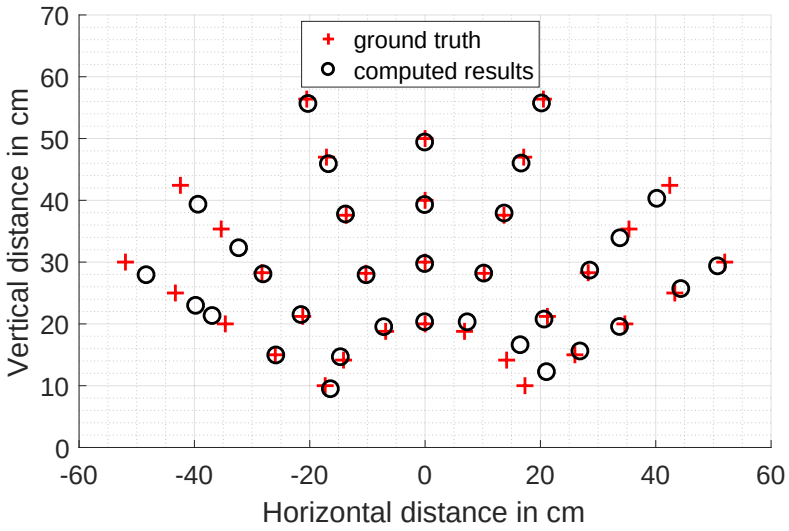

Figure 5.13: Accuracy results for distance changes.

Table 5.4: Comparing previous works with this study.

Study	Receivers	TXs	Evaluation	Calibration	Mobility
Lipro [132]	Single PD	Single LED	Empirical	Yes	Static (TX fixed)
Epsilon [130]	Single PD	Multiple LEDs	Empirical	Yes	Static (TX fixed)
Yin [133]	Single PD	Multiple LEDs	Theoretical	No	Static (RX fixed)
Yasir [131]	Single PD	Single/ Multiple LEDs	Empirical	No	Static (TX fixed)
Luxapose [33]	Single Image Sensor	Multiple LEDs	Empirical	No	Static (TX fixed)
This work	Single PD	Single LED	Empirical	No	Dynamic (TX and RX can be mobile)

5.6.4. Full mobile case

In this last scenario both nodes (i.e., the and) are free to move and orient themselves in any direction (within the FoV of each other). We conduct experiments by moving both nodes on tracks as shown in Fig. 5.14(a-b) to simulate the movements of robots or vehicles on a road. The localization is performed whenever the is at a reference position (marked in red color). We note that for straight paths like Track 1, ideally there are no changes in orientation of the nodes. However, even minor changes in the orientation are useful and exploited by our system. The localization results are shown in Fig. 5.14(c-d). The estimated localized positions of the are marked with blue squares. We observe that the results are quite accurate, with a maximal error of around 2 cm.

This approach works because the sampling rate and data communication are sufficiently fast to assume that the distance does not change between two consecutive samples.

(a) Setup of Track 1. (b) Setup of Track 2.

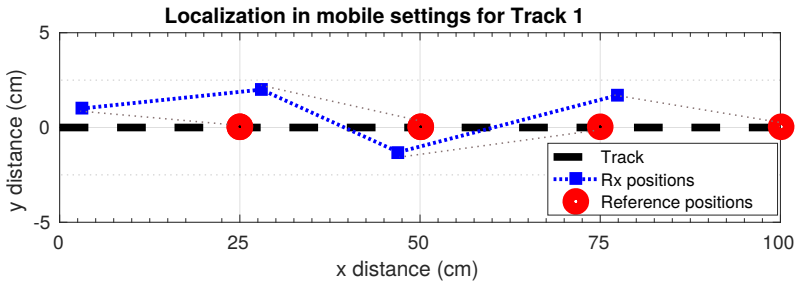

(c) Localization results (Track 1).

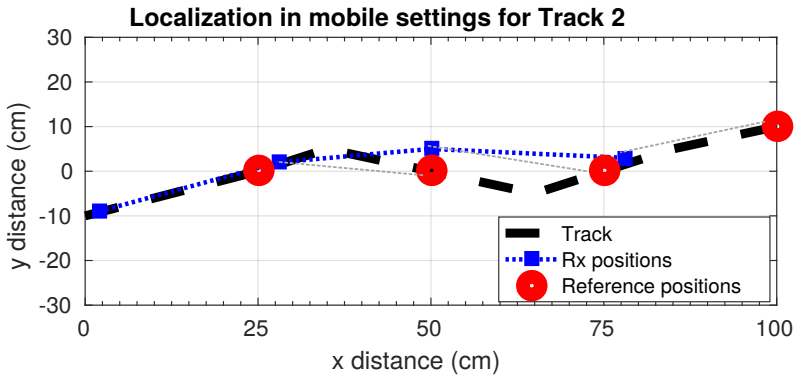

(d) Localization results (Track 2).

Figure 5.14: Accuracy results for fully mobile case.

5.7. Related Work

Localization has been studied widely [137]. Here we focus our discussion on studies leveraging methods based on visible light. Table 5.4 positions the novelty of our work within the most relevant studies in the SoA.

Initially, most visible light positioning solutions were built based on existing radio-frequency techniques [124], borrowing concepts such as proximity [138] and fingerprinting [126]. A second generation of methods then started to exploit the unique properties of light, such as the free-space optical propagation property and the limited multi-path [122]. In this category of works, at least three Lambertian LED sources are required to achieve high accuracy using multilateration algorithms [130]. Other schemes exploit the difference in attenuation between signals received from LED pairs to estimate distances [139], but they require a high modulation rate as well as an increased installation cost to synchronize the LEDs. Camera/vision-based positioning has also been studied in the context of smartphones acting as receivers [33, 34]. Most

studies rely on intensity modulation, but polarization-based modulation has been also considered for wearable applications [140]. Other works use legacy light fixtures such as fluorescent lights [34, 141] to provide indoor localization. These methods exploit existing infrastructure, but fluorescent lights are energy inefficient, contain mercury and have reduced life cycle. All the methods above share the common idea that there exist a few static light fixtures, usually in the ceiling. Our work removes the need of static light fixtures and reduce the required LED luminaries to one. Our method also assumes *full mobility* with arbitrary Lambertian orders of the LED luminary, and no pre-recorded environmental/fingerprinting information is needed.

There has been some related work exploiting angle information. LiPro [132] is a system that uses a single light for localization. A smartphone, acting as receiver, has to be manually rotated around three axes, in a way similar to the calibration of a magnetic sensor. Another localization system is proposed in [133], describing a method using both angle-of-arrival principles and free-space attenuation of light signals to estimate the position of a mobile device. The method adopts multiple PDs installed on a ceiling acting as receivers, extending the proofs in [131]. Localization is done for a mobile device made of multiple LEDs fixed at different angles acting as a transmitter. The method is refereed to as *angle diversity transmitter* (ADT). Yasir et. al. [131] also exploit the Lambertian model of luminaries for localization. Both single or multiple reference points can be used to localize an object. The case of single fixed LED and multiple adjacent PDs has also been covered in [129], which estimates position using the difference of power received by multiple tilted optical receivers. However, in all the works described above there always exists reference points that are stationary. In our approach, we build on the works [131–133] by proposing, modeling and demonstrating a localization method with *nodes that can move freely* and with a single reference point (light). We proved that fully mobile systems require their own math and that, contrary to the simpler setup faced in indoor scenarios, the uniqueness of a solution is not guaranteed.

For fully mobile contexts, visible light that uses laser-ranging could be applied as well for distance computation. In particular, time-of-flight ranging sensors are gaining attention as a means to measure distances and being also available with small form factors. For instance, the recent VL53L1 [142] can accurately measure distances up to 4 meters. Economical laser lights have stringent safety regulations (Europe: EN 207, US: ANSI Z136) that limit the output power to only a few mW. In this work, we compute 2D relative position using standard LED sources, which are widely available.

5.8. Conclusion

We proposed a framework to compute the relative position of objects using LEDs. The method allows nodes to move freely in any direction and it works independently

of the surrounding environment. Our approach requires a single light source and we derive close-form expressions to obtain unique localization solutions in most cases. The model is validated in different scenarios, and simulations are performed to obtain additional insights about the performance of our positioning method. Experimental results demonstrate the high accuracy of our model and also highlight the key factors affecting this accuracy. Currently, our model can compute the relative locations in a two-dimensional space and we envision that it can be applied in applications for manned and unmanned vehicles such as motorbikes and robots.

"Never limit your vision based on your current resources"

Michael Hyatt (former chairman and CEO of Thomas Nelson)

6 Filtering Visible Light Reflections with a Single-Pixel Photodetector

Artificial lighting is everywhere, from the light bulbs on the ceilings to car headlights. It is expected that, by 2025, Light Emitting Diode (LED)s will account as 98% of the lighting [143]. This trend makes Visible Light Communication (VLC) an attractive technology for data transmission and, more recently, for locating objects and people using Light-based Positioning Systems (LPS). LPS are gaining significant attention from industry and the scientific community due to its high accuracy. Broadly speaking, LPS can be divided into two categories depending on the type of optical receiver they use: LPS relying on Photodiode (PD) and LPS relying on image sensor (camera) [33,130,144–146]. PDs provide higher throughput and energy-efficiency than cameras, and thus, they are a better choice for wearable devices (low energy) and vehicular networks (high throughput). Image sensors are popular due to their widespread availability in smartphones, but they have a reduced data rate (only kb/s rather than Mb/s or more [147]) and a higher energy cost (in the order of 300 mW rather than a few mW [148]).

The area of LPS for smartphones is relatively mature and there are products already in the market [149], whereas LPS in wearable devices and vehicular networks have encountered additional difficulties. Compared to cameras, the main constraint of PDs is its sensitivity to interference caused by optical reflections coming from the surroundings made of different type of surfaces and materials. All the incoming optical rays sum up at the PD due to operating as a single-pixel [150]. This implies that reflections affect LPS that work with a single-pixel PD as receiver, because the majority of LPS relies on the Received Signal Strength (RSS) resulting from all the light components [151]. Therefore, LPS relying on PD are affected by location errors (from a few cm to 1 m) in real deployments [150, 152–154]. Image sensors can instead solve optical interference by exploiting multiple pixels [155].

Discerning the Line-Of-Sight (LOS) path for positioning using a single-pixel PD receiver is not easy. Estimating the Channel Impulse Response (CIR) to infer the LOS and reflected Non Line-Of-Sight (NLOS) paths requires a powerful Analog-to-

Digital Converter (ADC) and high processing capabilities on the receiver side, which may not be available or desired [156]. Reflected paths have been ignored in recent experimental works for LPS relying on PD [5, 130, 144, 157], and reflections have been characterized only from a theoretical point of view [158–160]. However most of works are limited to simulations of simple static indoor scenarios [150, 152, 153, 161, 162] and the detection/discrimination of NLOS components has been done exploiting quite complex schemes as multiple sources [163] or multiple receivers [157, 164].

Against this background, this work presents a technique to discern and filter reflected light paths using a single-pixel PD. Our key idea is to avoid the need of the CIR knowledge and to propose a low-complexity practical approach based on time series of the RSS collected with a low-cost PD. This trade-off is possible in our scenarios of interest (vehicular or wearable networks) because (i) nodes are mobile, and thus, changes in reflection occur over a short period of time, (ii) the reflections from materials have unique statistical properties that can be exploited to filter out NLOS components, and (iii) our scenarios can accept some marginal delay, which allows us to work with short time series.

All in all, the main contributions of our work are:

- Section 6.1: system model containing the motivation and a dedicated experimental testbed that allows making controllable and realistic down-scaled characterization of visible light multipath in low-cost receivers.
- Section 6.2: a statistical method to identify and discriminate NLOS components with limited sampling rate of the receiver.
- Section 6.3: a decision tree algorithm that can run in low-cost receivers and that uses only two features, yet it can differentiate between different types of reflections.
- Section 6.4: a new algorithm for removing NLOS components and an experimental evaluation that shows that our method can improve the accuracy of LPS up to 93%.

6.1. System Model

6.1.1. Accuracy of positioning techniques with reflected light

As stated before, positioning techniques with single-pixel PDs typically rely on RSS-based estimation algorithms. These solutions exploit the so-called *inverse-square law*, which characterizes the relation between distance and RSS in pure LOS scenarios as:

$$\mathrm{RSS} \propto \frac{1}{d^2}, \tag{6.1}$$

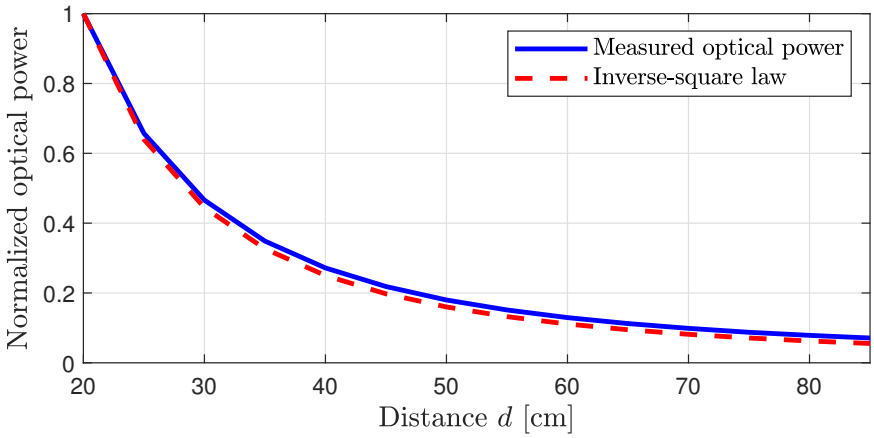

Figure 6.1: Comparison of inverse-square law with respect to RSS measurements using low-cost PD receivers in ideal LOS links.

where d is the distance between Transmitter (TX) and Receiver (RX). Figure 6.1 validates the application of the inverse-square law in VLC by comparing (6.1) to a simple practical experiment where the optical received power is measured for multiple d values. This good match has been the foundation behind the uptake of LPS in the last few years [151]. However, the monotonic curve of Fig. 6.1 becomes less predictable and more noisy if a fraction of the light does not reach the RX through the LOS link, i.e., it comes from reflected paths.

Since reflections are different for every reflector material, in order to quantify the impact of reflections, we first consider the ideal environment (considering only the LOS component), and subsequently, we include the presence of light reflected by different materials: shiny glass (glass), grey satin metal (metal), chipboard (rwood), plywood (swood) and foam core (wfoam). We use our experimental setup presented in Fig. 6.2 where the TX and the RX are at a fixed distance, and then we estimate, by using (6.1), the TX-RX distance based on the power received.

Figure 6.3 evaluates the TX-RX distance error with respect to the distance of the reflective material from the LOS link. As expected, the smallest error is obtained when there is only a LOS component (i.e. no reflective material). Depending on the reflective material and its distance, the accuracy of the LPS may change.

6.1.2. Basic intuition behind the proposed approach

The detection of NLOS components could be performed in systems with an expensive receiver with high gain-bandwidth product that samples quickly and have enough processing capabilities to compute the CIR [165]. This requires a very high sampling rate at the receiver, which is not feasible in low-end systems.

Instead of computing the CIR, we exploit the fact that, in a dynamic environment, a

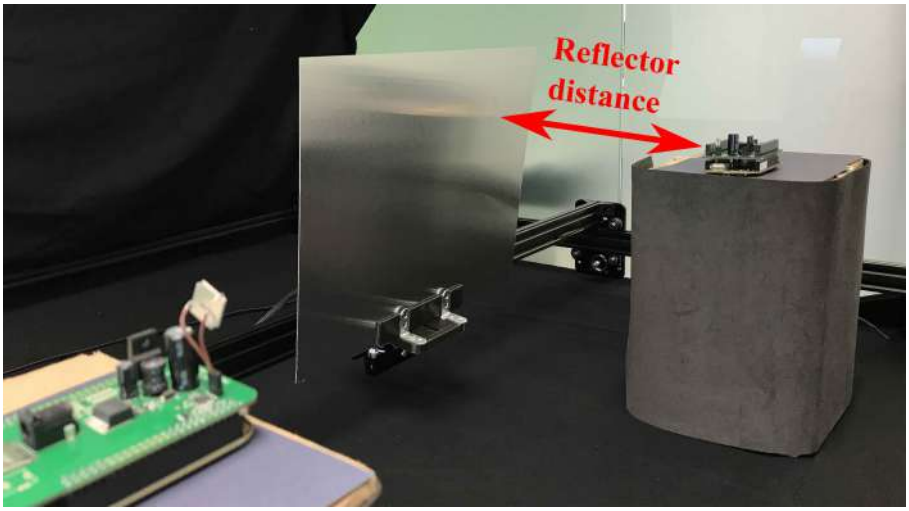

Figure 6.2: Testbed.

reflector does not appear suddenly, but enters the illuminated area progressively and the reflections are received by the PD at a certain speed (creating the NLOS component). This phenomenon is represented in Fig. 6.4[1] that, indeed, is the setup configured in this work for addressing LPS problems in mobile scenarios. Using past data, the receiver first observes an RSS variation due to the transition between the LOS and "LOS+NLOS" conditions and, subsequently, a transition between the "LOS+NLOS" and LOS conditions (when the reflective material is moving away). The shape of the time sequence depends on the material reflecting the light.

6.1.3. Testbed

Moving from the realm of theory and simulations into empirical evaluations under controlled circumstances requires a design and construction of a testbed. We require a testbed with modules capable of providing light-based positioning and mobility. On the one hand, the positioning requirement is addressed by the use of OpenVLC boards for LPS [5]. On the other hand, the mobility requirement is addressed by mounting these boards onto a structure that was originally designed for medium-scale laser cutting: the OpenBuilds ACRO movement structure [73] (see Fig. 6.2). We bundle these two systems together (mechanically, electrically and with software) to provide a precisely controlled environment. Our testbed can reproduce fully customizable down-scaled mobile scenarios and automatize the collection of raw data.

Without loss of generality, we move the reflecting materials and keep the positions of the TX and RX fixed (as the movement can be considered relative). The reflective

[1]We consider a mobile reflector and static TX and RX, but the same concepts applies to other cases as long as there is a relative movement among TX, RX and the reflector. For instance, mobile TX and RX, and a static reflector; or mobile TX and RX and a mobile reflector at different speeds.

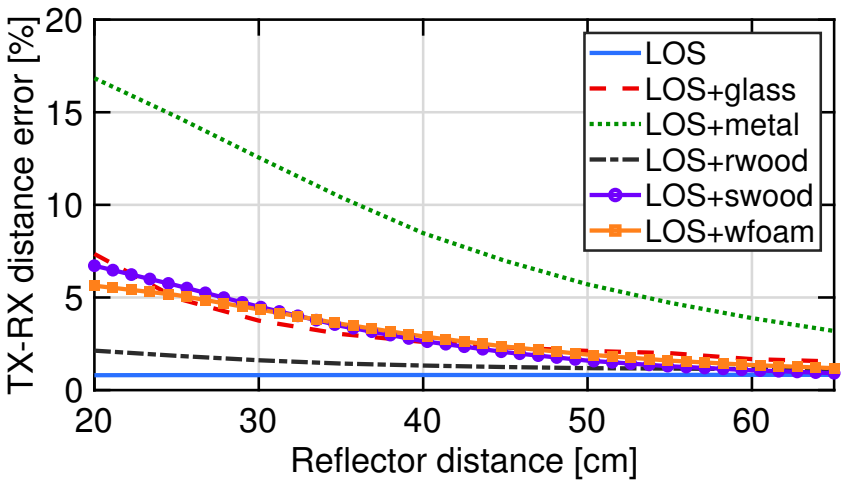

Figure 6.3: Experimental TX-RX distance error against reflector distance for multiple materials.

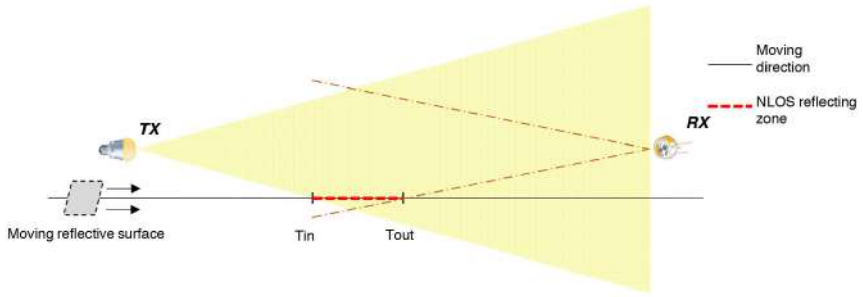

Figure 6.4: Schematic representation of reflections in the presence of mobility. The red dashed line depicts the reflecting zone.

surface is fixed on top of the mobile unit of the system (see Fig. 6.2), and it can move at a maximum speed of 42 cm/s and a maximum acceleration of 2 m/s^2 per axis. Table 6.1 reports the parameters of our testbed.

All of our experiments use the following setup. The TX and the RX are placed at a distance of 70 cm. The TX emits a fixed light intensity. The RX acquires a trace of 7.5 seconds (i.e. 1.5 million samples) for each movement. To consider various reflector distances, i.e., distances between the LOS path and the reflecting material, we start from a distance of 20 cm, and then move the reflective surface away in steps of 5 cm up to a distance of 65 cm (10 different distances). The same measurement is repeated 30 times in order to have statistical relevance. Thus 300 traces are collected for each material, plus another 300 for the "pure LOS" condition.

6.2. Identifying Reflections

As stated in Section 6.1, removing the effects of reflections in LPS implies identifying those reflections and discerning valid RSS changes. In this section, we tackle the identification problem for a series of RSS measurements.

Reflections cause peaks in light intensity, but these peaks can take widely different shapes depending on the properties of the reflecting material. Our first task is to define a minimal set of features to identify all such peaks. When light impinges upon a material, all reflected components are summed-up at the receiver. The final received intensity depends on two key properties: First, the reflection coefficient, i.e., the more reflective the material is (e.g., a mirror), the higher the light intensity reflected; Second, the material's smoothness, i.e., a very diffuse material (e.g., white paper) has a wide contribution in space because it reflects light in all directions. This type of materials lead to small but wide reflection peaks, as shown in Fig. 6.5a. On the other hand, a specular material (e.g. a smooth metallic plate) will only reflect light near the Snell angle [166] that leads to high and narrow peaks as represented in Fig. 6.5b. Furthermore, the shape of the peaks is also affected by the size of the reflective material, and by their relative speed (bigger materials and slower speeds lead to wider peaks). The number of materials and the variety of sizes and speeds would lead to a large number of peak shapes. To generalize the solution, we need an approach for identifying reflections that is material-, size- and speed-independent:

- Material-independent: The common property in all reflection peaks is that the power received by the NLOS component increases when the reflective object enters the illuminated area, reaches a maximum when the whole material is illuminated and starts decreasing when the material gets out of that area, action that can be seen in Fig. 6.4. This effect is represented in Fig. 6.5, where the received power is plotted for two different reflectors. Note that the maximum peak and the length of the transitions can change depending on the material, but the shape is similar and useful to identify a NLOS path.

- Size-independent: The plateau of the peak in Fig. 6.5 is determined by the size of the reflective material. Thus, if the study focuses only on the intrinsic properties of the received signal transitions, that is the upward and downward slopes highlighted in red in Fig. 6.5, the proposed approach may be considered as size-independent.

- Speed-independent: As long as the sampling rate is high enough compared to the object speed, speed variations will not modify the statistics of the received signal.

Each material and each surface type, due to its unique properties, reflects light differently. This is why this work analyzes different materials, trying to identify the reflection properties that are similar (and different) among them.

Table 6.1: Testbed parameters.

Parameter	Value
Nodes	2 OpenVLC boards
LED luminous flux	956 lm
LED half-power semi-angle	55°
Single-pixel PD area	7.02 mm^2
Single-pixel PD Field-of-View (FoV)	120°
Single-pixel PD spectral sensitivity	80 nA/lx
Sampling	12 bits @200 kHz
Movement grid	130 cm x 130 cm

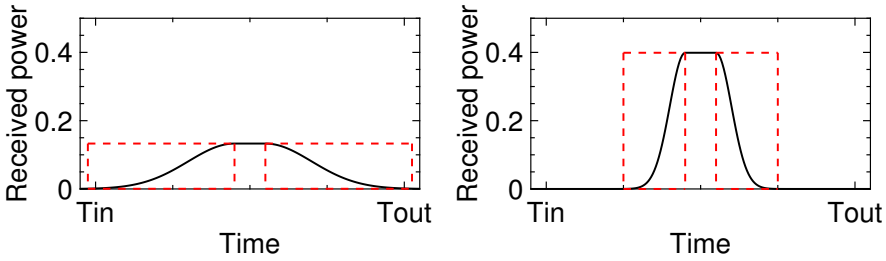

(a) Diffuse object. (b) Specular object.

Figure 6.5: Example of power received over time when diffuse (a) and specular (b) materials moves into an illuminated area.

6.2.1. Statistics considered

Let us consider a “Pure LOS” scenario and four “LOS+NLOS" scenarios composed by materials with distinct reflective features (glass, metal, plywood and chipboard). Using our experimental setup introduced in Section 6.1.3, Fig. 6.6 shows the raw traces and distributions for these cases. For the “Pure LOS” case (Fig. 6.6a and 6.6b), we observe, as expected, a normal distribution with $\mu = 238$ and $\sigma = 2.91$. For the “LOS+NLOS" scenarios (Fig. 6.6c-6.6j), we observe peaks that would lead to localization errors. To highlight the behavior hidden underneath the shot and thermal noises, a moving average filtering (white line) is shown. Except for the chipboard, which has a minimal effect due to its weak and diffuse reflection behavior, all the other materials have bimodal distributions. These bimodal distributions have different trends –for instance, glass produces a sharper reflection than metal or plywood (two well separated groups of values)– , but all distributions are clearly distinguishable from the “Pure LOS” case.

Our aim is to exploit the peculiar characteristics of the received signal variation

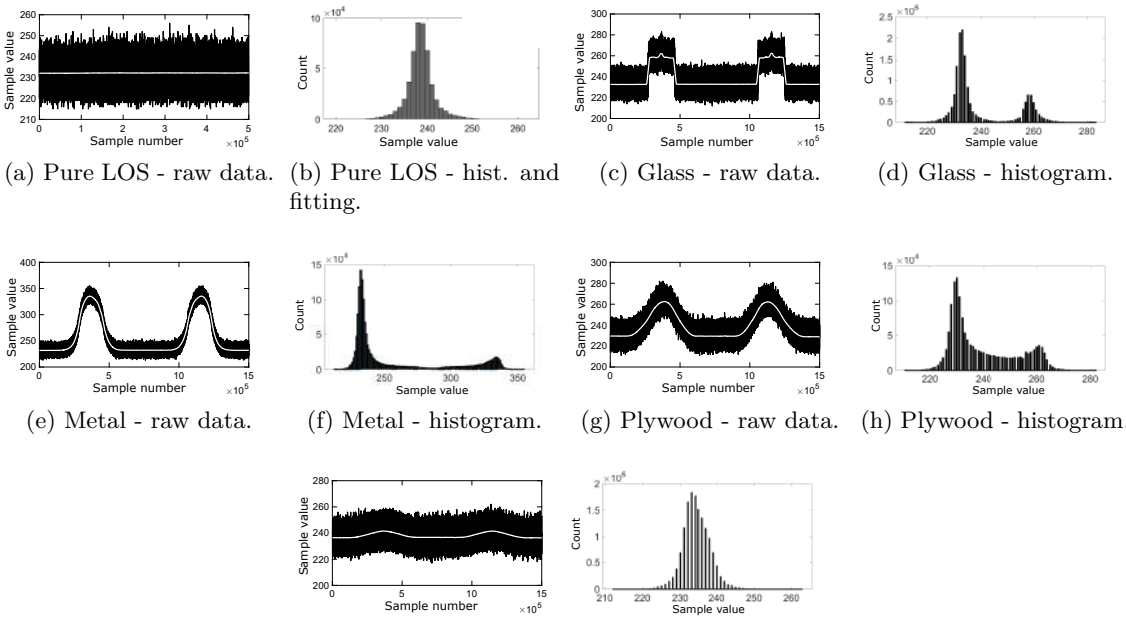

(a) Pure LOS - raw data. (b) Pure LOS - hist. and fitting. (c) Glass - raw data. (d) Glass - histogram.

(e) Metal - raw data. (f) Metal - histogram. (g) Plywood - raw data. (h) Plywood - histogram.

(i) Chipboard - raw data. (j) Chipboard - histogram.

Figure 6.6: Typical signal for "LOS+NLOS" condition with different reflector materials. Left side: raw data. Right side: histograms.

over time when light bounces on a given surface. With the constraint of using a given sequence of RSS samples collected with low-cost receivers, here we propose to look at the statistical moments up to the fourth-order as candidate observables to monitor the channel transitions of the NLOS component, and identify those material properties. Apart from the mean μ, and the standard deviation σ, we then consider:

- skewness (*skew*): third-order statistic that gives the amount and direction of skew (departure from horizontal symmetry).
- kurtosis (k): fourth-order statistic that defines how tall and sharp the central peak is, relative to a standard bell curve (i.e. normal distribution where $k = 3$) [167].

Since we are considering low-computational algorithms, Higher-Order Statistics (HOS) are not taken into account because the higher the moment, the harder it is to estimate. Moreover, larger sample intervals are required in order to obtain estimates of similar quality. Finally, due to the high noise power produced when using off-the-shelf devices, HOS are significantly less robust than lower-order statistics.

6.2.2. RoI of the received signal

The detection of the NLOS component in the received signal depends on the statistical properties of the signal itself. As observed previously, the information we want to exploit

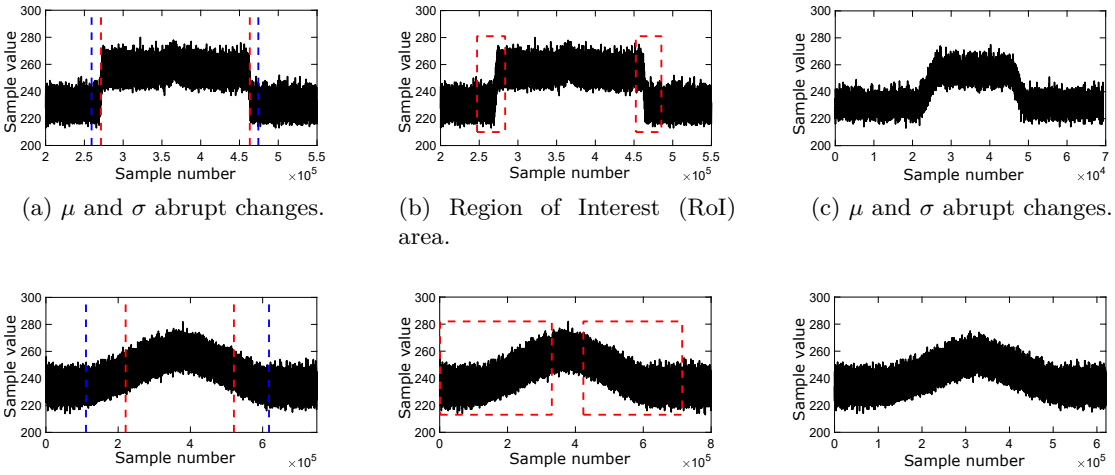

(a) μ and σ abrupt changes.

(b) Region of Interest (RoI) area.

(c) μ and σ abrupt changes.

(d) Cropped signal.

(e) Cropped signal.

(f) RoI area.

Figure 6.7: Typical RoI selection steps for glass (a-c) and white foam (d-f).

lies in the transition between LOS and "LOS+NLOS". Therefore, we compute the RoI to select the subset of samples in which we are interested by searching for changes in the signal [168]. Here, a change in the signal is given by a significant variation of the first statistical moment (i.e. the mean value μ). If μ is considered as a parameter for detecting changes, then a proper algorithm would return the index at which μ changes most significantly (e.g. index corresponding to the blue dashed line of Fig. 6.7a and Fig. 6.7d). Instead, if the statistical parameter is the standard deviation σ, the algorithm would return the index relative to the beginning/ending of the rising/falling edge, which is the area where σ changes the most (red dashed line of Fig. 6.7a and Fig. 6.7d). As examples of this procedure, the RoI selection steps for two different materials are shown in Fig. 6.7. Starting from the whole traces of Fig. 6.6, the signal is processed for detecting regions of interest [2]. Our algorithm works as follows. First, the signal is filtered with a smoothed filter for removing the noise, allowing a sharp RoI selection. Then the regions between the μ and σ changes, named "transition intervals", are selected. As a final step, in order to not alter the statistical properties of the signal and make the selection size-and-speed-independent, the RoI is extracted by taking the "transition interval" plus the same amount of this interval before and after it as highlighted in Fig. 6.7b and Fig. 6.7e. The result is the signals of Fig. 6.7c and Fig. 6.7f for glass and white foam, respectively, that will be used for the statistical analysis.

6.2.3. Assessment of observables to identify reflections

As we perform an statistical analysis, we show in Fig. 6.8 the trend of each observable (mean, standard deviation, skewness and kurtosis) over step numbers, which represent the

[2] With a double passage back and forth, there are two RoIs.

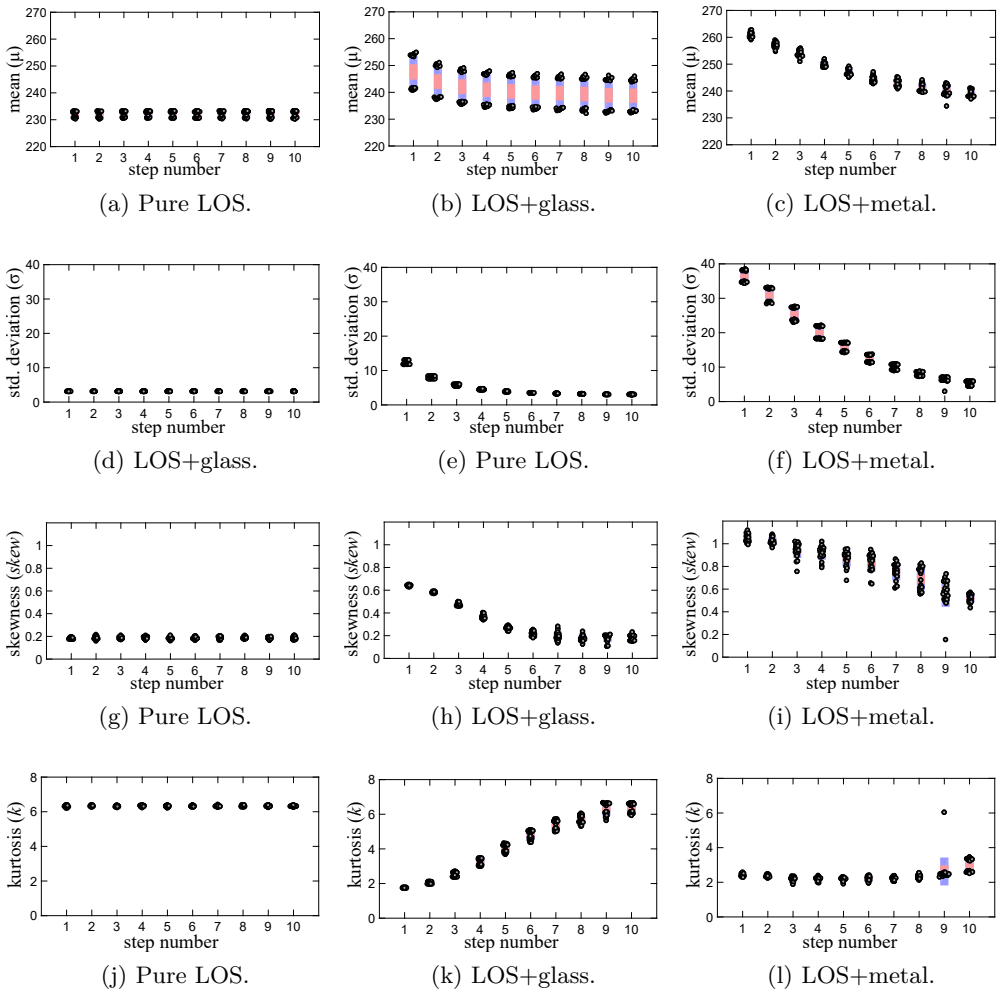

(a) Pure LOS. (b) LOS+glass. (c) LOS+metal.

(d) LOS+glass. (e) Pure LOS. (f) LOS+metal.

(g) Pure LOS. (h) LOS+glass. (i) LOS+metal.

(j) Pure LOS. (k) LOS+glass. (l) LOS+metal.

Figure 6.8: Statistical parameters μ (a-c), σ (d-f), skewness (g-i), kurtosis (j-l) against step numbers (i.e., distance to the reflector that goes from 20 to 65 cm in steps of 5 cm). Comparison among "Pure LOS", "LOS+glass" and "LOS+metal" are also represented.

reflector distances from 20 to 65 cm in steps of 5 cm. Note that only values for glass and metal reflectors are discussed, as they are the most representative cases of reflection types and typical materials in vehicular scenarios. By analyzing the results, we note that μ and σ follow the same trend for both sample materials. Their values decay exponentially along the distance and converge to the pure LOS case.

Something more peculiar can be stated for skewness and kurtosis. In the experiment involving glass, skewness tends to the LOS case value more quickly than the values corresponding to the metal case. Glass produces a more directional reflection than metal and only in a determined position along the movement. This is why the NLOS due to glass (specular-like) is lost after a few steps, becoming indistinguishable from noise.

Skewness and kurtosis values for all the scenarios acquired at step number 1 are

reported in Table 6.2. Note that the "Pure LOS" condition can be clearly identified by only exploiting skewness and kurtosis. Also, a high-level estimation of different types of reflector materials can be given, leaving space to an in-depth analysis for a material-type classification (i.e. specular-like and diffuse-like material groups), but this is out the scope of this paper and a possible future research line. Besides, note that all the "LOS+NLOS" scenarios tend to look the same and approach the "Pure LOS" condition for larger distances. Therefore, after a certain distance, the identification cannot be performed anymore. However, in those cases, the reflection due to a specific material or large distance leads to a very weak contribution and the generated noise into the system can be considered negligible (refer to Section 6.1.1).

6.3. Classification problem proposed

We start this section by solving the problem of discerning valid RSS changes, and then introducing a low-cost decision tree algorithm to learn which RSS sequences are caused by reflections.

6.3.1. Relative movement of devices

As seen until now, a quick object passage produces a reflection which can lead to a significant amplitude signal variation. As shown in Section 6.1.2 a relative linear movement between devices cannot be distinguished easily from the presence of a reflector object. It is necessary to investigate if the statistical methodology proposed in the previous section to detect reflections is still valid in the presence of relative movements between the TX and the RX. We want our method to detect reflections when they are present, but not to remove relative movement information, as it will lead to a higher error on the computed location, our final goal. The statistical properties of the received signal are related to the relative speed between TX and RX, and the sampling rate. Thus, if we want to make our metric speed independent, we should look into the number of samples obtained per unit of distance. Two acquisition systems will have the same behavior if the $\frac{speed}{sampling\ rate}$ ratio is the same. Therefore, keeping the number of samples per unit of area constant, the movement of fast objects can be studied by decreasing the sampling rate. In fact, increasing the number of samples per unit of distance does not necessarily lead to a better detector accuracy. A system needs sufficient samples to detect a change in the signal, but sampling the signal more will not improve the detection. Taking this into account, we perform tests at low speeds, so that the experimentation can be correctly performed with a low-sampling-rate device and safely carried out also in an indoor controlled environment, such as the one presented in Section 6.1.3.

We acquire a set of relative movements (approaching or moving away TX and RX) at five different speeds. Figure 6.9 shows four tests performed at different speeds (v1

Table 6.2: Example values of skew and k for step number 1.

Condition	skew	k
Pure LOS	0.18	6.32
LOS+metal	1.05	2.41
LOS+swood	0.67	2.08
LOS+glass	0.64	1.77
LOS+wfoam	0.55	2.20
LOS+rwood	0.17	4.92

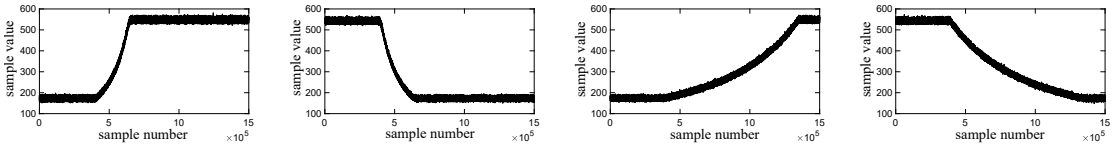

(a) Towards at speed v1 = 42 cm/s: $skew$ = -0.50, k = 1.34.

(b) Backwards at speed v1 = 42 cm/s: $skew$ = 0.78, k = 1.72.

(c) Towards at speed v2 = 10.6 cm/s: $skew$ = 0.82, k = 2.24.

(d) Backwards at speed v2 = 10.6 cm/s: $skew$ = 0.21, k = 1.43.

Figure 6.9: Typical received signal for a linear movement towards and backwards the transmitter.

= 42 cm/s and v2 = 10.6 cm/s). We see that raw data look specular as the reflective object approaches or moves away at the same speed. Skewness assumes positive and negative values with respect to the direction of the movement, but overall both statistical parameters we consider have very similar values to those related to "LOS+NLOS" conditions (Fig. 6.8). Therefore a way for discriminating the relative movement of devices is needed. By supposing that dynamics change fast in mobile environments, a certain NLOS component may last over a short period of time. An option would be to introduce a time metric, but this would come at the cost of a speed-dependent measure, which is undesired. We instead propose a metric that is the sum of the difference between consecutive samples, denoted by the variable

$$isMov = \sum_{i=2}^{N} (x_i - x_{i-1}), \tag{6.2}$$

where x_i is the i-th sample of the received signal and N is the window signal length. Equation (6.2) gives us an intuition on how the system is at the end of the measuring window compared to the beginning. If they are very different, the system has changed, otherwise the system relative distance did not change significantly. This let us distinguish

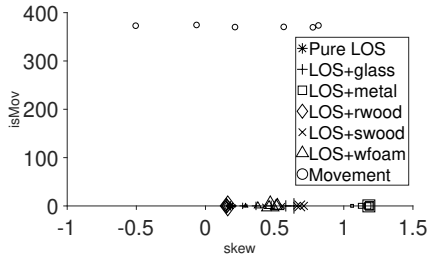

Figure 6.10: "isMov" parameter against skewness for different conditions along distances (mean values for each step).

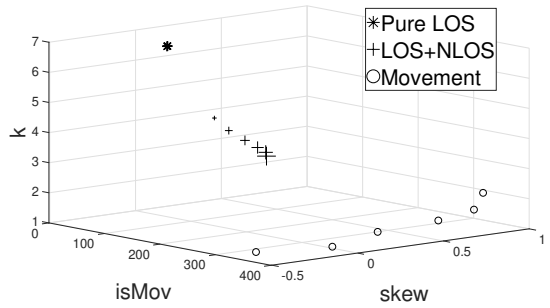

Figure 6.11: Scatter plot representing kurtosis, skewness and isMov parameter in case of "Pure LOS", "LOS+NLOS" and "movement".

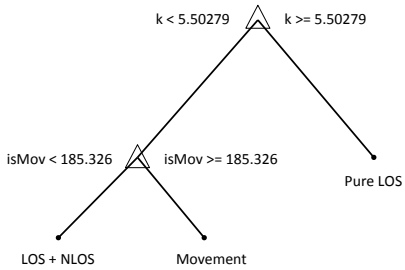

Figure 6.12: Coarse tree classification model.

between a signal variation which returns to the same value of the beginning after a fast perturbation (like the one coming from a reflection) and an increased final value. Typical values of "isMov" are high (always above 100) only for relative movements. This avoids the fact that relative movements of devices are confused with a LOS to "LOS+NLOS" transition (see Fig. 6.10).

6.3.2. Detection of NLOS reception

Contrary to the first and second moments, skewness and kurtosis are independent of signal intensity, and they can be efficiently used as features in a classification model where three classes are discerned: "Pure LOS", "LOS+ NLOS" and "Movement". Figure 6.11 represents a scatter plot with axes that are skewness, kurtosis and the "isMov" variable. The "LOS+ NLOS" class aggregates all the different materials in our tests. We observe that classes are well separated as the "Pure LOS" and "Movement" persist in a very limited area.

Therefore, we can introduce the algorithm to detectNLOS components in fully dynamic scenarios. In order to classify different scenarios, we choose a supervised machine learning classifier. We introduce a Decision Tree (DT) algorithm to let low-cost devices take advantage of NLOS recognition. In particular, here a coarse tree is used as it has a simple structure, fast prediction speed, small memory usage and low

Table 6.3: Confusion matrix, model trained on data set A and tested on data set B.

		Predicted condition		
		Pure LOS	LOS+NLOS	Movement
True condition	Pure LOS	100%	0%	0%
	LOS+NLOS	26.8%	73.2%	0%
	Movement	0%	0%	100%

computational cost [169, 170].

The classification model we exploit is shown in Fig. 6.12. As mentioned before, it is a coarse tree algorithm with only two decision nodes. It uses only kurtosis and "isMov" as predictors since, as demonstrated in Fig. 6.11, they are sufficient for discriminating the three labels (or classes): "Pure LOS", "LOS+NLOS" and "Movement". Skewness would be useful in distinguishing one material type from another, but is not the objective of this study.

For training the classification model, we take a data set (A) of 1320 collected observations (200 traces for each LOS and NLOS scenario and 120 for movement). 5-fold cross-validation is employed for avoiding an overfitted training. It has an accuracy of 100%. Finally, for testing the trained classification model, another data set (B) of 660 collected observations (100 for each LOS and NLOS condition, 60 for movement) is used.

Table 6.3 shows how the model performs in discerning among classes. In particular, "Pure LOS" and "Movement" are always detected correctly. Differently, the "LOS+NLOS" class suffers the condition in which the NLOS contribution is very poor, and the classification model fails to detect it, with an error of up to 26.8%. Therefore, as expected, there are conditions with low predictive accuracy. Nevertheless, it must be taken into account that the final objective is not simply to detect when there is an object reflecting, but to also remove that reflection.

6.4. Localization accuracy improvement by removing the NLOS component

The final goal of this work is to design a method to remove the NLOS components from received VLC signals in order to improve the LPS accuracy. Using the results obtained in Section 6.3, we design a NLOS removal algorithm. This algorithm detects and corrects the NLOS components that appear in mobile environments.

As input, the algorithm takes the raw data from the ADC. This data is statistically analyzed and the NLOS is detected using the method introduced in Section 6.3. Then, the NLOS component is removed as follows.

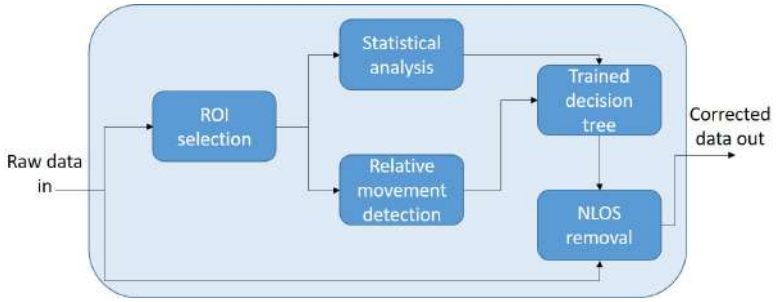

Figure 6.13: Block diagram of the algorithm implemented.

6.4.1. NLOS removal

After the signal is received, it is analyzed for detecting the presence of a NLOS component. In order to remove it, the algorithm performs the following steps that are also depicted in Fig. 6.13:

- First, it detects the RoI where the NLOS component is located. In order to do so, it computes the average value of the whole signal and it detects what part deviates from it. This is selected as the RoI (see Section 6.2.2).
- The RoI is processed by computing the variable "isMov" as (6.2) indicates in order to determine if there is a relative movement in it. Then it is statistically analyzed using the metrics explained in Section 6.3.
- The results of this analysis are passed to the trained decision tree, that decides if there is a NLOS component or not.
- If the decision tree decides that a NLOS component is present, it is filtered out from the received signal. This is done by subtracting the RoI average value μ_{ROI} from the RoI and adding the average value of the rest of the received signal μ_{RX}:

$$ROI' = ROI - \mu_{ROI} + \mu_{RX}. \tag{6.3}$$

The NLOS removal algorithm works only on the RoIs as the rest of the signal is not being filtered. An example of how the algorithm performs is shown in Fig. 6.14 for a "LOS+metal" condition. As can be seen, the incoming raw data (blue curve) with a significant RSS variation is corrected into a reflection-less signal (red curve) following (6.3).

6.4.2. Algorithm Evaluation

This section analyzes the performance of the NLOS removal algorithm. In particular, using the same data set and the received-power based localization algorithm [5] used in

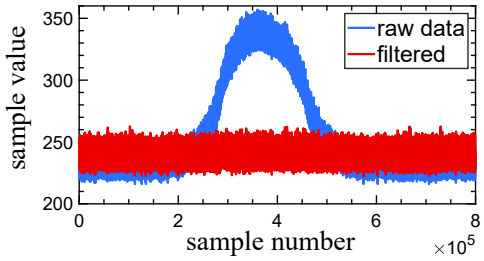

Figure 6.14: NLOS removal in a "LOS+metal" condition.

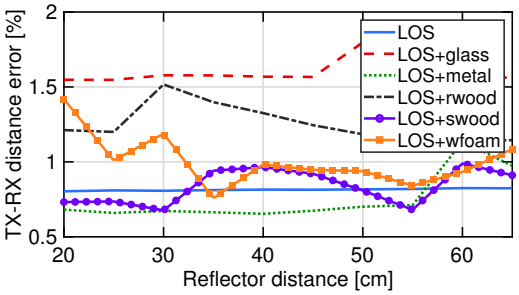

Figure 6.15: Distance calculation error after applying correction algorithm (glass = shiny glass, metal = grey satin metal, rwood = chipboard, swood = plywood, wfoam = foam core).

Table 6.4: Position accuracy (worst cases).

Condition	W/o correction [cm]	With correction [cm]	Improvement [%]
Pure LOS	0.58	0.58	0
LOS+glass	5.15	1.32	74.3
LOS+metal	11.79	0.8	93.2
LOS+rwood	1.5	1.06	28.8
LOS+swood	4.7	0.7	85.3
LOS+wfoam	3.94	0.99	74.8

Section 6.1.1, we compute the relative localization of an object using light. The position of the object was computed both with and without the correction algorithm.

As it can be seen in Fig. 6.15, the distance error[3] is always below 2% and the accuracy may increase up to 93% as shown in Table 6.4. The reason why the error stays low even if the relative power of the NLOS component increases, is that the removal algorithm works better with higher NLOS components. Indeed, when the NLOS component is weak (i.e. the reflective material is far), the RoI selection does not work very well, but the effect on the localization is, as seen in Section 6.1.1 (Fig. 6.3), negligible.

As an additional result, Fig. 6.16 shows the CDF of the estimated TX-RX distance error. Note that a significant improvement is achieved when the reflection correction algorithm is used. Indeed, the positioning error can reach a sub-centimeter accuracy with a probability of 80% when filtering is applied.

[3]Percent distance error means the positioning error w.r.t. the ground truth, that is fixed and equal to 70 cm.

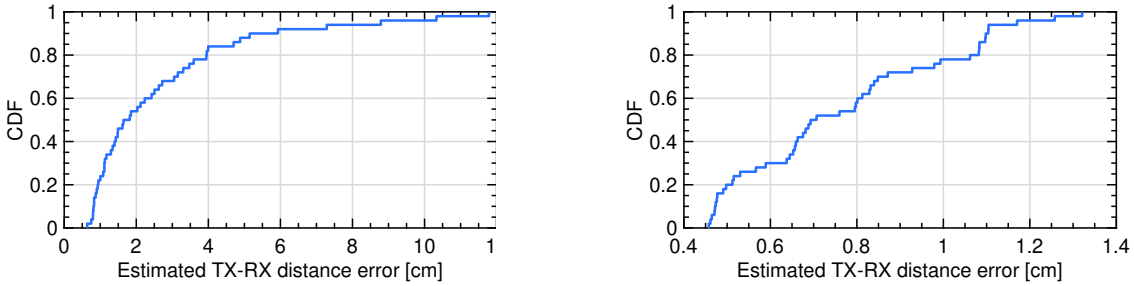

(a) Without correction algorithm.

(b) With correction algorithm.

Figure 6.16: Cumulative Distribution Function (CDF) of the estimated TX-RX distance error. Plots aggregate all the steps and all the conditions.

6.5. Related Work

Authors in [150, 156, 160, 161] perform an efficient characterization of the NLOS components to model the CIR. However these solutions come at the cost of high computational resources and the use of specific software even for outdoor applications [171]. Proper countermeasures are then needed for compensating the undesired effect of reflections in VLC [172] and for avoiding the positioning error in LPS [155, 157, 173, 174].

Among other solutions, LOS and NLOS identification algorithms make use of multiple LEDs and PDs [163,164,174], which is a condition that may not be available. Some others works opt to use image sensors for canceling the effect of reflections [145,146], but it comes with much higher energy consumption cost and lower data rate of communication.

Our work is an experimental-based analysis and the solution we introduce here exploits a low-cost system allowing a NLOS cancellation in both specular and diffuse reflection conditions. It implies a statistical study of time series of the received signal instead of a way-more-complex and costly CIR estimation. Our method, by using a single-pixel photodetector, offers a solution for overcoming the reflection effect and guaranteeing a better accuracy than more sophisticated systems do.

6.6. Conclusion

Motivated by the errors caused by reflections in LPS, this paper presented a new method to filter out multipath components that is suited to be implemented in low-cost positioning receivers that use a single-pixel photodetector. By analyzing the statistical properties of NLOS components and exploiting the results obtained from a custom-built and fully controllable testbed, we have performed extensive experimental campaign that drove us to design a low-computational cost algorithm to effectively remove the NLOS components, and then to improve the LPS accuracy in up to a 93 %. Going forward, we

envision that this approach can inspire follow-up system research towards this promising direction.

"Ends are not bad things, they just mean that something else is about to begin"

C. JoyBell C

7 Conclusions

In this thesis, we have explored the intersection between Visible Light Communication (VLC) and Internet-of-Things (IoT) applications. We recognize the synergies of using VLC for IoT, but also acknowledge that the combination of both technologies needs to be further studied to find when and how using VLC is beneficial for an IoT deployment. Moreover, we believe that to have the highest possible impact, the research work done needs to take into account practical aspects, with experimental implementations and evaluations. This work is a step towards that direction.

In Part II, we have presented the new OpenVLC, a open-source, software-based, low-cost VLC platform. We have shown the main goals that we wanted to reach and the design decisions that were taken along the process, and how Software (SW), Firmware (FW) and Hardware (HW) have been concurrently designed to improve the platform capabilities. The data pipeline scheme has been introduced, and we have demonstrated that offloading time-sensitive operations to the Programmable Real-time Unit (PRU) can significantly increase the data rate between VLC devices. Changes in the HW have reduced the form factor, and allowed faster physical layer symbol rates. OpenVLC has been the backbone of this thesis. One of the greatest challenges has been to make it flexible, with a general-purpose interface, which has proven to be the correct decision over time, as it has allowed us to speed up the design process for the rest of prototypes and deployments of this thesis. We tested the versatility of the platform by using it to overcome significant scientific and technical challenges in a wide range of scenarios.

Using OpenVLC as a framework, the main contributions of this thesis have been presented in Part III, where we showed challenges that VLC can help to overcome both in the infrastructure and the IoT devices.

Chapter 3 analyzed the behavior of VLC systems where luminaries have been densely deployed. To share the resources in an optimum way and minimizing mutual interference, a power budget optimization algorithm was proposed, studied, and implemented in this Multiple-Input-Single-Output (MISO) system. To coordinate the different transmitters, a method to synchronize them using the Non Line-Of-Sight (NLOS) component of

the light was proposed. In our opinion, this work was a step-forward into realistic dense deployments and showed that well-known problems as the synchronization between Transmitter (TX) could be leveraged using the visible light signal itself.

Chapter 4 showed the benefits of using VLC as a receiver in a battery-free device. Together with backscattering, both technologies allowed creating bidirectional communication, leveraging some of the problems that using a unique technology has, to achieve continuous communication at low transmission rates under realistic scenarios. This is probably one of the most exciting results of the whole thesis, as it shows the combination of two very different technologies that result in a system that leverages most of the problems that using a single technology has.

Chapter 5 showed how Visible Light Positioning System (VLPS) can be achieved by only having a pair or transmitter-receiver in a dynamic scenario, only adding a reference direction using an electronic compass. We mathematically proved that in all the cases, the correct result is within the set of results, and under some scenarios, the result is always unique. We simulated the system and implemented it, showing that such a configuration is not only possible but doable. This work is the most analytical one and shows simple models as the Lambertian model of the Light Emitting Diode (LED) can be used as tools to design innovative solutions.

Finally, in Chapter 6, we presented a solution to detect NLOS components of a VLC signal only using simple statistics and a Photodiode (PD). Although quite specific, this work shows that even easily computable values such as the statistics from first- to fourth-order are able to detect fast NLOS components of the signal. We collected received samples for different materials and movement conditions and trained a decision tree able to detect these components and remove then, increasing the localization accuracy.

In summary, this thesis showed the capabilities of VLC for IoT using custom HW, studied some critical aspects for realistic deployments and proposed future paths for continuing researching around such technology.

Appendices

A Proof of Propositions

A.1. Proof of Proposition 5.3.2

Proof: Since the TX does not rotate, we then have $\alpha_{\text{tx}}^1 = \alpha_{\text{tx}}^2$ and $\psi_1 = \psi_2$. From Eq. (5.8), we have

$$\theta_i = \psi_1 - \pi - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^i, \quad i \in [1, 2] \tag{A.1}$$

Substituting Eq. (A.1) into Eq. (5.12), we get

$$\begin{cases} P_r^1 = \frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1) \cdot \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1 - \pi) \cdot g_r + N \\ \quad = -\frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1) \cdot \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1) \cdot g_r + N \\ P_r^2 = \frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1) \cdot \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^2 - \pi) \cdot g_r + N \\ \quad = -\frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1) \cdot \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^2) \cdot g_r + N \end{cases}$$

which leads to

$$\begin{aligned} \frac{P_r^1 - N}{P_r^2 - N} &= \frac{\cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1)}{\cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^2)} \\ &= \frac{\cos(\psi_1 - \alpha_{\text{tx}}^1)\cos(\alpha_{\text{rx}}^1) - \sin(\psi_1 - \alpha_{\text{tx}}^1)\sin(\alpha_{\text{rx}}^1)}{\cos(\psi_1 - \alpha_{\text{tx}}^1)\cos(\alpha_{\text{rx}}^2) - \sin(\psi_1 - \alpha_{\text{tx}}^1)\sin(\alpha_{\text{rx}}^2)} \\ &= \frac{\cos(\alpha_{\text{rx}}^1) - \tan(\psi_1 - \alpha_{\text{tx}}^1)\sin(\alpha_{\text{rx}}^1)}{\cos(\alpha_{\text{rx}}^2) - \tan(\psi_1 - \alpha_{\text{tx}}^1)\sin(\alpha_{\text{rx}}^2)} \end{aligned} \tag{A.2}$$

From Eq. (A.2), we can easily derive the closed-form expression of ψ_1:

$$\psi_1 = \arctan\left(\frac{(P_r^1 - N)\cos(\alpha_{\text{rx}}^2) - (P_r^2 - N)\cos(\alpha_{\text{rx}}^1)}{(P_r^1 - N)\sin(\alpha_{\text{rx}}^2) - (P_r^2 - N)\sin(\alpha_{\text{rx}}^1)}\right) + \alpha_{\text{tx}}^1$$

Substituting the expression of ψ_1 into Eq. (5.8), we can easily obtain the expression of θ_2. ■

A.2. Proof of Proposition 5.3.2

Proof: Since the RX does not rotate, we then have $\alpha_{\text{rx}}^1 = \alpha_{\text{rx}}^2$ and $\theta_1 = \theta_2$. From Eq. (5.8), we now that

$$\theta_i = \psi_1 - \pi - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1, \quad i \in [1, 2] \tag{A.3}$$

Substituting Eqs. (5.10) and (A.3) into Eq. (5.12), we get

$$\begin{cases} P_r^1 = \frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1) \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1 - \pi) \cdot g_r + N \\ P_r^2 = \frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1 + \alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1) \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1 - \pi) \cdot g_r \\ \qquad +N \end{cases}$$

which leads to

$$\begin{aligned} \frac{P_r^1 - N}{P_r^2 - N} &= \frac{\cos^m(\psi_1)}{\cos^m\left(\psi_1 + \alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1\right)} \\ &= \left(\frac{\cos(\psi_1)}{\cos(\psi_1)\cos\left(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1\right) - \sin(\psi_1)\sin\left(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1\right)} \right)^m \\ &= \left(\frac{1}{\cos\left(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1\right) - \tan(\psi_1)\sin\left(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1\right)} \right)^m \end{aligned} \tag{A.4}$$

From Eq. (A.4), we can easily derive the closed-form expression of ψ_1:

$$\psi_1 = \arctan\left(\frac{\cos\left(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1\right) - \sqrt[m]{\frac{P_r^2 - N}{P_r^1 - N}}}{\sin\left(\alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1\right)} \right) \tag{A.5}$$

Substituting Eq. (A.5) into Eq. (5.8), we can easily obtain the expression of θ_2, as given in Eq. (5.14). ■

A.2.1. Proof of Proposition 5.3.3

Proof: In the full-rotation scenario, we have

$$\begin{cases} P_r^1 = -\frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1) \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1) \cdot g_r + N \\ P_r^2 = -\frac{(m+1)A_{\text{RX}}}{2\pi d^2} \cos^m(\psi_1 + \alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1) \\ \qquad\qquad \cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^2) \cdot g_r + N \end{cases}$$

which leads to

$$\frac{P_r^1 - N}{P_r^2 - N} = \frac{\cos^m(\psi_1)\cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^1)}{\cos^m(\psi_1 + \alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1)\cos(\psi_1 - \alpha_{\text{tx}}^1 + \alpha_{\text{rx}}^2)} \tag{A.6}$$

Let us first make some notations:

$$\begin{cases} b = \frac{P_r^1 - N}{P_r^2 - N} \\ c = \alpha_{\text{rx}}^1 - \alpha_{\text{tx}}^1 \\ e = \alpha_{\text{tx}}^2 - \alpha_{\text{tx}}^1 \\ h = \alpha_{\text{rx}}^2 - \alpha_{\text{tx}}^1 \\ k = \cos e - \tan(\psi_1) \cdot \sin e \end{cases} \tag{A.7}$$

Then Eq. (A.6) becomes

$$\begin{aligned} b &= \frac{\cos^m(\psi_1)}{\cos^m(\psi_1 + e)} \frac{\cos(\psi_1 + c)}{\cos(\psi_1 + h)} = \frac{1}{k^m} \frac{\cos c - \tan(\psi_1) \sin c}{\cos h - \tan(\psi_1) \sin h} \\ &= \frac{1}{k^m} \frac{\sin e \cos c - \cos e \sin c + k \sin c}{\sin e \cos h - \cos e \sin h + k \sin h} \\ &= \frac{1}{k^m} \frac{\sin(e - c) + k \sin c}{\sin(e - h) + k \sin h} = \frac{1}{k^m} \frac{\sin(\alpha_{\text{tx}}^2 - \alpha_{\text{rx}}^1) + k \sin c}{\sin(\alpha_{\text{tx}}^2 - \alpha_{\text{rx}}^2) + k \sin h} \end{aligned} \tag{A.8}$$

Let $x = \sin(\alpha_{\text{tx}}^2 - \alpha_{\text{rx}}^1)$ and $y = \sin(\alpha_{\text{tx}}^2 - \alpha_{\text{rx}}^2)$, then the above equation can be re-written as

$$b \cdot \sin(h) \cdot k^{m+1} + b \cdot y \cdot k^m - \sin(c) \cdot k - x = 0 \tag{A.9}$$

■

B Video demo with OpenVLC

Visible Light Communication (VLC) has been pointed in the last years as a new technology to connect to the Internet [27] [28]. In the last few years there has been several applications which use VLC: communication with toys [31], human sensing [30], indoor localization [33] [34], mobile interaction [32], and passive VLC [35] [36].

Nevertheless, the access to VLC technology is not trivial as communications using light in the free space have physical properties immensely different from other Radio Frequency (RF). This makes it difficult to start the research in VLC systems. This is why we introduced OpenVLC at the VLCS'14 workshop [26]. OpenVLC is a Open Source VLC platform that we have improved since then [1] [2].

Although there has been other projects [175] [176] that have transmitted video using light, this is, to the best of our knowledge, the first work that tries to minimize the cost of the system, while trying to have a throughput sufficiently high for video transmission.

In this demo, we show the capabilities of the latest OpenVLC version. Due to the increase in throughput compared to previous versions, we are able, for the first time, to transmit video through VLC using the OpenVLC platform.

B.1. System view

The system that has been designed and implemented has several components that are described here separately. Due to space constraints, no picture of the setup is shown but a video URL is provided in Sec. B.2.

B.1.1. VLC link

The VLC link is the main component whose performance is shown in this demo. The VLC link is created using two BeagleBone Blacks (BBB) with two OpenVLC capes on top. One of the BBBs is used as VLC transmitter (TX) and the other as VLC receiver (RX). In the current version, the maximum UDP throughput of this link 415 kb/s.

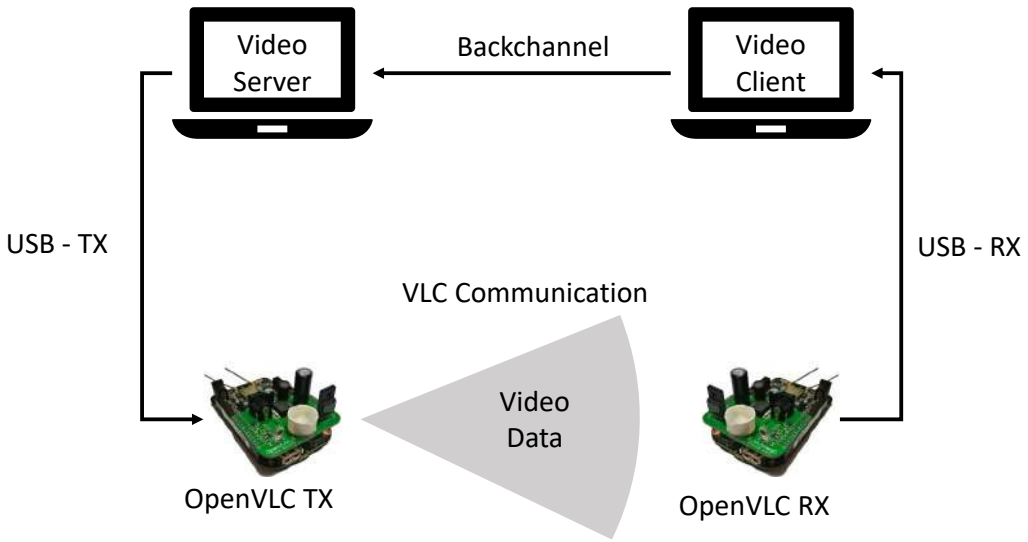

Figure B.1: System scheme. Each black line represents a different communication link.

B.1.2. Video transmission and playback

In order to transmit and reproduce the video we use an HTTP stream. One computer behaves as an HTTP video server and the second as an HTTP client. Both computers are running Ubuntu 18.0.4. For the server, we use an Apache 2.4.29. The video is stored as a Transport Stream file (.ts) as it is the best format for our case scenario. It is compressed in MPEG format using FFmpeg and ready to be transmitted, minimizing workload on the computers.

B.1.3. Network connections

The VLC interface that this demo uses is OpenVLC1.3. OpenVLC has increased the throughput to more than 400 kb/s. This allows users of OpenVLC to have a data-rate high enough to, for the first time, be able to transmit video in real-time.

The transmission between the computers and the BBBs where OpenVLC software is running is done through USB using TCP/IP. The OpenVLC software and firmware then handles the hardware control and packet TX and RX.

OpenVLC1.3 creates a VLC simplex connection. This follows the current trend where VLC hybrid system are usually deployed with VLC as downlink and RF as uplink [71]. Nevertheless, as in any communication system, there are sometimes transmission errors. Without a backchannel, the video shows errors and the performance of the system may drop dramatically. For this reason, the nodes of the system have been prepared to transmit the control messages necessary for the retransmission of lost video packets. Messages in the backchannel are sent using Ethernet. A scheme of the system can be seen in Fig. B.1.

B.2. Results

Once the system has been correctly configured and all the devices set up, the video client (which is responsible for playing the video) makes an HTTP request to the server through the backchannel and then the server replies through VLC.

The video player has a buffer that is used to store the video. In case of an error, the buffer should be large enough to allow the video player to request the missing packets without disrupting the transmission. Nevertheless, if we voluntarily interrupt the connection, for example putting our hand on top of the LED, it can be seen how the video stops as soon as the buffer is depleted.

Multiple video qualities have been tested, from 480p to 720p. When no VLC link blockage is present, the system is able to transmit the videos without perceivable errors. Nevertheless, having a fixed buffer size, what we see is that although the amount of video in bits is the same, the buffered time is higher when the quality is lower.

All the code for running this demo can be found in the Github repository of OpenVLC. https://github.com/openvlc/OpenVLC. The video showing the demo of our system running a can be found at https://youtu.be/DeEhffECR5k.

The organization should provide a long table of 2 meters and a power connection. The BBBs, capes and computers will be provided by IMDEA Networks Institute.

B.3. Conclusion

In this demo we show the capabilities of the latest OpenVLC, which is now able to transmit real-time video using low-end embedded system. This can be done due to the improved communication throughput that OpenVLC allows. We hope that the new capabilities and easy access of OpenVLC will attract new researchers and developers.

C Communication and Press appearences

During my research as a pre-doctoral student, I had the chance to participate in several dissemination activities related to spreading the scientific world to the general population, primarily focusing on children in the school age. I also had several press appearances related to the same topic.

C.1. Week of Science Madrid (2017)

Every year, the region of Madrid organizes the week of science, a week where science is exalted, and scientists present interesting scientific topics to the general population. In 2017, I gave a lecture about gambling statistics in front of the general public.

Figure C.1: Talk at week of Science in Madrid.

C.2. Ciencia en Acción (2018)

Ciencia en Acción (Science in Action) is a dissemination contest that takes place in different locations of Spain. The main objective is to try to present scientific concepts to school kids in an easily understandable way. I presented a demo where Visible Light

Communication (VLC) was used to transmit data from a controller to a computer. The computer was running a game similar to Pang in it, and in that way, the kids could understand how and when the information was transmitted.

Figure C.2: Demo at Ciencia en Acción.

C.3. Demo in Aula Ciencia (2018-2019)

Every year, the region of Madrid presents its future career alternatives to anyone interested to know about them, although the main objective is to inform to last year highschool students. There, IMDEA Networks was part of the initiative, and we prepared 1 demo related with light in 2018 and two demos in 2019, one related with ElectroSense where we created and decoded an FM signal and another one with Virtual Reality where we show a demo of Non Line-Of-Sight (NLOS) component calculation in a 3D environment. I collaborated with the design, deployment and explanation of the experiments in Aula.

Figure C.3: Me presenting the demo "Play through light" to the IFEMA Director, the Director of the Madri+d Foundation for Knowledge and other authorities.

C.4. Talk in Aula Ciencia

In the same event, they set up a scenario where researchers could present any of their recent discoveries to the general audience. I lectured a talk about VLC, the evolution through history, and where we are today. I also presented my most recent work, and I did a small demo to show how it works.

Figure C.4: Talk at Aula Ciencia.

C.5. Short exposition in RTVE (spanish public national television)

Lab24 is a TV program where new developments in research and technology are shown and explained. For one of the programs, they came to IMDEA and recorded OpenVLC working and interviewed me about how VLC works and what it could be used for.

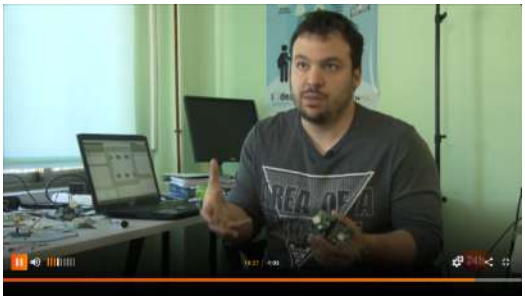

Figure C.5: Appearance at TV.

References

[1] A. Galisteo, D. Juara, Q. Wang, and D. Giustiniano, "OpenVLC1.2: Achieving higher throughput in low-end visible light communication networks," in *2018 14th Annual Conference on Wireless On-Demand Network Systems and Services, WONS 2018 - Proceedings*, vol. 2018-Janua, 2018.

[2] A. Galisteo, D. Juara, and D. Giustiniano, "Research in Visible Light Communication Systems with OpenVLC1.3," in *IEEE 5th World Forum on Internet of Things, WF-IoT 2019 - Conference Proceedings*, 2019.

[3] J. Beysens, A. Galisteo, Q. Wang, D. Juara, D. Giustiniano, and S. Pollin, "DenseVLC: A cell-free massive MIMO system with distributed LEDs," in *CoNEXT 2018 - Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies*, 2018.

[4] J. Beysens, Q. Wang, A. Galisteo, D. Giustiniano, and S. Pollin, "A cell-free networking system with visible light," *IEEE/ACM Transactions on Networking*, 2020.

[5] A. Galisteo, Q. Wang, A. Deshpande, M. Zuniga, and D. Giustiniano, "Follow that light: Leveraging LEDs for relative two-dimensional localization," in *CoNEXT 2017 - Proceedings of the 2017 13th International Conference on emerging Networking EXperiments and Technologies*, 2017.

[6] A. Galisteo, D. Juara, H. Cordobes, and D. Giustiniano, "Demo: Video transmission using low-cost visible light communication," in *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2019.

[7] A. Galisteo, H. Wu, Q. Wang, D. Juara, M. Zuniga, and D. Giustiniano, "Openvlc1. 2 for increased data rate with embedded systems," in *Proceedings of the 4th ACM Workshop on Visible Light Communication Systems*, 2017, pp. 33–33.

[8] R. Aversa, R. V. Petrescu, A. Apicella, and F. I. Petrescu, "The basic elements of life's," *American Journal of Engineering and Applied Sciences*, vol. 9, no. 4, pp. 1189–1197, 2016.

[9] R. J. Rapf and V. Vaida, "Sunlight as an energetic driver in the synthesis of molecules necessary for life," *Physical Chemistry Chemical Physics*, vol. 18, no. 30, pp. 20 067–20 084, 2016.

[10] E. N. Borza, "Alexanderâs communications," *Laourdas et Makarras*, pp. 295–303, 1977.

[11] E. T. ALVA, "Electric lamp," Nov. 4 1879, uS Patent App. 223,898.

[12] H. Riechert, "Lighting the 21st century: Isamu akasaki, hiroshi amano, and shuji nakamura received the 2014 nobel prize in physics for the development of the blue light-emitting diode," *physica status solidi (a)*, vol. 212, no. 5, pp. 893–896, 2015.

[13] "How Energy-Efficient Light Bulbs Compare with Traditional Incandescents." [Online]. Available: https://www.energy.gov/energysaver/save-electricity-and-fuel/lighting-choices-save-you-money/how-energy-efficient-light

[14] N. Sklavos, M. Hübner, D. Goehringer, and P. Kitsos, *System-level design methodologies for telecommunication.* Springer, 2013.

[15] Cisco, "https://www.cisco.com/c/dam/en/us/products/collateral/se/internet-of-things/at-a-glance-c45-731471.pdf."

[16] H. Burchardt, N. Serafimovski, D. Tsonev, S. Videv, and H. Haas, "Vlc: Beyond point-to-point communication," *IEEE Communications Magazine*, vol. 52, no. 7, pp. 98–105, 2014.

[17] "The free open source radio ecosystem Â · gnu radio." [Online]. Available: https://www.gnuradio.org/

[18] D. Tsonev, H. Chun, S. Rajbhandari, J. J. McKendry, S. Videv, E. Gu, M. Haji, S. Watson, A. E. Kelly, G. Faulkner *et al.*, "A 3-gb/s single-led ofdm-based wireless vlc link using a gallium nitride μ *led*,$''$ *IEEE Photonics Technology Letters*, *vol.* 26, *no.* 7, *pp.*637 − −640, 2014.

[19] Y. Wang, L. Tao, X. Huang, J. Shi, and N. Chi, "8-gb/s rgby led-based wdm vlc system employing high-order cap modulation and hybrid post equalizer," *IEEE Photonics Journal*, vol. 7, no. 6, pp. 1–7, 2015.

[20] H. M. Oubei, J. R. Duran, B. Janjua, H.-Y. Wang, C.-T. Tsai, Y.-C. Chi, T. K. Ng, H.-C. Kuo, J.-H. He, M.-S. Alouini *et al.*, "4.8 gbit/s 16-qam-ofdm transmission based on compact 450-nm laser for underwater wireless optical communication," *Optics express*, vol. 23, no. 18, pp. 23 302–23 309, 2015.

[21] D. Tsonev, S. Videv, and H. Haas, "Towards a 100 gb/s visible light wireless access network," *Optics express*, vol. 23, no. 2, pp. 1627–1637, 2015.

[22] G. Margelis, R. Piechocki, D. Kaleshi, and P. Thomas, "Low throughput networks for the iot: Lessons learned from industrial implementations," in *2015 IEEE 2nd world forum on internet of things (WF-IoT)*. IEEE, 2015, pp. 181–186.

[23] A. Ageev, E. Luci, C. Petrioli, and N. Thakker, "Vulcan: A low-cost, low-power embedded visible light communication and networking platform," in *Proceedings of the 22nd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2019, pp. 127–134.

[24] S. Yin, N. Smaoui, M. Heydariaan, and O. Gnawali, "Purple vlc: Accelerating visible light communication in room-area through pru offloading." in *EWSN*, 2018, pp. 67–78.

[25] Y. Yang, J. Luo, C. Chen, W.-D. Zhong, and L. Chen, "Synlight: Synthetic light emission for fast transmission in cots device-enabled vlc," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 1297–1305.

[26] Q. Wang, D. Giustiniano, and D. Puccinelli, "OpenVLC: Software-Defined Visible Light Embedded Networks," in *ACM VLCS*, 2014.

[27] "pureLiFi," *https://purelifi.com/*, 2018.

[28] J. Zhang, X. Zhang, and G. Wu, "Dancing with light: Predictive in-frame rate selection for visible light networks," in *INFOCOM*, 2015.

[29] C. B. Liu, B. Sadeghi, and E. W. Knightly, "Enabling Vehicular Visible Light Communication {(V2LC)} Networks," in *Proc. VANET*, 2011.

[30] T. Li, C. An, Z. Tian, A. T. Campbell, and X. Zhou, "Human Sensing Using Visible Light Communication," in *Proc. MobiCom*, 2015.

[31] N. O. Tippenhauer, D. Giustiniano, and S. Mangold, "Toys communicating with leds: Enabling toy cars interaction," in *IEEE CCNC*, 2012.

[32] C. Zhang, J. Tabor, J. Zhang, and X. Zhang, "Extending Mobile Interaction Through Near-Field Visible Light Sensing," in *Proc. MobiCom*, 2015.

[33] Y.-S. Kuo, P. Pannuto, K.-J. Hsiao, and P. Dutta, "Luxapose: Indoor Positioning with Mobile Phones and Visible Light," in *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2014, pp. 447–458.

[34] C. Zhang and X. Zhang, "LiTell: Robust Indoor Localization Using Unmodified Light Fixtures," in *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2016, pp. 230–242.

[35] Q. Wang, M. Zuniga, and D. Giustiniano, "Passive communication with ambient light," in *ACM CoNEXT*, 2016.

[36] X. Xu, Y. Shen, J. Yang, C. Xu, G. Shen, G. Chen, and Y. 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*, ser. MobiCom '17. New York, NY, USA: ACM, 2017, pp. 180–192. [Online]. Available: http://doi.acm.org/10.1145/3117811.3117843

[37] "802.11bb," *https://standards.ieee.org/develop/project/802.11bb.html*, 2018.

[38] A. Galisteo, H. Wu, Q. Wang, D. Juara, M. Zuniga, and D. Giustiniano, "Demo: OpenVLC1.2 for increased data rate with embedded systems," in *VLCS 2017 - Proceedings of the 4th ACM Workshop on Visible Light Communication Systems, co-located with MobiCom 2017*, 2017.

[39] "BeagleBone Black," *http://beagleboard.org/Products/BeagleBone+Black*.

[40] "{IEEE} Standard for Local and Metropolitan Area Networks–Part 15.7: Short-Range Wireless Optical Communication Using Visible Light," *IEEE Std 802.15.7-2011*, 2011.

[41] I. S. Association *et al.*, "Ieee std. for local and metropolitan area networksâpart 15.7: Short-range wireless optical communication using visible light," *IEEE computer Society*, 2011.

[42] O. Jung and C. Ruland, "Analysis of the Statistical Self-Synchronization Mode of Operation," *ITG FACHBERICHT*, pp. 119–126, 2004.

[43] "No Title," *http://www.earth-policy.org/data_highlights/2011/highlights15*.

[44] M. Yamada and D. Chwastyk, "Adoption of light-emitting diodes in common lighting applications," Tech. Rep., 2015.

[45] S. Dimitrov and H. Haas, *Principles of LED Light Communications: Towards Networked Li-Fi*. Cambridge University Press, 2015.

[46] T. Komine and M. Nakagawa, "Fundamental analysis for visible-light communication system using LED lights," *IEEE transactions on Consumer Electronics*, 2004.

[47] Z. Su, D. Xue, and Z. Ji, "Designing LED array for uniform illumination distribution by simulated annealing algorithm," *Optics express*, 2012.

[48] I. Moreno, M. Avenda, and R. Tzonchev, "Designing light-emitting diode arrays for uniform near-field irradiance," *Appl. Opt.*, 2006.

[49] D. Ramane and A. Shaligram, "Optimization of multi-element LED source for uniform illumination of plane surface," *Optics express*, 2011.

[50] G. Varma and \emphet. al., "Power allocation for uniform illumination with stochastic LED arrays," *Optics Express*, 2017.

[51] H. Yang and E. al., "Uniform illumination rendering using an array of LEDs: a signal processing perspective," *IEEE trans. on signal processing*, 2009.

[52] A. Sewaiwar and \emphet. al., "Smart LED allocation scheme for efficient multiuser visible light communication networks," *Optics express*, 2015.

[53] Z. Wang and \emphet. al., "Performance of a novel {LED} lamp arrangement to reduce {SNR} fluctuation for multi-user {VLC} systems," *Optics express*, 2012.

[54] L. Azizan and \emphet. al., "Optimization of SNR for wireless light-emitting diode communication in modern lighting layouts," *Optical Engineering*, 2014.

[55] Y. Zhang, H. Yu, and J. Zhang, "Block Precoding for Peak-Limited MISO Broadcast VLC: Constellation-Optimal Structure and Addition-Unique Designs," in *IEEE Journal on Selected Areas in Communications*, 2018.

[56] Z. Sun, H. Yu, W. Li, Z. Tian, and Y. Zhu, "Power-Efficient Linear Precoding for MU-MISO VLC Systems With Channel Uncertainty," in *IEEE Photonics Technology Letters*, 2018.

[57] E. Nayebi, A. Ashikhmin, T. Marzetta, and H. Yang, "Cell-Free Massive MIMO systems," in *Asilomar Conference on Signals, Systems and Computers*, 2015.

[58] R. Jiang, Q. Wang, H. Haas, and Z. Wang, "Joint User Association and Power Allocation for Cell-Free Visible Light Communication Networks," *IEEE Journal on Selected Areas in Communications*, 2018.

[59] H. Ngo, A. Ashikhmin, H. Yang, E. Larsson, and T. Marzetta, "Cell-Free Massive MIMO Versus Small Cells," in *IEEE Transactions on Wireless Communications*, 2017.

[60] L. Nguyen, T. Duong, H. Ngo, and K. Tourki, "Energy Efficiency in Cell-Free Massive MIMO with Zero-Forcing Precoding Design," in *IEEE Communications Letters*, 2017.

[61] E. Nayebi, A. Ashikhmin, T. Marzetta, H. Yang, and B. Rao, "Precoding and Power Optimization in Cell-Free Massive MIMO Systems," in *IEEE Transactions on Wireless Communications*, 2017.

[62] H. Ngo, L. Tran, T. Duong, M. Matthaiou, and E. Larsson, "On the Total Energy Efficiency of Cell-Free Massive MIMO," in *IEEE Transactions on Green Communications and Networking*.

[63] M. A. Lombardi, L. M. Nelson, A. N. Novick, and V. S. Zhang, "Time and frequency measurements using the global positioning system," *Cal Lab: International Journal of Metrology*, 2001.

[64] D. Giustiniano, E. Goma, A. Lopez Toledo, I. Dangerfield, J. Morillo, and P. Rodriguez, "Fair WLAN Backhaul Aggregation," in *MobiCom*, 2010.

[65] A. Tsiatmas˜ \emphet al, "Joint illumination and visible-Light Communication systems: Data rates and extra power consumption," *IEEE ICC*, 2015.

[66] E. Schubert, *Light-Emitting Diodes*. Cambridge University Press, 2006.

[67] "CREE XT-E LED," *https://goo.gl/ocs7bP*.

[68] "ISO 8995-1:2002 - Lighting of work places – Part 1: Indoor." [Online]. Available: https://www.iso.org/standard/28857.html

[69] "RFC 5905," *Network Time Protocol Version 4: Protocol and Algorithms Specification*.

[70] "IEEE Std 1588-2008," *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*.

[71] S. Naribole, S. Chen, E. Heng, and E. W. Knightly, "LiRa: A WLAN Architecture for Visible Light Communication with a Wi-Fi Uplink," *IEEE SECON*, 2017.

[72] D. Pauluzzi and N. Beaulieu, "A comparison of SNR estimation techniques for the AWGN channel," *IEEE Trans. on Communications*, 2000.

[73] "ACRO," *https://openbuilds.com/builds/openbuilds-acro-system.5416/*.

[74] Z. Wang, C. Guo, Y. Yang, and Q. Li, "Antenna Selection Based Dimming Scheme for Indoor MIMO Visible Light Communication Systems Utilizing Multiple Lamps," *IEEE Vehicular Technology Conference*, 2016.

[75] K. Ying and \emphet. al., "{MIMO} Transceiver Design in Dynamic-Range-Limited VLC Systems," *IEEE Photonics Technology Letters*, 2016.

[76] P. Saengudomlert, "Transmit beamforming for line-of-sight MIMO VLC with IM/DD under illumination constraints," in *IEEE ECTI-CON*, 2015.

[77] I. Din and H. Kim, "Energy-Efficient Brightness Control and Data Transmission for Visible Light Communication," *IEEE Photonics Technology Letters*, 2014.

[78] A. Jalajakumari and \emphet. al., "An energy efficient high-speed digital LED driver for visible light communications," in *ICC*, 2015.

[79] A. Nordrum, "The Internet of Fewer Things," *IEEE Spectrum*, 2016.

[80] I. Analytics, "State of the IoT 2018: Number of IoT devices now at 7B," 2018.

[81] M. M. Tentzeris, A. Georgiadis, and L. Roselli, "Energy harvesting and scavenging," *Proc. IEEE*, vol. 102, no. 11, 2014.

[82] M. Anderson, "Potential Hazards at Both Ends of the Lithium-Ion Life Cycle," *IEEE Spectrum*, 2013.

[83] T. Li and X. Zhou, "Battery-Free Eye Tracker on Glasses," in *ACM MobiCom '18*, New York, NY, USA, 2018.

[84] H. Truong, S. Zhang, U. Muncuk, P. Nguyen, N. Bui, A. Nguyen, Q. Lv, K. Chowdhury, T. Dinh, and T. Vu, "CapBand: Battery-free Successive Capacitance Sensing Wristband for Hand Gesture Recognition," in *ACM SenSys '18*, New York, NY, USA, 2018.

[85] V. Talla, B. Kellogg, S. Gollakota, and J. R. Smith, "Battery-Free Cellphone," *IMWUT*, vol. 1, no. 2, 2017.

[86] S. Naderiparizi, M. Hessar, V. Talla, S. Gollakota, and J. R. Smith, "Towards Battery-Free {HD} Video Streaming," in *{NSDI} 18*. USENIX, 2018.

[87] X. Li, B. Hussain, L. Wang, J. Jiang, and C. P. Yue, *Journal of Lightwave Technology*.

[88] Z. Tian, K. Wright, and X. Zhou, "The Darklight Rises: Visible Light Communication in the Dark: Demo," in *Proceedings of the 22Nd Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '16. New York, NY, USA: ACM, 2016, pp. 495–496. [Online]. Available: http://doi.acm.org/10.1145/2973750.2987384

[89] J. Li, A. Liu, G. Shen, L. Li, C. Sun, and F. 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.

[90] H. Haas, "High-speed wireless networking using visible light," *SPIE Newsroom*, vol. 1, no. 1, 2013.

[91] A. Varshney, A. Soleiman, L. Mottola, and T. Voigt, "Battery-free Visible Light Sensing," in *Proceedings of the 4th ACM Workshop on Visible Light Communication Systems*, ser. VLCS '17. New York, NY, USA: ACM, 2017, pp. 3–8. [Online]. Available: http://doi.acm.org/10.1145/3129881.3129890

[92] A. Varshney, A. Soleiman, and T. Voigt, "TunnelScatter: Low Power Communication for Sensor Tags Using Tunnel Diodes," in *The 25th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '19. New York, NY, USA: ACM, 2019, pp. 50:1—-50:17. [Online]. Available: http://doi.acm.org/10.1145/3300061.3345451

[93] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith, "Ambient Backscatter: Wireless Communication out of Thin Air," in *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 39–50. [Online]. Available: http://doi.acm.org/10.1145/2486001.2486015

[94] M. Rostami, J. Gummeson, A. Kiaghadi, and D. Ganesan, "Polymorphic Radios: A New Design Paradigm for Ultra-low Power Communication," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '18. New York, NY, USA: ACM, 2018, pp. 446–460. [Online]. Available: http://doi.acm.org/10.1145/3230543.3230571

[95] C. Pérez-Penichet, C. Noda, A. Varshney, and T. Voigt, "Battery-free 802.15.4 Receiver," in *ACM/IEEE IPSN '18*, Piscataway, NJ, USA, 2018.

[96] E. C. for Standardization, "Light and Lighting-Lighting of work places-Part 1: Indoor work places," *European Std. EN 12 464-1*, 2011.

[97] IXYS, "Slmd121h04l. http://ixapps.ixys.com/datasheet/slmd121h04l-data-sheet.pdf."

[98] A. Varshney, O. Harms, C. Pérez-Penichet, C. Rohner, F. Hermans, and T. Voigt, "LoRea: A Backscatter Architecture That Achieves a Long Communication Range," in *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys '17. New York, NY, USA: ACM, 2017, pp. 18:1—-18:14. [Online]. Available: http://doi.acm.org/10.1145/3131672.3131691

[99] B. Kellogg, V. Talla, S. Gollakota, and J. R. Smith, "Passive Wi-Fi: Bringing Low Power to Wi-Fi Transmissions," in *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, ser. NSDI'16. Berkeley, CA, USA: USENIX Association, 2016, pp. 151–164. [Online]. Available: http://dl.acm.org/citation.cfm?id=2930611.2930622

[100] V. Talla, M. Hessar, B. Kellogg, A. Najafi, J. R. Smith, and S. Gollakota, "LoRa Backscatter: Enabling The Vision of Ubiquitous Connectivity," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 1, no. 3, pp. 105:1—-105:24, 2017. [Online]. Available: http://doi.acm.org/10.1145/3130970

[101] Texas Instruments. bq25570, "bq25570. http://www.ti.com/lit/ds/symlink/bq25570.pdf."

[102] PowerFilm. MP3-37, "No Title."

[103] On Semiconductor, "https://www.onsemi.com/pub/Collateral/NCS2200-D.PDF."

[104] Texas Instruments, "MSP430FR5949. http://www.ti.com/lit/ds/symlink/msp430fr5949.pdf."

[105] Linear Technology, "https://www.analog.com/media/en/technical-documentation/data-sheets/6906fc.pdf."

[106] Analog Devices, "https://www.analog.com/media/en/technical-documentation/data-sheets/ADG904.pdf."

[107] Ablic, "https://www.ablic.com/en/doc/datasheet /voltage_regulator/S1313_E.pdf."

[108] BeagleBoard, "https://cdn-shop.adafruit.com/datasheets/BBB_SRM.pdf."

[109] Texas Instruments, "http://www.ti.com/lit/ds/symlink/cc1310.pdf."

[110] Ettus Research, "https://www.ettus.com/wp-content/uploads/2019/01/b200-b210_spec_sheet.pdf."

[111] K. D. C. P. P. S. E36313A, "No Title," https://literature.cdn.keysight.com/litweb/pdf/E36311-90001.pdf, 2019.

[112] M. Kamel, W. Hamouda, and A. Youssef, "Ultra-Dense Networks: A Survey," *IEEE Communications Surveys Tutorials*, vol. 18, no. 4, 2016.

[113] S. Han, C. l. I, Z. Xu, and C. Rowell, "Large-scale antenna systems with hybrid analog and digital beamforming for millimeter wave 5G," *IEEE Communications Magazine*, vol. 53, no. 1, 2015.

[114] B. Kellogg, A. Parks, S. Gollakota, J. Smith, and D. Wetherall, "{Wi-Fi} Backscatter: Internet Connectivity for {RF}-powered Devices," in *ACM SIGCOMM*, 2014.

[115] D. Bharadia, K. R. Joshi, M. Kotaru, and S. Katti, "BackFi: High Throughput WiFi Backscatter," *SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 4, 2015.

[116] S. P. Y. F. X. T. H. Y. X. W. Renjie Zhao Fengyuan Zhu, "OFDMA-Enabled Wi-Fi Backscatter," in *ACM MobiCom '19*, 2019.

[117] R. Bloom, M. Zuniga, Q. Wang, and D. Giustiniano, "Tweeting with Sunlight: Encoding Data on Mobile Objects," in *IEEE INFOCOM*, 2019.

[118] F. Tonolini and F. Adib, "Networking Across Boundaries: Enabling Wireless Communication Through the Water-air Interface," in *ACM SIGCOMM '18*, New York, NY, USA, 2018.

[119] Y. Wang, D. A. Basnayaka, X. Wu, and H. Haas, "Optimization of Load Balancing in Hybrid LiFi/RF Networks," *IEEE Transactions on Communications*, vol. 65, no. 4, 2017.

[120] J. Gummeson, J. Mccann, C. Yang, D. Ranasinghe, S. Hudson, and A. Sample, "RFID Light Bulb: Enabling Ubiquitous Deployment of Interactive RFID Systems," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 1, no. 2, 2017.

[121] D. Giustiniano, A. Varshney, and T. Voigt, "Connecting Battery-free IoT Tags Using LED Bulbs," in *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*, ser. HotNets '18. New York, NY, USA: ACM, 2018, pp. 99–105. [Online]. Available: http://doi.acm.org/10.1145/3286062.3286077

[122] T.-H. Do and M. Yoo, "An in-depth survey of visible light communication based positioning systems," *Sensors*, vol. 16, no. 5, p. 678, 2016.

[123] Z. Ghassemlooy, W. Popoola, and S. Rajbhandari, *Optical wireless communications: system and channel modelling with Matlab®*. CRC press, 2012.

[124] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of wireless indoor positioning techniques and systems," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 6, pp. 1067–1080, 2007.

[125] A. Arafa, S. Dalmiya, R. Klukas, and J. F. Holzman, "Angle-of-arrival reception for optical wireless location technology," *Optics express*, vol. 23, no. 6, pp. 7755–7766, 2015.

[126] J. Vongkulbhisal, B. Chantaramolee, Y. Zhao, and W. S. Mohammed, "A fingerprinting-based indoor localization system using intensity modulation of light emitting diodes," *Microwave and Optical Technology Letters*, vol. 54, no. 5, pp. 1218–1227, 2012.

[127] G. Kail, P. Maechler, N. Preyss, and A. Burg, "Robust asynchronous indoor localization using LED lighting," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014, pp. 1866–1870.

[128] C. Danakis, M. Afgani, G. Povey, I. Underwood, and H. Haas, "Using a CMOS camera sensor for visible light communication," *Proceedings of the IEEE Globecom Workshops*, 2012.

[129] S. H. Yang, E. M. Jeong, and S. K. Han, "Indoor positioning based on received optical power difference by angle of arrival," *Electronics Letters*, vol. 50, no. 1, pp. 49–51, 2014.

[130] L. Li, P. Hu, C. Peng, G. Shen, and F. Zhao, "Epsilon: A Visible Light Based Positioning System." in *Proceedings of the USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2014, pp. 331–343.

[131] M. Yasir, S.-W. Ho, and B. N. Vellambi, "Indoor positioning system using visible light and accelerometer," *Journal of Lightwave Technology*, vol. 32, no. 19, pp. 3306–3316, 2014.

[132] B. Xie, S. Gong, and G. Tan, "LiPro: light-based indoor positioning with rotating handheld devices," *Wireless Networks*, pp. 1–11, 2016.

[133] L. Yin, X. Wu, and H. Haas, "Indoor visible light positioning with angle diversity transmitter," in *Proceedings of the IEEE Vehicular Technology Conference (VTC Fall)*, 2015, pp. 1–5.

[134] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen, "Zee: Zero-effort Crowdsourcing for Indoor Localization," in *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2012, pp. 293–304.

[135] Q. Wang, D. Giustiniano, and D. Puccinelli, "An Open-Source Research Platform for Embedded Visible Light Networking," *IEEE Wireless Communication*, 2015.

[136] Q. Wang, D. Giustiniano, and M. Zuniga, "In Light and In Darkness, In Motion and In Stillness: A Reliable and Adaptive Receiver for the Internet of Lights," *IEEE Journal on Selected Areas in Communications (JSAC)*, 2017.

[137] D. Lymberopoulos, J. Liu, X. Yang, R. R. Choudhury, V. Handziski, and S. Sen, "A Realistic Evaluation and Comparison of Indoor Location Technologies: Experiences and Lessons Learned," in *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2015.

[138] J. Armstrong, Y. Sekercioglu, and A. Neild, "Visible light positioning: A roadmap for international standardization," *IEEE Communications Magazine*, vol. 51, no. 12, pp. 68–73, 2013.

[139] K. Panta and J. Armstrong, "Indoor localisation using white LEDs," *Electronics letters*, vol. 48, no. 4, pp. 228–230, 2012.

[140] Z. Wang, Z. Yang, J. Zhang, C. Huang, and Q. Zhang, "Wearables Can Afford: Light-weight Indoor Positioning with Visible Light," in *Proceedings of the ACM Annual International Conference on Mobile Systems, Applications, and Services (Mobisys)*, 2015.

[141] S. Zhu and X. Zhang, "Enabling High-Precision Visible Light Localization in Today's Buildings," in *Proceedings of the ACM Annual International Conference on Mobile Systems, Applications, and Services (Mobisys)*, 2017.

[142] "STmicroelectronics - VL53L1," 2017. [Online]. Available: http://www.st.com/content/st{_}com/en/products/imaging-and-photonics-solutions/proximity-sensors/vl53l1.html

[143] (2017) LEDs to account for 98% of all lighting by 2025, Frost & Sullivan forecasts. [Online]. Available: https://www.eenewsled.com/news/leds-account-98-all-lighting-2025-frost-sullivan-forecasts

[144] Q. Wang, H. Luo, A. Men, F. Zhao, X. Gao, J. Wei, Y. Zhang, and Y. Huang, "Light positioning: A high-accuracy visible light indoor positioning system based on attitude identification and propagation model," *International Journal of Distributed Sensor Networks*, vol. 14, no. 2, p. 1550147718758263, 2018.

[145] T. Do and M. Yoo, "Visible light communication-based vehicle-to-vehicle tracking using CMOS camera," *IEEE Access*, vol. 7, pp. 7218–7227, 2019.

[146] M. Ratosi and G. Simon, "Real-time localization and tracking using visible light communication," in *Proc. International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2018.

[147] P. Hu *et al.*, "ColorBars: Increasing data rate of LED-to-camera communication using color shift keying," in *Proc. ACM CoNEXT*, 2015.

[148] R. LiKamWa *et al.*, "Energy characterization and optimization of image sensing toward continuous mobile vision," in *Proc. ACM MobiSys*, 2013.

[149] (2018) Philips indoor positioning. [Online]. Available: http://www.lighting.philips.com/main/systems/lighting-systems/indoor-positioning

[150] W. Gu, M. Aminikashani, P. Deng, and M. Kavehrad, "Impact of multipath reflections on the performance of indoor visible light positioning systems," *J. Lightw. Technol.*, vol. 34, no. 10, pp. 2578–2587, May 2016.

[151] Y. Zhuang, L. Hua, L. Qi, J. Yang, P. Cao, Y. Cao, Y. Wu, J. Thompson, and H. Haas, "A survey of positioning systems using visible LED lights," *IEEE Commun. Surveys Tuts*, vol. 20, no. 3, pp. 1963–1988, thirdquarter 2018.

[152] W. Gu, M. Aminikashani, and M. Kavehrad, "Indoor visible light positioning system with multipath reflection analysis," in *Proc. IEEE International Conference on Consumer Electronics (ICCE)*, Jan. 2016, pp. 89–92.

[153] W. Tang, J. Zhang, B. Chen, Y. Liu, Y. Zuo, S. Liu, and Y. Dai, "Analysis of indoor VLC positioning system with multiple reflections," in *Proc. International Conference on Optical Communications and Networks (ICOCN)*, Aug. 2017, pp. 1–3.

[154] D. Plets, A. Eryildirim, S. Bastiaens, N. Stevens, L. Martens, and W. Joseph, "A performance comparison of different cost functions for RSS-based visible light positioning under the presence of reflections," in *Proc. ACM Workshop on Visible Light Communication Systems (VLCS)*, 2017, p. 37â41.

[155] W. Pan, Y. Hou, and S. Xiao, "Visible light indoor positioning based on camera with specular reflection cancellation," in *Proc. Conference on Lasers and Electro-Optics Pacific Rim (CLEO-PR)*, July 2017, pp. 1–4.

[156] C. Chen, D. A. Basnayaka, X. Wu, and H. Haas, "Efficient analytical calculation of non-line-of-sight channel impulse response in visible light communications," *J. Lightw. Technol.*, vol. 36, no. 9, pp. 1666–1682, May 2018.

[157] X. Yu *et al.*, "Indoor positioning system based on single LED using symmetrical optical receiver," in *Proc. Asia Communications and Photonics Conference (ACP)*, 2018.

[158] D.-q. Ding and X.-z. Ke, "A new indoor VLC channel model based on reflection," *Optoelectronics Letters*, p. 295â298, Oct. 2010.

[159] Y. Qiu, H.-H. Chen, and W.-X. Meng, "Channel modeling for visible light communicationsâa survey," *Wireless Communications and Mobile Computing*, vol. 16, no. 14, pp. 2016–2034, 2016.

[160] C. Chen, D. Basnayaka, and H. Haas, "Non-line-of-sight channel impulse response characterisation in visible light communications," in *Proc. IEEE International Conference on Communications (ICC)*, May 2016, pp. 1–6.

[161] N. A. Mohammed and M. A. Elkarim, "Exploring the effect of diffuse reflection on indoor localization systems based on RSSI-VLC," *Opt. Express*, vol. 23, no. 16, pp. 20 297–20 313, Aug. 2015.

[162] M. A. Elkarim, N. A. Mohammed, and M. H. Aly, "Exploring the performance of indoor localization systems based on VLC-RSSI, including the effect of NLOS components using two light-emitting diode lighting systems," *Optical Engineering*, vol. 54, no. 10, pp. 1 – 9, 2015.

[163] C. Huang and X. Zhang, "LOS-NLOS identification algorithm for indoor visible light positioning system," in *Proc. International Symposium on Wireless Personal Multimedia Communications (WPMC)*, Dec. 2017, pp. 575–578.

[164] J. Xu, H. Shen, W. Xu, H. Zhang, and X. You, "LED-assisted three-dimensional indoor positioning for multiphotodiode device interfered by multipath reflections," in *Proc. IEEE Vehicular Technology Conference (VTC Spring)*, June 2017, pp. 1–6.

[165] J. Chen and C. Yan, "A channel model for indoor visible light communication system with specular reflection," in *Proc. International Conference on Optical Communications and Networks (ICOCN)*, Aug. 2017, pp. 1–3.

[166] M. Fox, *Optical properties of solids (2nd ed.)*. Oxford: Oxford University Press, 2010.

[167] S. Brown, *Measures of Shape: Skewness and Kurtosis*. Oak Road Systems, 2018.

[168] R. Killick, P. Fearnhead, and I. A. Eckley, "Optimal detection of changepoints with a linear computational cost," *Journal of the American Statistical Association*, vol. 107, no. 500, pp. 1590–1598, 2012.

[169] A. Mannini and A. M. Sabatini, "Machine learning methods for classifying human physical activity from on-body accelerometers," *Sensors*, vol. 10, no. 2, pp. 1154–1175, 2010.

[170] D. B. M. Susi and G. Lachapelle, "Accelerometer signal features and classification algorithms for positioning applications," *International Technical Meeting of The Institute of Navigation*, 2011.

[171] S. J. Lee, J. K. Kwon, S. Y. Jung, and Y. H. Kwon, "Simulation modeling of visible light communication channel for automotive applications," in *Proc. International IEEE Conference on Intelligent Transportation Systems*, Sep. 2012, pp. 463–468.

[172] F. Miramirkhani, O. Narmanlioglu, M. Uysal, and E. Panayirci, "A mobile channel model for VLC and application to adaptive system design," *IEEE Commun. Lett.*, vol. 21, no. 5, pp. 1035–1038, May 2017.

[173] H. Huang, L. Feng, P. Guo, A. Yang, and G. Ni, "Iterative positioning algorithm to reduce the impact of diffuse reflection on an indoor visible light positioning system," *Opt. Eng.*, vol. 55, no. 6, pp. 1 – 7, 2016.

[174] S. Shawky *et al.*, "Improved vlc-based indoor positioning system using a regression approach with conventional rss techniques," in *2017 13th International Wireless Communications and Mobile Computing Conference (IWCMC)*, 2017.

[175] D. K. Son, E. B. Cho, and C. G. Lee, "Demonstration of visible light communication link for audio and video transmission," in *2010 Photonics Global Conference.* IEEE, 2010, pp. 1–4.

[176] Y. He, L. Ding, Y. Gong, Y. Wang *et al.*, "Real-time audio & video transmission system based on visible light communication," *Optics and Photonics J*, vol. 3, pp. 153–157, 2013.