Angular Diversity Based Visible Light Positioning

ANGULAR DIVERSITY BASED VISIBLE LIGHT POSITIONING

BY

MICHAEL THOMAS TAYLOR, B.Eng., (Electrical Engineering)

McMaster University, Hamilton Ontario, Canada

A THESIS

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING

AND THE SCHOOL OF GRADUATE STUDIES

OF MCMASTER UNIVERSITY

IN PARTIAL FULFILMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE

Master of Applied Science (2014) McMaster University

(Electrical & Computer Engineering) Hamilton, Ontario, Canada

TITLE:	Angular Diversity Based Visible Light Positioning
AUTHOR:	Michael Thomas Taylor B.Eng., (Electrical Engineering) McMaster University, Hamilton Ontario, Canada
SUPERVISOR:	Dr. Steve Hranilovic
NUMBER OF PAGES:	xxi, 136

This work is dedicated to my father Thomas for starting me down this road, my mother Elizabeth for steering me straight every time I hit a pothole, and my sister Sarah for traveling with me the whole way.

Abstract

Visible Light Positioning (VLP) uses modulated LED luminaries to help locate a receiver inside a building. Indoor positioning is becoming an increasingly important service, and visible light has many advantages over other technologies used in indoor positioning systems (IPS). However existing VLP approaches have major drawbacks in robustness that have hindered their ability to be commercially deployed.

This work proposes and demonstrates a new way of using light signals in an indoor localization system, titled Angular Diversity Visible Light Positioning (AD-VLP). AD-VLP uses optics at the transmitter to create a structured overlapping light pattern that can be used for positioning. This method is shown to have several advantages over existing VLP approaches, including increased robustness over intensity based techniques while still using a single element receiver.

This work also includes an experimental implementation of the proposed AD-VLP system using existing mobile device technology. The experiments prove that sub-meter accuracy is obtainable, even when the receiver is oriented away from the transmitter.

Acknowledgements

I would like to acknowledge the immense amount of guidance and support given to me by my supervisor, Dr. Steve Hranilovic. From my first summer internship, through my undergraduate capstone project, and now through this master's degree, he has constantly helped me achieve more than I ever thought possible. Without his mentorship along the way I would not have been able to accomplish what I have, nor been able to get to where I'm going next. It's been fun Sir!

Symbols and Notation

E_e	Irradiance
E_v	Illuminance
F	Lens focal point
I_e	Radiant Intensity
M	Image magnification
$R(\lambda)$	Receiver responsivity
R_{cen}	Projected light pattern seperation radius
R_{cir}	Projected light circle radius
$V(\lambda)$	Luminosity function
Φ_e	Radiometric Flux
$\Phi_e(\lambda)$	Source spectral power density
Φ_v	Luminous Flux
η	Region detection signal strength threshold
λ	Wavelength
ϕ	Transmitter to receiver azimuth angle
ψ	Receiver orientation angle
θ	Transmitter to receiver polar angle

ε	Position error
d	Transmitter to receiver distance
f	Lens focal length
h	Transmitter to receiver vertical seperation distance
m	Lambertian order
r	Lens radius
s	Object (LED) to lens seperation distance
s'	Image distance
y	Object height
y'	Image height

Acronyms

AD-VLP	Angular Diversity Visible Light Positioning
AOA	Angle of Arrival
CDMA	Code Division Multiple Access
DD	Direct Detection
DOA	Direction of Arrival
FDM	Frequency Division Multiplexing
FOV	Field of View
GPS	Global Positioning System
IM	Intensity Modulation
IMU	Inertial Measurement Unit
IPS	Indoor Positioning System
LBS	Location Based Services
LED	Light Emitting Diode
LOS	Line of Sight

MSE	Mean Squared Error
NLOS	Non-Line of Sight
PSD	Power Spectral Density
RSS	Received Signal Strength
SPD	Spectral Power Density
TOA	Time of Arrival
VLC	Visible Light Communications
VLP	Visible Light Positioning

Contents

List of Figures

Chapter 1

Introduction

The importance of knowing where you are does not stop when you walk inside a building, but conventional navigation technologies cease to work indoors. Global Positioning System (GPS) satellites can locate a receiver almost anywhere on the planet, but they cannot be used inside. Even a compass is unreliable indoors due to the many sources of local magnetic interference. Finding your way around inside an unfamiliar building becomes an orienteering adventure largely dependent on how good the posted signage is and if you can find a map. In our increasingly interconnected world where mobile computing is ubiquitous this is a major shortfall. There are many applications that cannot be realized without knowing the device's position. Whether it's a busy traveler trying to find a gate at an airport, or a hospital trying to page a doctor, there is a real need for a system that can provide accurate positioning information indoors the same way GPS does outside.

1.1 The Drive for Indoor Positioning

There are a number of applications that require accurate position information indoors, but by far the development of indoor positioning systems (IPS) is driven by commercial interests. Location based services (LBS) are projected to become a $4 - $10 billion dollar market over the next ten years [4, 5]. This market includes the variety of services that are made possible with accurate location information indoors. Position information can be used to help a customer find their way around a store, or direct them to a particular product. Businesses can use position information to provide customized advertisements directly to customers, or to offer reward programs for visiting their stores. Different IPS pilot projects have already been deployed in many malls across the United States. Wal-Mart and Home Depot have developed apps for navigating their stores [6], Apple recently purchased indoor localization start-up WiFiSLAM for $20 million [7], and Google now provides indoor maps at several locations [8].

Indoor position information can also be used for secure communications, asset tracking and management, network planning, and virtual reality, among other applications [9]. Given how useful it would be to navigate an unfamiliar airport, connect with a customer in a store, or locate equipment in a hospital, it is surprising such systems are only now starting to become available. There have been a number of systems proposed over the years based on different technologies, but they have all had drawbacks that have kept them from widespread adoption. It is only recently that the technology has matured enough to realize a complete indoor positioning system ready for adoption. The current front runner is indoor positioning using WiFi signals, but the accuracy of these approaches is limited. A promising new approach

developed recently is visible light positioning (VLP), which uses light signals in the visible wavelengths to locate a receiver indoors. Start-up ByteLight has raised over $3 million dollars to develop a light based IPS [10], and recently lighting giant Philips has begun marketing its own VLP system [11].

1.2 Previous Approaches to Indoor Positioning

Indoor positioning has been an active topic of research for the last fifteen years. A brief overview of previous technologies used in indoor positioning is presented in this section. The material in this section was gathered from the following indoor localization surveys [12–18].

Computer Vision (CV) Techniques

Cameras and image processing techniques can be used to locate and track objects in their field-of-view [17]. These techniques require a large number of cameras to be installed to cover a given area which must be calibrated and synchronized. Providing coverage over multiple rooms requires a prohibitively expensive number of cameras, and positioning performance suffers without some known reference markers on the objects being tracked. It is also possible for a camera on the receiver to locate itself through scene recognition. Nearly every mobile device is equipped with a camera but the complexity of the image processing task, lack of known markers in the scene, movement of the receiver in the users hand and heavy drain on the battery make these approaches inaccurate and unattractive.

Ultrasonic (US) Techniques

Ultrasonic systems use echolocation to form a location estimate. A pulse outside the human audible range is emitted by a speaker tag attached to the user and received by a series of microphones installed in the ceiling [19]. Measurements allow distance to be estimated from the speed of the traveling sound wave, while multiple receivers allow angle to be determined. Ultrasonic systems are effective at room level accuracy, but suffer from reflections around the room which limit their absolute accuracy. A high number of receivers are required to achieve centimeter level accuracy for a given coverage area, increasing the system cost. Also these systems are limited in the number of tags they can track in a given area due to signal interference.

Infrared (IR) Techniques

Infrared (IR) systems transmit IR light pulses from tags on a user which are tracked by a series of sensors in the ceiling, or by IR cameras [14]. The ability to use a camera increases the possible positioning accuracy of the system but also increases the cost. Position can be determined to centimeter level accuracy if enough cameras and tags are used, but many cameras are needed and they must be well calibrated. IR systems also suffer from interference from background radiation.

Ultra Wide Band (UWB) Techniques

Ultra Wide Band (UWB) systems use very wide bandwidth, short time duration pulses of electromagnetic radiation. This allows for better rejection of multi-path signal components and more accurate timing measurements than IR or US approaches

with fewer base stations. However UWB systems require accurate timing synchronization, are more expensive, and have not yet been able to realize accuracy improvements over other simpler systems [13].

Radio Frequency Identification (RFID) Techniques

Radio Frequency Identification (RFID) uses passive or active RFID tags that emit low power RF signals which can be used for positioning [15]. The signals are gathered by multiple antennas and combined to form an estimate of the receiver position. These systems have very limited range and accuracy due to the low power of the transmit tags and the low frequency signals used. Accuracy is typically on the order of meters.

Inertial Measurement Unit (IMU) Navigation

Inertial Measurement Unit (IMU) based navigation uses a combination of accelerometer, gyroscope and compass readings to estimate the orientation and position of a receiver. If a starting position and orientation is provided these measurements can be integrated to calculate the displacement in time, which provides an estimate of the receivers current location. These methods use no external points of reference, and are known as "dead reckoning" navigation from their absolute confidence in (and dependence on) internal calculations. The major drawback of this approach is that any error in measurement will propagate forward in time for ever. Any small deviation in angle or acceleration will result in an infinitely growing displacement error, or a positional drift. These systems are still attractive however, as they require no external infrastructure to be installed. Also since many mobile devices now have multi-axis accelerometers and gyroscopes built in there is the possibility of enabling

mobile positioning for a variety of existing consumers. There are different approaches that attempt to mitigate drift by using step counts or body motion models [18], but ultimately an external reference is required eventually to update the systems actual position and reset the drift.

WiFi Techniques

Indoor localization using WiFi signals has been an active topic of research for the last decade [20]. WiFi is nearly ubiquitous indoors, and mobile devices with WiFi connectivity have been widespread for several years. This presents an attractive possibility of being able to use already deployed infrastructure and receivers to realize a low cost IPS. Unfortunately WiFi positioning has drawbacks due to the type of signal being used and the primary function of the infrastructure. WiFi positioning uses received signal strength measurements to estimate the distance from the router to the receiver. As the RF signals go through walls and floors they suffer from fading and multi-path distortion. The receiver cannot accurately recover position information from the signal measurements if it does not precisely know the environment the signal traveled through. Typical positioning accuracy for WiFi based IPS are still 1 to 3 m, even with high node density [21]. Indoor localization systems need as many overlapping nodes as possible for more accurate position estimation, and requires a minimum amount of overlap to provide any estimate at all. WiFi networks are not usually designed with large amounts of overlap as it is unnecessary for providing network connectivity, and can increase interference.

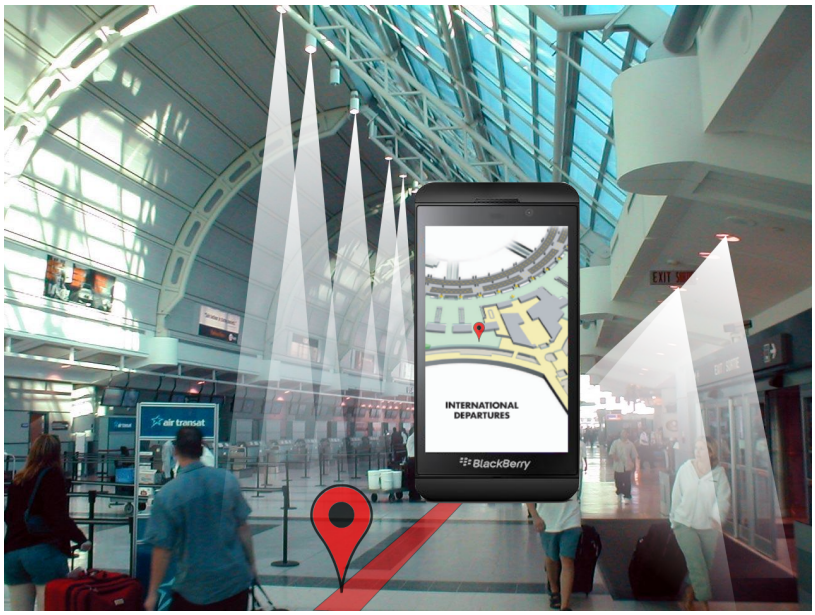

Figure 1.1: Depiction of Visible Light Positioning (VLP) in Toronto's Pearson airport. Modified from [1] by author.

1.3 Visible Light Positioning

Visible light has several characteristics that make it suited for indoor positioning applications. Buildings already have a large number of light fixtures that cover the entire structure, since people need lighting to see. These fixtures make for a large number of potential transmitter locations which allows a much higher transmitter density than other technologies. Light does not travel through opaque boundaries such as walls or floors, so the signal is localized to the room it was transmitted in. This also eliminates interference between transmitters in adjacent rooms allowing high spatial reuse of bandwidth. Light based positioning raises fewer privacy concerns as well since covering the receiver guarantees the system is not in use. Finally, light is also highly directional and can be shaped with inexpensive optics. These advantages make visible light an obvious choice for designing an indoor positioning system.

Visible light positioning (VLP) uses visible light signals to locate a receiver. It was originally conceived in the mid 2000's when researchers considered enabling communications through fluorescent lighting [22]. The development of the bright white LED transformed the field of visible light communications (VLC) and was an enabling technology for VLP. VLC is the research field that studies communications using the visible spectrum of light [23]. It has recently been standardized in IEEE 802.15.7 for short range networks [24].

In addition to being longer lasting and more energy efficient than existing lighting technologies, LED lights can be modulated to communicate information [25,26]. VLC enabled luminaries can be used to realize an indoor positioning system by turning luminaries in light fixtures into transmitters at known locations. Each newly installed VLC bulb can be used as a measurement point in the building by recording which light fixture it was mounted in. The bulbs can identify themselves using VLC so that each measurement is unique. There is an unprecedented opportunity to implement indoor positioning functionality in existing buildings as they are refit with more energy efficient lighting.

While VLP has many inherit advantages over other indoor positioning approaches, and has begun to see commercial development, there are still fundamental issues to be addressed. As will be shown in Chapter 2 the particulars of transmitting and detecting light make using optical measurements in indoor positioning difficult, and existing approaches to VLP are not yet able to fully realize its potential.

1.4 Thesis Contribution

This thesis describes and develops a new method of indoor localization using visible light based on the principle of angular diversity. The method was proposed by the author in [3, 27]. This new approach proposes the use of optical elements at the luminary transmitter in conjunction with individual modulation of the luminary LEDs to create a structured overlapping light pattern. Each region in the overlapping light pattern can be uniquely identified by the signals that form it, which allows a bound on the receiver position to be formed using simple fingerprinting techniques.

To the author's knowledge, this is the first time such a system has been proposed in the context of indoor localization, and the first time it has been implemented experimentally. The concepts behind this new approach are explained in detail, and it is shown that the proposed system addresses several of the drawbacks of previous VLP systems while maintaining the main benefits of light based indoor positioning approaches.

The second contribution is the development of a thresholding method for identifying the light regions. The illumination characteristics of the system are presented and tied to the positioning performance through the choice of a region detection threshold.

The final contribution of this work is the experimental validation of the proposed VLP system. A positioning enabled luminary is built and characterized according to the methods proposed. The luminary is used in experiments with a mobile receiver to test the validity of the new approach, and the positioning performance of the system.

1.5 Thesis Outline

Chapter 2 presents relevant background information and a literature review of existing approaches to indoor positioning using visible light. Basic radiometric quantities are reviewed and the VLC channel is described. Both intensity and imaging approaches to indoor positioning using visible light are reviewed and their strengths and weaknesses discussed.

Chapter 3 presents a new angular diversity approach to indoor positioning using visible light. Observations taken from imaging based VLP are extended to realize the new VLP method, termed *Angular diversity visible light positioning* (AD-VLP). The use of optical elements at the transmitter to create a structured, overlapping light pattern is proposed in detail. The use of this structured light pattern in indoor positioning is compared to previous approaches and its advantages and drawbacks are discussed. A basic design example using an ideal thin lens is presented to develop the theory of AD-VLP, which is followed by simulation results of the systems theoretical positioning performance.

Chapter 4 further develops the theory of AD-VLP to include the illumination characteristics of the luminary, and presents a method of identifying the individual overlapping light regions. A method of setting an intensity threshold for detecting light regions is presented, and the effect of this threshold on the geometry of the pattern is examined. Simulation results of positioning performance using non-ideal transmitters and lenses is presented. Finally the problem of illumination and positioning co-design is examined, and the advantage of more intricate optical elements is presented.

Chapter 5 presents experimental work done to validate this new approach to VLP.

The design of an AD-VLP transmitter and single element receiver is detailed, followed by experimental illumination and positioning results.

Finally Chapter 6 is a short concluding section that summarizes the key points of this work, and presents some of the many potential future research directions.

Chapter 2

Light and Indoor Positioning

This chapter presents an overview of the concepts relevant to indoor positioning using visible light followed by a review of existing approaches to VLP. The chapter includes a review of localization techniques, measurement methods, and performance metrics before describing important factors in indoor localization. A review of light measurement systems, the VLC channel model, and perspective projection describes the theory underlying existing approaches to VLP techniques. A review of previous approaches to VLP is broken into two sections based on the class of technique used, intensity and imaging. The chapter concludes with a summary of the benefits and challenges of existing VLP approaches.

2.1 Principles of Localization

Localization is the act of confining or restricting something to a particular place. In the context of indoor positioning it is used to describe narrowing down the possible

locations of an object to determine its actual position. Before explaining the specific factors impacting indoor localization it is useful to provide a brief discussion of localization techniques, measurement systems and metrics. A good review is found in [28, Ch.3] and the following surveys [13, 14, 29].

2.1.1 Localization Techniques

Localization techniques use measurements of distance or angles from a known point to solve for the location of the unknown one. Each measurement creates a set of possible points the receiver could possibly be at to produce that measurement. As more measurements are taken to different known points, the receiver must be located at the intersection of the sets. The type of measurement taken determines the shape of the boundary surfaces, and the technique used to solve for their intersection.

Trilateration

Trilateration uses distance measurements between the unknown point and several known ones to solve for its location. Figure 2.1 depicts this problem in two dimensions. Each distance measurement produces a circle of possible locations around the fixed point extending out to the unknown point. Each new point and distance measurement produces a new circle, and the unknown point must be located at the intersection of all three circles. The same principle works in three dimensions with each measurement producing a sphere, and at least four known points are required.

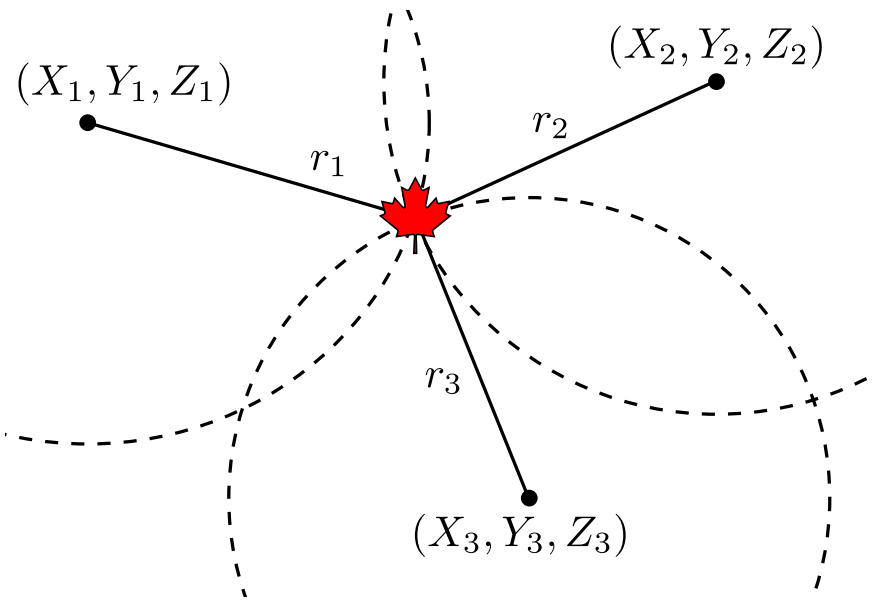

Figure 2.1: Geometry for Trilateration. The maple leaf is the unknown point. Each distance measurement (r_1, r_2, r_3) forms a new circle. The intersection of all three circles is the unknown location.

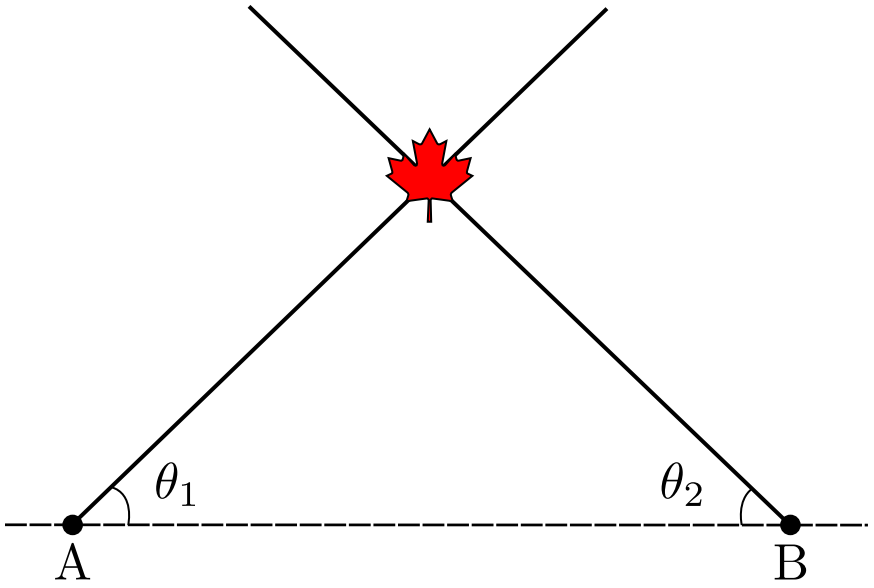

Figure 2.2: Geometry for Triangulation. The maple leaf is the unknown point. Each angle measurement θ_1, θ_2 combined with the known points A, B form a line. The unknown point is located where the lines intersect.

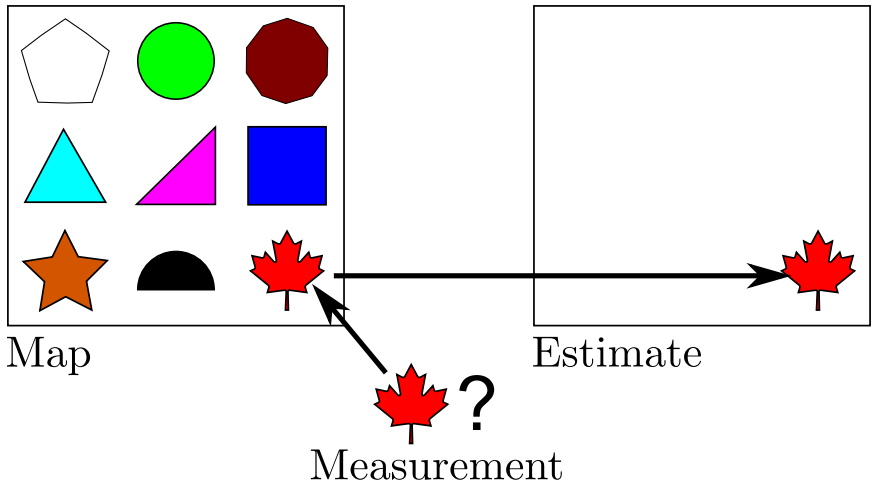

Figure 2.3: Example of a fingerprinting localization system. The closest matching measurement point (maple leaf) is used as the position estimate.

Triangulation

Triangulation uses angle measurements to determine the location of an unknown point. Figure 2.2 illustrates the geometry of triangulation in 2D. Each angle θ_1, θ_2 from known points A, B creates a line on the plane. The unknown point must be at the intersection of these lines. In two dimensions only two measurements are needed, but the point being measured can not be in the same line as the known points. In three dimensions orientation is often measured in spherical coordinates using two angles θ and ϕ, and two pairs of measurements are enough to locate the object.

Scene Analysis

Scene analysis is a different class of localization technique that does not use measurements to solve for the unknown location geometrically. Instead it uses measurements of features such as landmarks in images or the signal strength of radio waves to identify a location based on an existing map. The area must be measured beforehand, and

relevant features are recorded at several locations. The features are used to identify a location in the same way as a fingerprint identifies a person. This analogy leads scene analysis techniques to sometimes be referred to as "fingerprinting" techniques. The receiver uses the point on the map that matches its current measurement the closest in the position estimate, either directly or with some interpolation between points.

A depiction of a fingerprinting scheme is shown in Figure 2.3 where the features at each measurement point are shapes. The object measures a maple leaf, and checks to see that this measurement matches the map in the lower right corner. The object then assumes that since that location is where the same measurement was recorded on the map that is where it must be.

These methods can be very useful when accurate modeling and measurement is not possible or high positioning accuracy is not required. Even the best WiFi signal strength fingerprinting techniques are limited to about 2m accuracy [30]. These techniques require the feature map to be made beforehand, and they are limited by how accurate and distinct the features are. If the features change without the map being updated the estimate will be incorrect. If the features are not distinct, say every shape in the Figure 2.3 was a maple leaf, then the object cannot distinguish one location from another and will produce a poor estimate.

2.1.2 Measurements in Localization Systems

Many indoor localization systems use measurements of an incoming signal to estimate the information needed for the algorithms presented in Section 2.1.1. There are several different methods of extracting position information from a signal, the most common of which are described briefly below.

Angle of Arrival (AOA)

Angle of Arrival (AOA) (sometimes called Direction of Arrival (DOA)) measures the angle between the receiver and an incoming signal for use in triangulation. Detecting the incident angle of a signal can be accomplished with multiple antennas at known distances apart measuring the difference in signal phase, or by exploiting the directionality of a signal and receiver.

Time of Arrival (TOA)

Time of Arrival (TOA) uses prior knowledge of the propagation speed of a signal to estimate how far it has traveled. In ultrasonic systems the time of flight of a sound pulse is compared to the speed of sound, while in RF or optical systems electromagnetic waves are measured against the speed of light. TOA provides an estimate of the distance between the transmitter and receiver that can be used in trilateration. TOA requires precise timing and synchronization between the transmitter and receiver to accurately estimate the distance traveled.

Received Signal Strength (RSS)

Received Signal Strength (RSS) uses knowledge of the emission and propagation of a signal to compare the amount of power received with a known amount of power sent by the transmitter. Many signals decay with distance, such as $1/d^2$ for free-space RF propagation, which accounts for how the signal power spreads out as it gets further from the transmitter. This dependence can be used to estimate the distance between the transmitter and receiver. There are many other factors that affect the received signal energy, including attenuation, absorption, scattering and directionality of the

receiver. It is challenging to develop an accurate propagation model that accounts for all of these effects and can produce an accurate distance measurement [12].

2.1.3 Signals in Indoor Environments

Indoor environments have very complex geometry. There are walls, ceilings, windows and furniture all made of different materials and placed at different locations. The surfaces of these objects can reflect and scatter signals so that they bounce around the room before being detected at the receiver. The path of the bounces can be very complex and practically impossible to predict. This affects the measurement accuracy as signals that bounce have longer path lengths and arrive at different angles than they were transmitted. The effect of combined reflections is called multi path distortion.

Another problem that arises in indoor localization is random deviations in the attenuation of a signal, known as signal fading. Measurements are often based on assumptions about the propagation path of the signal. The receiver does not know about the objects in the way or how they affect the properties of the signal it is trying to measure. In indoor environments doors can open, people can walk between the transmitter and receiver, and a number of other factors can cause the signal strength to change for the same location. These changes are difficult to model and can be mis-interpreted by the receiver as changes in position.

2.1.4 Metrics for Localization Systems

There are several important metrics for measuring and comparing localization systems. Positioning performance is the most obvious one as it quantifies how good the system is at actually locating an object, but it is not enough to judge the suitability of

a system by its accuracy alone. Cost, complexity, scalability, robustness, and privacy are all important factors to consider when assessing an indoor localization system.

Positioning Performance

Positioning performance quantifies the ability of the system to produce location estimates that are close to the actual location of the object. The accuracy of the system is the difference between the location estimate $\hat{\xi} = [\hat{x}, \hat{y}, \hat{z}]$ and the actual object position $\xi = [x, y, z]$. Accuracy of a localization system is usually reported as average error, or mean squared error (MSE) which is the expected value of the position error squared as shown in equation (2.1). Many works suggest that sub-meter accuracy is required for indoor positioning [31, P. 32], but there has been little work done to define a target accuracy for a given application. Systems with 15-30cm accuracy are among the best reported, and considered to be high accuracy [14].

$$MSE(\hat{\xi}) = E[(\hat{\xi} - \xi)^2] \tag{2.1}$$

The estimate variance $\sigma^2_{\hat{\xi}}$ describe the precision of estimates produced by the localization system. It intuitive use standard deviations, $\sigma_{\hat{\xi}}$, as they can be used to describe the percent chance that an estimate will fall within a given distance. A lower variance means a lower standard deviation and a higher precision. Variance is not reported in the literature as well as accuracy. Meter level accuracy systems report between 50-90% of estimates within 1-6 m, while some systems report an impressive 99% of estimates within 30 cm [13, Table 1]. Likewise a target variance has not presented previously.

Coverage Area, Capacity, and Scalability

The coverage area of the system is the size of the region it can successfully track targets in. Coverage area is just as important as the positioning accuracy, as it determines how many sensors or additional hardware need to be installed for a given area. High accuracy over a few square feet is not useful for a store-wide positioning system, but neither is a system that covers the whole store with very low accuracy.

The maximum number of objects a system can track is also important depending on the application. Tag based indoor positioning systems designed for small offices with few employees do not need to handle many separate position estimates simultaneously, while a system for navigating an airport would need to be able to service hundreds or even thousands of users at a time.

Scalability describes how easy it is to expand the capabilities of a system once it has been installed. Either to cover more area, provide better positioning performance, or track more users. A system that is flexible and expandable represents a more attractive investment than one that is not.

Update Frequency

The update frequency of the system is how often it can produce a position estimate. This sets the limit on how fast an object the system can track. For a user walking through a building the update frequency should at least match their step frequency, usually 2-3Hz. There is little benefit to updating faster than 10 Hz [18].

Cost

The cost of the system is a major deciding factor in whether it is suitable for commercial application or not. Cost includes the price of hardware for the system, its installation and maintenance, and its power consumption. In general installing new infrastructure specifically for positioning makes the system more expensive. In certain applications, such as industrial factories or research labs this cost may be viewed as installing any other piece of equipment, but for businesses and public deployment it is preferable to make use of existing infrastructure to reduce cost.

Complexity

Complexity is related to cost through the amount of hardware, set up, and maintenance required but is also represents considerations of the reliability, ease of implementation, and convenience of the system. Systems that require a large amount of processing power to run the required algorithms might be unsuited for implementation on mobile devices, while systems that require additional tags or set up for the user are much less attractive in general purpose applications than ones that do not.

Privacy and Security

Privacy and security are especially relevant in localization systems. Position information is a large component of surveillance which raises a number of privacy concerns, especially if the user of the system is unaware his or her location is being broadcast. WiFi systems are especially prone to abuse as the user does not know what the radio is doing when they think they have put the device away. Google has faced legal action in the US and Europe over the mapping and collection of data from unsecured

wireless networks without the owner's consent [32]. Security of the system is also important to ensure the position information is not obtained by an outside party, or otherwise misused.

2.2 Principles of Light

This section describes some of the basic measurement quantities and properties of light that are needed to discuss its application to indoor positioning.

2.2.1 Radiometric and Photometric Quantities

Visible light is electromagnetic radiation composed of wavelengths perceived by humans, typically between 380 and 780 nanometers. There are two measurement systems used to describe the properties of light. Radiometry measures the power propagating in the electromagnetic wave, while photometry measures the same radiation in terms of its physiological response in the human eye. Photometry accounts for the wavelength dependent response of the human eye by weighting the power over the visible wavelengths using a luminosity function $V(\lambda)$ [33, Ch. 10]. This relates the physical spectral power density (SPD) of the light to its physiological perception by the eye.

The important radiometric quantities, their photometric counterpart and the conversion between the two are explained below. This information was mainly gathered from references: [34–36].

Figure 2.4 shows the geometry of the source element $\mathrm{d}S$, the surface element $\mathrm{d}A$, and the solid angle subtended by the surface element $\mathrm{d}\Omega$. The angle between $\mathrm{d}A$ and

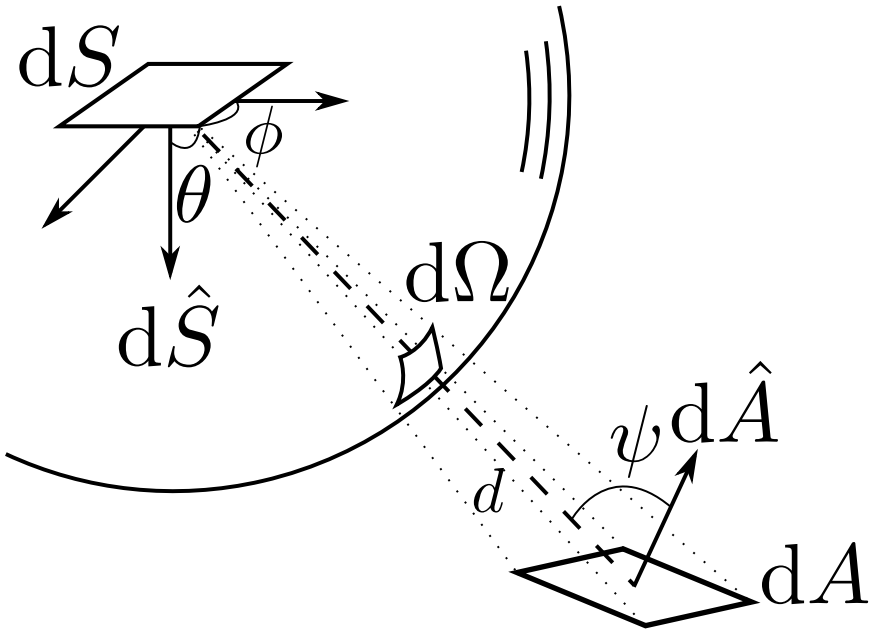

Figure 2.4: Geometry of surface elements dS, dA, and solid angle dΩ used in photometric and radiometric quantities.

the source normal d$\hat{S}$ is θ, while the angle between the source normal and the surface normal d$\hat{A}$ is ψ. The distance between the source and surface is d. These parameters will be used to describe the optical systems in this work.

Radiant and Luminous Power or Flux

The radiometric measure of power is radiant power or radiant flux Φ_e. It is the total amount of power emitted by a light source over all wavelengths in units of Watts W. The spectral power density (SPD) of the source,$\Phi_e(\lambda)$, describes how much power is at each wavelength.

$$\Phi_e = \int_{\lambda_{min}}^{\lambda_{max}} \Phi_e(\lambda) \mathrm{d}\lambda \tag{2.2}$$

The photometric equivalent is the luminous flux Φ_v. It is the total power weighted

by wavelength according to the typical response of a human eye, measured in units of lumens lm. To obtain the luminous flux from the radiant flux the radiometric SPD is multiplied by the luminosity function and integrated over the visible wavelengths.

$$\Phi_v = K_m \int_{380nm}^{780nm} \Phi_e(\lambda) V(\lambda) \mathrm{d}\lambda \tag{2.3}$$

Where $V(\lambda)$ is typically the CIE Photopic luminosity function and $K_m = 683\,\mathrm{lm/W}$ is a constant that fixes the relationship between the physical SPD and physiological luminosity curve [37]. Any radiometric quantity can be converted to its photometric equivalent using this relationship between radiant and luminous flux.

Radiant and Luminous Intensity

Radiant intensity I_e is the amount of power emitted per unit solid angle measured in Watts per steradian W sr^{-1}. The photometric equivalent is luminous intensity I_v measured in lumens per steradian lm/sr which is also known as candela cd.

$$I_e = \frac{\mathrm{d}\Phi_e}{\mathrm{d}\Omega} \tag{2.4}$$

Intensity describes the directionality of the emission of light. Radiant intensity describes how much power is traveling in a certain direction from the source, while luminous intensity relates this to the perceived brightness. The eye collects power over a small angle subtended by the pupil, so a source that emits a lot of power in that direction will appear extremely bright. There are important standards for eye safety with optical sources including IEC-60825 for coherent sources [38], and more recently IEC-62471 for high brightness LEDs [39].

If the distance from the light source is large enough that the radiant intensity

appears to only change with angle $I_e(\theta, \phi, d) \approx I_e(\theta, \phi)$ then the source can be approximated as a point. This is known as the far-field condition and leads to simplifying assumptions when describing optical systems. The exact distance for the start of the far-field region depends on the geometry of the source, but a general rule is to use five times the largest source dimension [34].

Irradiance and Illuminance

Irradiance E_e is the amount of power that falls on a unit area. It has units of Watts per meter squared W m^{-2}. The photometric equivalent is illuminance which has units of Lumens per meter squared lm m^{-2}, also known as lux. Illuminance E_v is the primary measure of how well a room is lit. There are standards that specify the minimum amount of lux required for different working conditions [40].

$$E_e = \frac{\mathrm{d}\Phi_e}{\mathrm{d}A} \tag{2.5}$$

Irradiance describes how much power will be collected by a detector over its entire surface area. If the far-field condition is met, a surface element $\mathrm{d}A$ at distance d from the source will subtend a solid angle $\mathrm{d}\Omega$ according to equation (2.6).

$$\mathrm{d}\Omega = \frac{\cos(\psi)\mathrm{d}A}{d^2} \tag{2.6}$$

In Figure 2.4 the angle ψ determines the effective area of the surface seen by the source in the projection. This changes the solid angle subtended by the surface. For the same surface orientation the solid angle also gets smaller with the distance squared. This is known as the inverse square law and is an important assumption in modeling optical systems [34]. Using equation (2.6), a relationship between radiant

intensity and irradiance can be formed to give equation (2.7).

$$E_e = \frac{\mathrm{d}\Phi_e}{\mathrm{d}A} = \frac{cos(\phi)}{d^2}\frac{\mathrm{d}\Phi_e}{\mathrm{d}\Omega} = \frac{I_e \cos(\phi)}{d^2} \tag{2.7}$$

Irradiance and radiant intensity are related by the ratio of the surface element to the solid angle it subtends.

$$E_e = I_e \frac{\mathrm{d}\Omega}{\mathrm{d}A} \tag{2.8}$$

Radiance or Luminance

If the distance to the source is on the same order of magnitude as the dimensions of the source it can no longer be assumed to be a point. This is called the near-field condition [41]. In the near field the distribution of radiant intensity over the surface of the emitter will cause the radiation pattern to look different at different distances. Radiance L_e is used to measure the radiant intensity for extended (non-point) sources. Radiance is defined as the radiant power $\mathrm{d}\Phi_e$ emitted from the projected source area $\cos(\theta)\mathrm{d}S$ per unit solid angle $\mathrm{d}\Omega$. It has units of Watts per steradian meter squared W sr^{-1} m^{-2}. Projected source area is used due to the interpretation of luminance as the apparent brightness of an object.

$$L_e = \frac{\mathrm{d}^2\Phi_e}{\cos(\theta)\mathrm{d}S\mathrm{d}\Omega} \tag{2.9}$$

The photometric equivalent is Luminance L_v measured in candela per meter square cd m^{-2}. Luminance is the luminous intensity per projected source area:.

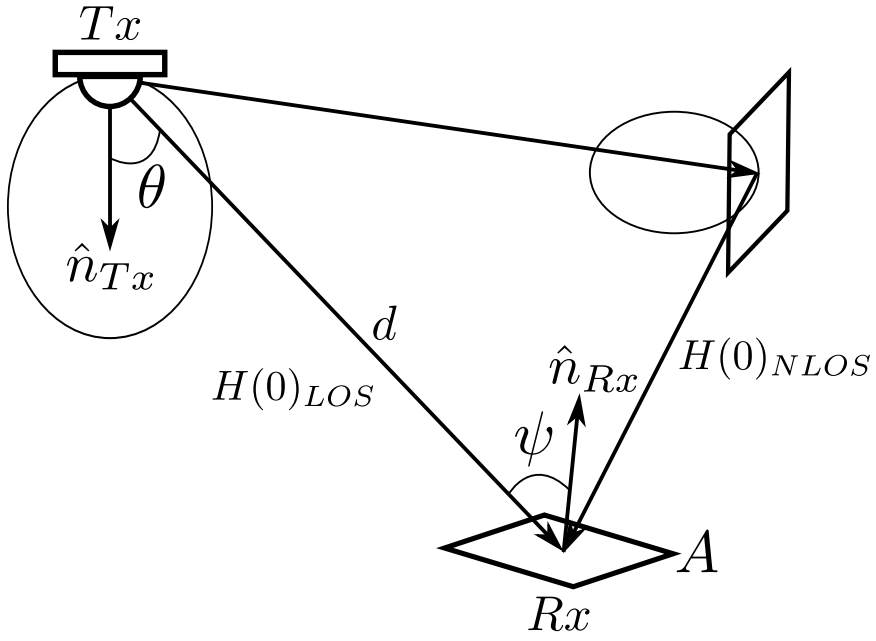

Figure 2.5: Geometry of the VLC Channel. (following definitions used in Figure 2.4)

$$L_v = \frac{\mathrm{d}^2\Phi_v}{\cos(\theta)\mathrm{d}S\mathrm{d}\Omega} = \frac{\mathrm{d}I_v}{\cos(\theta)dS} \tag{2.10}$$

In the same way as a tilted surface $\mathrm{d}A$ appears smaller to the source, the source area $\mathrm{d}S$ appears smaller to an observer at different angles. The combination of how much power is emitted (or reflected) in a particular direction and the apparent surface area that power is spread over determines how bright an object appears to a human observer. In older texts luminance and brightness were used interchangeably.

2.3 The VLC Channel Model

The VLC channel model is a mathematical description of the propagation of an optical signal from transmission through to detection at the receiver. In indoor positioning

the channel model can be used to estimate geometric positioning information by comparing the received signal to a what was sent by the transmitter. A description of the channel model that was originally developed for infrared communications, but is directly applicable to VLC, is presented in [42]. The channel geometry is shown in Figure 2.5 and each component is briefly explained below.

The most common VLC channel model is the intensity modulation / direct detection (IM/DD) channel. Since the transmitters are LEDs which are incoherent sources intensity modulation (IM) is used to modulate information onto the instantaneous intensity, or radiant power of the optical signal. This optical signal passes through the channel and is then collected by the receiver. The receiver produces an electrical signal proportional to the radiant power collected over its surface area. This is direct detection (DD) of the incident signal power. The channel is typically free-space where the medium between the transmitter and receiver is air, however the channel also includes any other reflected paths a signal may take to the receiver.

2.3.1 The Transmitter

The largest potential application of VLC is in adding communications functionality to LED illumination products. Therefore the most important VLC transmitter is the high intensity white LED designed for indoor illumination, shown in Figure 2.5 as Tx.

LEDs are solid state semiconductor devices that convert an input current to an incoherent light emission. They are identical in principle to regular diodes, with the distinction that the recombination of electrons and holes in the depletion region releases photons with energy corresponding to frequencies of visible light (IR and

UV LEDs are also possible). This is the inverse of the photoelectric effect, where the frequency of light produced is dependent on the band-gap energy of the semiconductor used. Since LEDs do not produce a very wide spectrum of light, there are two main ways to realize a white LED. The first is to optically combine the outputs of a red, green and blue LED. The human eye perceives the combination as white. The other option is to use a bright blue LED encapsulated in a phosphor coating. Photons from the blue LED are absorbed by the coating and re-emitted in the yellow region of the spectrum. The result is a spreading of power from blue wavelengths into the rest of the spectrum which the human eye observes as white.

Many LEDs can be modeled in the far-field as a generalized Lambertian emitters. The radiant intensity of the LED is approximated by a rotationally symmetric Lambertian radiation pattern:

$$I_e(\theta) = I_0 \frac{(m+1)}{2\pi} \cos^m(\theta) \tag{2.11}$$

Where $I_0 = I_e(\theta = 0)$ is the maximum radiant intensity of the source. The above equation is normalized such that the integral over a full hemisphere is the radiant flux of the source Φ_e. The value m is the Lambertian index which is related to the transmitter half power semi-angle $\Phi_{1/2}$ by equation (2.12) .

$$m = -\frac{1}{\log_2(\cos(\Phi_{1/2}))} \tag{2.12}$$

Usually $m = 1$ is used, which corresponds to a half power semi-angle of 60°. Larger values of m correspond to a more tightly focused radiant intensity pattern.

2.3.2 The Receiver

The receiver element Rx in Figure 2.5 is typically a photodiode. Photodiodes are solid state semiconductor devices that produce a current proportional to the optical flux incident on their surface area. The material band gap and structure of the device determines which wavelengths it is most sensitive to. The photodiode responsivity curve $R(\lambda)$ is a measure of how much current is produced for a given optical power and wavelength, usually given in Amps per Watt A/W. This is a similar weighting function as the luminosity function, $V(\lambda)$, but specific to the physical response of the semiconductor material rather than the human eye. The total current produced i_r is:

$$i_r = \int R(\lambda)\Phi_{e,receiver}(\lambda)\mathrm{d}\lambda \tag{2.13}$$

Where $\Phi_{e,receiver}(\lambda)$ is the SPD of the signal incident on the receiver. The amount of flux collected by the receiver depends on the source radiant intensity and the surface area of the detector seen by the source. Using the inverse square law this is given by:

$$\Phi_{e,receiver} = \frac{A}{d^2}\cos(\psi)I_e(\theta,\phi) \tag{2.14}$$

The maximum signal is produced when the receiver is oriented towards the source, maximizing the projection of its surface area. To increase the total amount of signal output the surface area of the photodiode can be increased, but this comes at the price of increased capacitance which limits the device bandwidth. Other possibilities are to use a concentrator to focus more light onto the detector, effectively increasing its surface area [42].

2.3.3 The Channel Model

The VLC IM/DD channel can be modeled as a linear channel with additive white Gaussian noise:

$$Y(t) = RX(t) \otimes h(t) + N(t) \tag{2.15}$$

Where the output electrical signal $Y(t)$ is the result of the detector responsivity R multiplied by the input optical signal $X(t)$ and convolved with the channel impulse response $h(t)$. The signal is corrupted by additive white Gaussian noise $N(t)$. $X(t)$ is the transmitted instantaneous optical power which is the convolution of the input electrical signal with the response curve of the LED.

Channel Response

For VLC channels it is convenient to separate the channel into two components as shown in Figure 2.5; the line-of-sight (LOS) component and non-line-of-sight (NLOS) component [42]. The total channel response $h(t)$ is the sum of these two components as shown in equation (2.16).

$$h(t) = h_{LOS}(t) + h_{NLOS}(t) \tag{2.16}$$

The LOS channel response is the effect of the channel on the signal that travels straight from the transmitter to the receiver. Unlike outdoor free-space optics, indoors there is no appreciable scattering, distortion or absorption from the air because the distances are small. The only effect is free-space propagation loss and a time delay Δt_{LOS} due to the propagation speed of light, i.e,

$$h_{LOS}(t) = H(0)_{LOS}\delta(t - \Delta t_{LOS}). \tag{2.17}$$

In this case the LOS channel given by equation (2.17) can be characterized by the DC gain $H(0)_{LOS}$ which directly relates the transmitted and received power. Combining equations (2.11) and (2.14), the DC gain is given by

$$H(0)_{LOS} = \frac{\Phi_{Receiver}}{\Phi_{Source}} = \frac{(m+1)A}{2\pi d^2}\cos(\psi)\cos^m(\theta). \tag{2.18}$$

The NLOS channel response contains the portion of the signal that arrives at the receiver after reflecting and scattering off objects in the environment. It is more difficult to model accurately due to the complex multiple scattering and reflections from different surfaces. It can be simulated recursively for a particular environment by treating the reflection off each surface element in the room as a new Lambertian transmitter that emits the percentage of light reflected by that surface [43]. The channel impulse response for the k^{th} bounce is $h^{(k)}(t; Tx, Rx)$ and the NLOS response is the sum of any path with one or more bounces.

$$h_{NLOS}(t) = \sum_{k=1}^{\infty} h^{(k)}(t; Tx, Rx) \tag{2.19}$$

The total channel impulse response for the k^{th} bounce depends on all previous bounces, and so it is calculated by a recursive convolution that starts with the 0^{th} bounce (the LOS path) using equation (2.20).

$$h^{(k)}(t; Tx, Rx) = \sum_{i=1}^{\infty} h^{(0)}(t; Tx, \epsilon_i) \otimes h^{(k-1)}(t; \epsilon_i, Rx) \tag{2.20}$$

Equation (2.20) shows that the reflecting element ϵ_i starts as a receiving element that the light from the transmitter lands on, then in the next bounce it is treated as a new transmitter that is emitting the light it reflected. This method requires knowledge of all of the geometry and materials in the room being simulated, and has a high computational complexity.

The NLOS channel response is often ignored in VLP techniques because it is much harder to work with and there is no practical way of being able to know the geometry for every single room beforehand. However, if the receiver is located near even a diffuse reflecting surface (light is scattered equally in every direction) the first bounce can still contain approximately 10% of the power seen by the LOS path [43]. This is enough power to affect a distance measurement by nearly 5% using the inverse square law of equation (2.6), or skew an angle measurement by several degrees. If the reflecting surface is specular (reflections are directional like a mirror) the effect is far more pronounced. It is important for a system to be robust to the NLOS channel response.

2.4 Intensity Based Visible Light Positioning Techniques

Intensity based VLP techniques use measurements of a modulated visible light signal and knowledge of the channel model to estimate the receiver position. Each VLC luminary broadcasts information about its position and transmit power that the receiver can detect. The receiver makes a position estimate by comparing the received signal strength of the transmission to the values contained in the message, and attempting

to infer the distances or angles needed to apply one of the localization techniques presented.

Solving for Distance Using the Channel Model

From Figure 2.5 it can be seen that the visible light signal measurements depend not only on the distance d but also the angle between the transmitter and receiver normal ψ. This extra degree of freedom makes locating the receiver difficult because θ is an angle in the global coordinate frame, while ψ is an angle in the receivers local coordinate frame. The measurement being used to extract position information depends on the orientation of the receiver, not just it's position. Without knowledge of the receiver's orientation six separate measurements are required to determine its location, and the equations are non-linear and sensitive to measurement errors.

This orientation dependence is problem common to all intensity based VLP approaches. Many works avoid the problem altogether by restricting the receiver axis to be parallel to the transmitter axis as shown in Figure 2.6. In this case, $\theta = \psi$ and the channel LOS DC gain reduces to equation (2.21).

$$H(0)_{LOSAligned} = \frac{\Phi_{Receiver}}{\Phi_{Source}} = \frac{(m+1)A}{2\pi d^2}\cos^{m+1}(\theta) \tag{2.21}$$

By supposing the receiver is at a known height h, and using $\cos(\theta) = h/d$ the distance to each luminary d can be solved for directly using equation (2.22).

$$d = \sqrt[m+3]{\frac{\Phi_{Source}}{\Phi_{Receiver}}K} \tag{2.22}$$

In equation (2.22) K is the combined effect of the height, surface area, Lambertian order and any collector gain or filter gain G [44]. It can be determined from

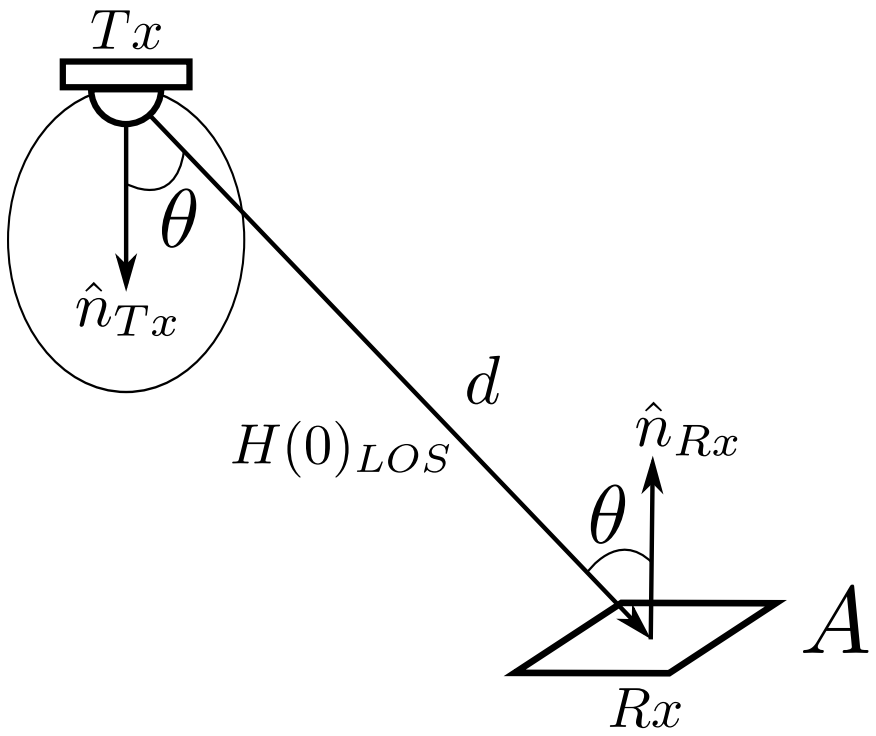

Figure 2.6: Special case of Figure 2.5 where the receiver is perpendicular to the transmitter so that the transmitter and receiver normals are aligned. In this case $\theta = \psi$ which reduces the degrees of freedom in the channel model.

equation (2.23).

$$K = \frac{(m+1)}{2\pi} AGh^{m+1} \tag{2.23}$$

In this case the distance estimate depends only on the ratio of sent to received signal power $\Phi_{Source}/\Phi_{Receiver}$. This is an overly optimistic simplification. If the receiver is no longer perfectly aligned with the transmitters, or $\theta \neq \psi$, the channel model cannot be reduced in this manner. Also this method requires the receiver height h to be known. This assumption may hold for a robot or cart that guarantees the receiver will remain parallel to the luminaries in the ceiling and at a fixed height, but it is unrealistic for a user carrying a mobile device.

2.4.1 Prior Works

Working directly with the channel model is not possible without knowing the orientation angle ψ. Previous work in intensity based VLP have devised different solutions for working around the orientation dependence of the channel model. Some of the techniques proposed and demonstrated are surveyed below.

IMU Assisted

An early attempt at positioning using intensity measurements from LED lights was a single element receiver combined with an IMU accelerometer to detect changes in orientation [45]. This so called "Switching Element Receive" detected the visible light ID (VLID) from different transmitter luminaries and estimated their distance from the receiver. It made use of a 6 axis accelerometer to provide information on the changes in receiver orientation, and recognized that a narrow receiver FOV could help

isolate the orientation before making the distance measurement. The key observation was that if a light was visible the receiver orientation ψ must be bounded to within the FOV angle. This helped remove the ambiguity in orientation and allowed the distance to be estimated. The reported simulation accuracy of this system was 30 cm to 1 m depending on the value of ψ.

Several other works also propose using accelerometer and gyroscope measurements to help estimate the orientation of the receiver [46–48]. A recent approach used multiple measurements at the same position but with the receiver tilted to different orientations [49]. An accelerometer provided information on the angular separation between measurements, and the difference in intensity and angle measurements was used to solve for the receiver location in three dimensions. The position accuracy in simulation was reported to be 20 cm for high enough SNR (>35 dB), but combining the measurements in practice would be challenging as the receiver will move between measurements.

Differential Photosensors

Another proposed method is to use measurements from multiple photodetectors arranged with different orientations. In one implementation a corner cube was constructed from three separate photodetectors arranged orthogonally [50]. The signal on each face could be compared to make AOA measurements of the incoming light. The reported accuracy was 4 cm, but the receiver had to be aligned with the global coordinate frame. This approach also had a limited effective FOV, since the lights needed to fall on all three faces of the detector.

A more recent implementation used six photodetectors oriented in a half dodecahedron [51]. Intensity measurements from each sensing face were used to produce a position estimate with an average error of <40 cm, and a standard deviation of 20 cm. However this approach still required an IMU and calibration to provide orientation information. With a smart phone implementation of trilateration the authors were able to achieve 45 cm to 1 m accuracy while the receiver was held straight up and had its heading aligned.

Fingerprinting Techniques

Scene analysis approaches can be used with VLP as well. The light intensity can be used to make a map of measurement points over the entire area of interest. The receiver compares the measurements it takes to the database of previous measurement points and chooses the nearest match as its position estimate. One such system used multiple photodetectors arranged in different directions and used the power spectral density of the received signal as the measurement fingerprint [52]. Another system that used the received signal strength from multiple QPSK transmitters was demonstrated in [53]. Calibrated measurements over the $60 \times 60 \times 60$ cm^3 test area were be used to locate the receiver to within 4 cm. In both cases the receiver was positioned at a fixed distance from the transmitter and perpendicularly aligned.

Hybrid Approach

An interesting hybrid approach was proposed that used relative signal strength of visible luminaries to form a weighted average coarse position estimate [54]. This method is more akin to fingerprinting than channel model approaches as it only uses

the relative strength of the visible luminaries to form a weighted average position estimate. This method is still dependent on the receiver orientation, however, as changes in orientation will alter the relative intensities of the visible light signals. Transmitters the receiver is pointing towards will be more favorably weighted. The simulation accuracy of the coarse positioning algorithm was reported as 35 cm.

The second part of the approach proposed in [54] was to use the coarse positioning algorithm as the starting point for a more accurate algorithm using signal strength measurements and the channel model. By having a rough idea of the location of the receiver some of the ambiguity in orientation can be removed. Simulations of the fine positioning algorithm reported an average error of 15 cm, but the approach also assumed the transmitter and receiver were aligned, and that the receiver could measure the angle of incidence of the light signal.

2.5 Camera Projection Theory and Imaging Based VLP

The same theory that allows a camera to capture an image can be used for indoor positioning. There have been several optical indoor positioning systems proposed [17], but VLC has enabled a new approach. Imaging based approaches to VLP seek to avoid the problems with channel model measurements entirely by recasting the positioning problem as a perspective projection using an imaging receiver. The VLC luminaries act as reference points in the image that the receiver can use to solve for its position.

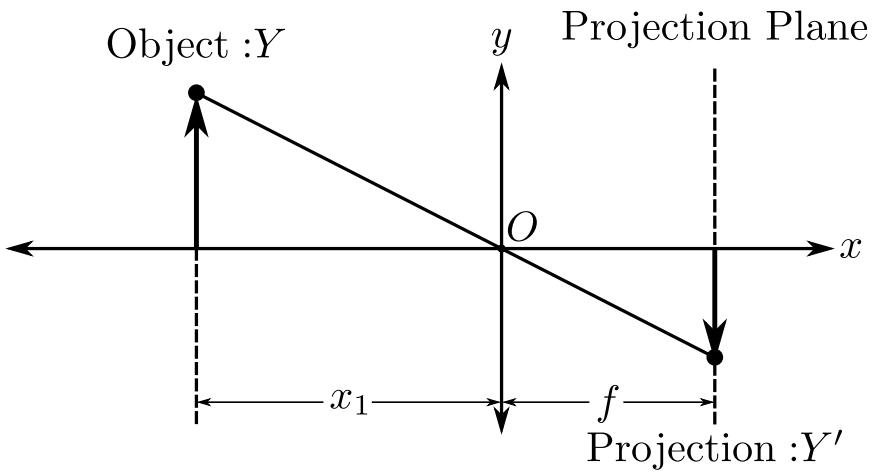

Figure 2.7: Perspective projection in two dimensions.

2.5.1 Camera Perspective Projection

Cameras capture images by performing a perspective projection of light from the scene onto their image sensor. A perspective projection is a projection of points onto a lower dimensional space such that they are scaled by how far away they are [55, Sec. 1.4]. This is the same type of projection our eyes perform which makes objects further away appear smaller and closer objects appear larger. The relationship between distance to the object and its size in the image can be used for positioning.

Figure 2.7 shows the perspective projection of the object point Y onto the projection plane at $x = f$. The projection passes through the origin O to create point Y'. The relationship between the coordinates of point Y and its projection Y' can be determined by similar triangles using the line segment connecting them as the hypotenuse:

$$Y' = -\frac{f}{x_1}Y \tag{2.24}$$

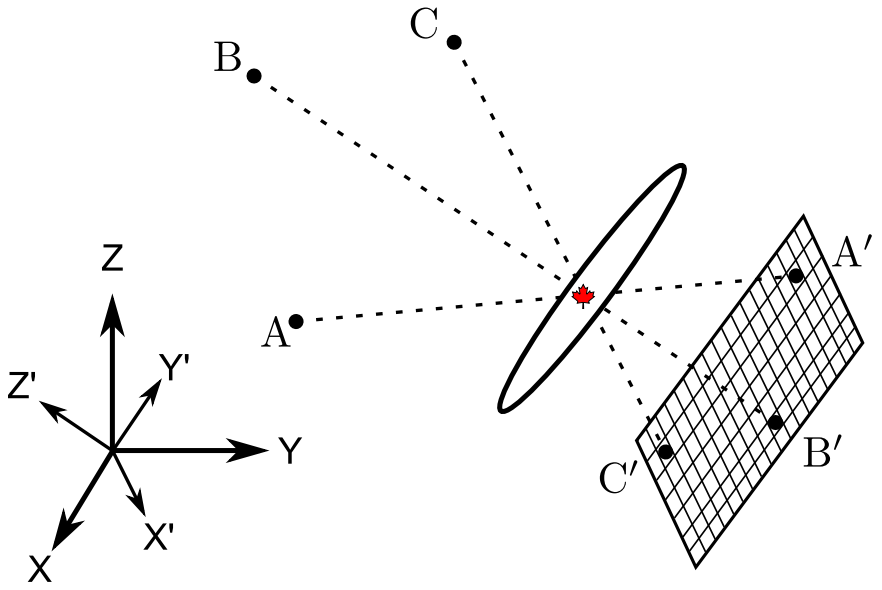

Figure 2.8: Perspective projection of the points A,B and C with an imaging receiver. The maple leaf is the vanishing point of the system, which is used as the position estimate.

The projection appears on the opposite side of the x axis, which accounts for the negative sign. The ratio of f/x_1 determines the scaling factor of the projection where x_1 is the distance from the aperture to the object along the common axis (x axis in the figure) and f is the distance from the aperture to the projection plane. This scaling based on distance is the basis for using perspective projection in indoor positioning. If the object height, projection height, and distance f are known the distance can be determined.

Figure 2.8 shows a 3D perspective projection of three known points A, B, C onto a multi-element receiver (camera CCD or CMOS sensor). The projection is performed by lenses which capture and focus the light, and will be explained in more detail in the next chapter. If the projection is ideal each point is mapped directly to its projection. The location of each point is known, and their projected coordinates

A', B', C' can be determined by identifying which pixel the spot was projected on to. Note that the coordinates of points A', B', C' are in the cameras local coordinate frame, which depends on how the camera is positioned and oriented relative to the global coordinate frame. Since points A, B, C are known in the global coordinated frame rotation matrices must be used to align the coordinate systems. These six points can be used to identify the location of the camera. The trigonometry of solving for the vanishing point (the maple leaf where the rays cross) is cumbersome when the axes are not aligned, but is presented in [55, Sec. 1.7].

2.5.2 Imaging Based VLP Techniques

Imaging based approaches to VLP use VLC transmitters in the ceiling as the known points A, B, C to solve for the receiver location. The projection coordinates A', B', C' are provided by addressing the pixels on the image sensor corresponding to their projection. An early work on using imaging techniques for VLP presented a simulation analysis of the factors affecting the position estimate including the number and size of pixels at the receiver, orientation of the receiver plane, and location of the receiver relative to the positioning lights [56]. It showed that imaging approaches have the potential to produce accurate estimates, but still suffer from variations in orientation as well as limited field of view. The density of luminaries needs to be very high to ensure enough lights can be detected in each frame.

There is potential in applying this theory to indoor localization using visible light, however little progress has been made on imaging VLP. The biggest challenge is the ability to correctly identify and map the projection points A', B', C'. While there have been methods proposed and implemented performing low-rate visible light

communications using camera image sensors [57, 58] they are not suited for spatially isolating multiple signals to the correct bright spot on the image plane when the receiver is mobile and measurements need to be correlated across multiple frames. Without knowing which projection came from which transmitter the system cannot solve for its position. The communications capability required to realize imaging VLP is not present in current imaging receiver technology.

2.6 Conclusions

This chapter has presented the basic concepts in indoor localization and VLC required to review existing approaches to VLP. It has been shown that despite the advantages visible light has over other signals used in indoor positioning it also has unique challenges to be overcome. The orientation dependence of the VLC channel model prevents position information from being estimated directly from signal strength measurements, and intensity based approaches have struggled to come up with a means of addressing this dependence. Techniques that worked well for RF localization are not suited for visible light due to the directionality of the receiver. Imaging based VLP techniques cannot be implemented with conventional imaging sensors, but they make an important observation that the directionality of light can be used for indoor positioning. This reformulation of the positioning problem from intensity measurements to geometric calculations motivates a new approach to VLP proposed in the next chapter.

Chapter 3

Angular Diversity Approach to Indoor Positioning Using Visible Light

This chapter presents a new approach to indoor positioning using visible light based on the principle of angular diversity. The proposed method combines elements of imaging and intensity based VLP techniques in a new way that compliments the strengths of both, while simultaneously addressing many of the drawbacks. The directionality of light is examined before proposing angular diversity visible light positioning (AD-VLP). The idea behind AD-VLP is presented and compared to existing VLP approaches, citing the advantages and drawbacks of this new method. A simple design example based on the thin lens equation is presented to develop the concepts of AD-VLP, followed by simulation results of the positioning performance.

3.1 Using the Directionality of Light

Chapter 2 showed that imaging receivers could be used for indoor positioning with visible light, and that they did not use intensity measurements or the channel model. This reformulation of the problem as a geometric projection is possible due to the directionality of light, and specifically its ability to be re-directed using optics. Cameras use systems of lenses and apertures to collect and focus light onto their image sensors, approximating as closely as possible a true perspective projection of the scene they are capturing. The ability to change the direction of visible light using inexpensive materials like glass or clear acrylic is one of the most useful properties of the visible spectrum, and one that is under-utilized in indoor positioning applications.

3.1.1 Geometric Optics

The earliest system developed to describe the properties of light, and the one most relevant to indoor positioning, is geometric optics. Geometric optics treats light as a series of rays that travel in straight lines [59]. The ability to describe the behavior of light using individual rays and geometric equations leads to the interchangeable names "geometric optics" and "ray optics". A thorough description of the properties of geometric optics can be found in [60, Ch. 1-4], but an important example presented next is that of the ideal thin lens.

The Thin Lens Equation

Lenses use the index of refraction of their material and the shape of their surface to focus light that falls on them. With certain simplifying assumptions a simple equation can be derived to describe the optical properties of a lens, known as the thin lens

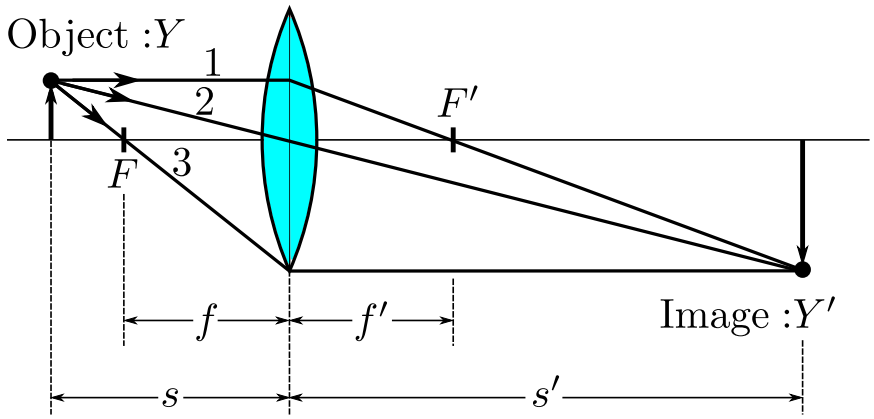

Figure 3.1: Geometry of a simple biconvex lens, modeled using the thin lens equation.

equation.

Figure 3.1 shows rays of light from a point Y traveling a distance s to pass through a biconvex lens. The lens redirects the rays and forms an inverted image of the point Y' at a distance s'. Under the thin lens approximation the shape and thickness of the lens is ignored, and the only quantity needed to describe its optical properties is the focal length f. For a thin lens the primary and secondary focal lengths are identical, so $f = f'$.

The focal length is the distance along the optical axis from the lens to its focal point F. The focal point is the theoretical point all rays parallel to the optical axis will be focused to by the lens. Ray 1 from object Y is traveling parallel to the optical axis, and it is redirected by the lens to pass through point F'. By the same token any ray that appears to originate from the focal point will be redirected to travel parallel to the optical axis. Ray 3 passes through F on its way to the lens, where it is redirected to travel parallel to the optical axis. Where ray 1 and ray 3 intersect is where the image Y' will be formed. It can be shown that *any* ray originating from the object point Y passing through this ideal lens will be redirected towards the image

point Y'. Ray 2 is a special case known as the *center ray* or *chief ray* that passes through the center of the lens and is not redirected.

The thin lens equation, equation (3.1), describes the relationship between the distance of the object s, the focal length of the lens f, and the distance of the image formed s'.

$$\frac{1}{f} = \frac{1}{s} + \frac{1}{s'} \tag{3.1}$$

It can then be shown from similar triangles that the height of the image point Y' is related to the height of the object point Y by equation (3.2). Where the ratio $M = Y'/Y$ is known as the magnification of the lens.

$$Y' = MY = -\frac{s'}{s}Y = \frac{f}{f-s}Y \tag{3.2}$$

3.1.2 Optics at the Receiver: Imaging Receiver

An ideal perspective projection requires an infinitesimally small aperture that allows only a single ray from each point in the scene through. This is not practical since the film or image sensor capturing the image responds to the irradiance on its surface. If very few rays of light pass through the aperture a very long exposure time is needed to form an image. Imaging receivers use lenses to capture much more light over a larger aperture, while still closely approximating a true perspective projection.

It is not difficult to see how lenses can be used to form images. The geometry of Figure 3.1 bears a strong resemblance to the 2D projection shown in Figure 2.7, and equation (3.2) looks very similar to the equation for 2D projection (2.24) presented in Section 2.5. Both take a single point and map it to another point on the image

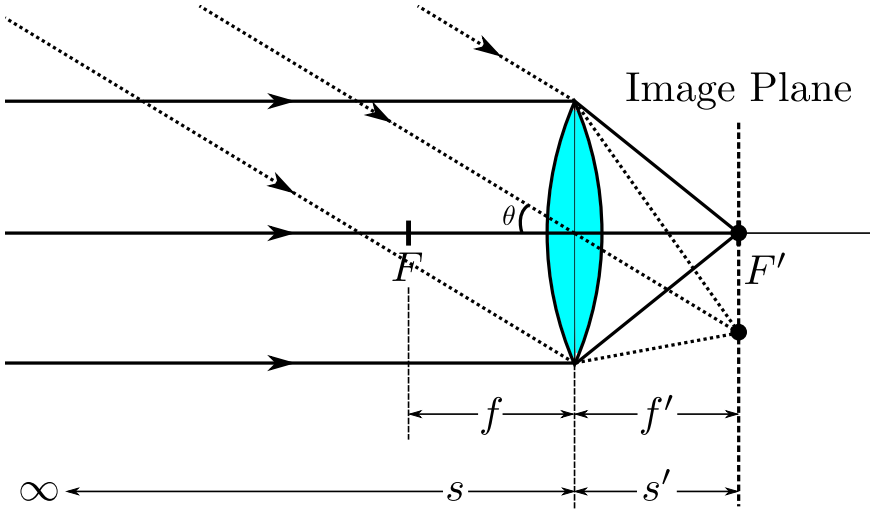

Figure 3.2: Light from far away objects ($s \to \infty$) focused onto the focal plane by the lens.

plane. An image is just the superposition of an infinite many points over the whole scene being captured.

Angle-to-Space Mapping

If the distance from the object to the lens becomes very large($s \to \infty$), then equation (3.1) will reduce to $s' = f$ and all object points will be projected onto the focal plane. This is a special case where all the rays of light from the object point can be considered approximately parallel (equivalent to a planar wavefront), and the lens will correctly image them all to the focal plane as shown in Figure 3.2.

Figure 3.2 demonstrates a very important property of lenses that forms the basis of imaging VLP techniques. Lenses at the receiver perform angle-to-space mapping on the incident light. All the parallel rays of light are imaged to a unique point on the focal plane that depends on their angle. The lens is performing a mapping from angle to height on the focal plane.

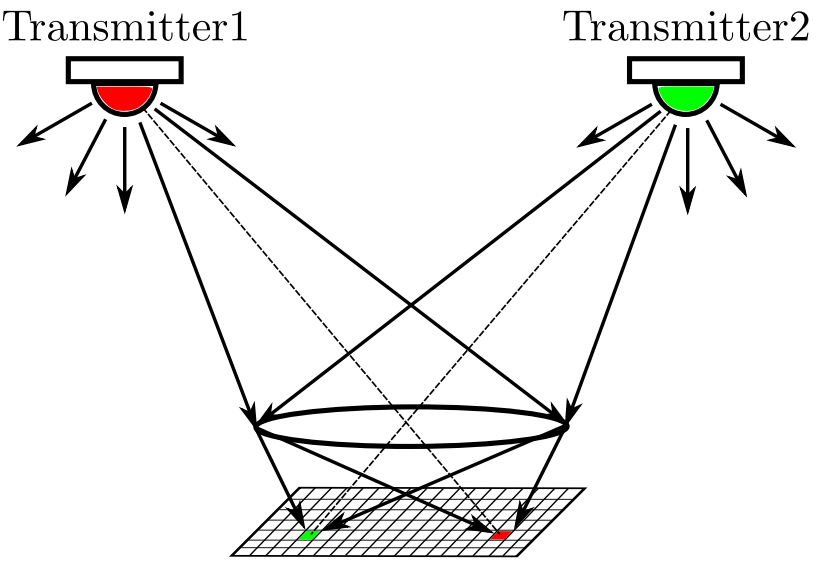

Figure 3.3: Lens in an imaging receiver performing angle-to-space mapping on the light from VLC luminaries in the ceiling.

Figure 3.3 shows how this property is used in imaging approaches to VLP. The light from each transmitter is emitted over an entire hemisphere, but the lens at the receiver subtends only a small set of angles of that light. The light that the lens collects is focused onto the multi-element receiver at a point in space. Since all the light collected came from roughly the same angle, the image point is known thanks to the pixel address it fell on, and the optical properties of the lens is known, the angle the light arrived at can be determined. This is the geometric information required to solve for location of the imaging receiver.

3.1.3 Optics at the Transmitter: Angular Diversity

A basic but often overlooked property of geometric optics is that ray tracing works both ways. In Figure 3.1 the object point Y and the image point Y' could be swapped with nothing more than a change of direction of the rays. This property reveals an

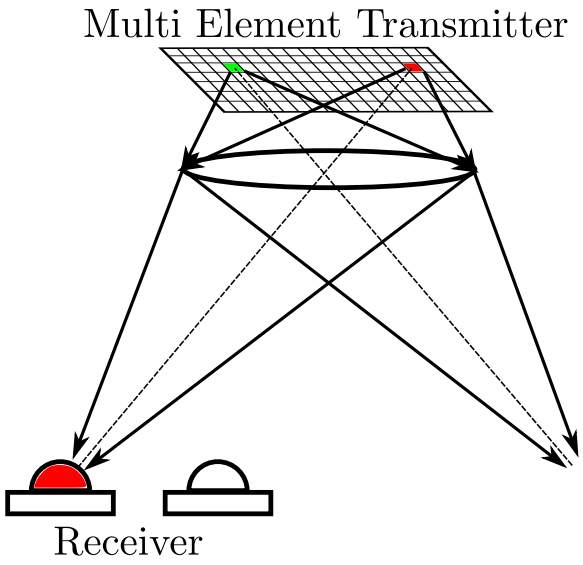

Figure 3.4: Lens performing space-to-angle mapping at the transmitter. The light from each transmitter element is sent in different directions by the lens.

opportunity to use the directionality of light in a new way. Figure 3.4 shows the imaging receiver from Figure 3.3 reversed. The lens is placed in-front of a multi element transmitter, and the receiver is the single element. Now the lens is performing space-to-angle mapping at the transmitter. Since each element in the transmitter is at a different location in space the light from each will be mapped to a different set of angles. This is the principle of angular diversity, and it forms the basis of this works new approach to indoor positioning using visible light.

3.2 Angular Diversity Visible Light Positioning: AD-VLP

Angular diversity visible light positioning (AD-VLP) is proposed as a new method of indoor positioning using visible light based on the principle of angular diversity. It

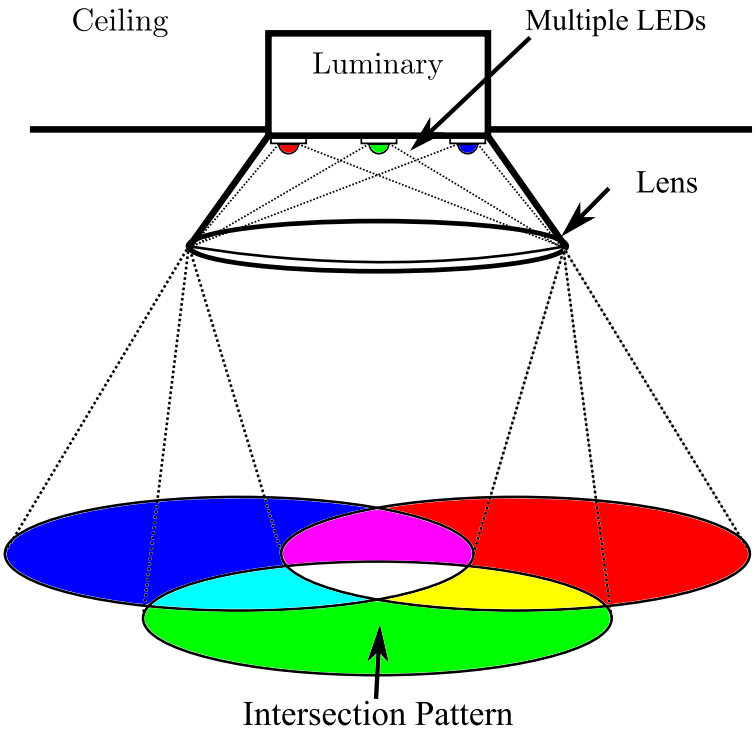

Figure 3.5: Angular Diversity Visible Light Positioning

uses the space to angle mapping shown in Figure 3.4 to encode position information in the structure of the light emitted by the transmitter. The idea is that if each element in the transmitter could be uniquely identified, and the light was sent in different directions by the lens, then there would be a set of directions the receiver could identify just from which lights it could see. This approach allows a single element receiver to identify position information without using the channel model.

3.2.1 Structured Overlapping Light Pattern

The key to AD-VLP is the creation of a structured light pattern at the transmitter that can be detected by a single element receiver. The geometric information is contained in the structure of the light pattern itself, so the receiver only needs to be

able to identify which part of the pattern it is in to form a bound on its position.

Figure 3.5 depicts this core concept implemented in a single transmitter. Commercial LED luminaries contain multiple LEDs to achieve the required illumination levels [61]. Using VLC, each LED can be independently modulated and used as a separate transmitter. This is shown in the figure as different shades for clarity, but in the actual system it would all be white light identified by different signals. Optical versions of multiple access techniques such as frequency division multiple access (FDMA) or code division multiple access (CDMA) could be used to ensure the signals from each LED are orthogonal and uniquely identifiable [62].

The lens captures the light from each LED and sends it in a different direction. In this case the LEDs have been placed slightly inside the focal length of the lens so that the light diverges into cones. This causes the light from each LED to cover a wider area, and it also causes some overlap between the lights. This overlap is extremely useful because it creates more regions that can be uniquely identified by the binary combination of the signals that form them. If a receiver sees the signal from the left LED light it knows it must be in the right circle, but if it also sees the signal from the middle LED light it must be in the overlapping region formed by both. If it only sees the one LED signal it can only be in the outside region where there is no overlap. Each unique region forms a new bound on the receiver position.

Each luminary is part of the greater AD-VLP system that would cover the desired area in a building with many uniquely identifiable regions. Each transmitter luminary must have its own ID so that the regions it produces can be distinguished from other luminaries in the building. A map of the structure where the system is installed must be annotated with the ID of each light and its location. The receiver identifies which

region it is in, and which transmitter formed that region from the optical signals it receives. Then it can use the system map to look up where that region is in the entire building to determine its position. This map could be stored on the mobile device itself, possibly communicated through VLC when a user enters the building, or it could be stored on an external server and accessed through the traditional wireless network.

3.2.2 Advantages of AD-VLP

Positioning using the structured overlapping light pattern described above has several advantages over previously proposed methods of indoor positioning using visible light.

The primary advantage of AD-VLP is its ability to use a simple single element receiver while avoiding the orientation dependence of intensity measurement based VLP techniques. The receiver only needs to identify which light signals are visible, not where they came from or how strong they are. This is largely independent of the receiver orientation since as long as the transmitter is within the receiver's field of view it will see the signal regardless of whether it is pointing towards or away from it. Separating the position information from intensity measurements greatly increases the robustness of the system.

Using optics at the transmitter gives much more control over the position information provided to the receiver. The size and shape of the overlapping regions can be controlled by the design of the optical elements in the transmitter. There is even the possibility of creating luminaries with a custom illumination pattern for particular areas of a structure, such as hallways, cubicle farms, or large conference rooms. This approach to positioning is very similar to scene analysis techniques, with the

advantage that the layout of the scene is under the system designer's control. Other scene analysis approaches relied on whatever signals were available for measurement, while AD-VLP uses its own pattern layout. Compared to other scene analysis VLP techniques AD-VLP does not fingerprint on the intensity measurement of the signal, but on the geometry, again avoiding the orientation dependence. It also creates many unique fingerprinting locations from the overlapping light pattern. An additional advantage shared with scene analysis techniques is that the receiver only needs to see one signal from one LED to form a position estimate.

Another advantage is the modularity and scalability of a system implemented using AD-VLP. Each AD-VLP transmitter does not require external communications infrastructure or synchronization with other AD-VLP luminaries. To upgrade the system a user only needs to install a new luminary and record its location. Each individual luminary guarantees a certain worst case positioning error over the area it illuminates, but as will be shown when multiple AD-VLP luminaries are used together in a room this error drops significantly.

An AD-VLP system does not need to replace every luminary in the structure. The AD-VLP luminary can be designed to cover a certain area with a certain positioning performance, so that only as many luminaries as needed to cover the desired area are required. If it is later decided that better positioning performance is required additional AD-VLP luminaries can be installed, replacing the non-positioning luminaries. An update of the positioning map with the new luminaries installed is all that is needed to refit the system.

3.2.3 Drawbacks of AD-VLP

AD-VLP introduces it own trade-offs between how the light signal is used and how accurate the position estimate is. While these trade-offs are largely to the benefit of the indoor positioning system there are still drawbacks that must be addressed.

The most obvious drawback is that AD-VLP does not provide a point estimate of the receiver, but a region that bounds its location. The receiver has no idea where it is located inside the region. If the regions are large the position estimate could be very poor. By reducing the signal measurement to a binary detection the receiver is throwing away any other information contained in the intensity of the signal. This information could still be useful in forming a more accurate position estimate similar to the coarse / fine example presented in Chapter 2.

Another major drawback is the trust placed in the model of the geometric structure of the overlapping light pattern. Any un-modeled geometry or reflections will distort the pattern and cause the receiver to estimate incorrect position bounds, even if the right region is identified. Methods of setting a region decision threshold to mitigate this are discussed in Chapter 4.

Finally the AD-VLP transmitter is more complex than a regular LED luminary. Modulating each LED in parallel requires more components which increases cost, and reduces the energy efficiency of the luminary [63, P.67]. The cost of adding optics to the system is marginal however as LED lights already use collimating lenses and reflectors to create more uniform illumination distributions [64].

3.3 AD-VLP Design With an Ideal Thin Lens

The theory behind AD-VLP can be demonstrated via a simple design example using a single ideal thin lens at the transmitter. This example is an idealized version of the experimental setup presented in Chapter 5. It is shown here and will be carried through the remaining chapters to further develop the presented theory.

3.3.1 Luminary Geometry

Transmitter LEDs

The pattern design is based on the geometry of the LEDs in the luminary. For this example six LEDs are arranged in a circle evenly spaced 60° apart, with one additional LED at the center as shown in Figure 3.6. This geometry corresponds the LED disc used later in the experimental transmitter. For the purpose of this example the LEDs in the luminary will be considered ideal point source Lambertian emitters.

Since the disc is circularly symmetric and so is the lens, only the middle and top LED need to be considered in the design. The regions formed by the outer LEDs will be rotated versions of the region from the top LED.

Center LED

Figure 3.7 shows the resulting projection of the light from the center LED. According to the thin lens equation, the distance from the lens to the focus point of the point source is given by equation (3.3).

$$s' = \frac{sf}{s - f} \tag{3.3}$$

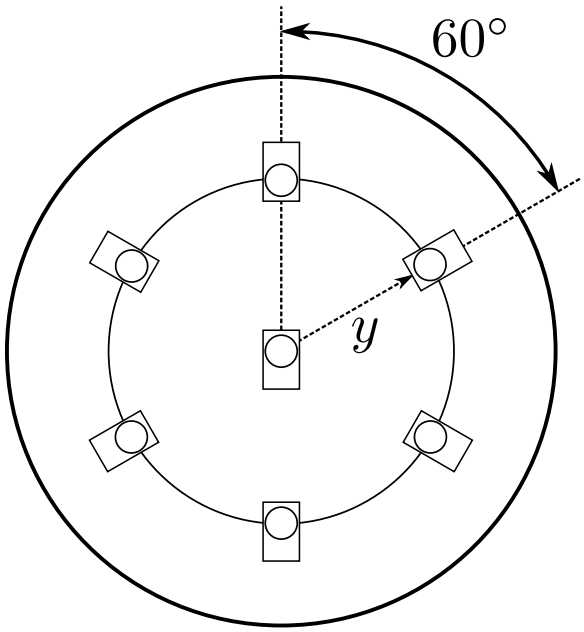

Figure 3.6: Geometry of the LED disc used in the transmitter.

In this case $s < f$ and s' will be negative. This forms a virtual image point behind the lens where all the rays appear to originate from. This has the advantage of causing the rays to diverge which allows overlapping regions to be formed. It can be seen that the furthest rays form the edges of a circular cone with the image of the LED as the apex.

This lends itself to the edge ray principle of non-imaging optical systems. The edge ray principle is used to design non-imaging optical systems by only modeling the furthest rays (or edge rays) in the transmitter [65]. The rays that strike the edges of the lens will determine the outer boundary of the light from each LED, and hence the geometry of the structured overlapping light pattern.

Using the edge ray principle the entire geometry of the light pattern can be described in terms of the cones of light produced by the lens. The center cone can be parameterized by its apex and its cross sectional radius r_{cir}. For a given height h the

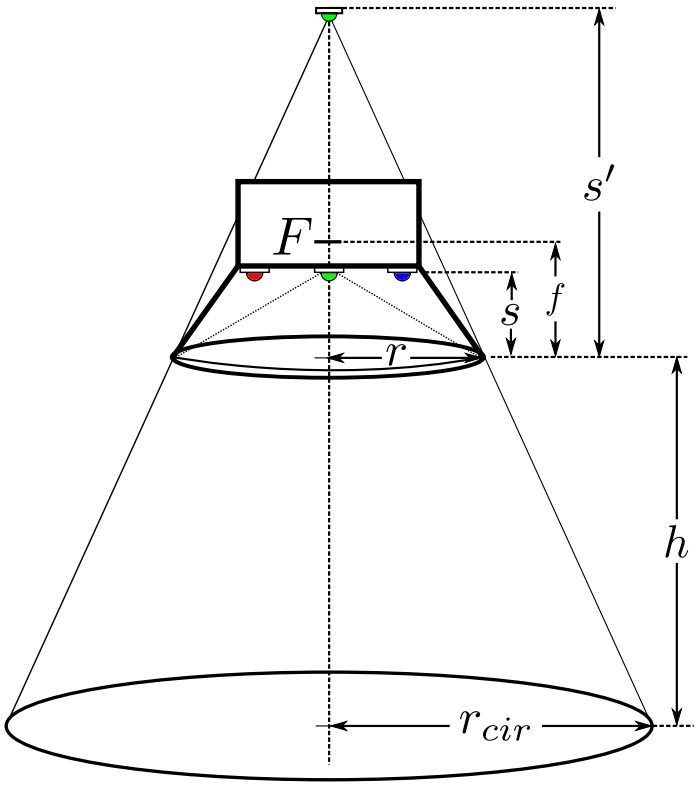

Figure 3.7: On-axis LED placed inside the focal length of the lens. A virtual image is formed behind the LED by the lens.

radius of the cone is given by similar triangles as equation (3.4).

$$r_{cir} = \frac{r(h + s')}{s'} \tag{3.4}$$

Equation (3.4) can also be expressed in terms of the lens variables using equation (3.3) to give equation (3.5).

$$r_{cir} = r\left(\frac{h(f - s)}{sf} + 1\right) \tag{3.5}$$

Equation (3.5) shows that as $s \to f^-$, $r_{cir} \to r$, or that as the LEDs are moved to the focal plane all the rays of light become parallel and r_{cir} becomes the lens radius. The radius r serves as a scaling factor for the regions. It does not affect the virtual image point but it determines how far out the edge of the cone will be. The focal length f and LED distance s both change the virtual image point, which in turn affects r_{cir}.

Off-Axis LED

Figure 3.8 shows the projection of the light from any one of the off axis LEDs. In this case an oblique circular cone is formed along the center ray. This directionality along the center ray creates the angular diversity between the light from each LED. For a given height an identical circle is formed to the one in the on-axis case, but centered along the center ray. The distance from the optical axis to the center of the projection r_{cen} can be determined by similar triangles which gives equation (3.6).

$$r_{cen} = \frac{hy}{s} \tag{3.6}$$

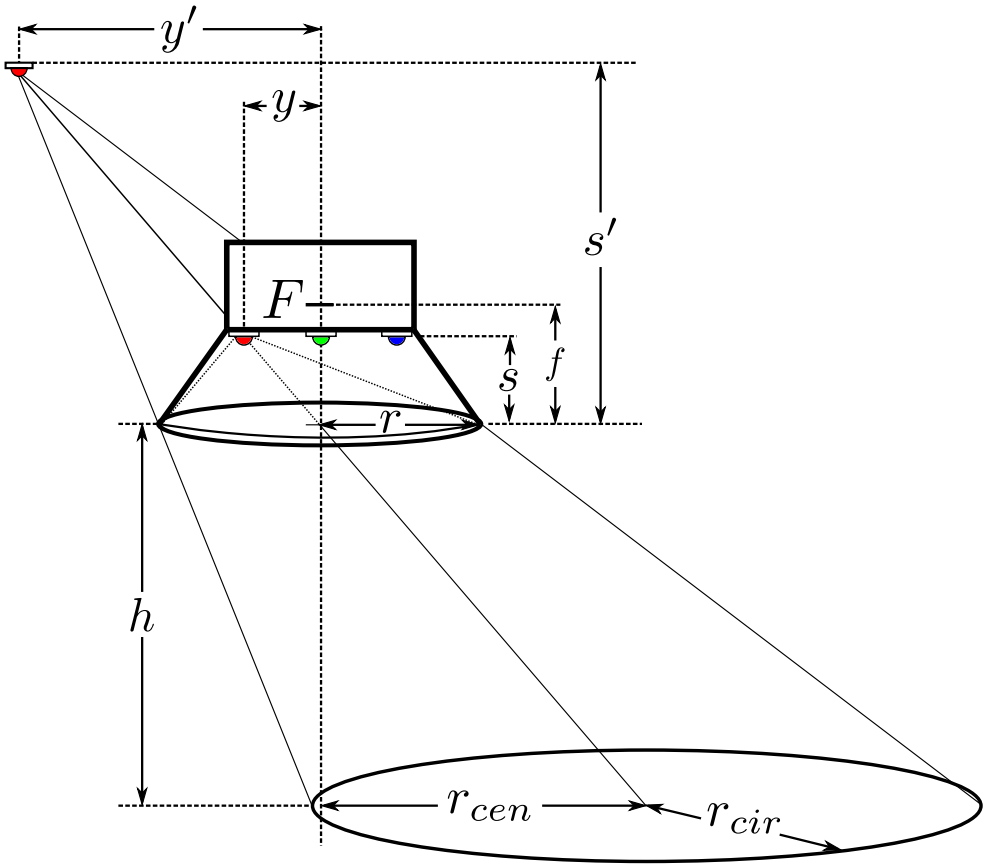

Figure 3.8: Off-axis LED placed inside the focal length of the lens.

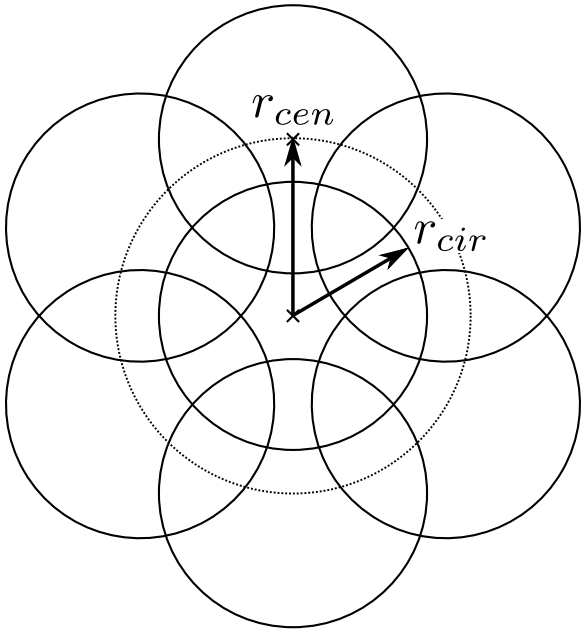

Figure 3.9: Overlapping light pattern for a given height h

Both equations (3.6) and (3.4) depend on the LED-lens distance s. The relationship between the size of the circles and how far apart they are is fixed for a given lens by these equations.

3.3.2 Overlapping Light Pattern

The geometry of the light pattern is determined by the intersection of the cones of light from each LED. For a given height h the cross section of the light pattern will be a series of overlapping circles. Figure 3.9 shows the shape of the overlapping pattern in two dimensions, and how r_{cen} and r_{cir} describe the dimensions of the pattern.

Lenses with large values of f create a poor overlap pattern when the LEDs are near the lens (s is small). They cause r_{cir} to be much larger than r_{cen} which clusters all the overlap directly under the luminary. The LEDs cannot be moved too far from

the lens though or the pattern will cover too little area and the luminary will be too large. For example if a lens was used with radius $r = 2.5$ cm, and focal length $f = 20$ cm the LEDs would need to be placed $s = 10$ cm from the lens to make the pattern look like Figure 3.9. The whole pattern would only cover < 1 m^2 with $r_{cen} = 24$ cm and $r_{cir} = 27$ cm. On the other hand if the lens had a focal length $f = 5$ cm the LEDs could be placed $s = 3$ cm from the lens and form the same shape pattern. In this case $r_{cen} = 70$ cm and $r_{cir} = 90$ cm which covers > 5 m^2. Lenses with shorter focal lengths, and especially small ratios of focal length to radius (f/r) are better for the pattern design.

3.3.3 2D Positioning Using the Light Pattern

With no a priori knowledge of the receiver location it is assumed to be uniformly distributed under the light. In this case it makes sense to use the center of mass of each region as the 2D position estimate. If the receiver identifies it is in a region, it assumes it is at the center.

Figure 3.10 shows a portion of the 2D overlap pattern with the center of mass for each region marked. The odd geometry of many overlapping circles means the center of mass must be calculated using numerical integration. The position error ε is the Euclidean distance between the center of mass of the region, which is the position estimate $\hat{\xi}$, and the actual receiver location ξ.

$$\varepsilon = |\hat{\xi} - \xi|_2 \tag{3.7}$$

The worst case error for the region ε_{max} is the distance from the center of mass to the furthest edge. The average error over the entire region will be lower for smaller

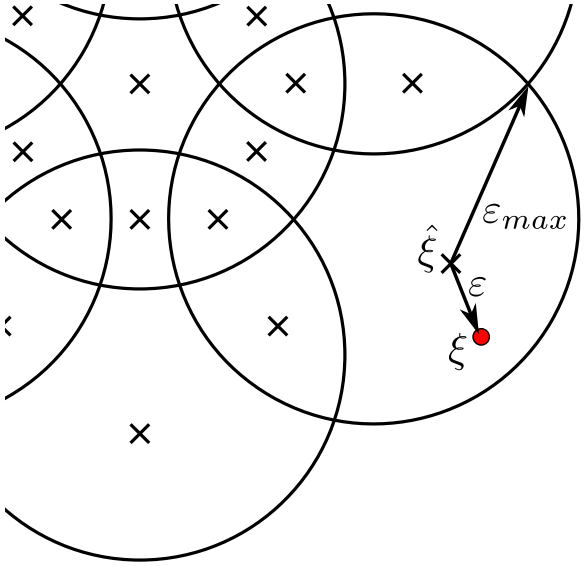

Figure 3.10: 2D positioning using the region center of mass.

closely packed regions, which encourages creating as many small regions as possible in the pattern.

3.4 Simulation Results

This section presents simulation results of the positioning error using the proposed thin lens AD-VLP transmitter. Each simulation calculates the overlapping light pattern from the transmitter parameters f, r, s at a fixed height h and determines the position error using the center of mass of each region. The luminaries are all located in the ceiling facing straight down, and have the following parameters;

$$\begin{aligned} f &= 50\text{mm} \\ r &= 25\text{mm} \\ s &= 33\text{mm} \end{aligned} \tag{3.8}$$

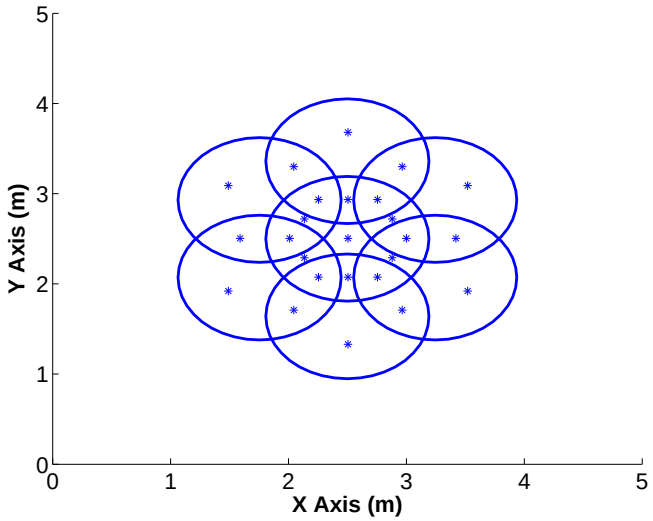

Figure 3.11: Simulated light overlap pattern using an ideal thin lens. The center of mass of each overlapping region is marked with an asterisk.

The values of f and r correspond to one of the lenses used later in the experimental setup, while the LED-lens distance s was chosen to clearly demonstrate the overlap concept in the simulation figures. Each simulation is carried out in a $5 \times 5 \times 3$ m^3 room with a height $h = 2$ m from the transmitter (or $h = 1$ m from the floor). Results for one luminary and multiple luminaries in the room are presented, followed by the effect of changing the receiver height h. Finally a simulation of the effect of changing the LED-lens distance s is shown.

3.4.1 Single Luminary

Using the given luminary parameters and equations (3.5) and (3.6) the 2D overlap pattern for one light can be drawn with $r_{cir} = 0.86$ m and $r_{cen} = 0.69$ m. Figure 3.11 shows the overlapping pattern with the centers of mass for each region.

Figure 3.12 shows a map of the position error using one luminary located in the

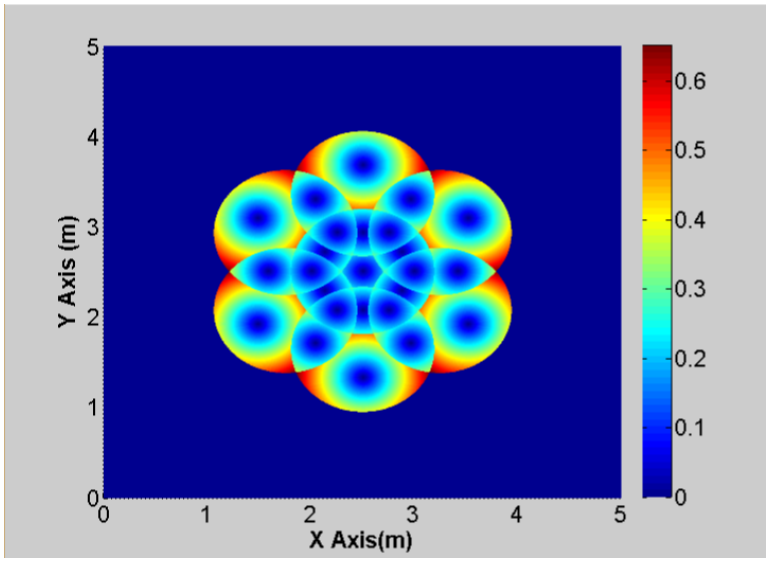

Figure 3.12: 2D positioning error (cm) using a single AD-VLP luminary.

# of Regions	Coverage Area [m^2]	Average Region Size [m^2]	Average Position Error [cm]	Std. Dev. [cm]	Max Error [cm]
25	6.75	0.27	25.2	12.7	65.0

Table 3.1: Simulation results for one AD-VLP luminary.

center of the room. The distinct ring pattern is from using the center of mass as the position estimate. The position error increases radially as the receiver moves from the center of the regions, with the largest position error at the furthest edge of the largest region. The simulation results are summarized in Table 3.1.

One luminary has an average position error of only 25.2 cm over a 6.75 m^2 coverage area, which already places AD-VLP in the sub-meter category of indoor positioning systems. This is an excellent result for an unoptimized choice of luminary parameters. In addition to the low average positioning error this pattern shows a worst-case position error of 65 cm, which shows that even the worst possible position estimate

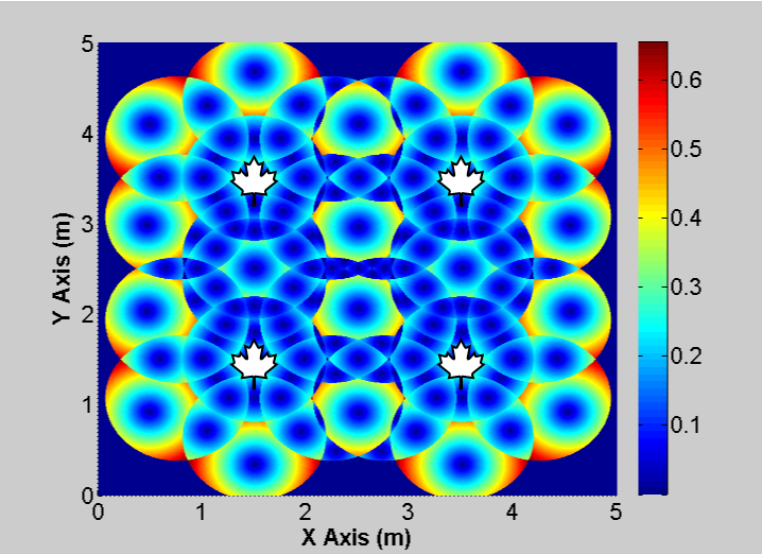

Figure 3.13: 2D positioning error (cm) using four AD-VLP luminaries. The maple leafs represent the locations of the luminaries.

formed in the area covered is less than a meter. This holds up well to the 35 cm average error and 55 cm maximum error in the coarse phase positioning algorithm from [54].

3.4.2 Multiple Luminaries

Since one luminary is typically insufficient to cover the whole room, multiple luminaries can be installed to cover more area. Figure 3.13 shows the positioning error map when four of the same luminaries are placed in the room together. The luminaries are all identical to the single luminary simulation and their locations are indicated in the figure. The same pattern can be seen around the outside of the room, but in the middle there appears to be many more regions than expected. There are 177 unique regions in this pattern, which is much higher than the $4 \times 25 = 100$ contributed from each individual luminary. The additional regions are from the overlap between

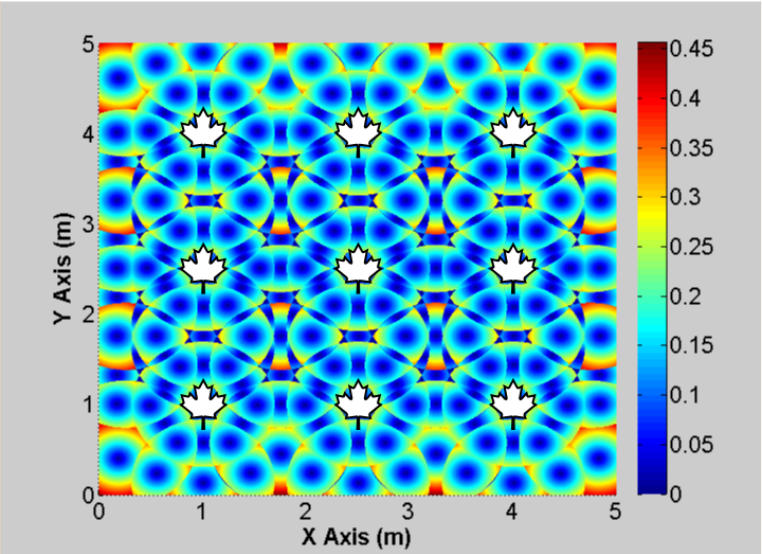

Figure 3.14: 2D positioning error (cm) using nine AD-VLP luminaries. The maple leafs represent the locations of the luminaries.

# of Luminaries	# of Regions	Coverage Area [m^2]	Avg Region Size [m^2]	Avg Pos Error [cm]	Std. Dev. [cm]	Max Error [cm]
2	51	13.4	0.263	24.9	12.7	65.0
4	177	21.4	0.121	21.2	12.1	65.0
9	675	25.0	0.037	15.1	9.3	45.6

Table 3.2: Simulation results for multiple AD-VLP luminaries.

adjacent luminaries which can also be used for positioning.

Figure 3.14 shows the positioning error map when nine luminaries are used. In this case 675 regions are formed and the average position error drops to 15.1 cm over the entire 25 m^2 room, with a standard deviation of 9.4 cm and a worst case error of 45.6 cm. The results of using two, four and nine luminaries are summarized in Table 3.2.

The increase in positioning performance from using multiple luminaries shows that AD-VLP is a very scalable technique. If an application finds the need for more coverage area or better performance, new luminaries can always be installed later

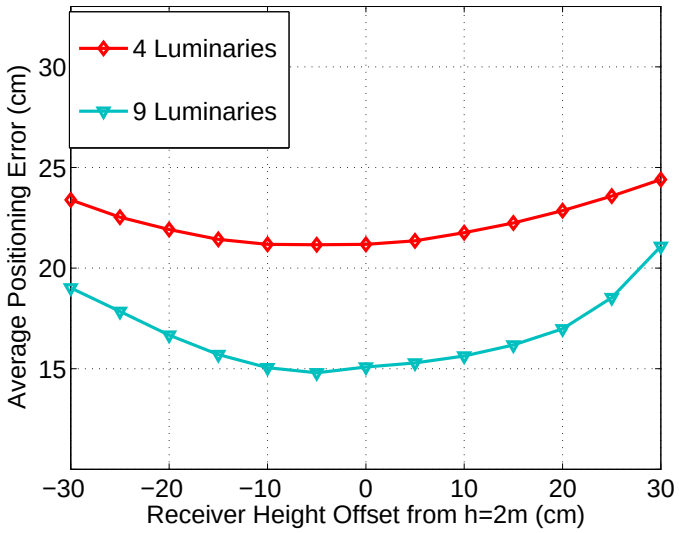

Figure 3.15: Average position error vs height offset from the expected $h = 2$ m.

replacing non-positioning lights.

3.4.3 Effect of Changing Receiver Height

The previous simulations assumed the receiver was held at a known fixed height from the transmitter. If this height is not correct the receiver will see a different map of the overlapping light than the one it is using for positioning and produce greater positioning error. A simulation was run to test the penalty for using an incorrect receiver height.

Figure 3.15 shows the effect of receiver height offset on the average positioning error for the four and nine luminary cases. The assumed receiver height was at $h = 2$ m, while the receiver was simulated between $h = 1.7$ m and $h = 2.3$ m. This ± 30 cm offset could be caused by people of different heights using the system, or periodic height variations as a user walks along. It is shown that the position error

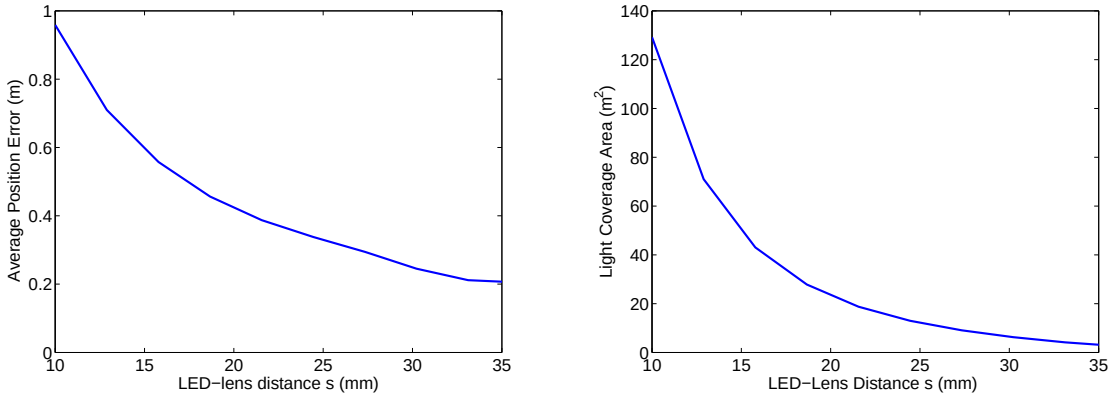

(a) (b)

Figure 3.16: Positioning error (a) and light coverage area (b) for a single luminary with changing s.

increases as the magnitude of the height offset increases. Even with a 30 cm height offset the average position error increases by less than 7 cm, which suggests AD-VLP is robust to variations in the receiver height.

3.4.4 Effect of Changing LED-Lens Distance

For a given lens with focal length f and radius r the only design parameter is the LED-lens separation distance s. From equations (3.5) and (3.6) it can be seen that both r_{cir} and r_{cen} depend on s, so the size of the projected circles will change along with their distance from the center axis.

It is difficult to compare the effect of changing s on the positioning performance because it also has a major impact on the coverage area of the light. Figure 3.16 shows that as the LED-lens distance is increased the average position error decreases, but so does the area covered. Since the average error depends on the distance from the center of mass from each region, as the regions get smaller the average error will

also drop. It is difficult to distinguish what benefit in positioning error is due to changes in the overlap pattern vs how much is just from the reduction in coverage area. There is limited design freedom only changing one lens parameter. If a different trade-off between coverage area and position performance is desired a new lens with different values of f and r is needed.

3.5 Conclusions

This chapter introduced a new method of indoor positioning using visible light: angular diversity visible light positioning. It was shown that by using optical elements at the transmitter to create a structured overlapping light pattern a single element receiver can be used for positioning without relying on signal strength measurements and the VLC channel model. This method largely separates the position information from the intensity measurements and makes the system more robust to changes in receiver height or orientation.

A design example using an ideal thin lens as the optical element at the transmitter has demonstrated a basic framework to describe the overlapping light pattern. A method for using the center of mass of each overlapping region to locate the receiver in two dimensions was presented and explored via simulation. The simulation results showed the average and worst-case theoretical performance of the system to be highly competitive with other VLP systems. It was also shown that adjacent AD-VLP luminaries produced many more overlapping regions that can be used to achieve more accurate position estimates.

In the next chapter the illumination characteristics of the proposed system will be investigated and their connection to the positioning performance is examined. It will

be shown that a minimum signal threshold must be set when identifying regions, and a method of setting this threshold is presented. The theory presented in the design example will also be further developed to include more realistic lens models to help build towards realizing AD-VLP with an experimental setup.

Chapter 4

Illumination and Positioning in AD-VLP

This chapter presents the illumination characteristics of the proposed AD-VLP system, and shows how they are related to the overlapping pattern used for positioning. The illuminance distribution of the luminary is presented with an ideal and more realistic lens before discussing the importance of positioning-illumination co-design. It is shown that detecting regions must be done through setting a signal strength threshold. Methods of setting the threshold and its effect on the positioning regions are also discussed.

4.1 Illumination in AD-VLP

Chapter 3 presented the positioning characteristics of AD-VLP but the luminary must also provide adequate illumination. The key to VLP being adopted lies in its ability to provide positioning functionality while also being a more energy efficient

LED lighting solution.

4.1.1 Lighting Requirements

There are several constraints placed on VLC systems to ensure they provide adequate illumination and that their modulation will not produce any adverse health effects. A description of these constraints and the health considerations, with more references, is presented in [66, Sec. 2.2]. For the purposes of this work the two main factors to be considered are the total amount of illumination provided by the luminary, and the uniformity of that illumination.

Illuminance

As presented in section 2.2.1 the illuminance is the amount of luminous flux that falls on a surface. It is a measurement of how well a room is lit, and there are standards in place that specify how much illuminance is required for different types of working conditions. A house or theater may have as low as 150 lux, while an operating room requires at least 1000 lux [40].

For the AD-VLP transmitter to provide the required level of illuminance the luminary must have powerful enough LEDs to provide the required luminous flux over the coverage area of the light. Practically this means that a single AD-VLP luminary cannot be spread to cover too large an area while still being the primary source of illumination. When combined with other regular luminaries the AD-VLP luminary does not need to provide all of the required illumination by itself.

Illumination Uniformity

Another important characteristic is the uniformity of illumination provided by the luminary. Ideally a luminary would have completely uniform illuminance over the entire area it covers, or no one spot under the light would appear any brighter or darker than another. The challenge with AD-VLP is that the light from each LED is being separately directed with a certain amount of overlap. If the light from each LED was uniform by itself the overlap would be much brighter where multiple lights intersect, and dimmer in regions with few lights.

The illuminance distribution changes with height. This is unavoidable since the same surface area at different heights subtends different solid angles to the transmitter. Therefore uniform illuminance can only be specified at a fixed height, which is usually taken to be the height of the work surface such as a desk. Uniform illuminance can then be measured by looking at the variance of illuminance over the desired coverage area, the peak-to-average ratio, or the peak-to-peak illuminance [67].

4.1.2 Illumination with LEDs and Optical Elements

The illumination distribution of a luminary depends on its luminous intensity. As presented in Section 2.2.1 the luminous intensity $I_v(\theta, \phi)$ describes the directionality of a point source. By knowing how the light from the source is directed its projection on a surface can be determined. For an LED luminary using optics such as lenses the luminous intensity is re-shaped by the optical elements. The analytical description of the light propagation through the optical elements is extremely involved using Fourier optics, and a large part of the advantage of AD-VLP was to avoid precisely modeling the optical intensity in the system. Instead it can be approximated using ray-optics

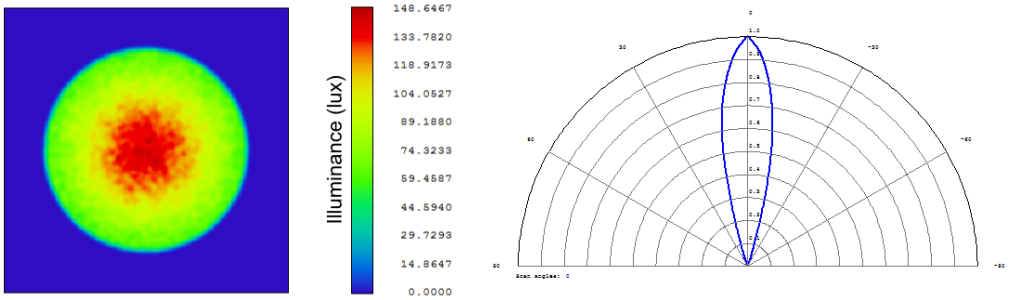

(a) (b)

Figure 4.1: Illuminance (a) and normalized luminous intensity (b) for a single on-axis point source through a thin lens. The area of the illuminance plot is 1.5m × 1.5m.

and a process known as ray tracing.

Ray Tracing

Ray tracing is a powerful tool for numerically simulating the response of optical systems. Rather than manipulating the complex wave propagation equations used in wave optics, ray tracing represents the wavefront as many individual rays that can each be traced through the system using geometric optics. When many rays are generated according a given source distribution, their sum is a close approximation to how the light from that source will actually propagate. The result of all of the rays traveling through an optical system is a sampled version of the wavefront. It is less precise than working with the full wave propagation equations, but it is much easier to use and sufficient for simulating the new luminous intensity distribution produced by the optics in the AD-VLP transmitter.

Figure 4.1 shows the results of a ray tracing simulation run on the on-axis LED with the same ideal thin lens as presented in Section 3.4. The simulation was run using an academic license of the commercially available optical simulation program

Zemax [68], which includes a powerful ray tracer. The simulation was run with 1 million rays, and both the luminous intensity of the entire transmitter and illuminance at $h = 2$ m were recorded. Figure 4.1a shows the illuminance at $h = 2$ m which provides an average of 100 lux over its nearly 1 m^2 surface area.

Figure 4.1b shows the luminous intensity. It can be seen that the lens truncated the Lambertian emission pattern of the LED and focused it to be more directional. The half power angle $\Phi_{1/2}$ of the light leaving the lens is approximately 10°, compared to 60° from the Lambertian source.

The results of ray tracing simulations on the transmitter can be used to represent the source in later simulations. The simulated luminous intensity can be saved and re-sampled, or an equation can be fit to the numerical results and used in further analysis. A method for fitting Gaussian and raised cosine functions to rotationally symmetric radiant intensity profiles was proposed in [69]. However, once the light source is not aligned along the optical axis it becomes much more difficult to find a functional fit.

4.1.3 Illumination Pattern with an Ideal Thin Lens

Figure 4.2 shows the simulated illumination pattern for the thin lens AD-VLP transmitter described in Section 3.4. The pattern can be clearly seen as the summation of light from each LED, and it strongly resembles the overlapping regions used for positioning. The bright spots are where multiple lights overlap while the lower intensity spots are the regions that only have one light.

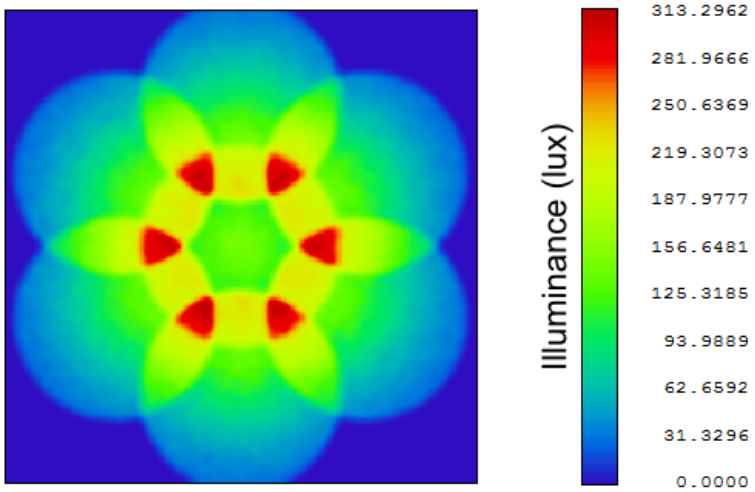

Figure 4.2: Illuminance distribution for the 7 LED thin lens AD-VLP transmitter presented in Chapter 3. Area of the plot is 2.5m $\times$ 2.5m.

4.1.4 Illumination Pattern with a Non-Ideal Lens and Transmitter

There are many distortions and aberrations that arise from using real spherical lenses, and from simulating the actual LED die instead of using point sources. An excellent description of the different kinds of lens aberrations is presented in [60, Ch.9], while a description of the difficulties modeling illumination in the near field is presented in [70]. Ray tracing can also be used to simulate non-ideal optics by modeling the surface geometry of the lens. This requires more computation, but it captures the effect of the non-idealities by modeling the ray interactions more completely.

Figure 4.3 shows the simulated illumination pattern for a glass biconvex glass lens EO32978 from Edmund Optics [71]. This lens is made of N-SF11 optical glass with a focal length $f = 50$ mm , radius $r = 25$ mm, and a total thickness of 16mm. The same ideal point source transmitters were used, at the same spacing as before. The

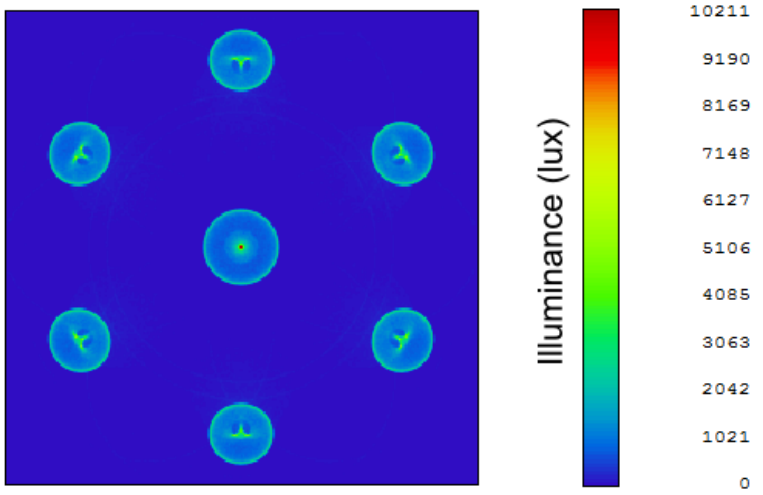

Figure 4.3: Illuminance distribution for the 7 LED AD-VLP transmitter with a thick glass lens ($s = 33$ mm, $f = 50$ mm). Area of the plot is 1.5m $\times$ 1.5m.

pattern produced is significantly different from the thin lens, even though the focal length and radius are identical. This caused by the effect of lens non-idealities, or deviations from the ideal thin lens behavior. Specifically it is the result of spherical aberration, coma, distortion and the lens thickness.

Spherical aberration causes the focal length of the lens to decrease with radius, which bends rays near the edge of the lens more than rays near the middle. This both decreases r_{cir} for each LED and causes the rings to form around the outside of each light region. It also makes the outer light regions slightly smaller than the middle one. Coma is the effect of different magnification seen across the aperture of the lens by light that is away from the optical axis. It causes the pinched bright spot in the middle of the outer light regions. The middle light region does not have this shape since the center LED is along the optical axis. The final major distortion is caused by both the thickness and shape of the lens which prevents it from performing perfect rectilinear projection. The most noticeable effect is the location of the center of each

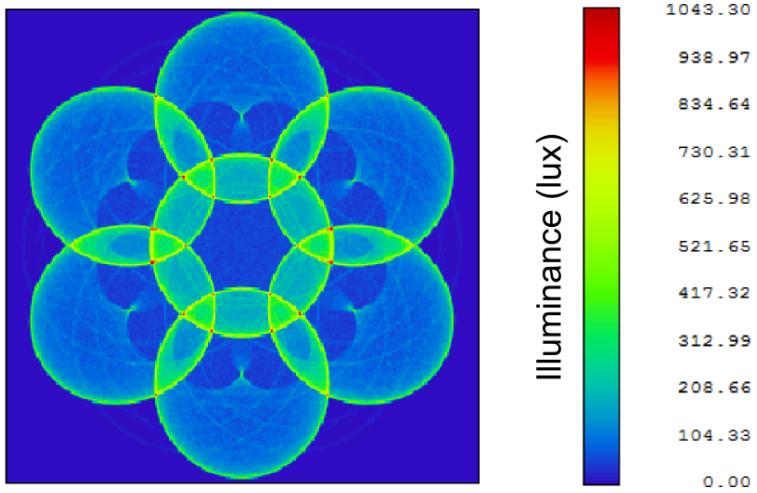

Figure 4.4: Illuminance distribution for the 7 LED AD-VLP transmitter with the thick glass lens, ($s = 18$ mm, $f = 50$ mm). Area of the plot is 2.5m $\times$ 2.5m.

light region r_{cen} is much smaller than predicted by the ideal thin lens.

By moving the lens closer to the LEDs the overlap pattern can be recovered. Figure 4.4 shows the same lens with the LED-lens distance changed to $s = 18$ mm. In this case the light from each LED is overlapping again, and the pattern is similar in size and shape to the one formed by the thin lens. The outer lights become more oval in shape than circular, which is caused by the lens aberrations. The illuminance distribution is unsuitable for general illumination because the edges of the regions are nearly ten times brighter than the center, and will be clearly visible when projected on objects or the floor.

Finally, since real LEDs are not point sources the same simulation is run using a LED source file from Phillips [2,72]. These specific LEDs are used in the experimental transmitter presented in Chapter 5. Figure 4.5 shows the illumination distribution with the non-ideal LED model. It is the same as the point source case, but more spread out due to the light being emitted over the LED die surface area.

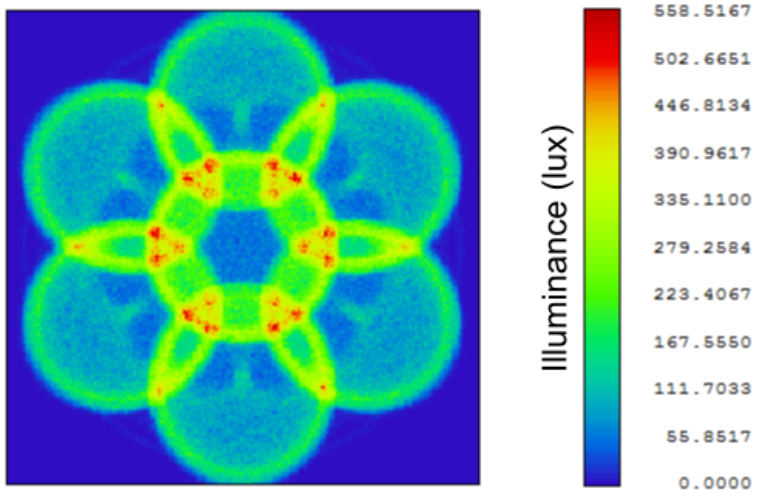

Figure 4.5: Illuminance distribution for the 7 LED AD-VLP transmitter with the thick glass lens, (s = 18 mm, f = 50 mm), and physical LED model. Area of the plot is 2.5m $\times$ 2.5m.

4.1.5 Optics Design for AD-VLP Luminaries

The illumination distributions presented in the previous section are not ideal for lighting a room or office due to their non-uniformity. However, there is a body of work on designing optical elements for LED luminaries that can be adapted to develop structured overlapping light patterns that are also suited for illumination.

Using optics in LED luminaries is not a new idea. Almost every LED luminary available has a plastic lens or diffuser on it that converts the light from the individual LEDs into a more suitable illumination distribution. Different numerical and ray tracing methods can be used to design the target illumination distribution, some of which are reviewed in the following reference [73]. It has been shown that these methods can be used to create a uniform illumination distribution over an arbitrary defined shape in the target plane, such as a rectangle in [74].

Another approach is to use microlens arrays, which are an array of many tiny lenses

that can be individually tailored to produce the desired illumination distribution [75]. These methods have even more design freedom than a single surface custom lens, and can create uniform illumination over odd shapes even when projected off the optical axis [76].

These techniques could be used to create uniformly illuminated overlapping light patterns, allowing the positioning regions to be customized while still producing adequate illumination. Since many LED luminaries already use custom optical elements there would be little added cost to the AD-VLP transmitter.

4.2 Identifying Positioning Regions

Detecting the signals that form a positioning region has a lot in common with the classical problem of detecting a known signal in noise [77, Ch.8]. However, it is not enough to only detect the presence of a signal. Unlike in Chapter 3 where the light signals were perfectly confined to the direction they were transmitted in, the signals will bounce and scatter outside of the desired positioning regions. Signals detected by the receiver outside of their intended region could lead to incorrect region estimates and positioning errors. The system needs to detect the signal *and* decide if that signal is strong enough to be part of the positioning region.

4.2.1 Setting a Region Detection Threshold

The receiver can make a better decision about whether it is actually inside a region or not by comparing the received signal strength from each LED against a set threshold η. This threshold helps separate the intended LOS signal from the unwanted NLOS

portions by using the fact that the LOS portion should be higher in intensity. It should be set high enough for the receiver to ignore reflections and background noise, but not so high that it becomes less robust to changes in orientation or fails to recognize when it is actually inside a region.

Using a fixed threshold is rigid and will produce poor results in situations different from those considered when the threshold was set. It is likely that elements of radar detection theory [78] could be applied to developing an adaptive threshold that changes with the total amount of signal received, and better detects regions under different conditions. However, for the purpose of illustrating the effect of the threshold and why it is needed only a fixed threshold will be considered.

Since the signaling method has not been discussed, the illumination distribution is an appropriate starting point for setting the region threshold. The simulated illumination distributions can be examined for certain best and worst case scenarios to help choose an appropriate value for η.

Line-of-Sight Illumination Distribution

The best case scenario for the system would be a luminary in the open with a direct LOS path to the receiver everywhere covered by the light. The threshold η should be set so that the receiver can correctly identify each region everywhere under the light, which in this case means being lower than the lowest illuminance level in the pattern.

Figure 4.6 shows a cross section of a hypothetical illumination distribution from a properly designed optical element, such as the type presented in section 4.1.5 [75]. It bears a resemblance to the magnitude response in filter design. The goal is to keep as much of the light contained in the defined region, with as little ripple as possible

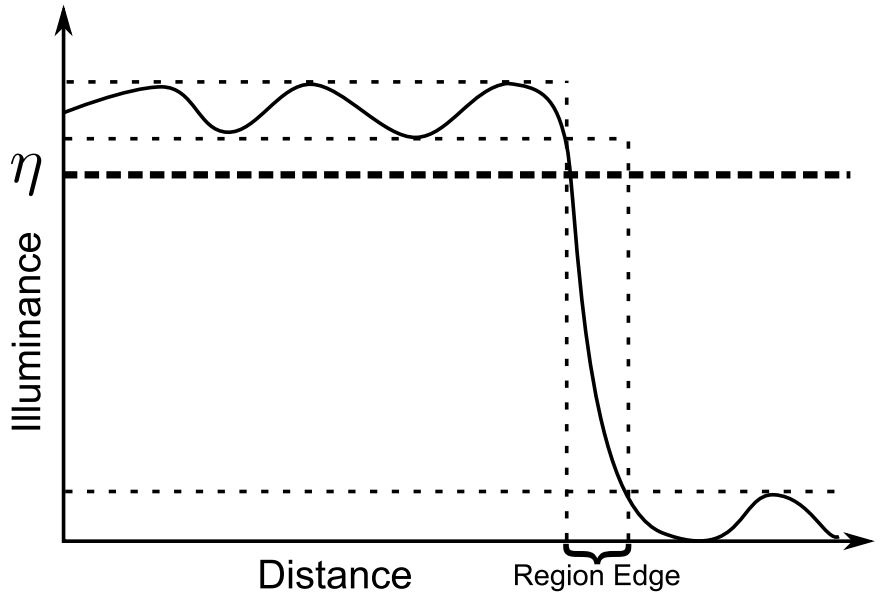

Figure 4.6: Cross section of a hypothetical illuminance distribution inside one circular region.

for uniform illumination. In this case the threshold η is set below the lowest point of ripple in the main part of the light region. The edge of the region should have a sharp cutoff, since the choice of η will determine the edge of the region as seen by the receiver.

Figure 4.7a shows the horizontal and vertical cross sections of the illuminance distribution from the middle and top regions from Figure 4.4. Since the on axis LED maintains the circular symmetry, the middle region is symmetric. Unfortunately this lens creates an illumination distribution that decreases towards the center. In order to properly detect this region the threshold can not be set any higher than $\eta = 45$ lux or the receiver will fail to identify the center of the region.

Figure 4.7b shows the same cross sections for the off axis regions. The same decrease in intensity towards the center of the lens can be seen and the threshold of $\eta = 45$ lux is below this minimum as well. These examples show how important

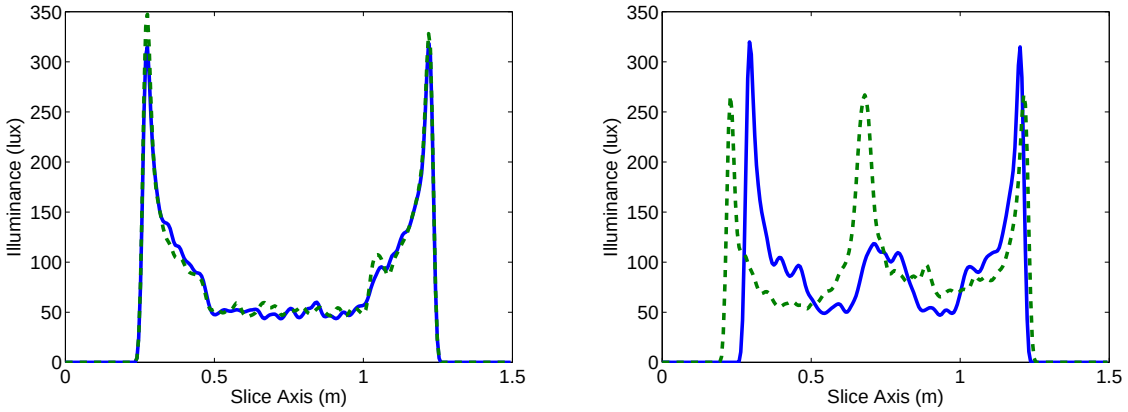

(a) (b)

Figure 4.7: Horisontal (solid) and vertical (dashed) cross sections of the illumination distribution from the on (a) and off (b) axis LEDs.

proper illumination design is to the positioning aspects of the system, in addition to the lighting requirements.

Guarding Against Reflections and Scattering

No matter how well the light is directed by the optics, it will scatter and reflect around the room once it hits a surface. Light that reflects can be detected by the receiver outside the original region in the structured overlapping light pattern, leading it to estimate an incorrect position. The signal threshold can be set high enough to help distinguish signals that arrive at the receiver from reflections. As was presented in Section 2.3.3 the NLOS portion of the channel is much harder to model because it requires knowledge of the geometry around the transmitter. However, there are some observations that can be used to investigate the possible effects of reflections on AD-VLP.

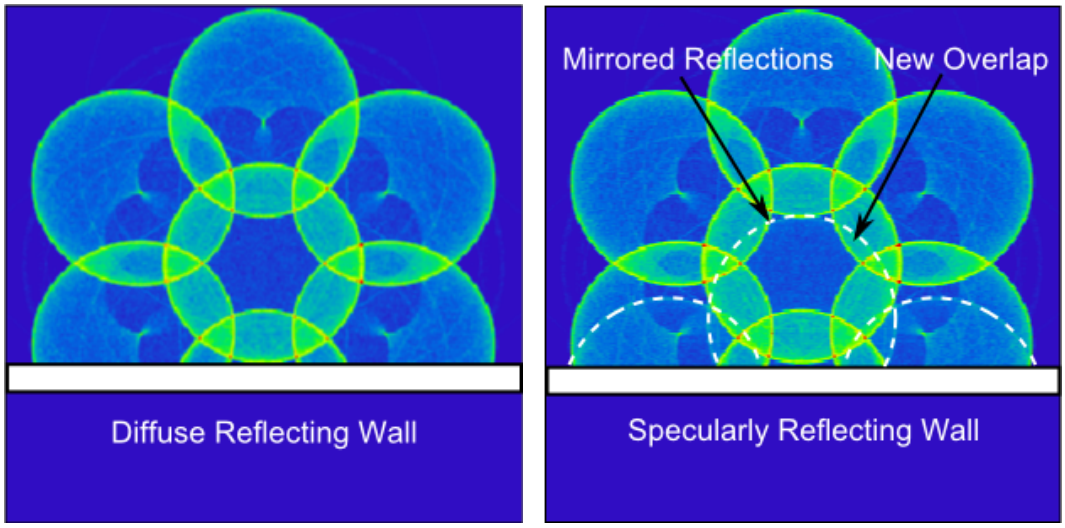

(a) (b)

Figure 4.8: Illumination distribution produced near a diffuse reflecting (a) and specular reflecting (b) wall.

Many materials found indoors can be approximated as diffuse reflectors with reflectivity between 0.1 and 0.8 [79]. These materials scatter light in approximately a Lambertian distribution regardless of its angle of incidence. Since the light is scattered over a wide range of angles the radiant intensity in any one direction is very low. The first bounce contains the most reflected power, and is the most likely to affect the region detection.

A common scenario where reflections could affect the system is when the luminary is near a wall. In this case the light from the outside LEDs can reflect back into the pattern. Figure 4.8a shows the illumination distribution at a distance of 2 m, with the luminary located 0.5 m from an 80% diffuse reflecting wall. There is no noticeable effect on the illumination pattern.

There are also specular reflectors such as glass, high gloss painted walls, and mirrors. Many of the materials typically assumed to be diffuse reflectors can also have

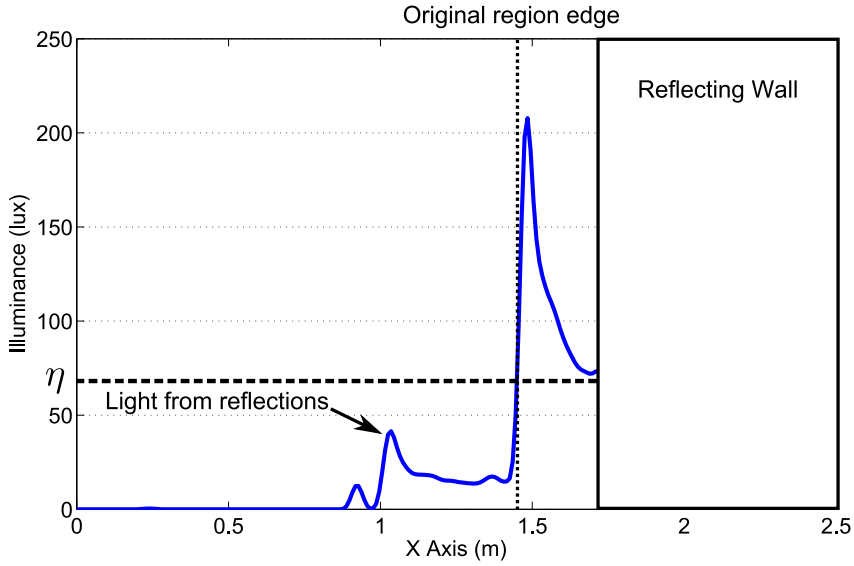

Figure 4.9: Cross section of one LED region reflected off a specular reflecting wall.

a specular component [80]. These materials produce some directional reflections with significant power. Figure 4.8b show the same example with a 20% specular reflecting wall. There is clearly a mirroring of the light pattern which will be detectable by the receiver.

This mirroring distorts the shape of regions near the wall, and also produces new overlap that is not expected. The highlighted region in Figure 4.8a shows where the light from the outer region reflected back over onto one of the opposite light regions. This can actually be used to the systems advantage. If the receiver sees a combination of light signals that do not exist according to the region map, it can guess that this region might be caused by reflections. The receiver could look at the stronger signal to decide which region it should really be in, or in more advanced systems it could even attempt to update the region map with the new region.

Figure 4.9 shows a cross section of the illuminance reflected off of a wall. It can

be seen that the original part of the pattern formed by the line-of-sight signal is still present, but the reflections off the wall have mirrored the rest of the pattern over to the other side. In this case the threshold η must be set higher than the highest reflection at 42 lux , but lower than the lowest illuminance in the original region which is 67 lux. As will be seen in Section 4.2.3 there is a trade-off between setting the threshold higher to avoid reflections, and lower to increase robustness to changes in orientation.

Based on these observations AD-VLP can be made robust to reflections and scattering with proper selection of the threshold η. Purely diffuse reflections have little effect on the system since they are so much lower in intensity than the LOS components, and specular reflections can be guarded against by keeping the threshold sufficiently high. The level to set η at will depend on the luminary, the target height, and the expected worst case illuminance situation.

4.2.2 Effect of Signal Threshold on Region Shapes

Although the AD-VLP luminary optics were designed to produce a particular shape for the light pattern, the illumination distribution combined with the receiver threshold will determine what the edges of the regions look like to the receiver.

The selection of the signal threshold η can have a large impact on the appearance of positioning regions to the receiver. Figure 4.10a shows the region map for the illumination distribution presented in Figure 4.4 with a threshold of $\eta = 25$ lux. The regions are quite distinct and match the expected overlap pattern. Figure 4.10b shows the effect of setting the threshold too high at $\eta = 100$ lux. The pattern has large gaps in the regions because the receiver threshold is higher than the illuminance level

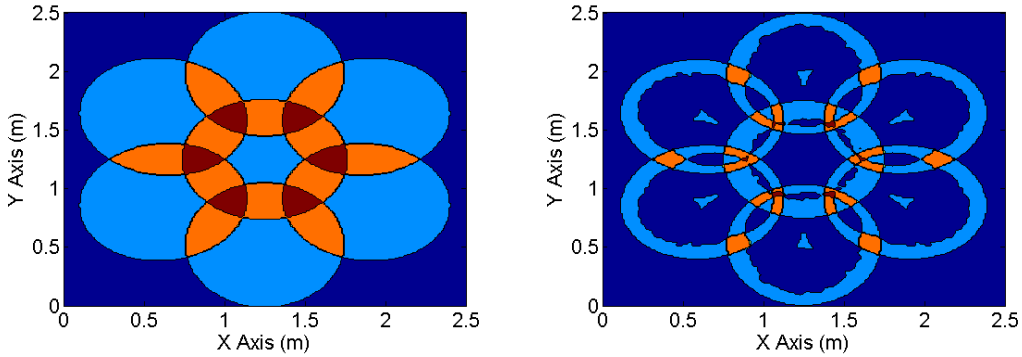

Figure 4.10: Region map produced by setting the illuminance threshold to $\eta = 25$ lux (a) and $\eta = 100$ lux (b).

inside the region, and causes the receiver to fail to recognize when it is inside.

Changing Height and Region Thresholds

Just as the shape of the regions changes with height, so too does their illuminance distributions. The further the distance from the luminary the lower the illuminance will be, as predicted by the inverse square law. In terms of setting a threshold this means the receiver will detect less power inside a region the further its vertical separation from the luminary. A fixed threshold could be set based on the lowest expected operating height of the receiver, which would guarantee that it would be met at any higher elevation since the receiver can only see more power the higher it is held. However this may make it more sensitive to reflections closer to the luminary.

4.2.3 Orientation and Region Thresholds

The received signal intensity depends on the receiver orientation, and so will the detection of regions using a threshold. When the receiver is facing away from the

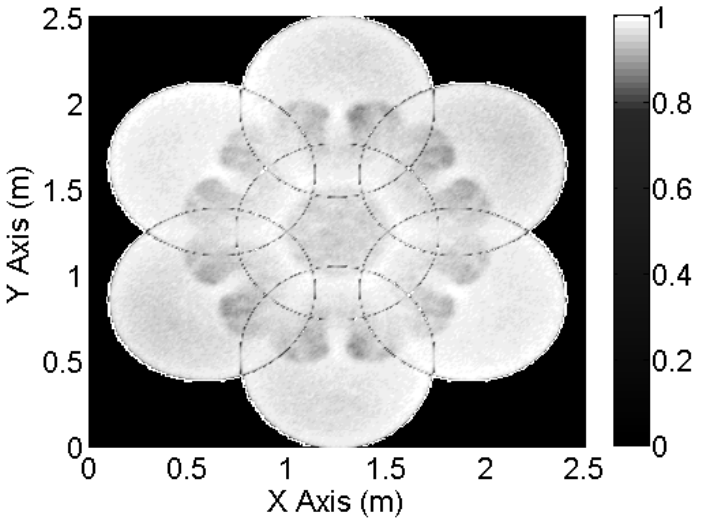

Figure 4.11: Probability of receiver making a correct decision for $\eta = 25$ with random orientation.

luminary it will need more luminous intensity to reach the signal threshold, while when it is facing towards the light it will need less. This leads to the possibility of setting a worst case decision threshold that is the lowest amount of power seen by the receiver when it is inside a region and facing a certain amount away.

If the receiver orientation ψ is modeled as a Gaussian random variable with mean $\mu = 0$ (receiver held straight up) and standard deviation $\sigma = \pi/6$ (30 degrees from center) then a simulation can be run over the illumination distribution to see how well the receiver can still correctly identify regions. The choice of tilt distribution covers the range of orientations considered in [45], while keeping about half of the tilt angles between $\pm 30°$.

Figure 4.11 shows the probability of the receiver making a correct decision over the area covered by the luminary with the threshold kept at $\eta = 25$ lux. The average probability of correct detection was 94.5% over 10,000 trials. This suggests AD-VLP

can be made very robust to changes in receiver orientation. The figure shows that the probability of correct detection is proportional to the illuminance of the region. Areas where the light was dim and the edges of regions are hardest to detect correctly. The lower the threshold is set the more the receiver can be tilted away before the measurement becomes too low to meet it, and it can better detect the dim parts of the light regions at different orientations. However, setting the threshold lower also increases the chance of detecting reflections or stray light.

4.3 Conclusions

This chapter has presented the basic illumination characteristics of the proposed AD-VLP system, and shown how they relate to the positioning function of the system. It was shown that simple lenses produce poor illumination distributions, but advanced optical design techniques could be used to realize a structured overlapping light pattern with uniform illumination characteristics.

Furthermore it was shown that by using a signal threshold when identifying regions AD-VLP is very robust to reflections and changes in orientation. Methods of setting a fixed decision threshold were presented based on the minimum expected region illuminance, and worst case reflection scenarios.

The examples of setting the region threshold showed that optical design in AD-VLP is important for both the illumination and positioning performance of the system. A uniform illumination distribution with sharp cutoff at the region edge is easier to detect than an oddly illuminated region with fuzzy edges. Optical design for AD-VLP will likely be an important future direction.

The next chapter uses the presented theory and simulation data to implement

an experimental AD-VLP setup, and demonstrates the validity of the concept with several tests.

Chapter 5

Experimental Validation of AD-VLP

This chapter presents an experimental implementation of the proposed AD-VLP indoor localization system based on the theory presented in Chapters 3 and 4. The design of a simple AD-VLP transmitter is presented, and the structured overlapping light patterns it produces are examined. A single element receiver based on a mobile device ambient light sensor is then used to perform positioning experiments with the constructed luminary. The experimental setup is briefly discussed before presenting the results of the experiments, and comparing them to simulation results.

5.1 AD-VLP Transmitter

This section presents the design of an AD-VLP luminary / transmitter using a single lens as the optical element.

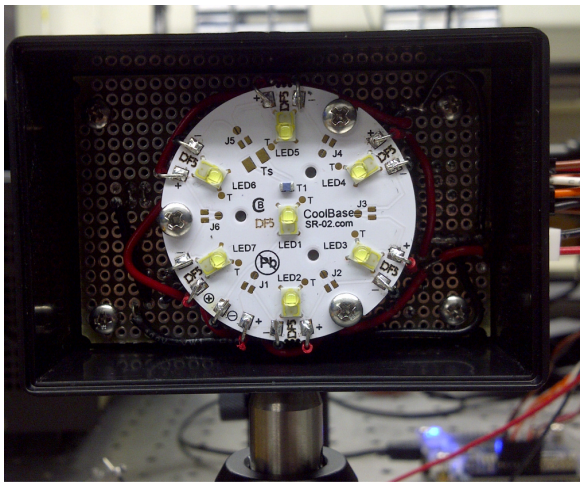

Figure 5.1: AD-VLP transmitter and housing. The seven Phillips LUXEON® Rebel LEDs [2] can be seen arranged in a hexapolar configuration.

5.1.1 Transmitter Design

The experimental transmitter was designed using a LuxeonStar CoolBase 7 LED disc [81], which features seven Phillips LUXEON® Rebel LXML-PWC1-0120 bright white LEDs [2]. Each LED is capable of outputting 220 lm of luminous flux at 700 mA drive current. Each LED is driven independently by a PN2222 NPN BJT [82] through a current limiting resistor. This enables a logic level signal to drive each LED. The modulation current was approximately 440 mA while on for each LED, and nearly 0 mA while off. The completed transmitter and housing are shown in Figure 5.1.

5.1.2 Optics

For this work a single lens was chosen as the optical element at the transmitter. The basic theory of AD-VLP presented in Chapter 3 was built up from a single thin lens, and while this was shown to be a non-ideal approach in Chapter 4 it still represents the most straightforward way to demonstrate the basic principles of AD-VLP.

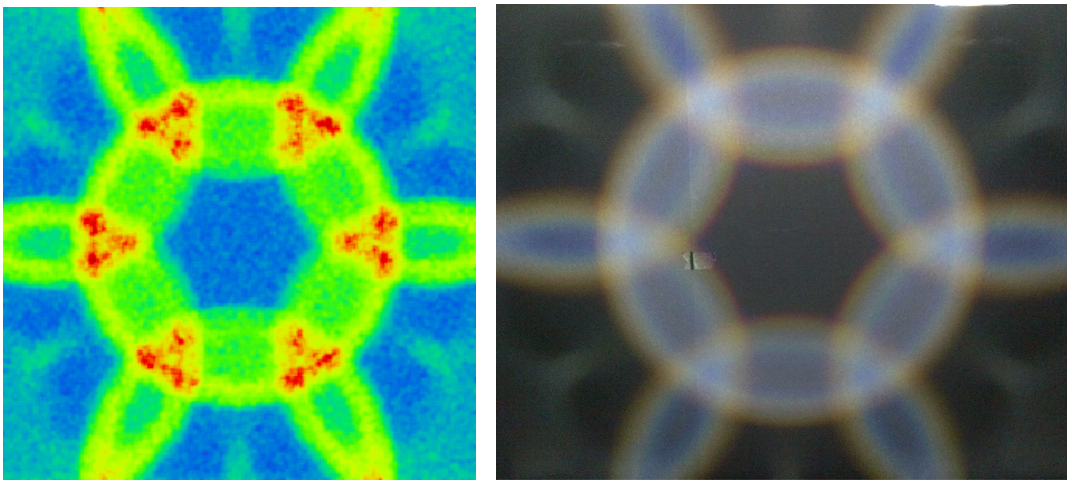

(a) (b)

Figure 5.2: Structured overlapping light pattern produced by the luminary with the glass biconvex lens ($s = 18$ mm, $f = 50$ mm, $r = 25$ mm). Simulation (a) and experiment (b).

The lens used in the experiments was a thick spherical glass biconvex lens from Edmund Optics [71]. The lens has a radius of $r = 25$ mm and an effective focal length $f = 50$ mm. This is a short focal length for the lens radius. It was mentioned in section 3.3.2 a short focal length is desirable for creating the overlapping light pattern. This lens was the same used in the simulations in Chapter 4, where it was modeled from the manufacturer supplied lens data [71].

5.1.3 Structured Overlapping Light Pattern

Figure 5.2a shows the simulated illumination distribution at a distance of $h = 1$ m using the thick glass lens at $s = 18$ mm from the LEDs. Beside it Figure 5.2b is a picture of the illumination distribution produced by the experimental luminary with the same spacing. The observed distribution matches what was expected from simulation and the overlapping regions can clearly be distinguished in the image.

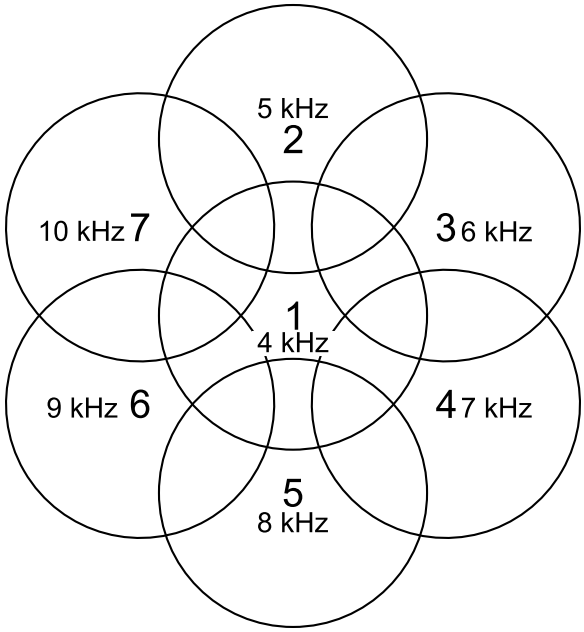

Figure 5.3: Numbering for the light projections from each LED. The drive frequencies used in [3] for identifying each LED are displayed in each region.

5.1.4 Signaling and Region Identification

Frequency division multiplexing (FDM) was used to identify the light from each LED independently. By modulating each LED with a different frequency square wave the signals can be identified by looking at the FFT of the received optical signal. Figure 5.3 shows the numbering for the projections from each LED along with their drive frequencies. The center projection is from LED1, and the outer regions are numbered clockwise. Since each LED is modulated with a square wave the average current is half the on current, or 220 mA per LED. The average luminous intensity at this current was 65 lm per LED.

In a previously published conference work, a THORLABS photodetector [83] was used to detect the modulated light signals [3]. Figure 5.4 shows this receiver placed inside one of the overlapping light regions with the FFT of the received signal shown

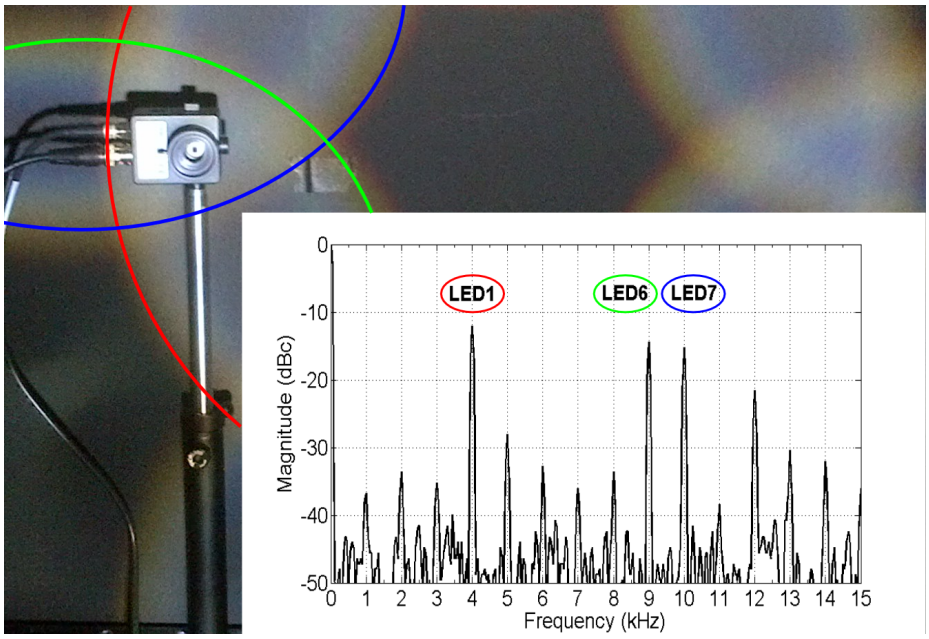

Figure 5.4: Single element photodetector placed inside one of the overlapping light regions, with highlighted FFT shown in the inset.

in the inset. The largest spikes in the FFT are the LOS signals visible to the receiver. They appear at 4, 9 and 10 kHz, which corresponds to the frequencies that were assigned to LEDs 1, 6 and 7. The overlapping illumination area from each of these LEDs is highlighted in the figure. If the receiver sees all three of these lights, it must be in the region they all overlap, and this is in fact the case. This was one of the earliest proof of concept experiments performed using the AD-VLP idea.

The other spikes in the FFT come from different sources. The spike at 12 kHz is the first harmonic of the signal from LED1. As the LEDs are modulated with square waves harmonics will appear at odd multiples of the fundamental frequency $3f_0, 5f_0, 7f_0$...etc. The spikes at 5,6,7 and 8 kHz are reflections from the other modulated LEDs. In this case the luminary was resting on a highly reflective optical bench, and there was also a low shelf above. The signals from the other LEDs reflected off

of these surfaces and still arrived at the transmitter, even though it did not have a direct line-of-sight to these regions. These signals arrived nearly 20 dB below the line of sight signals. This observation motivated setting a signal strength threshold above the expected reflected signals, as presented in section 4.2.1.

5.2 AD-VLP Mobile Receiver

The goal of AD-VLP is to realize a method of indoor positioning that can be implemented in a mobile device such as a smart phone. To that end a mobile receiver was developed that could be used for testing.

Ambient Light Sensor as an Optical Receiver

Previous works on enabling optical communications on a mobile device have used the camera. While this has been demonstrated to be possible, it is very limited in its capability and has a major power consumption penalty both in running the camera and the intensive image processing afterwards. One of the biggest advantages of AD-VLP is the ability to avoid using the camera and only use a single element receiver. There is another optical sensor in most modern mobile devices that could be adapted for communications, the Ambient Light Sensor (ALS).

Mobile devices use the ALS to detect the background illumination level and control the brightness of the screen. These devices are single element photodetectors sensing light intensity. If these parts could be sampled fast enough they could be used for low rate optical communications. Currently mobile devices only sample the light sensor at around 10 Hz, but the device itself can be polled as fast as 2.5 kHz if the sample precision is reduced [84]. This is enough for very low rate VLC, but with some

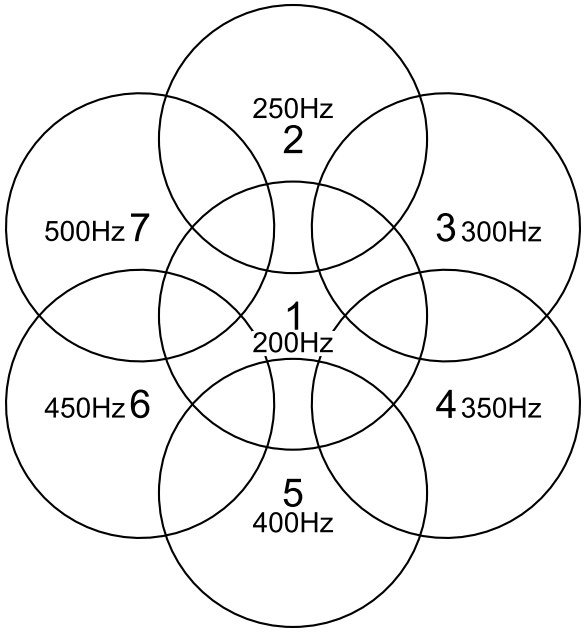

Figure 5.5: Numbering for the light projections from each LED with the new drive frequencies used with the TIVA mobile receiver.

adjustments to the part it could be made much faster. The limitation on the sampling rate is not the photodetector itself, but the integrating ADC it is packaged with. With a faster ADC it could be possible to enable higher rate optical communications on mobile devices without installing a new aperture or expensive sensor.

5.2.1 TI Microcontroller Based Mobile Receiver

Current mobile device API's do not allow access to the low level functions required to increase the ALS sampling rate inside the device. However a sensor development board was found with the same type of light sensor as is currently available inside mobile devices, but open to access.

The Texas Instruments TIVA Launchpad development was chosen for the mobile receiver [85]. One of the add-on hardware modules for the TIVA platform is the

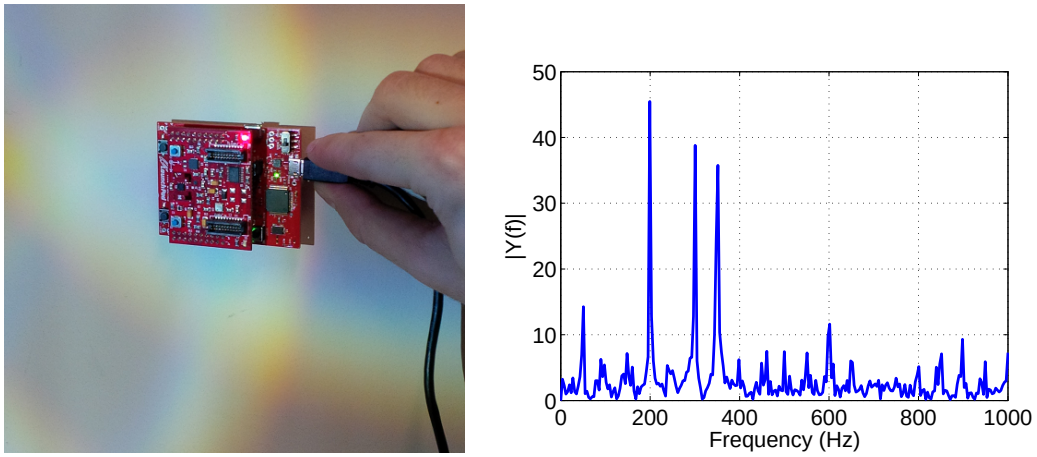

(a) (b)

Figure 5.6: TIVA-SENSHUB based receiver placed inside the overlapping light pattern (a) and the FFT of the detected signals (b).

TIVA-SENSHUB sensor development board [86]. This sensor board includes an ISL-29023 ALS [84], and a MPU-9150 [87] IMU. It was chosen for use in testing as both of these sensors are used in current mobile devices.

A program was developed to increase the sampling rate of the ALS to a potential 2.5 kHz by reducing the resolution from 16 to 8 bits. The TIVA samples the ALS at 2 kHz and the IMU at 50 Hz and passes the values to a PC for processing via a serial connection. Because the sampling rate of the ALS low, the frequencies of each LED were assigned from 200 to 500 Hz spaced 50 Hz apart as shown in Figure 5.5.

Figure 5.6 shows the mobile receiver placed inside one of the three light overlapping regions alongside the single sided FFT magnitude of the received ALS sensor values. The highest spikes at 200 Hz, 300 Hz and 350 Hz correspond to LEDs 1, 3 and 4, which are the lights that form that region.

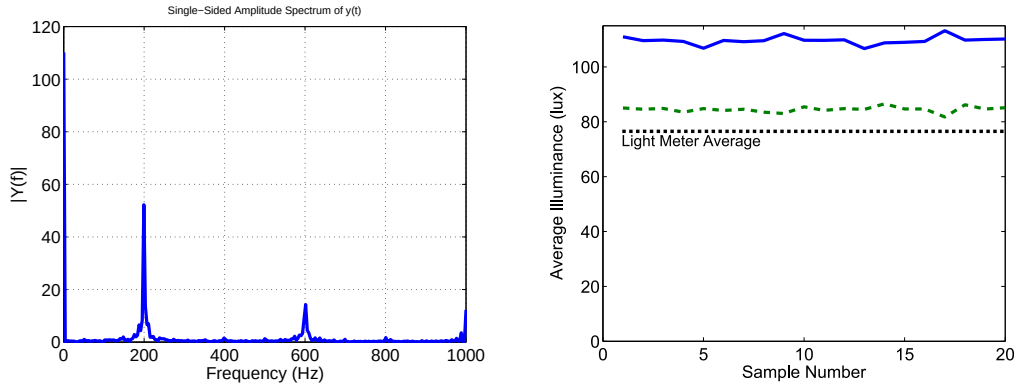

(a) (b)

Figure 5.7: FFT of the received light signal from one LED (a) and the average illuminance over a few seconds of measurement (b).

5.2.2 Measuring the Received Illuminance

Since the ALS has a physical filter designed to match the human eye luminosity function it produces measurements in lux only. This is an odd unit to work with for signal strength, but the same principles apply using average illuminance as the measurement instead. This also allows the measurements to be verified with a commercial light meter (Extech Instruments HD400 [88]).

In the Fourier transform of a square wave the magnitude of the fundamental frequency component is the amplitude of the signal. This can be used to quickly estimate the amount of light seen from each LED by looking at the appropriate frequency bin in the FFT. The average illuminance from each LED should be the magnitude at the fundamental frequency divided by two since the driver duty cycle is 50%. The total average illuminance seen by the receiver will appear as the DC component of the FFT.

Figure 5.7a shows the FFT of the received signal for one measurement of the

center LED region, while Figure 5.7b shows the extracted average illuminance for ten measurements. Using a 512 sample FFT with a sampling rate f_s of 2 kHz the bin resolution is $f_s/n = 2\ kHz/512 = 3.906$ Hz/bin. Since the frequencies selected for each LED do not exactly divide by this value, and because the FFT window causes spreading of power to adjacent bins, the measurements were taken over a 5 bin window. This also sets the maximum update frequency of the system at 3.906 Hz because of how long it takes to collect the 512 samples at 2 kHz. A lower point FFT would increase the sample rate but reduce the frequency resolution.

In Figure 5.7b the blue (solid) curve is the total average illuminance seen by the receiver estimated from the DC component of the FFT, while the green (dashed) curve is the average illuminance from LED1 estimated from the fundamental frequency component. The annotated dotted line is the average illuminance as measured by the commercial light meter. It is apparent that there are two problems, first the average illuminance measured by the ALS does not match the commercial light meter, and second the average illuminance estimated from the LED is does not equal the total average. The test was run in a dark room with the LED as the only light source, so the average from the LED should match the total average.

The first problem is explained by the spectral response of the sensor being more biased towards the blue wavelengths than the typical human vision luminosity function (see spectral response in [84]). Since the LEDs are the bright blue-phosphor type they have a very strong blue component in their PSD [2]. This extra blue power is incorrectly weighted and artificially inflates the illuminance reported by the receiver by nearly 40%. To compensate for this the lux values from the sensor were scaled by $1/1.4 = 0.714$ which caused the average to properly match the light meter.

	Light Meter Avg (lux)	ALS Avg (lux)	ALS LED Avg (lux)	Scaling Factor
LED1	70.4	72.5	57.8	1.26
LED2	130.5	128.1	76.4	1.68
LED3	136.3	139.3	119.5	1.17
LED4	135.1	140.4	139.8	1.00
LED5	166.4	165.9	158.9	1.04
LED6	142.5	143.9	117.2	1.23
LED7	139.0	136.2	82.4	1.65

Table 5.1: Measurements of illuminance in each single LED region. The scaling factor is the ALS average divided by the ALS LED average.

The second problem was due to the drive circuitry in the luminary and the power supply used. Driving the LEDs through a single BJT hard on/off produced a poor square wave which pushed a sizable fraction of the power seen by the receiver to frequencies outside the fundamental component. This made the estimate taken from the fundamental frequency too low. The problem became worse as more LEDs were turned on since the current draw on the power supply varied with each turning on and off, causing noticeable extra harmonics in the light signals. Since the error manifested as a transfer of power away from the fundamental frequency, a measured constant was used to scale the measurement to what was expected. The average illuminance values were still in agreement.

Table 5.1 shows the results of measuring the center of each single LED region with the light meter, the calibrated ALS measurement, the LED average extracted from the fundamental frequency component, and the resulting scaling factor. The signal from LED7 was the most distorted and had a much weaker fundamental frequency component. LED4 however had very little distortion. These scaling factors agreed well with measurements taken at other points inside the regions.

Noise Source	Measured Noise Power (Brightness, lux)
Receiver dark current offset	2.789
Receiver thermal noise (std. dev.)	0.108
Ambient light noise (std. dev.)	0.576
Receiver std. dev. at 85 lux total signal	1.183
Receiver std. dev. at 300 lux total signal	2.643
Reflected signals	1-5
Stray light signals	5-42

Table 5.2: Measurements of different noise sources at the receiver.

5.2.3 Receiver Noise Measurements

There are several sources of noise that can impact the receiver illuminance measurement. Measurements of different noise sources for the experimental set up are presented in Table 5.2. Since the receiver is calibrated to lux, all measurements are taken as the illuminance reported due to that noise source.

The receiver dark current and thermal noise were measured by completely covering the photodetector. The dark current produced a 2.789 lux offset, and the receiver thermal noise had a standard deviation of 0.108 lux. Next the ambient light noise from the fluorescent sources in the room was measured. The average illuminance in the experimental area was 300.15 lux, and the standard deviation was 0.576 lux. The ballast driving the fluorescent lights introduced noise around 120 Hz, but this was outside the LED drive frequencies. Reflections of the LED light signals were between 1 and 5 lux outside their intended regions and were measured by turning on one LED at a time.

The receiver had a large amount of measurement variance that depended on the amount of incident light signal. The receiver was placed at different locations inside

the pattern and recorded measurements for 30 seconds. The measurements were taken continuously as in Section 5.2.2, extracted from the 512 point FFT at almost 4 Hz. When the average illuminance from the LED1 signal measured 85 lux the receiver had a measurement standard deviation of 1.183 lux, but when it was moved to a brighter part of the LED1 region it had a measurement standard deviation of 2.643 lux. This measurement variance also appeared in all the other LED frequency bins regardless of whether the receiver was inside a region formed by that LED. This was likely caused by receiver shot noise. Averaging over multiple measurements could reduce the effect of this measurement variance at the expense of the system update frequency.

Stray light is any light that escaped the edge of the lens and was not directed towards the intended region. There was a slight gap between the lens and the cardboard used as an aperture stop that light could pass through. A maximum of 42 lux was measured outside the main region from LED7 by turning on each light individually and measuring outside the intended light region.

5.3 Experimental Setup

Figure 5.8 shows the experimental setup used for testing. The luminary was elevated to reduce reflections off of the table and the lens was placed on a stand to allow the distance to be adjusted. The FPGA [89] used to drive the LEDs can be seen beside the luminary stands. The luminary was placed facing the wall 1 m away so that the light pattern would be projected at a known distance. The receiver can be moved across the wall to test the pattern at different points while maintaining a fixed distance.

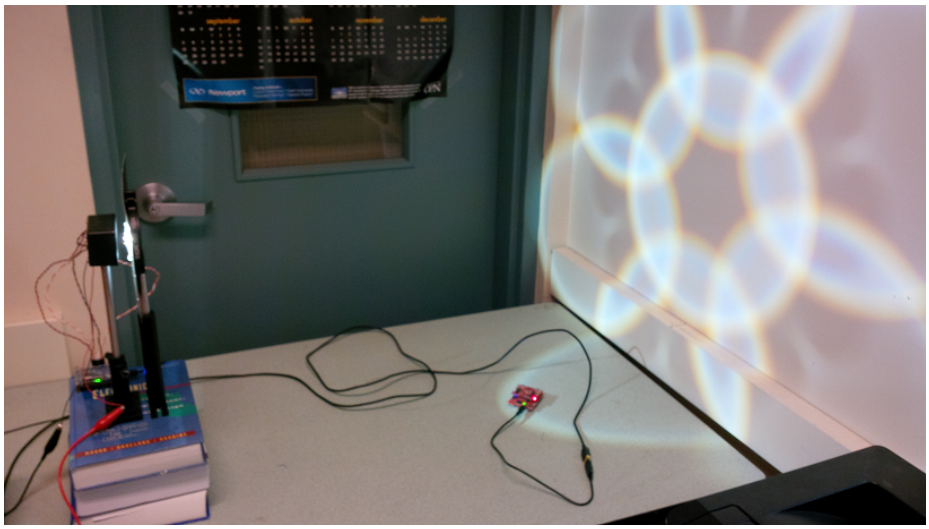

Figure 5.8: Experimental setup for verifying the illumination pattern and positioning performance.

5.3.1 Light Pattern Illuminance

Figure 5.9 shows a surface plot of the expected illuminance distribution for the experimental setup. A small script was written to extract the simulation values from Zemax and import them into matlab to enable easier visualization and post processing. The lowest expected illuminance is 68lux in the middle of the center region, and the highest was 600 lux where three lights overlap. The lowest illuminance measured with the light meter was 70 lux, and the highest was 510 lux. The minimum values were in good agreement with the simulation, but the peak values were slightly lower than expected. It is likely that the simplified PSD of the source used in the simulation underestimated chromatic dispersion from the lens. Chromatic dispersion causes the power from different wavelengths to spread out more around the edges.

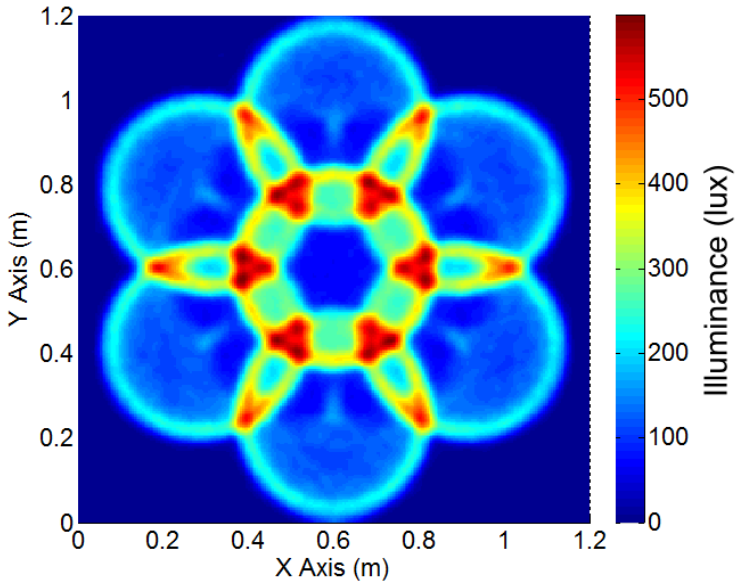

Figure 5.9: Matlab visualization of the calibrated simulation illuminance distribution, colour map is in units of lux.

5.3.2 Dimensions of the Projected Light Regions

The geometry of the on-axis light region is circular due to the symmetry of the lens, and the off axis regions can be reasonably well approximated by an ellipse. An ellipse was fit directly to the simulated illuminance data via linear regression, using the technique outlined in [90]. A general bivariate conic section was fit to the convex hull of the data with constraints that restrict the solutions to ellipses. This direct fitting technique was implemented in a matlab function EllipseDirectFit(XY) by Nikolai Chernov [91].

Figure 5.10 shows a single off axis light region (a) and the resulting ellipse fit to its boundary (b). The ellipse is parameterized by its center (x_c, y_c), semi-major axis a, semi-minor axis b, and rotation angle ϕ. The rotation angle is $\pi/3$, which matches the rotation of the LEDs on the disc.

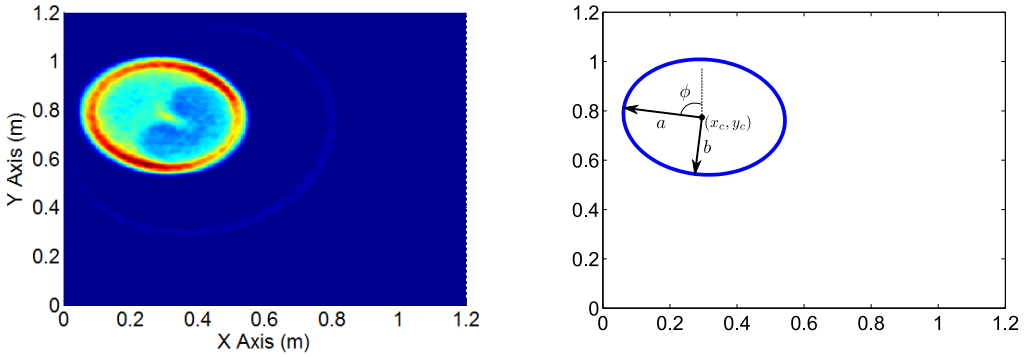

(a) (b)

Figure 5.10: Projected illumination distribution for single off axis LED (a) and the least squares fit elliptical boundary (b).

	Simulation (m)	Measurement (m)
Radius of center region (m)	0.240	0.245
Off axis region semi-major axis (a)	0.245	0.250
Off axis region semi-minor axis (b)	0.230	0.235
Region 2 center (x,y)	(0,0.34)	(0,0.35)
Region 3 center (x,y)	(0.28,0.17)	(0.3,0.18)
Region 4 center (x,y)	(0.28,-0.17)	(0.3,-0.16)
Region 5 center (x,y)	(0,-0.34)	(0,-0.34)
Region 6 center (x,y)	(-0.28,-0.17)	(-0.29,-0.16)
Region 7 center (x,y)	(-0.28,0.17)	(-0.29,0.18)

Table 5.3: Simulated and measured dimensions of the overlapping light pattern.

Table 5.3.2 shows the expected dimensions for the projected light pattern at a distance of 1 m and a LED-lens separation of 19 mm, and the experimentally measured values. The centers of the outer regions were taken as the bright 'pinch' point, as it was easier to measure to. The results are in good agreement though the pattern is slightly shifted up and a bit wider than expected from simulation. The main sources of error are from alignment of the system. Slight deviations in the lens position and orientation can change the edges of the pattern quite drastically. These observations could be important to follow up on in the investigation of manufacturing tolerances on the system. For the purposes of this work however the fit is close enough to allow comparison between the measurements and simulation.

5.3.3 Expected Positioning Error

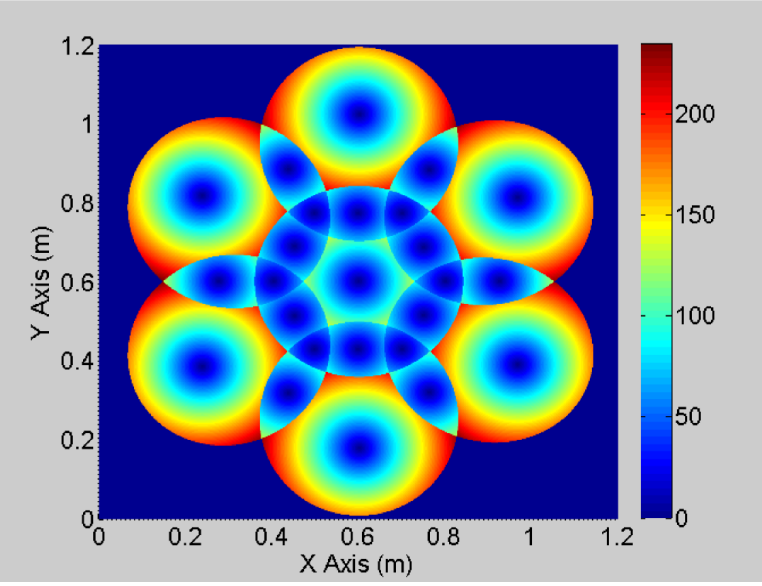

Figure 5.11: Expected error map from the dimensions of the experimental pattern. Z axis is error in (mm).

Figure 5.11 shows the expected error map based on the dimensions of the experimental pattern from Section 5.3.2. Using the center of mass of each region as the position estimate the expected average positioning error over the whole pattern is 10.17 cm.

5.4 Results

This section presents the results of experiments performed with the described set up. First a test to set the region threshold is presented, followed by a series of measurements taken over the overlap pattern to form a rough estimate of positioning performance. Finally measurements taken with the receiver orientation changed are presented.

5.4.1 Choosing a Region Threshold

The experimental setup can be used to investigate choosing a fixed region threshold. As was mentioned in section 4.2.1 the threshold needs to be set lower than the minimum expected illumination level for an LED inside the region, but high enough to avoid reflections and stray light. The receiver was passed through the regions as shown in Figure 5.12 and measurements were recorded of the illuminance from each separate LED.

Figure 5.13 shows the average illuminance levels measured by the mobile receiver as it was passed through the regions. The thick curve represents the total average illuminance seen by the receiver, while the other curves are the individual contribution from each LED extracted from the fundamental frequency component in the FFT.

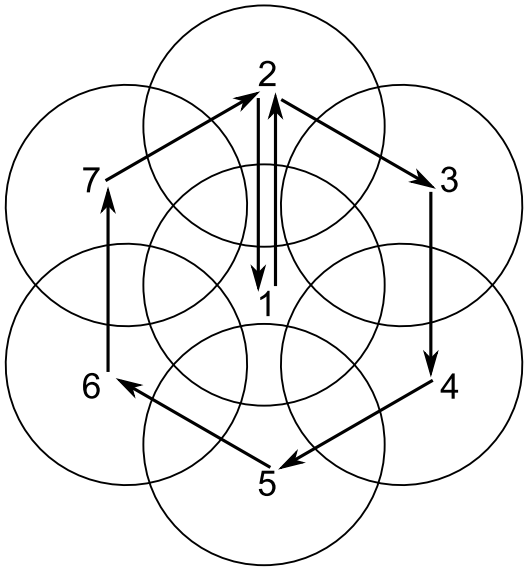

Figure 5.12: Path of the receiver for the region threshold experiment.

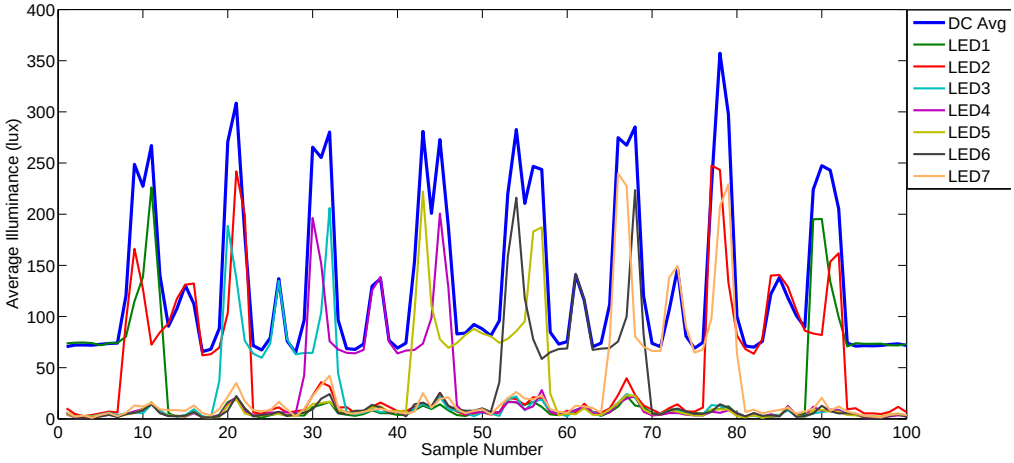

Figure 5.13: Average illuminance levels seen by the mobile receiver passing through the outer regions.

As the receiver was moved from region to region the LEDs that make up each region appear brighter due to the LOS signal path. In regions formed by only one LED the total average illuminance matches the illuminance of the single LED, while in regions with two LEDs the average is the sum of light from both. The transitions from region to region can clearly be seen when the total average illuminance rises or falls sharply, and likewise the LEDs that make up the region can be seen from the curves that rise and fall with it. The rest of the LEDs are only detected by reflections, stray light, or noise and they all cluster near the bottom of the graph. The gap between the bottom of the higher curves that make up the regions and the top of the lower curves that are stray light is the margin for setting the region threshold.

The minimum measured illuminance for a single LED region was 58.5 lux from LED6 at sample 57. The brightest signal visible outside of its intended region was 42.06 lux from LED7 at sample 32. This measurement was due to stray light escaping the edge of the lens, which can also be seen in the simulation in Figure 5.10a as the dim outer ring outside the main region. There is a trade-off between setting the region higher to avoid reflections or stray light, and lower to be more robust to changes in orientation. The threshold was chosen halfway between the two limits at $\eta = 50$ lux without attempting to balance the effect of one type of error against another.

5.4.2 Measured Positioning Error

To test the positioning performance of the system the receiver was placed at different locations inside the light pattern and the received light signals were recorded. The measurement data was used to determine which region the receiver would identify at that location based on the set threshold. The positioning error was then taken as the

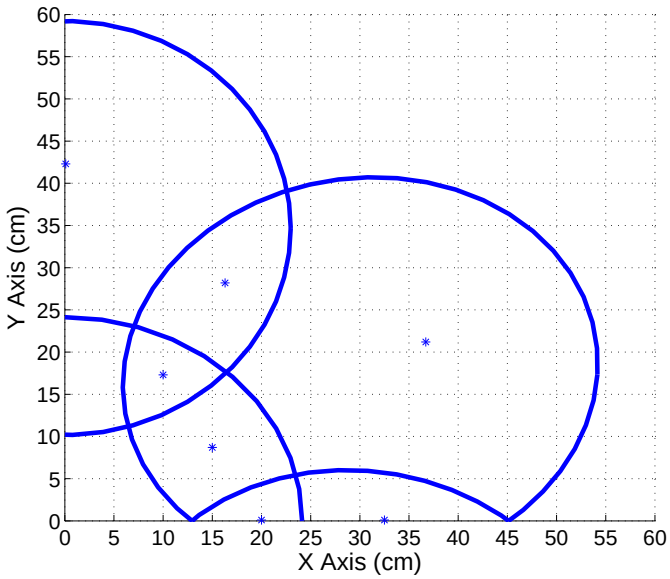

Figure 5.14: Measurement grid for testing the positioning performance of the system. The center of mass for the visible regions are marked by asterisk.

distance from the measurement point to the simulated center of mass of the identified region.

Since the pattern is symmetric only the top right quadrant needs to be measured. A 60cm × 60cm grid was drawn out with grid lines every 5 cm for a total of 169 measurement points. Figure 5.14 shows the simulated region boundaries overlaid with the grid used for the measurements. The calculated center of mass for each region is also plotted.

Measured Illumination and Region Identification

Figure 5.15a shows the measured illuminance distribution over the grid with the receiver held perpendicular to the luminary. The resolution is fairly low due to the wide spacing between measurement points, but the shape and illumination levels

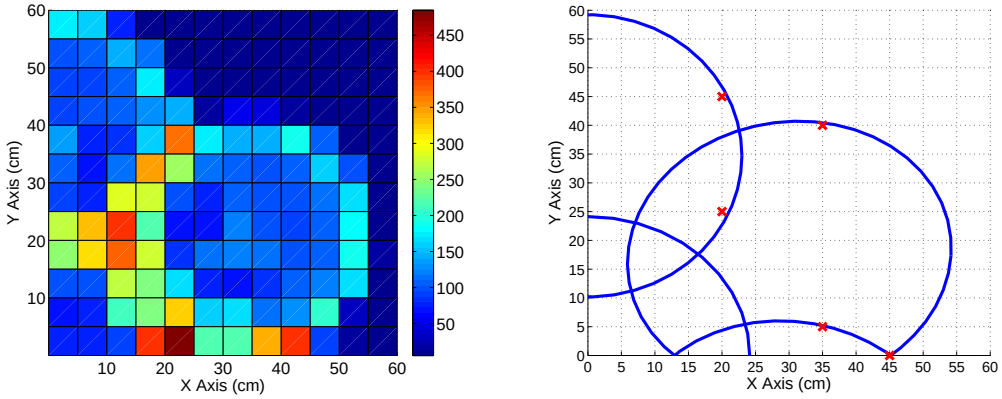

(a) (b)

Figure 5.15: Measured average illuminance (a) and region estimation errors with $\eta = 50$ lux (b).

Measurement Point (cm)	Expected LEDs	Measured LEDs
(45,0)	3,4	3
(35,5)	3,4	3
(20,25)	2,3	3
(35,40)	3	0
(20,45)	2	0

Table 5.4: Summary of measurement points that produced an incorrect region estimate when the receiver was perpendicular to the luminary.

resemble the simulation results. Figure 5.15b shows where the receiver incorrectly identified which region it was in with the threshold set at $\eta = 50$ lux, marked with a red $\times$. A region was considered incorrectly identified if the signals that passed the threshold did not match the region expected from simulation.

Only five of the 169 measurement points produced an incorrect region estimate. Each point is presented in Table 5.4 along with the LEDs expected from the simulation and the LEDs the receiver identified. Since each error was a failure to identify an

LED and not a false detection of an LED outside its region the threshold successfully ignored stray light and reflections. It can be seen from the figure that the receiver only made mistakes near the edge of light regions. The far edges of the regions were less bright and less defined than expected from simulation. At $(35, 5)$ for example, the receiver measured 31.85lux from LED4, which was too low to include it in the region estimate with the threshold at $\eta = 50$ lux.

The error at $(35, 40)$ is more damaging because the receiver failed to detect the only LED that makes up the outer region. Not seeing a light tells the system much less information than seeing one, so when no signals meet the threshold it cannot form a position estimate. The receiver recorded 46.86 lux from LED3, which was just short of the threshold. It may be a good idea to dynamically adjust the threshold to help handle cases where the receiver has trouble forming a region estimate.

Positioning Error

The average position error for the experiment was 10.71 cm, which was very close to the 10.17 cm that the simulation predicted. The difference is due to the receiver incorrectly estimating regions at those few highlighted points and using the wrong center of mass for its position estimate. The coverage area was only 0.235 m^2 in this experiment, but it shows that the values claimed from simulation are practically possible and that AD-VLP can realize sub-meter accuracy in practice. For using nearly the same distance and measurement area as the approach presented in [53], this experiment has demonstrated results on the same order of magnitude (10cm vs 4cm average position error) with a far simpler and easier to implement approach.

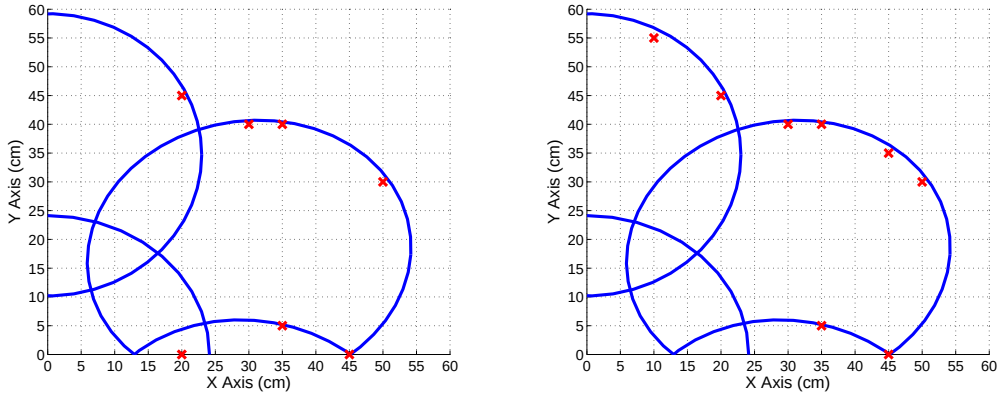

(a) (b)

Figure 5.16: Region estimation errors when the receiver is tilted 30° left (a) and 30° right (b), $\eta = 50$.

5.4.3 Changing Orientation

To investigate the effect of changing orientation the receiver was mounted at fixed angles of $+30°$ and $-30°$ in the horizontal plane and the measurements were repeated. These new measurements can be used with the threshold to check if the receiver can correctly identify the regions at different orientations.

When the receiver was tilted to the left it was facing more towards the luminary

Measurement Point (cm)	Expected LEDs	Measured LEDs
(20,0)	1,3,4	1,3,4,7
(35,5)	3,4	3
(50,30)	3	0
(30,40)	3	0
(30,45)	3	0
(20,45)	2	0

Table 5.5: Summary of measurement points that produced an incorrect region estimate when the receiver was angled 30° towards the luminary.

and will see higher than expected illuminance levels. Figure 5.16a shows the measurement points where the receiver did not correctly identify the region while it was facing 30° left. There were only six incorrect measurement points which means the receiver still correctly identified which region it was in 96% of the time. Many of the errors occurred at the region edges, just as in the case where the receiver was aligned. Table 5.5 presents the expected and identified LED signals at each error point.

The error at (20,0) is of a different type than the rest. Instead of failing to identify an expected signal the receiver detected one that was not part of the expected region. It detected LED7 in the region formed by LEDs 1, 3, and 4. This was caused by stray light from LED7 escaping around the edges of the lens, as mentioned in Section 5.4.1. The receiver measured 53.07 lux from LED7 at this point, which was higher than the 45.40 lux measured when the receiver was perpendicular to the light, and high enough to meet the detection threshold. Because the receiver was facing the light, the stray signal appeared brighter and was enough to pass the threshold. This type of error could also occur for reflections, making them appear stronger if the receiver was near the reflecting surface and facing it.

The errors where the receiver incorrectly detects a signal in a region it does not belong can be corrected by observing the signal strength of the other LEDs that form the region. The measured illuminance from the correct LEDs was 213.21 lux from LED1, 88.10 lux from LED3, and 113.82 lux from LED4, which are all higher than was observed from LED7. Since there is no region where LEDs 1,3,4, and 7 intersect on the region map the receiver can be certain that there has been an error. It can then correct the error by removing the weakest signals until a valid region estimate is formed.

Measurement Point (cm)	Expected LEDs	Measured LEDs
(35,5)	3,4	3
(50,30)	3	0
(45,35)	3	0
(30,40)	3	0
(30,45)	3	0
(20,45)	2	0
(10,55)	2	0

Table 5.6: Summary of measurement points that produced an incorrect region estimate when the receiver was angled 30° away from the luminary.

When the receiver was tilted to the right it was facing away from the luminary and sees less average illuminance than expected. Figure 5.16b shows there were seven incorrect measurement points all clustered at the edges of the regions. Table 5.6 presents each error point, and shows that all the errors were a failure to identify an expected signal. In this experiment some measurements came very close to being below the threshold in the middle of the single LED regions. The minimum illuminance measured in the center of the LED3 region was 52.02 lux. Much lower and it would not meet the detection threshold causing gaps to form in the position estimate.

Positioning Error

The average position error for the left tilted experiment was 12.24 cm which is only 15% higher than the aligned experiment. The average position error for the right tilted experiment was 12.76 cm. The performance was worse when the receiver was tilted away from the luminary because it measured less illuminance and failed to meet the region threshold in more locations. These results prove that AD-VLP is capable of high positioning accuracy even when the receiver orientation changes.

5.5 Conclusions

This chapter has demonstrated the validity of the proposed angular diversity visible light positioning system through multiple experiments. The design of the AD-VLP transmitter was presented and the light pattern was characterized. An ambient light sensor was used as the optical element in a mobile receiver and it was successful in detecting the individual LED signals. These sensors are already present on many devices, and with a faster sampling ADC they could be used for optical communications and VLP without any radical changes to existing mobile device designs.

Experiments at a fixed distance produced an average position error of 10.71 cm when the receiver was held perpendicular to the luminary, 12.24 cm when it was tilted towards, and 12.76 cm when it was tilted away. The positioning accuracy is tied to the ability of the receiver to correctly identify regions, which depends on the choice of signal detection threshold. The threshold used in the experiments was chosen from measurements of the light signal, and proved successful in detecting the correct regions. If the stray light could be eliminated from the transmitter the threshold could be set even lower and the detection at the edges of regions and when the receiver was facing away would improve. Ultimately an adaptive threshold that adjusts based on the relative intensity of all the visible signals could improve the region identification.

Hopefully these experiments have given some insight into the many interrelated components of AD-VLP design, and proven that it is a viable new approach to VLP. There are still many factors to consider before a complete AD-VLP system can be realized, and the next chapter will present some of these future directions.

Chapter 6

Conclusions and Future Work

This thesis has proposed and demonstrated a new method of indoor positioning using visible light based on the principle of angular diversity. It was shown in Chapter 3 that by applying imaging techniques at the transmitter a structured overlapping light pattern could be created and used for positioning. This approach enabled a single element receiver to be used while overcoming the orientation dependence of the channel model. A simple example of designing the light pattern and using it for positioning was presented and used in simulations, which showed an average position error of 15.1 cm was possible over a 5m $\times$ 5m area.

The illumination characteristics of AD-VLP were examined in Chapter 4, where it was proposed that a luminary could be designed with good illumination and positioning performance using more advanced optical design techniques. The most important observation was the need to set a signal strength threshold for identifying positioning regions. It was shown that the threshold needs to be set high enough to ignore reflections and stray light, but lower than the minimum LOS signal for a region. A method for setting the threshold based on simulations of the AD-VLP luminary was

proposed.

An experimental AD-VLP set up was developed and used for testing in Chapter 5. Experiments produced an average positioning error of 10.71cm over the 60cm $\times$ 60cm measurement area, proving AD-VLP can be practically implemented. A simple receiver based on a mobile device ambient light sensor proved effective at detecting the signals from the transmitter and identifying the correct region in the structured overlapping light pattern. A fixed region threshold was set from measurements of the illuminance of the light pattern. This threshold proved effective in separating directed LOS signals from stray light and reflections, allowing the receiver to correctly identify which region it was in 95% of the time even with changing orientation.

6.1 Future Work

This work is a first step in the development of AD-VLP, and it is hoped that the material presented will provide insight for future development of the idea. To that end some proposed future directions are presented.

6.1.1 Mobile Device Integration

Ambient Light Sensor based VLC

It was shown that an ambient light sensor could be used as a low rate optical receiver, and that many devices have this sensor already. It would be worthwhile to investigate the possibility of replacing these sensors with a similar, but higher sampling rate component to enable optical communications on mobile devices with minimal design changes. In the absence of a new component, device manufacturers should allow

faster sampling of the existing parts in their devices. It was shown that the current sensor could be sampled at over 2kHz, which is still an order of magnitude better sampling rate than the camera can provide.

Inertial Measurement Unit Integration

It was mentioned that the mobile receiver also had an IMU on the board. Most mobile devices also have a six or nine axis IMU integrated already. These sensors can be used to provide three dimensional relative motion information that can assist positioning and tracking. Integration of AD-VLP with IMU based indoor positioning could be an extremely promising avenue of research, as existing IMU approaches suffer from drift, but AD-VLP could provide bounding regions to limit this error.

6.1.2 Positioning Algorithm Development

This work has largely focused on being able to define and correctly identify the positioning regions, but not how to translate that information into an updated position estimate. The center of mass estimate used does not consider any other sensor inputs or previous measurements. AD-VLP enables a host of new overlapping region based positioning algorithms to be developed. Information could be extracted from the region transitions, the frequency and direction of region changes, or even looking at the intensity information that was previously discarded in thresholding. The possibility of combining AD-VLP with other techniques also lends itself to new algorithm design. How to most effectively use the positioning information given by AD-VLP should be investigated.

6.1.3 Systems and Communications Development

AD-VLP requires a backbone server system to store the region position map and communicate with the devices being located. How the regions are mapped, stored, and updated is an important topic to be addressed. Advanced possibilities like updating the map based on repeated user measurements could also be useful.

Luminary Signaling and Identification

The biggest part of the system design is the selection of a suitable signaling and identification scheme for the luminaries. FDM is impractical once multiple luminaries are installed, but perhaps CDMA could be used. The scheme should be able to work without synchronization between the luminary and the receiver, or even between adjacent luminaries. Ideally it needs to be able to handle overlap between luminaries as well to take full advantage of the new regions.

6.1.4 Optics Design and AD-VLP Luminaries

A challenging and exciting area of future research for AD-VLP is the design of the structured overlapping light pattern. The simple lens used proved to be effective at demonstrating the idea of AD-VLP, but in order to realize a real AD-VLP system a custom optical element should be considered. Given the ability to design custom illumination distributions with advanced optical design techniques and micro lens arrays there are a number of options to explore designing the overlapping light pattern.

References

[1] Mortadelo2005. The grand hall of terminal 3. Online. Liscensed under CC BY 2.0. [Online]. Available: http://upload.wikimedia.org/wikipedia/commons/7/77/International_airport_toronto_pearson.jpg

[2] Philips, "LUXEON Rebel General Purpose White Portfolio," 2012.

[3] M. Taylor and S. Hranilovic, "Angular diversity approach to indoor positioning using visible light," in *2013 Proceedings IEEE GLOBECOM*, Atlanta, GA, Dec 2013.

[4] F. Donovan. (2013, Oct) Indoor location market to reach $4 billion in 2018. Online. FierceMobileIT. [Online]. Available: http://www.fiercemobileit.com/story/indoor-location-market-reach-4-billion-2018-predicts-abi/2013-10-18

[5] P. Harrop. (2014) Online. IDTechEx. [Online]. Available: http://www.idtechex.com/research/articles/mobile-phone-indoor-positioning-systems-create-10bn-market-00006207.asp

[6] A. Oursler. (2012, Aug) Lost at Wal-Mart? There's an app for that. InvestorPlace. [Online]. Available: http://money.msn.com/top-stocks/post.aspx?post=1c4300dc-0e9e-43af-bc42-11c7dcebb180

[7] T. Costa. (2013, March) Apple + WiFiSLAM = Game on for Indoor Location. Frrester Research. [Online]. Available: http://www.forbes.com/sites/forrester/2013/03/29/apple-wifislam-game-on-for-indoor-location/

[8] Google Maps: Indoor. Online. Google Inc. [Online]. Available: http://maps.google.com/help/maps/indoormaps/index.html

[9] J. Schiller and A. Voisard, *Location-Based Services*, ser. The Morgan Kaufmann Series in Data Management Systems. Elsevier Science, 2004. [Online]. Available: http://books.google.ca/books?id=wj19b5wVfXAC

[10] Bytelight. [Online]. Available: http://www.bytelight.com/

[11] Philips. (2014, March) Philips creates shopping assistant with LEDs and smart phone. IEEE Spectrum online. IEEE. [Online]. Available: http://spectrum.ieee.org/img/philips%20infog%20store%20light-1392764161845.jpg

[12] K. Pahlavan, X. Li, and J.-P. Makela, "Indoor geolocation science and technology," *Communications Magazine, IEEE*, vol. 40, no. 2, pp. 112–118, 2002.

[13] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of Wireless Indoor Positioning Techniques and Systems," *IEEE Transactions ON Systems, Man, and Cybernetics*, vol. 37, no. 6, pp. 1067–1080, November 2007.

[14] Y. Gu, A. Lo, and I. Niemegeers, "A survey of indoor positioning systems for wireless personal networks," *IEEE Communications Letters Surveys & Tutorials*, vol. 11, no. 1, pp. 13–32, 2009.

[15] H. Koyuncu and S. H. Yang, "A survey of indoor positioning and object locating systems," *IJCSNS International Journal of Computer Science and Network Security*, vol. 10, no. 5, pp. 121–128, 2010.

[16] K. Al Nuaimi and H. Kamel, "A survey of indoor positioning systems and algorithms," in *Innovations in Information Technology (IIT), 2011 International Conference on*, April 2011, pp. 185–190.

[17] R. Mautz and S. Tilch, "Survey of optical indoor positioning systems," in *Indoor Positioning and Indoor Navigation (IPIN), 2011 International Conference on*, Sept 2011, pp. 1–7.

[18] R. Harle, "A survey of indoor inertial positioning systems for pedestrians," *Communications Surveys Tutorials, IEEE*, vol. 15, no. 3, pp. 1281–1293, Third 2013.

[19] M. Hazas and A. Hopper, "Broadband ultrasonic location systems for improved indoor positioning," *Mobile Computing, IEEE Transactions on*, vol. 5, no. 5, pp. 536–547, May 2006.

[20] Y. Wang, X. Jia, H. Lee, and G. Li, "An indoors wireless positioning system based on wireless local area network infrastructure," in *6th Int. Symp. on Satellite Navigation Technology Including Mobile Positioning & Location Services*, no. 54, 2003.

[21] I. Bisio, F. Lavagetto, M. Marchese, and A. Sciarrone, "Performance comparison of a probabilistic fingerprint-based indoor positioning system over different smartphones," in *Performance Evaluation of Computer and Telecommunication Systems (SPECTS), 2013 International Symposium on*, July 2013, pp. 161–166.

[22] X. Liu, H. Makino, and Y. Maeda, "Basic study on indoor location estimation using visible light communication platform," in *30th IEEE EMBS Conference*, Vancouver, BC, August 2008, pp. 2377–2380.

[23] T. Komine and M. Nakagawa, "Fundamental analysis for visible-light communication system using LED lights," *IEEE Transactions on Consumer Electronics*, vol. 50, no. 1, pp. 100–107, February 2004.

[24] *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.

[25] T. Komine and M. Nakagawa, "Integrated system of white LED visible-light communication and power-line communication," *Consumer Electronics, IEEE Transactions on*, vol. 49, no. 1, pp. 71–79, Feb 2003.

[26] D. O'brien, L. Zeng, H. Le-Minh, G. Faulkner, J. Walewski, and S. Randel, "Visible light communications: Challenges and possibilities," in *IEEE. Int. Symp. Personal, Indoor and Mobile Radio Communications*, Sept 2008, pp. 1–5.

[27] M. Taylor, "Indoor light positioning system," US Provisional 61/914,089, 2013.

[28] E. C. Chan and G. Baciu, *Introduction to Wireless Localization: With IPhone SDK Examples*. John Wiley & Sons, 2012.

[29] J. Hightower and G. Borriello, "A survey and taxonomy of location systems for ubiquitous computing," *IEEE computer*, vol. 34, no. 8, pp. 57–66, 2001.

[30] L. Del Mundo, R. Ansay, C. Festin, and R. Ocampo, "A comparison of wireless fidelity (wi-fi) fingerprinting techniques," in *ICT Convergence (ICTC), 2011 International Conference on*, Sept 2011, pp. 20–25.

[31] S. Goswami, *Indoor Location Technologies*. Springer, 2013.

[32] Google faces streetview Wi-Fi snooping action. BBC. [Online]. Available: http://www.bbc.com/news/technology-24047235

[33] M. Bass, C. DeCusatis, J. Enoch, V. Lakshminarayanan, G. Li, C. Macdonald, V. Mahajan, and E. Van Stryland, *Handbook of optics. Volume III, Vision and vision optics*, 3rd ed. McGraw-Hill, Inc., 2010.

[34] A. Ryer, "Light measurement handbook," International Light Inc, Technical publication, 1998.

[35] "Handbook of LED metrology," Instrument Systems GmbH, Technical publication.

[36] D. B. Long. (2007) Photometry. [Online]. Available: http://www.drdrbill.com/optics.html

[37] *Colorimetry*, International Organization for Standardization ISO/ International Commission on Illumination CIE Std. ISO 11 664/CIE S 014. [Online]. Available: http://www.cie.co.at/Publications/Standards

[38] *Safety of laser products*, International Electrotechnical Commission Std. IEC 60 825, 2007.

[39] *Photobiological safety of lamps and lamp systems*, International Electrotechnical Commission Std. IEC 62 471, 2006.

[40] *Artificial lighting of interiors*, Deutsches Institut für Normung, DIN Std. 5035, 1990.

[41] I. Ashdown and M. Salsbury, "A near-field goniospectroradiometer for led measurements," in *Contract Proceedings 2006*. International Society for Optics and Photonics, 2007, pp. 634 215–634 215.

[42] J. Kahn and J. Barry, "Wireless Infrared Communications," *Proceedings of the IEEE*, vol. 85, no. 2, pp. 265–298, 1997.

[43] J. Barry, J. Kahn, W. Krause, E. Lee, and D. Messerschmitt, "Simulation of multipath impulse response for indoor wireless channels," *IEEE Journal on Selected Areas in Communications*, vol. 11, no. 3, pp. 367–379, April 1993.

[44] W. Zhang and M. Kavehrad, "A 2-D indoor localization system based on visible light LED," *Photonics Society Summer Topical Meeting Series*, pp. 80–81, 2012.

[45] C. Sertthin, E. Tsuji, S. Kuwano, K. Watanabe, and M. Nakagawa, "A switching estimated receiver position scheme for visible light based indoor positioning system," in *ISWPC*, Feb 2009.

[46] M. Rahaim, G. Prince, and T. Little, "State estimation and motion tracking for spatially diverse vlc networks," in *Globecom Workshops (GC Wkshps), 2012 IEEE*, Dec 2012, pp. 1249–1253.

[47] Md. Shariful Islam, "Indoor positioning through integration of optical angles of arrival with an inertial measurement unit," Master's thesis, UBC, 2012.

[48] C. Sertthin, T. Fujii, O. Takyu, Y. Umeda, and T. Ohtsuki, "On physical layer simulation model for 6-axis sensor assisted vlc based positioning system," *Proc. Globecom 2011*, pp. 1–5, Dec 2011.

[49] M. Yasir, S.-W. Ho, and B. N. Vellambi, "Indoor localization using visible light and accelerometer," in *Global Communications Conference (GLOBECOM), 2013 IEEE*, Atlanta GA, Dec 2013, pp. 3363–3368.

[50] A. Arafa, X. Jin, and R. Klukas, "Wireless indoor optical positioning with a differential photosensor," *Photonics Technology Letters, IEEE*, vol. 24, no. 12, pp. 1027–1029, June 2012.

[51] B. Xie, G. Tan, Y. Liu, M. Lu, K. Chen, and T. He, "LIPS: A light intensity based positioning system for indoor environments," *arXiv preprint arXiv:1403.2331*, 2014.

[52] 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. [Online]. Available: http://dx.doi.org/10.1002/mop.26763

[53] H.-S. Kim, D.-R. Kim, S.-H. Yang, Y.-H. Son, and S.-K. Han, "Inter-cell interference mitigation and indoor positioning system based on carrier allocation visible light communication," in *ICSPCS*, 2011.

[54] G. B. Prince and T. D. Little, "A two phase hybrid RSS/AoA algorithm for indoor device localization using visible light," in *Globecom*, Anaheim, CA, Dec 2012, pp. 3347–3352.

[55] M. Shah, "Fundamentals of computer vision," 1997.

[56] M. Yoshino, S. Haruyama, and M. Nakagawa, "High-accuracy positioning system using visible LED lights and image sensor," in *IEEE Radio and Wireless Symposium*, Orlando, FL, January 2008, pp. 439 – 442.

[57] A. Ganick and D. Ryan, "Light positioning system using digital pulse recognition," Jan. 31 2013, US Patent App. 13/591,014. [Online]. Available: https://www.google.com/patents/US20130028475

[58] R. Roberts, "Space-time forward error correction for dimmable undersampled frequency shift on-off keying camera communications (camcom)," in *Ubiquitous and Future Networks (ICUFN), 2013 Fifth International Conference on*, July 2013, pp. 459–464.

[59] B. Saleh and M. Teich, *Fundamentals of Photonics*, ser. Wiley Series in Pure and Applied Optics. Wiley, 2013. [Online]. Available: http://books.google.ca/books?id=Qfeosgu08u8C

[60] F. A. Jenkins and H. E. White, *Fundamentals of optics*, 4th ed. McGraw-Hill, 2001.

[61] CREE INC., "LED luminarie design guide," 2012, Appl. Note CLD-AP15 rev 0A.

[62] T. C. W. Schenk, L. Feri, H. Yang, and J. P. M. G. Linnartz, "Optical wireless CDMA employing solid state lighting LEDs," in *Summer Topical Meeting, 2009. LEOSST '09. IEEE/LEOS*, July 2009, pp. 23–24.

[63] R. Lenk and C. Lenk, *Practical Lighting Design with LEDs*, ser. IEEE Press Series on Power Engineering. Wiley, 2011. [Online]. Available: http://books.google.ca/books?id=Fb2eSQAACAAJ

[64] Y.-Y. C. Allen Jong-Woei Whang and Y.-T. Teng, "Designing uniform illumination systems by surface-tailored lens and configurations of LED arrays," *Journal of Display Technology*, vol. 5, no. 3, pp. 94–103, March 2009.

[65] H. Ries and A. Rabl, "Edge-ray principle of nonimaging optics," *JOSA A*, vol. 11, no. 10, pp. 2627–2632, 1994.

[66] E. Monteiro, "Design and implementation of color-shift keying for visible light communications," Master's thesis, McMaster University, 2013.

[67] H. Yang, J. Bergmans, T. C. W. Schenk, J. Linnartz, and R. Rietman, "Uniform illumination rendering using an array of LEDs a signal processing perspective," in *IEEE Transactions on Signal Processing*, vol. 57, no. 3, March 2009.

[68] Zemax13, release 2. Zemax, LLC. [Online]. Available: http://www.zemax.com/

[69] I. Moreno and C.-C. Sun, "Modeling the radiation pattern of LEDs," *Optics express*, vol. 16, no. 3, pp. 1808–1819, 2008.

[70] I. Moreno, C.-C. Sun, and R. Ivanov, "Far-field condition for light-emitting diode arrays," *Applied optics*, vol. 48, no. 6, pp. 1190–1197, 2009.

[71] 50mm Dia. x 50mm FL Uncoated, Double-Convex Lens EO-32978. Edmund Optics. [Online]. Available: http://www.edmundoptics.com/optics/optical-lenses/double-convex-dcx-spherical-singlet-lenses/uncoated-double-convex-dcx-lenses/32978

[72] Optical LED design resources. Philips. [Online]. Available: http://www.philipslumileds.com/support/design-resources/optical

[73] F. R. Fournier, "A review of beam shaping strategies for LED lighting," *Illumination Optics II, Proc. SPIE*, vol. 8170, pp. 1–11, 2011.

[74] K. Wang, F. Chen, Z. Liu, X. Luo, and S. Liu, "Design of compact freeform lens for application specific light-emitting diode packaging," *Optics express*, vol. 18, no. 2, pp. 413–425, 2010.

[75] S. C. Sun Liwei, Shangzhong Jin, "Design of free-form microlens for LED general illumination," in *Communications and Photonics Conference and Exhibition*, Shanghai, China, November 2009, pp. 1–2.

[76] G. Wang, L. Wang, F. Li, and D. Kong, "Design of optical element combining fresnel lens with microlens array for uniform light-emitting diode lighting," *Journal of the Optical Society of America*, vol. 29, no. 9, pp. 1877–1834, September 2012.

[77] B. Levy, *Principles of Signal Detection and Parameter Estimation*. Springer, 2008.

[78] S. Kay, *Fundamentals Of Statistical Signal Processing*, ser. Prentice-Hall signal processing series. Prentice Hall, 2001, no. v. 1.

[79] K. Lee, H. Park, and J. R. Barry, "Indoor channel characteristics for visible light communications," *IEEE Communicatinons Letters*, vol. 15, no. 2, pp. 217–219, January 2011.

[80] J. Kahn, W. Krause, and J. Carruthers, "Experimental characterization of non-directed indoor infrared channels," *IEEE Transactions on Communications*, vol. 43, no. 2/3/4, pp. 1613–1623, Feb 1995.

[81] Luxeon star homepage. Luxeon Star LEDs. [Online]. Available: http://www.luxeonstar.com/

[82] PN2222A. Fairchild Semiconductor. [Online]. Available: http://www.fairchildsemi.com/ds/PN/PN2222A.pdf

[83] PDA36A - Si Switchable Gain Detector. Thorlabs Inc. [Online]. Available: http://www.thorlabs.com/thorProduct.cfm?partNumber=PDA36A

[84] Integrated digital light sensor with interrupt, ISL29023 datasheet. Intersil™. [Online]. Available: http://www.intersil.com/content/dam/Intersil/documents/isl2/isl29023.pdf

[85] Tiva™c series launchpad evaluation kit. Texas Instruments. [Online]. Available: http://www.ti.com/tool/ek-tm4c123gxl

[86] Sensor hub boosterpack. Texas Instruments. [Online]. Available: http://www.ti.com/tool/boostxl-senshub

[87] MPU-9150 nine-axis MEMS motion tracking device. InvenSense. [Online]. Available: http://www.invensense.com/mems/gyro/mpu9150.html

[88] Hd400: Heavy duty light meter. Extech Instruments. [Online]. Available: http://www.extech.com/instruments/product.asp?catid=10&prodid=56

[89] Terasic Technologies Inc. DE0-Nano development and education board. [Online]. Available: http://www.terasic.com.tw/cgi-bin/page/archive.pl?No=593

[90] A. Fitzgibbon, M. Pilu, and R. Fisher, "Direct least square fitting of ellipses," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 21, no. 5, pp. 476–480, May 1999.

[91] N. Chernov. (2009, Jan) Ellipse fit (direct method). [Online]. Available: http://www.mathworks.com/matlabcentral/fileexchange/22684-ellipse-fit-direct-method/content/EllipseDirectFit.m