INSTITUTO SUPERIOR DE ENGENHARIA DO PORTO
MESTRADO EM ENGENHARIA ELECTROTÉCNICA E DE COMPUTADORES



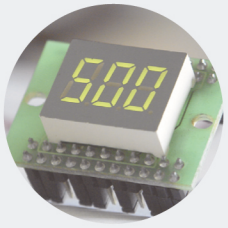


Enhancing Indoor Localisation: a Bluetooth 5.1 Beacon Placement approach

JOÃO PEDRO DA SILVA DIAS
novembro de 2023

POLITÉCNICO
DO PORTO

Politécnico do Porto
Instituto Superior de Engenharia do Porto

Enhancing Indoor Localisation: a Bluetooth Low Energy (BLE) Beacon Placement approach

João Pedro da Silva Dias

Master in Electrical and Computer Engineering
Specialization Area of Automation and Systems

Departamento de Engenharia Eletrotécnica
Instituto Superior de Engenharia do Porto

November, 2023

This dissertation partially satisfies the requirements of the Thesis/Dissertation course of the program Master in Electrical and Computer Engineering, Specialization Area of Automation and Systems.

Candidate: João Pedro da Silva Dias, No. 1181617, 1181617@isep.ipp.pt

Supervisor: Paula Maria Marques Moura Gomes Viana, pmv@isep.ipp.pt

Supervisor: Duarte Pereira, duarte.pereira@fraunhofer.pt

Supervisor: Miguel Roque, miguel.roque@fraunhofer.pt

Company: Fraunhofer Portugal Research

Departamento de Engenharia Eletrotécnica
Instituto Superior de Engenharia do Porto
Rua Dr. António Bernardino de Almeida, 431, 4200–072 Porto

November, 2023

Abstract

Indoor location-based services have become increasingly vital in various sectors, including industries, healthcare, airports, and crowded infrastructures, facilitating asset tracking and user navigation. This project addresses the critical challenge of optimising beacon placement for indoor location, employing Bluetooth technology as the communication protocol. The significance of this research lies in the efficiency and accuracy that an optimised beacon layout can provide, enhancing the effectiveness of indoor positioning systems. The algorithm developed takes into consideration materials attenuation, coverage and Line of Sight (LOS) conditions to optimise its layouts. Experimental validation of the algorithm's performance was conducted by comparing two beacon layouts: one optimised by the algorithm and the other manually arranged by individuals with empirical knowledge in the field. The experiment considered three distinct positions within the schematic, allowing for a comprehensive assessment of the optimised layout's superior performance. The results of this research offer insights into the potential of the algorithm to revolutionise indoor location services, providing a more reliable and cost-effective solution for a multitude of applications.

Keywords: Indoor localisation, Beacon placement, Bluetooth, Asset tracking.

Resumo

Os serviços de localização em ambientes internos tornaram-se cada vez mais essenciais em vários setores, incluindo indústrias, cuidados de saúde, aeroportos e infraestruturas movimentadas, facilitando o rastreamento de objetos e a navegação de utilizadores. Este projeto aborda o desafio crítico da otimização da colocação de beacons para localização em ambientes internos, utilizando a tecnologia Bluetooth como protocolo de comunicação. A importância desta pesquisa reside na eficiência e precisão que uma disposição otimizada de beacons pode proporcionar, melhorando a eficácia de sistemas de posicionamento em ambientes internos. O algoritmo desenvolvido leva em consideração a atenuação de materiais, a cobertura e as condições de visão direta para otimizar as suas disposições. A validação experimental do desempenho do algoritmo foi realizada ao comparar duas disposições de beacons: uma otimizada pelo algoritmo e outra organizada manualmente por indivíduos com conhecimento empírico na área. A experiência considerou três posições distintas no esquema, permitindo uma avaliação abrangente do desempenho superior da disposição otimizada. Os resultados desta pesquisa oferecem descobertas importantes sobre o potencial do algoritmo para revolucionar os serviços de localização em ambientes internos, proporcionando uma solução mais confiável e econômica para uma variedade de aplicações.

Palavras-Chave: Localização interna, Posicionamento de beacons, Bluetooth, Rastreamento de objetos.

Contents

List of Figures

List of Tables

List of Acronyms

G_{RX}	Receiver Antenna Gain
G_{TX}	Transmitter Antenna Gain
L	Losses
P_{RX}	Receiver Sensibility
P_{TX}	Transmitted Power
AoA	Angle of Arrival
AP	Access Point
ATE	Asset Tracking Engine
BLE	Bluetooth Low Energy
BS	Base Station
CSP	Constraint Satisfaction Problem
CTE	Constant Tone Extension
DEE	Departamento de Engenharia Electrotécnica
GPS	Global Positioning System
HPrecision	High Precision
IFFT	Inverse Fast Fourier Transform
ILS	Indoor Location Systems
IMU	Inertial Measurement Units
IoT	Internet of Things
IR	Infrared Radiation
ISEP	Instituto Superior de Engenharia do Porto
ISM	Industrial, Scientific and Medical

LED	Light Emitting Diode
LOS	Line of Sight
MEEC	Mestrado em Engenharia Electrotécnica e Computadores
NLOS	Non Line of Sight
OpenCV	Open Source Computer Vision Library
RGB	Red, Green and Blue
RS	Random Sampling
RSS	Received Signal Strength
RToF	Round-Trip Time of Flight
TDoA	Time Difference of Arrival
ToA	Time of Arrival
ToF	Time of Flight
UD	Unlocated Device
UL	Unique Location
UWB	Ultra-Wideband
WLAN	Wireless Local Area Networks

List of Symbols

Symbol	Description	Units
f	frequency	Hz
G	antenna gain	dBi
P	signal power level	dBm
P	power	W
t	time	s
x	displacement	m

Chapter 1

Introduction

In this introductory chapter, we present a contextualisation of the project, define the problem it addresses, outline its objectives and expected results, and provide an overview of the document's organisation.

1.1 Contextualisation

The presented project is part of the scope of the final work for the Master's thesis of the Mestrado em Engenharia Electrotécnica e Computadores (MEEC), from the Departamento de Engenharia Electrotécnica (DEE) at the Instituto Superior de Engenharia do Porto (ISEP).

Indoor localisation is a critical component across various industries where precise tracking and navigation within enclosed spaces are imperative. The relevance of this research extends to sectors such as:

- Retail and Hospitality: Improving customer experiences through targeted advertising, personalised promotions, and efficient indoor navigation within shopping malls or large establishments;
- Healthcare: Facilitating asset tracking, ensuring patient safety through real-time location monitoring, and optimising the workflow in medical facilities;
- Logistics and Warehousing: Enhancing inventory management, enabling efficient tracking of goods, and improving overall operational efficiency within warehouses;

- Smart Buildings: Contributing to energy conservation, optimising space utilisation, and implementing advanced security measures in modern smart buildings.

However, the application of indoor localisation systems faces notable challenges intrinsic to indoor environments, including:

- Signal Attenuation: Indoor spaces often contain materials that absorb or obstruct signals, leading to signal attenuation and reducing the accuracy of localisation;
- Multipath Effects: Signal reflections from surfaces can result in multipath effects, causing interference and inaccuracies in localisation estimates;
- Structural Complexity: Diverse architectural structures introduce complexities such as signal blockages and irregular signal propagation, further complicating accurate localisation.

In the realm of indoor localisation, a diverse array of technologies is employed to tackle these challenges. Technologies include WiFi-based positioning, Bluetooth Low Energy (BLE), Ultra-Wideband (UWB), Infrared Radiation (IR), and more. Each technology brings unique advantages and trade-offs, reflecting the need for tailored solutions based on specific use cases and environmental characteristics.

A distinctive characteristic of indoor localisation systems is the absence of a standardised approach. Unlike outdoor positioning systems with Global Positioning System (GPS) as a universal standard, each indoor environment presents its unique set of challenges, making it difficult to establish a one-size-fits-all solution. The lack of a standard in Indoor Location Systems (ILS) is attributed to the nuanced characteristics of individual indoor positioning problems.

1.2 Problem Definition

The effectiveness of indoor localisation systems is intricately linked to the strategic placement of beacons within the environment. The limitations of human-driven beacon deployment, whether through over-provisioning or under-provisioning, underscore the need for an advanced algorithmic approach to address these challenges.

Over-provisioning, involving the excessive placement of beacons, leads to increased costs and potential interference. Conversely, under-provisioning compromises the precision of localisation, hindering the system's performance.

The proposed approach described in this work aims to enhance indoor localisation by proposing an algorithm-driven beacon placement approach. The approach introduces key concepts to overcome challenges associated with common deployments typically done by humans:

1. LOS Conditions:
    - Objective: Mitigate signal attenuation and multipath effects;
    - Rationale: By strategically placing beacons in locations with optimal visibility, the algorithm aims to enhance signal propagation and minimise interference, addressing challenges encountered in conventional deployments.
2. Full Coverage of the Schematic:
    - Objective: Ensure the ability to compute the asset's position from any point in the schematic;
    - Rationale: Full schematic coverage guarantees that, regardless of the asset's location within the indoor space, the localisation system can accurately determine its position, eliminating potential blind spots.
3. Lateration Using Two BSs:
    - Objective: Reduce over-provisioning, minimising costs while maintaining system accuracy;
    - Rationale: Leveraging lateration with two strategically placed base stations offers a more cost-effective solution, reducing the number of required beacons while preserving accuracy.
4. Adaptation Based on Building Schematic:
    - Objective: Utilise information about the environment for adaptive base station positions;
    - Rationale: Incorporating the building schematic allows the system to identify materials causing signal attenuation, enabling the algorithm to adjust base station positions in coherence with their range, and enhancing the accuracy of distance estimations.

Through the implementation of these concepts, this research seeks to optimise beacon deployment, ensuring a more accurate, cost-effective, and adaptable ILS. The proposed approach aims to overcome the limitations of human-centric strategies, advancing the capabilities of indoor localisation in diverse real-world applications.

1.2.1 Objectives

To address the proposed problem, it is necessary to achieve the following fundamental objectives:

- Use the building schematic to obtain information about the environment

- Develop an algorithm that outputs an optimised beacon layout for indoor location
- Ensure the layout has full coverage of the schematic
- Test and compare the output layout with a pre-implemented layout
- Validate the hardware and software solution for indoor localisation, developed in the scope of a previous work

1.3 Dissertation Structure

This work is structured into five distinct chapters, each contributing to a comprehensive understanding of the project's scope and outcomes.

The current chapter, the introduction, serves as the entry point to the research, providing contextualisation for the project while establishing the overarching objectives. It outlines the problem statement and sets the stage for the subsequent chapters, ensuring readers grasp the significance of the work.

In the second chapter, a thorough exploration of the technologies and techniques relevant to indoor location is presented. Additionally, this chapter offers insights into existing solutions for beacon placement algorithms. By reviewing the current state-of-the-art, it lays the foundation for the algorithm's development in the subsequent chapters.

Chapter three examines and explores the application of supporting technologies, with a primary focus on image processing for acquiring features from images. This chapter establishes the foundation for comprehending the crucial role that the schematic plays in the developed algorithm.

The fourth chapter delves into the methodology employed in crafting the beacon placement algorithm. It elucidates the step-by-step approach taken in developing the algorithm, from conceptualisation to implementation. This chapter is instrumental in elucidating the strategies and techniques that underpin the algorithm's design.

Chapter five represents the heart of the work, where the results of experimental testing are presented and analysed. It offers empirical evidence of the algorithm's performance, shedding light on its strengths, limitations, and practical implications. These results provide invaluable insights into the algorithm's real-world applicability.

The sixth and final chapter brings the work to a conclusion, encapsulating the key findings, summarising the achievements, and addressing the limitations encountered during the research process. It also offers a broader perspective on the implications and potential future directions for this area of study, solidifying the work's contribution to the field of indoor location and beacon placement algorithms.

Chapter 2

State of the Art in Indoor Location Systems (ILS)

ILS have emerged as a pivotal technology in recent years [1], offering a wide array of applications across various sectors. These systems are instrumental in tracking and monitoring objects or people within indoor environments and their application ranges from commercial buildings to healthcare facilities and beyond. The importance of ILS lies in its ability to address the unique challenges presented by indoor spaces, where GPS signals are often weak or unavailable due to signal attenuation [1, 2]. In the absence of a single standardised solution for indoor location, enterprises and innovators have been actively developing a diverse range of technologies and systems to meet this growing demand [3, 4, 5].

The main goal of indoor location is to estimate the distance between a Unlocated Device (UD) and a BS. This is usually achieved with the support of a communication protocol that enables establishing a connection between the two devices. Solutions may not be limited to a single technique or technology, and a hybrid solution can enable increasing the accuracy. [3, 4, 5] are some examples found on the market that build their solutions according to the demands of the project. Typically the decision points to evaluate an ILS system are coverage, accuracy, precision, power consumption, and scalability among others [1, 6].

Sections 2.1 and 2.2 will discuss in more detail all these aspects.

2.1 Localisation Technologies

Several approaches for localisation can be identified. The technologies mentioned in Figure 2.1 include infrared, ultrasound and radio frequency systems. Although not represented in Figure 2.1 inertial systems can also be used for localisation as it is explained in Subsection 2.1.1.

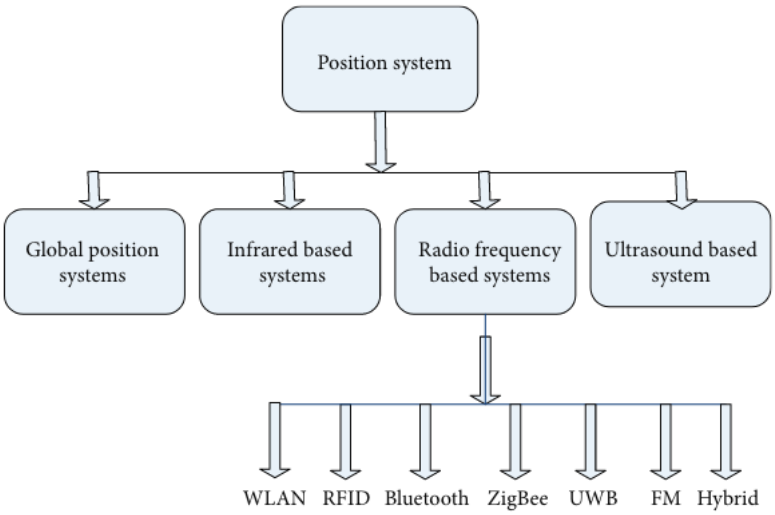

Figure 2.1: Position systems [7].

2.1.1 Inertial Systems

Inertial Measurement Units (IMU) are systems capable of calculating the current position of a moving object by using a previously determined position - this is known as Dead Reckoning [7]. For estimating those variations, the system uses a combination of an accelerometer, a gyroscope and, often, a magnetometer which measure linear accelerations, rotational speed, and magnetic field strength respectively, imparted on an object within a local frame of reference [8].

One of the advantages of this technology is the fact that it does not require anchors in the environment - only an IMU is required. Due to this and the low cost of the unit, this system is highly scalable [8].

In [8], it is pointed out that just like the GPS the use of a magnetometer for indoor situations is not viable as it suffers local disturbances mainly from the infrastructure and, as a result, it struggles to find the magnetic north, leading to inaccuracies.

The main limitations are the accumulation of errors over time and the fact that it only provides relative positioning. To obtain absolute positioning, it is necessary to implement supporting technology either through made devices or personal ones. The first one requires a costly approach as such devices need to be made for that application and are locked to it. The second approach takes advantage of a pre-built

infrastructure (smartphones, watch bands, smartwatches), although this method ensures an easier widespread of the technology, accessing the data from those devices raises privacy legal issues that make this option unavailable.

2.1.2 Infrared Light

Infrared light systems require an unobstructed LOS between the BS and the UD [7]. This type of system can be used as a very reliable room detector. The main advantages of a IR based system are its reduced dimensions and weight.

However, since light cannot traverse walls, a UD cannot detect light from an BS without being in the same room. For precise localisation, these solutions require installing many BSs and can struggle due to the low quality and strength of the signal. They are also susceptible to interference from fluorescent light or sunlight [9]. In [10] the importance of fluorescent lights is downplayed due to Light Emitting Diode (LED)s being the norm nowadays. Despite this downplay, in the described experiment each of the four ceiling lights used is modulated with a different frequency, to ensure each emits different signals and therefore can be distinguished by the algorithm used in the back end.

In [11], a system based on this technology, where the security cameras are the BSs, is proposed. The approach is to use the image captured from the camera, to detect the infrared light, and then compute its coordinates. Despite achieving position errors below 6 cm, the system implies the use of multiple cameras to ensure total area coverage. By using preexisting security cameras, costs can be reduced. However, total coverage of the infrastructure has to be guaranteed.

2.1.3 Ultrasound

Ultrasound technology uses acoustic signals with frequencies up to 25 KHz, delivering interference-free communication systems. By measuring the time of signals as they traverse from transmitters to receivers, ultrasound technology provides remarkably precise relative distance information [12].

Typically, an ultrasound system comprises ultrasound signal detectors (BSs) and ultrasound tags (UDs), enabling direct point-to-point communication between each tag and detector. An ultrasound tag consists of a wireless device equipped with its power supply, paired with a specific object. These tags periodically transmit their identifiers using ultrasonic waves, which are then captured by the nearest ultrasound detector [13].

Three ILS based on Ultrasound signals are described in [9]. The main drawbacks of these solutions are related to the fact that ultrasound waves cannot penetrate walls and the system performance is deeply held back due to interference from reflected

signals propagated by other sources [7, 9]. Another downside is the fact that Ultrasound requires a large density of devices deployed as detectors, therefore increasing the cost of implementation considerably [13].

2.1.4 Ultra-wideband (UWB)

The UWB is a radio technology for short-range and high-bandwidth communication (>500 MHz). It sends ultra short-pulses (<1 ns, time period), in the frequency range from 3.1 to 10.6 GHz, using a very low duty cycle which results in reduced power consumption [14].

The main advantage of this technology is the ability to propagate through solid materials (concrete, glass, and wood), which makes it appropriate in typical indoor environments where LOS is limited. It is worth mentioning that metals and liquids can interfere with the UWB spectrum [14]. Moreover, this technology is immune to interference from other signals, due to its unique radio spectrum, as shown in Figure 2.2. Another advantage comes from the use of short-duration pulses. This allows the system to be less sensitive to multipath, therefore providing a more accurate estimation of the Time of Arrival (ToA)/Time of Flight (ToF). Errors up to 10 cm can be expected with the use of this technology [15].

The limitations sit on costs. Since UWB is a short-range technology, it requires a high density of nodes to work properly. Thus, costs increase rapidly with the growth of the area to be covered, making it an ineffective solution on limited budgets.

In [16] it is argued that the introduction of 5 GHz networking will diminish the impact of UWB since it has similar wideband capabilities and cm level error (Figure 2.2). As of the current moment, 5 GHz networks have been extensively implemented, making the adoption of a system reliant on UWB technology an additional financial undertaking.

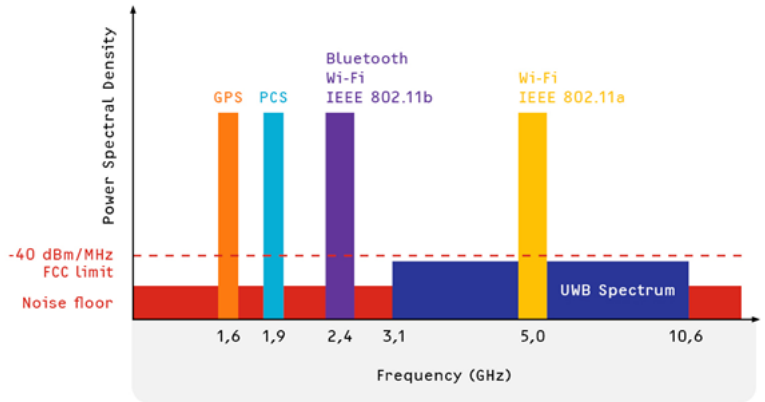

Figure 2.2: Radio spectrum comparison [17].

2.1.5 WiFi

The IEEE 802.11 standard, also known as WiFi Network or Wireless Local Area Networks (WLAN), has a unique aspect that sets it apart from the rest of the technologies - its availability. The technology is present in almost every device (mobile phones, laptops, tablets or other portable gadgets). Thus, developing an ILS using this communication standard, allows the developer to take advantage of an infrastructure already in place [14].

In many indoor environments, the network infrastructure includes Access Point (AP)s that have already been deployed for various purposes. Leveraging these APs as BSs for an ILS offers a cost-effective and efficient solution. These access points, originally intended for wireless communication, can be repurposed to support indoor positioning by utilising technologies like WiFi. This dual functionality not only reduces the additional hardware and infrastructure costs but also allows for seamless integration of indoor location capabilities without the need for extensive installation or disruption in the existing environment [18, 19].

The WiFi protocol operates across three distinct Industrial, Scientific and Medical (ISM) bands, namely 2.4 GHz, 5 GHz, and the relatively recent addition of 6 GHz [20]. The 2.4 GHz and 5 GHz bands were initially introduced with WiFi 2 in 1999, defined by the IEEE 802.11a standard. Despite its early introduction the 5 GHz band only gained relevance in later years when APs started to feature dual-band antennas, one for 2.4 GHz and another for 5 GHz. Nowadays, these deployments are the norm and with such an infrastructure in place, it is possible to work exclusively on this band and avoid interference with other protocols, Bluetooth and Zigbee (Figure 2.3), that operate in the 2.4 GHz [21, 22]. The 5 GHz also enables the attainment of centimetre-level precision [23], similar to what UWB technology achieves, in Subsection 2.1.4.

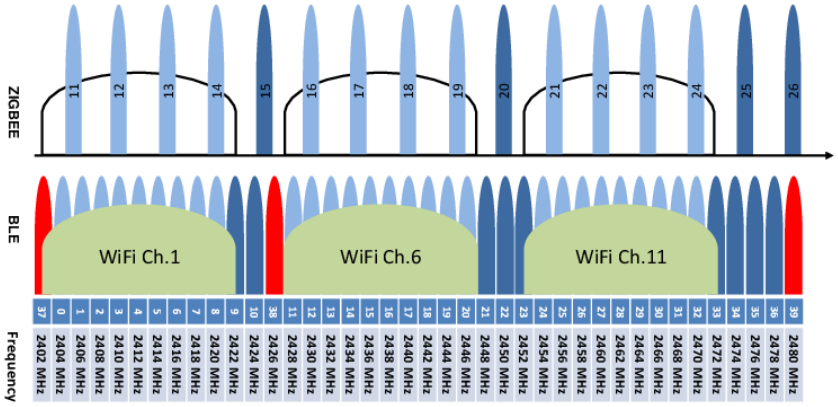

Figure 2.3: Spectrum of multiple wireless technologies in the 2.4 GHz band [24]: The colours indicate: light-green for WiFi bands; red for BLE advertising channels 37, 38 and 39; light-blue for BLE and ZigBee data channels which overlap with WiFi channels; and dark-blue for BLE and ZigBee data channels which do not overlap with WiFi channels (and are free from WiFi interference). [24].

WiFi presents several notable drawbacks, including the potential for device collisions in crowded environments and susceptibility to the multipath effect [25, 26]. Additionally, it's essential to recognise that WiFi was primarily designed for communication purposes rather than ILS, necessitating the development of innovative and efficient algorithms to enhance localisation accuracy while making use of the current infrastructure. Furthermore, WiFi technology tends to be power-hungry, consuming an average of approximately 216.71 mW, which can be sub-optimal for battery-operated devices [27]. Although WiFi 6 aims to improve this last problem, it is still relatively new (2021), and with little infrastructure in place, it is difficult to evaluate its performance [20]. Another challenge arises from the relatively slow scanning rate, as observed in smartphones, with intervals lasting between 3 to 4 seconds [25, 26, 27].

2.1.6 Bluetooth Low Energy (BLE)

BLE emerged in 2009 as a derivative of classic Bluetooth technology [28], introducing a streamlined and energy-efficient approach to wireless communication. While sharing foundational aspects with classic Bluetooth, BLE is tailored to address the specific needs of applications requiring low power consumption and periodic data transmissions. One key distinction is BLE's prevalence in the realm of Internet of Things (IoT) applications, where the demand for extended battery life and intermittent data exchange is paramount [29].

Operating in the 2.4 GHz ISM band, BLE utilises 40 channels with 2 MHz bandwidth ranging from 2402 MHz to 2480 MHz, Figure 2.3. The segmentation

into multiple channels allows for efficient utilisation of the spectrum and helps mitigate potential interference from other devices operating in the same frequency band [22, 24, 30].

The emphasis on low power consumption is a defining characteristic of BLE. The frequency-hopping mechanism not only contributes to interference resilience but also plays a vital role in reducing power consumption. By allowing devices to spend less time actively transmitting or receiving, BLE significantly extends the battery life of devices. This attribute makes BLE particularly well-suited for battery-powered devices, such as smartwatches, fitness trackers, and IoT devices that rely on efficient power management for sustained, long-term operation [28, 29, 31].

BLE is often used in ILS through the technique of RSS based fingerprinting [32]. This technique has some limitations that are mentioned in Subsection 2.2.5.

The introduction of Bluetooth 5.1, released in 2019, added a new field at the end of the packet structure called CTE (Figure 2.4). This structure can carry information such as angle calculation, enabling improving distance estimation based on AoA. Some manufacturers (*Texas Instruments*© [33] and *Nordic Semiconductor*© [34]) have already integrated this functionality. Easier deployment, smaller infrastructures and more robust scenarios when compared to raw RSS based methods can be expected [31, 35].

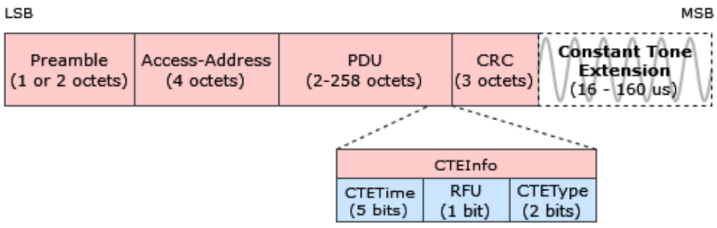

Figure 2.4: New Bluetooth 5.1 packet structure embracing the CTE appendices [31].

2.1.7 Discussion

During the study of technologies employed in ILS, the preceding analysis has scrutinised four primary methodologies: inertial, infrared, ultrasound, and radio frequency systems (Figure 2.1). Each of these approaches holds distinct advantages and limitations. The drawbacks of inertial systems in ILS applications include accumulated error over time, sensitivity to external forces, complex algorithm requirements, and limitations in providing absolute position information. While infrared and ultrasound systems exhibit accuracy in certain contexts, their susceptibility to LOS limitations poses challenges in scenarios where pervasive coverage and robust accuracy are paramount.

In contrast, radio frequency systems, characterised by their versatility and ability to penetrate obstacles, have emerged as a compelling alternative, offering the potential for more comprehensive and reliable ILS. Table 2.1 compares the radio frequency protocols presented earlier.

BLE often outperforms WiFi in ILS due to its lower power consumption, finer location granularity, and enhanced accuracy. BLE's ability to operate in low-power modes facilitates prolonged battery life in devices. Furthermore, the frequency-hoping mechanism implemented allows BLE devices to coexist and perform in the 2.4 GHz ISM band with other technologies without suffering from interference.

Moreover, the recent advancements in the BLE protocol, as detailed in Subsection 2.1.6, indicate a developmental trajectory aimed at enhancing its implementation and precision [36]. This evolution positions BLE as a cost-effective alternative to UWB technology.

In the initial section of Section 2.1, emphasis was placed on the prevailing approach in current solutions, where hybrid systems are pivotal in their development [3, 4, 5]. These solutions are finely tuned to the specific precision requirements of individual projects. Typically, the foundational layer of such solutions adopts a cost-effective alternative like BLE, which satisfies fundamental requirements. Subsequently, precision is augmented by integrating UWB capabilities, selectively applied in specific contexts. This strategy mirrors the approach employed in Apple's AirTags, which utilise BLE for general localisation and leverage UWB when users seek identification in close proximity contexts [37].

Due to the promising innovation in ILS capabilities, reduced costs, and ease of scalability, BLE was chosen in a previous study [31], utilising *Nordic Semiconductor*© hardware [38]. The technologies covered during this Section 2.1 serve as a contextualisation of the paradigms encountered when striving to implement an ILS, and their suitability depends on the specific situational requirements.

Table 2.1: Radio frequency technology comparison [7, 27, 6].

Technology	UWB	WiFi	BLE
Frequency band	3.1 - 10.6 GHz	2.4/5 GHz[1]	2.4 GHz
Channel bandwidth	>500 MHz	22 MHz	2 MHz
Techniques	ToA/ToF	ToA/ToF RSS AoA	ToA/ToF RSS AoA
Accuracy	10 cm	1 m - 5 m	0.9 m - 2 m
Power consumption[2]	Low TX: 750.1 mW RX: 750.1 mW	High TX: 722.7 mW RX: 709.5 mW	Low TX: 102.8 mW RX: 84.8 mW
Cost	High	Low	Low
Remarks	(1) Lack of infrastructure (2) No invasion of multipath (3) Difficult to scale	(1) Available infrastructure (2) Multipath susceptible (3) Easy to scale	(1) Available infrastructure (2) Upgrades aim to improve distance estimation (3) Easy to scale

[1] Depends on the WiFi version, as mentioned in Subsection 2.1.5;

[2] In [6], the study concludes that UWB consumes ± 7 times more energy than BLE when transmitting/receiving packages. On the other hand, if we take into consideration the bit rate and measure the mJ/Mb UWB is 10 times more efficient than BLE. When it comes to WiFi, although the transmission/reception numbers are very similar to UWB, it is worth noting that UWB had more than twice the bit rate than WiFi. This means that WiFi is not as efficient as UWB.

2.2 Localisation Techniques

Indoor environments represent a huge challenge to radio propagation. The multipath effects, frequent Non Line of Sight (NLOS) conditions, constantly changing environment, blocking objects and numerous reflective surfaces are just a few of the obstacles that contribute to a non-existent standard [39].

As shown in Figure 2.5, ILS can be employed with a variety of techniques. These techniques are mainly divided into three categories triangulation, scene analysis or proximity.

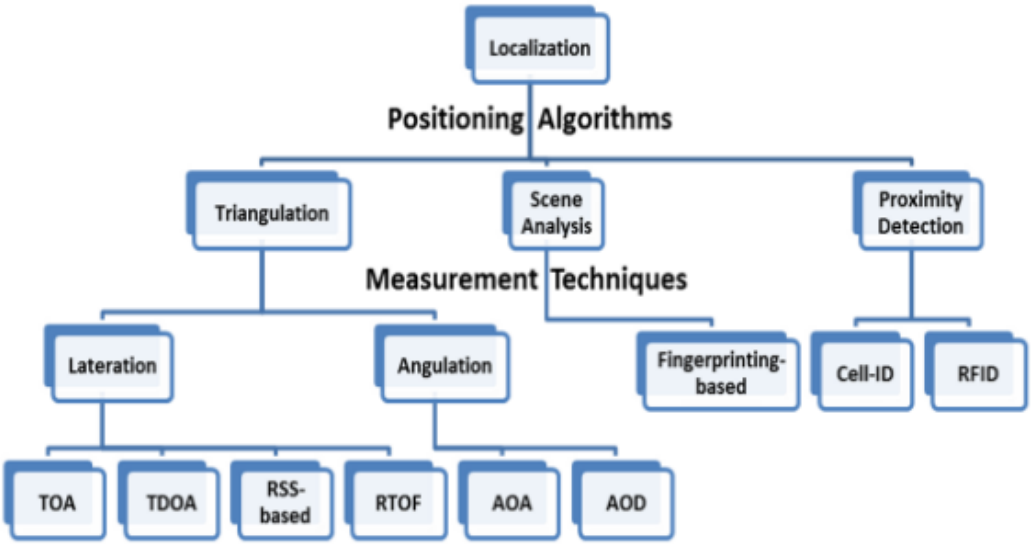

Figure 2.5: Classification of localisation methods [40].

2.2.1 Time of Arrival (ToA)

ToA measures the time that a signal takes to travel from the UD (transmitter) to the BS (receiver). The distance (d) is calculated indirectly by multiplying the measured time with the speed of light [41].

This calculated distance is then utilised to establish a circumference with a radius denoted as d, and the UD is expected to be positioned above it. Consequently, a minimum of three BSs is necessary for determining the exact location of an object, as illustrated in Figure 2.6 [40]. For the ToA technique to function effectively, it necessitates synchronised time references across all elements in the system, including both UDs and BSs. Additionally, as highlighted in [42], achieving this synchronisation can lead to an increased investment of development time and effort.

Time-based techniques are assumed to be working under LOS conditions [41]. When in NLOS conditions ToA an introduce inaccuracies in the measurement of the time taken for the signal to travel, leading to compromised distance estimations. In [43], RSS is fused with ToA in a hybrid-type technique to mitigate this problem. RSS data is leveraged to refine and adjust this distance estimation, therefore achieving more accurate and reliable distance estimations.

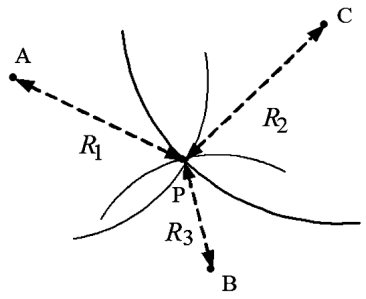

Figure 2.6: ToA Ranging [39].

2.2.2 Time Difference of Arrival (TDoA)

TDoA determines the position of a transmitter based on the time differences of signals received by multiple receivers. This difference in arrival times is then used to calculate the distance between the UD (transmitter) and each BS (receiver) [39].

TDoA utilises hyperbolic trilateration to compute the position of a transmitter based on the time differences of signals received by multiple receivers. The principle is rooted in the fact that the time difference between the arrivals of a signal at different receivers corresponds to a hyperbolic curve in the time-space domain, Equation 2.1 [41]. The intersection of these hyperbolic determine the transmitter's location [40], Figure 2.7.

In contrast with ToA, the synchronisation criteria only need to be addressed for the BSs. This alone guarantees a more robust system, less prone to errors [41].

$$c\Delta t_{ij} = [(x-x_i)^2+(y-y_i)^2+(z-z_i)^2]^{\frac{1}{2}} - [(x-x_j)^2+(y-y_j)^2+(z-z_j)^2]^{\frac{1}{2}} \quad (2.1)$$

where: (x, y, z) is the 3D position of the transmitter, (x_i, y_i, z_i) is the position of the i-th receiver, c is the speed of light, Δt_{ij} is the time difference between the i-th and j-th receivers.

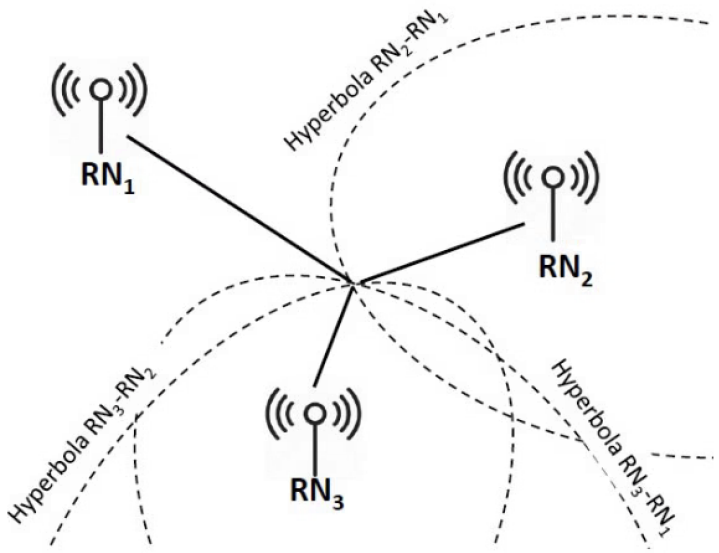

Figure 2.7: Positioning based on TDoA measurements [14].

2.2.3 Round-Trip Time of Flight (RToF)

Round-Trip Time of Flight (RToF) is a ranging technique that measures the total time taken for a signal to travel from a transmitter to a target and back to the receiver [14]. In RToF, the sender emits a signal, and the receiver measures the time it takes for the signal to travel to the target and for the reflected signal to return.

The distance between the UD and BS can be calculated using Equation 2.2 [14]. UD (transmitter) sends a message to a BS (receiver), the receiver then transmits a signal back to the transmitter, denoted as Δtr. Δd is the delay response time from the receiver. The outcome derived from this time difference signifies the duration of two trips (transmitter to receiver and receiver back to transmitter). Consequently, it should be divided by two to determine the time for a single trip and then multiplied by the speed of light.

RToF is commonly used in various ranging systems, including ultrasonic sensors, LIDAR, and some radar applications [44].

$$D_{ij} = \frac{\Delta tr - \Delta d}{2} \times c \tag{2.2}$$

2.2.4 Angle of Arrival (AoA)

The AoA approach measures the angle between the transmitted signal (UD) and the receiver (BS). This system features the antenna array on the receiver side so that by measuring the phase shift of the incoming signal, the receiver can determine the direction of the incoming signal [14, 45] (Figure 2.8).

The main advantage of AoA is its ability to distinguish between direct and reflected signals contributes to its effectiveness in mitigating multipath effects and enhancing the reliability of signal reception and transmission [46]. As it was mentioned in Subsection 2.1.7, the recent updates to the BLE protocol allow for the AoA to achieve better results than the typical RSS approach [35].

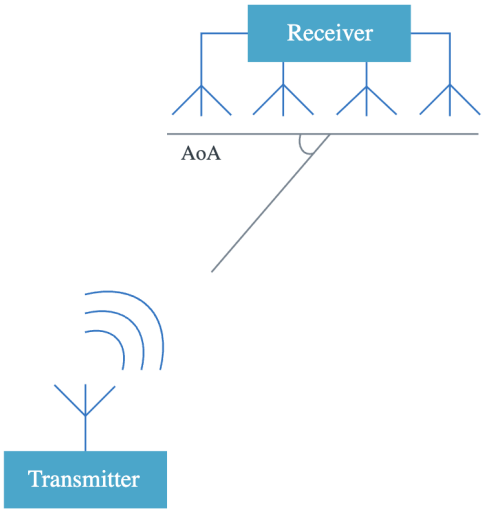

Figure 2.8: AoA method [45].

2.2.5 Received Signal Strength based Fingerprinting (RSS-based Fingerprinting)

RSS-based Fingerprinting is a localisation technique widely used in wireless communication systems to determine the location of a device based on the RSS from surrounding wireless APs or beacons [32]. The fundamental idea behind RSS-based fingerprinting is to create a database, or fingerprint map, of the signal strength at various locations in the environment. This database serves as a reference for estimating the location of a device when it is later detected based on its RSS.

In the process of building the fingerprint map, a device with known geographic coordinates collects RSS measurements from nearby access points or beacons at various locations. During localisation, when the device's location needs to be determined, it measures the RSS from the available APs. The system then compares these real-time RSS measurements with the pre-existing fingerprint map to identify the location that best matches the observed signal strengths [40].

In [47], the author addresses the main limitations when trying to build RSS-based Fingerprinting model, while highlighting the best practices for data acquisition. Some of the fundamental problems of this technique are determining the points to acquire beforehand, long setup periods (preparing the database) and ambiguity points. The last one is particularly important since the sample values stored in the database are only a fractional representation of the complete actual space. If the

system does not find the exact value for an object location, there is the possibility that more than one value falls within the threshold. Figure 2.9 demonstrates this limited representation.

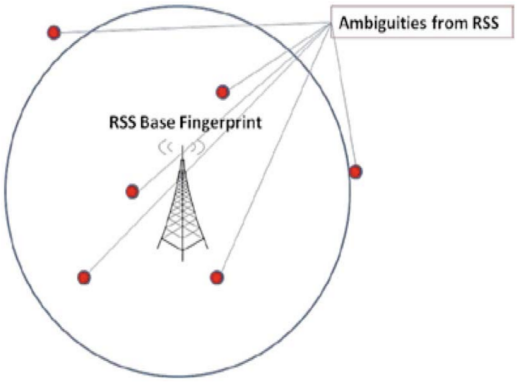

Figure 2.9: RSS-based Fingerprinting [40].

2.2.6 Proximity Detection

In proximity detection systems, the primary goal is to determine the relative closeness of a target asset to predefined reference points, rather than pinpointing the exact position of the asset. The system evaluates the RSS from multiple beacons and designates the asset's position as the location associated with the strongest RSS [7].

While it provides a straightforward solution for applications such as triggering actions based on proximity or context-aware services, it may not be suitable for use cases requiring high accuracy or detailed spatial information. The accuracy of the reported position is influenced by factors like signal strength variations, environmental interference, and the deployment density of reference points [48].

2.3 Beacon Placement

Although technologies and techniques used for ILS have been deeply studied and are currently well understood, an underlying problem, in the systems relying on BS/beacons to receive the advertised signal sent by any UD (asset), can still be identified. When building and optimising an ILS, the problem of beacon placement still imposes considerable challenges.

This issue is extremely relevant in terms of performance and overall costs. On one hand, if the beacon placement layout does not cover the entire floor plan, the system will not be capable of identifying the location of an object if it is positioned in a dead zone. On the other hand, when trying to achieve maximum coverage, system installers may over-provision indoor spaces with more beacons than required. As a

result, the deployment can be more expensive than it should be. For small deployments, the error margin could be small, or near-zero in some cases. However, as soon as the deployments start to expand to larger facilities such as airports, hospitals, and malls, an increase in the error margin and overall costs are expected. This is confirmed in [49], where a small experiment has proved that the error increases with the complexity of the layout.

Beacon placement is typically carried out by domain experts who follow a set of heuristics. The goal of this project is to automate this process by using an algorithm that receives as an input a floor plan and outputs the most efficient beacon placement layout, based on each floor plan's characteristics.

2.3.1 Constraint Satisfaction Problem

The reviewed literature suggests that the methodology in this area is quite varied, and no standard has been adopted yet. In [50], the authors use an image as an input for their experiment. As demonstrated in Figure 2.12, the first step consists of identifying the main materials that make up the floor plan and assigning them an attenuation value. This value differs between materials and will influence the signal propagation - the higher the attenuation value the greater the impact on the signal propagation, which reduces the beacon range. Then, two 2-dimensional arrays are created with the width and length of the map. The first one stores binary information about *go* (places where beacons can be deployed) and *no-go* (exterior spaces, walls and other obstacles) zones. The second one contains information on the attenuation factor for each material (Figure 2.10). Each pixel on the map will have two distinct values assigned to it - one for the zone type and another for the attenuation value.

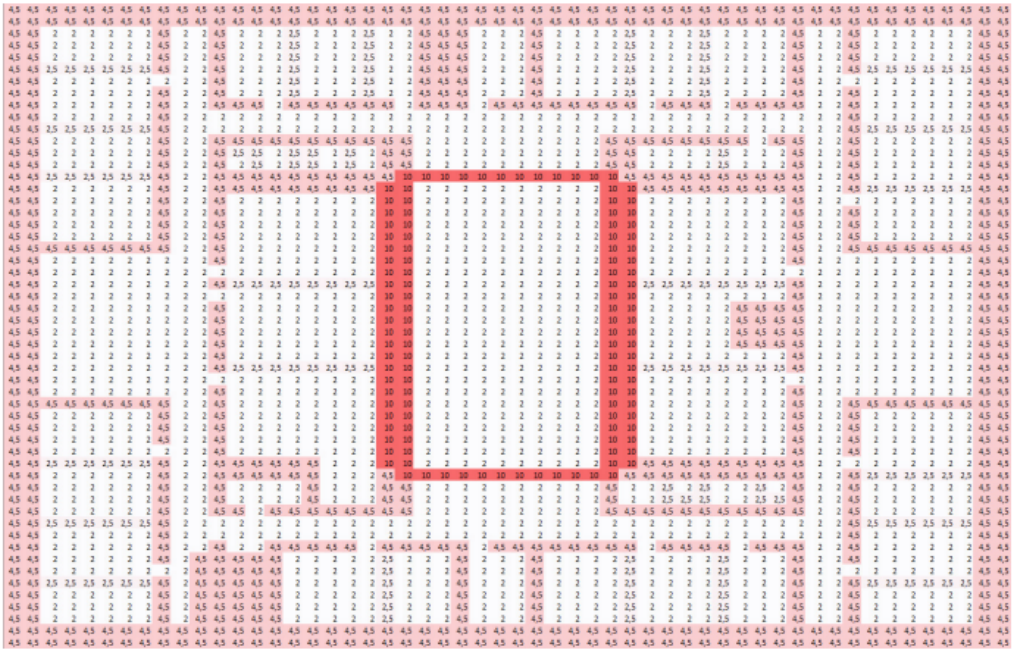

Figure 2.10: Floor plan represented with the different attenuation values [50].

To accelerate the process and use less computational resources, the image is down-scaled through the application of a sliding kernel. This kernel systematically moves across the image, identifying the maximum value within its defined region, Figure 2.11. This approach ensures that the algorithm consistently prioritises the highest attenuation value in the matrix, thereby guaranteeing performance in scenarios with sub-optimal conditions and maintaining robustness in more favourable ones.

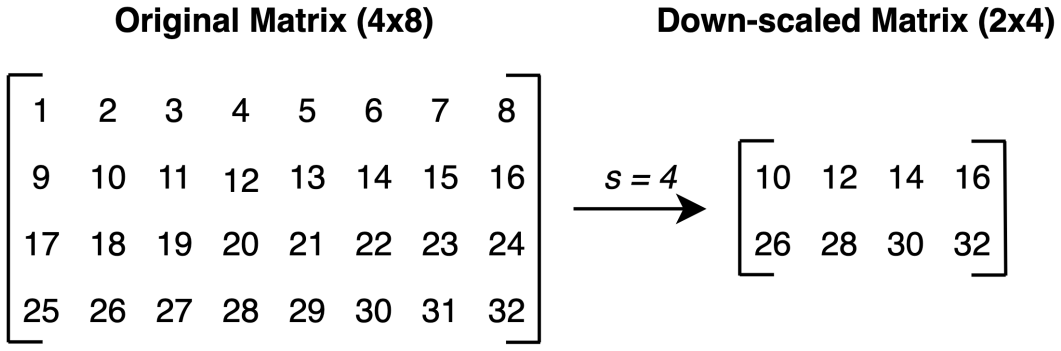

Figure 2.11: Representation of a down-scale in a matrix using a sliding kernel of 4×4, $s = 4$; s represents the down-scale factor $s = 1, 2, 3....$

Another important aspect to take into consideration is the beacon range. For this experiment, the range is calculated based on RSS, attenuation values, and thickness of the objects (walls). The process to calculate the position of the beacon is then

carried out by the Constraint Satisfaction Problem (CSP) algorithm. A CSP is a computational problem in which we have to find a solution that satisfies a set of constraints. The variables in the experiment are the location of the beacons, and the domain of each variable is the set of possible locations within the map. The constraints that are considered are the beacon range, each point being covered by three beacons, as well as a minimum distance of three metres between beacons. To optimise the solution, an additional constraint is introduced to minimise the number of beacons used. Figure 2.12 shows the final result.

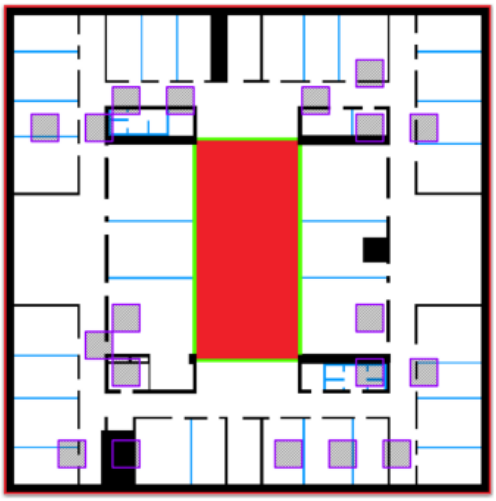

Figure 2.12: A graphic representation of the map [50]: Massive walls (Black), drywall (Blue), triple glazed windows (Green), free space covered by BSs (White), outdoor area no BSs required (Red), Beacons feasible location (Purple).

While studying the floor plan materials can lead to a more real testing scenario, this method can also bring other complexities. Firstly, the user is in charge of finding these materials and colouring the pixels accordingly. Additionally, the collected information cannot be interchangeable with other projects, so this initial setup has to be repeated for different floor plans. Another major downside is the difficulty in identifying every material on the floor plan since the output result will always lack detail.

2.3.2 Unique Location

The Unique Location (UL) of an asset typically refers to the precise and individualised geographical position of a specific object or device within a given environment. In the context ILS, this geographical position can be achieved using a variety of techniques as explained in Section 2.2. When employing the ToA technique in ILS the asset's locations are the computed result from the intersection of three BSs (multi-lateration). Three BSs is the minimum to guarantee a UL, Figure 2.13.

Both [51, 49] show that, in some scenarios, the location of an object can also be found using only two beacons. When a device is at a location p within the LOS of at least three or more beacons, using trilateration, it is possible to estimate the device location fairly accurately, as shown in Figure 2.13c. If it is only possible to have LOS with only one beacon, it is impossible to make any estimation as shown in Figure 2.13a.

The main problem lies in the two-beacon scenario, as the intersection of the circumferences will result in two symmetrical intersection points one mirror from the other. If no floor plan is considered, both points are plausible location choices, as evidenced in 2.13b.

Introducing the floor plan into the equation and taking into consideration its boundaries the results are analysed differently. In the scenarios presented in Figure 2.14, the two beacons' approach can provide sufficient coverage. For 2.14a, one of the intersection points can be left out, as its location does not belong to the floor plan. In Figure 2.14b, due to p_2 not being visible to b_2 this possibility is also excluded. The last case represented in Figure 2.14c, p_2 is visible to a third beacon, b_3, in this case, we are only interested in the points that only have LOS conditions for b_1 and b_2, therefore p_2 is excluded. In both Figures 2.14b and 2.14c, the deciding factor is the number of beacons visible to each point. There is also a fourth scenario not represented in any of the Figures 2.14, which is when the intersection of two beacons gives exactly one point $p_1 = p_2$. Although it does not happen as often, the result is straightforward and does not need any workarounds.

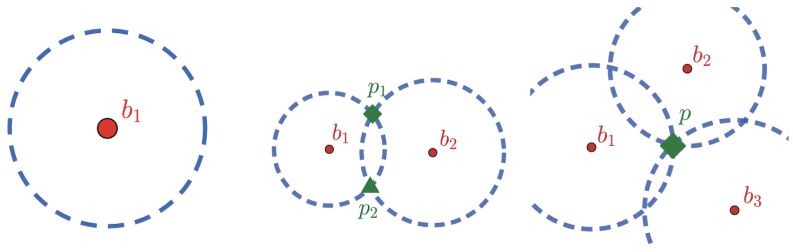

(a) (b) (c)

Figure 2.13: Localisation scenarios without floor plan: (a) Infinite number of solutions with a single beacon; (b) Two solutions for two beacons; (c) A unique solution with three beacons [49].

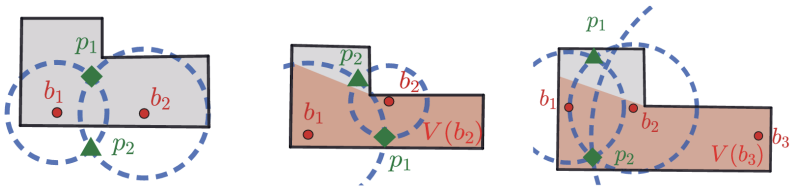

(a) (b) (c)

Figure 2.14: Localisation scenarios with floor plan: (a) p_2 is outside the floor plan; (b) p_2 is not visible to b_2; (c) p_2 is visible to a third beacon b_3 [49].

To test the UL of points, both [51, 49] established the following set of definitions and equation 2.3.

Definition 1 *Let P be the representation of the floor plan and B the set of candidate beacon locations, which guarantees a unique localisation, $B = \{b_i | 1 \leq i \leq n, b_i \in P\}$. The goal is to find the minimum set of beacons $D \subseteq B$ that can uniquely localise all the points in P.*

Definition 2 (Visibility) *Two points $p, q \in P$ are visible to each other if and only if the line segment $\overline{\text{pq}}$ is strictly inside P, i.e., does not intersect any point on the boundary of P.*

Definition 3 (Visible Region) *For a beacon bi, the set of points seen by b_i, is the visible region of b_i, denoted as $V(b_i)$.*

Definition 4 (Visible Beacon Set) *For a point $p \in P$, the visible beacon set of p, $V(p)$, is the set of candidate beacon locations that can see p. To represent the subset of beacons in the set D that are visible at p, we use $V_D(p) = b_i \in D | b_i canseep$.*

Definition 5 (Unique Localisation of a Point p) *Given a point p and a set of visible beacons $V_D(p)$, we say p can be uniquely localised, i.e., $UL_D(p) = 1$, if there is only one location consistent with the range measurements and the visibility information.*

$$UL_D(p_1) = \begin{cases} 1, |V_D(p_1)| \geq 3 \\ 1, |V_D(p_1)| = 2, p_1 = p_2 \\ 1, |V_D(p_1)| = 2, V_D(p_1) \neq V_D(p_2) \\ 0, |V_D(p_1)| = 2, V_D(p_1) = V_D(p_2), p_1 \neq p_2 \\ 0, |V_D(p_1)| \leq 1 \end{cases} \tag{2.3}$$

The work described in [49] provides also tests of two algorithms: the Greedy and Random Sampling algorithms. These two algorithms are based on the algorithms

presented in [51], with minor modifications aimed to improve performance. The process of comparing algorithms consists of evaluating their performance in various floor plan layouts. Both studies use the BLE technology for their beacons.

Before describing both algorithms, it is important to clarify some concepts. In [51, 49], the beacon placement is performed in a 2D representation, although in reality, they are going to be deployed in a 3D environment. Both studies claim that this problem can be mitigated by deploying the BSs close to the ceiling level and holding the UD at a regular height (at least 1 m from the ground). Nonetheless, having the UD 1 m from the ground is always a limitation due to the abundance of obstacles higher than this height. The ceiling mount option, allows the UD to have a higher chance of a clearer LOS with the BSs and therefore avoid being blocked by objects such as chairs or desks. Another point worth mentioning is that the algorithms are built to assure LOS conditions across the entire floor plan. Nonetheless, due to the reflections or signals penetration through walls, NLOS measurements will be registered. To avoid computing those values, a filter is applied to the data to eliminate their impact.

2.3.3 Greedy algorithm

As defined in [52], a greedy algorithm is an algorithm that makes a sequence of choices based on the best solution available at that time. Their main characteristics are simplicity, straightforwardness and efficiency. Although greedy algorithms can find an optimal solution, their output is usually a solution close to optimal.

In the exercise represented in Figure 2.15, the goal is to find the path for the largest sum. By following a greedy approach the path will be 7+12+6, which adds up to a total of 25. Looking at Figure 2.15, the optimal solution would be $7+3+99 = 109 > 25$. Due to its nature, the algorithm cannot find the best solution, and its decisions are based on the largest available number at each step.

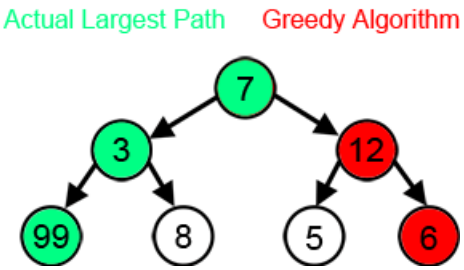

Figure 2.15: Greedy exercise [53]: Find the path for the largest sum.

Previous work into using greedy algorithms for beacon placement [51] aimed to divide the floor plan and then place beacons in the largest regions first. Only after this process would the layout be optimised. In [49], it is highlighted that this method could be sub-optimal as illustrated in Figure 2.16.

Figure 2.16 shows a rectangular-shaped schematic with spikes in the top and bottom parts. Beacons b_10 and b_11 can cover all the top spikes, whereas b_12 and b_13 cover the bottom ones. A greedy algorithm selects the red beacons by order of their index instead of these four blue beacons because once the rectangle is uniquely localised, the algorithm then tries to cover each tip that is seen once. When the number of tips is infinite, the greedy algorithm places infinite beacons while the optimal solution is to place the beacon b_1, b_2 and the four blue beacons (b_10, b_11, b_12 and b_13).

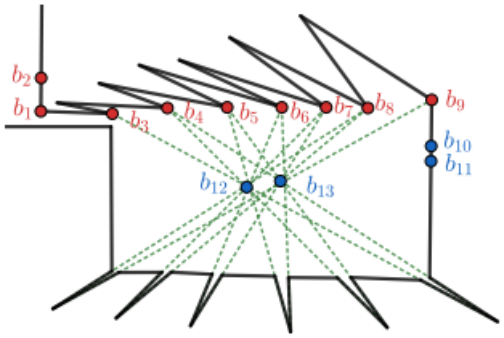

Figure 2.16: Sub-optimal solution using a greedy algorithm [49]: The algorithm selects the beacons in red resulting in a sub-optimal solution rather than choosing beacons b_1, b_2 and the four blue beacons.

2.3.4 Random Sampling algorithm

The Random Sampling (RS) algorithm can be better understood by taking into account the example in Figure 2.17. The initialisation process requires a set of beacons to be randomly generated in a given area. Initially, all generated beacons carry the weight of one (Figure 2.17a). As explained in Subsection 2.3.2, it is impossible to locate any point with just a single beacon. Therefore, the algorithm starts by selecting two candidate beacons. The higher the weight of a beacon, the higher its probability of being selected. The next step aims to find a point within the floor plan that is not uniquely localised. If no point is found, the algorithm terminates and returns the number of beacons and respective locations. Otherwise, it doubles the weight of the beacons in LOS with the point, as demonstrated in Figure 2.17b. The algorithm will then repeat the process of searching for not uniquely localised points and double the weight of the beacons seen by it. If the algorithm finds that the two beacons are not enough to localise the floor plan uniquely, then it will increase the number of beacons, reset the weight of the beacons, and restart the process of finding points until the coverage is maximized, Figure 2.17c. In [49], the process is explained in greater detail.

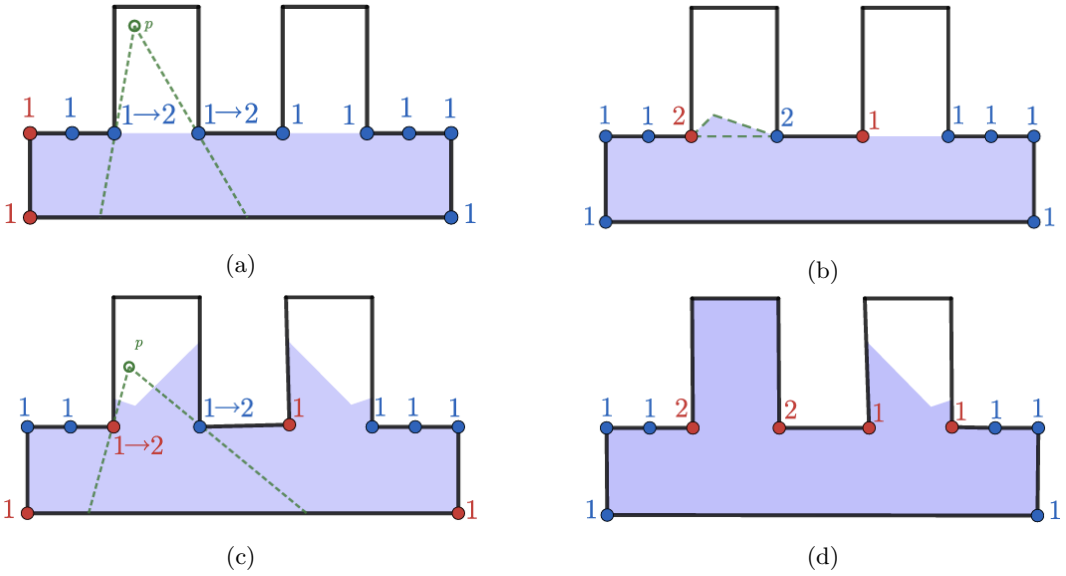

Figure 2.17: Random Sampling algorithm phase description: (a) Initialisation; (b) Cannot obtain a solution with 2 beacons; (c) Initialisation; (d) Feasible solution [49].

As shown in the performance analyses made in [49], the average number of beacons placed by the RS algorithm is lower than the Greedy one, especially with more complex floor plans. The main reason for this difference is that the greedy approach is built to optimise coverage by placing beacons sequentially without any feedback on past decisions. On the opposite side, the RS method by updating the beacon weights and restarting its process selection from the beginning can make more informed decisions, leading to fewer beacons.

Chapter 3

Supporting Technologies

The architectural blueprint's significance in the decision-making process for computer vision algorithms, as highlighted in Section 2.3, cannot be overstated. It serves as a foundational reference point, crucial for the analysis of key elements such as construction materials and the precise delineation of the schematic's boundaries. These aspects heavily influence algorithm performance and accuracy. Our research underscores the critical importance of the building schematic, which provides essential data for our analysis.

Open Source Computer Vision Library (OpenCV) [54] serves as a fundamental tool in the field of computer vision and image processing. This versatile framework offers a wide array of functions and algorithms, making it essential for various applications such as object detection and image recognition.

3.1 Image Conversion

The initial setup when working with image processing typically involves converting an image to binary format. This conversion serves two main purposes. Firstly, binary images are more memory-efficient as they use only one colour channel (pixels are either black or white) instead of the three in Red, Green and Blue (RGB) images. Secondly, many functions within the OpenCV library rely on binary images as input, making this conversion necessary. However, there are cases where coloured images are valuable, such as in the study by [50], which uses pixel colour to identify specific materials.

To achieve this, the image has first to be loaded by the algorithm using its source path and then proceed to transform the original image into a grayscale version, which reduces the image to shades of grey, allowing for efficient processing and analysis of image features without the complexity of colour information.

To obtain the binary image from the grayscale version it is necessary to apply a threshold. A threshold is a specific value or a range of values used to distinguish or separate pixels in an image into different categories, typically distinguishing between objects of interest (foreground) and the background. This categorisation is based on the pixel's intensity.

In OpenCV there are two types of binary thresholds, binary or inverse binary. The first one sets a pixel to 0 (black) if its value is smaller than the threshold, otherwise, it is set as 1 (white). The second one is the exact opposite of the first, setting a pixel to 0 (black) if this is above the threshold and to 1 (white) if it is below. In Figure 3.1 there is a representation of both techniques when applied to the same source image.

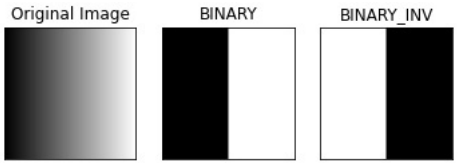

Figure 3.1: Binary threshold [54]

The threshold value is defined through Otsu's thresholding. This method is a valuable choice in image processing as it automatically calculates an optimal global threshold based on the image's colour histogram (Figure 3.2). This approach is advantageous because it maximises the separation between foreground and background, making it highly adaptable to various image inputs. Its automated nature also minimises the need for manual threshold adjustments, reducing the potential for errors in the threshold process.

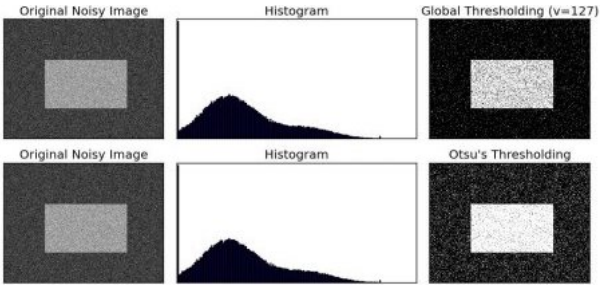

Figure 3.2: Global vs. Otsu's thresholding [54]: The first image the threshold value was defined by the user (127), the second one had its threshold determined by Otsu's method.

3.2 Morphological Transformations

Morphological transformations are multipurpose operations that allow to performance of tasks such as noise removal, image segmentation, orientation determination, or extracting size characteristics from objects within an image. This versatility was notably evident in Subsection 2.3.1, where morphological transformations were employed to gather data regarding wall thickness variations (size), identify the precise location of doors and windows (image segmentation), and distinguish between concave and convex walls (orientation).

These transformations constitute fundamental operations in image processing, especially in the context of binary or grayscale images. They encompass operations like dilation, erosion, opening, and closing, which are performed by applying a small, user-defined matrix known as a kernel to the input image. The kernel serves as a structural element that dictates the extent and shape of the transformation. By convolving this kernel with the image, morphological operations effectively eliminate noise, isolate objects of interest, and modify the image's structure.

OpenCV offers three kernel types (Figure 3.3): rectangular, elliptical and cross-shaped. These kernels provide flexibility in structuring the scope and shape of the transformations applied to images. Figure 3.4 demonstrates the different outputs when applying the different kernels to the same input image.

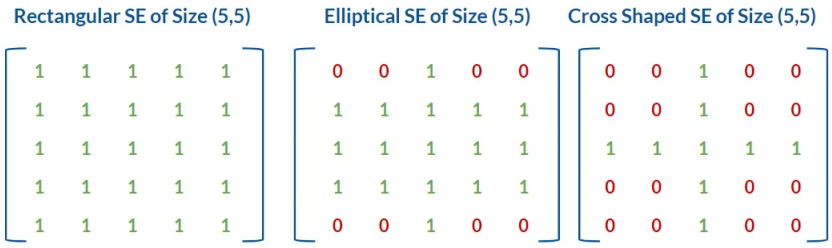

Figure 3.3: Kernel Shapes [55]: Representation of the $(5 * 5)$ matrices generated for each kernel shape.

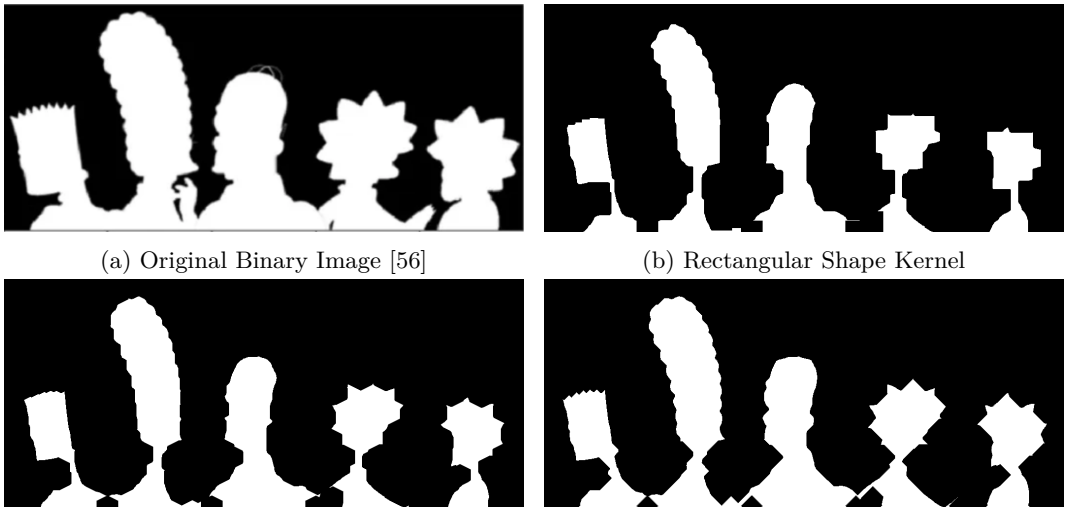

(a) Original Binary Image [56] (b) Rectangular Shape Kernel

(c) Elliptical Shape Kernel (d) Cross Shape Kernel

Figure 3.4: Different Kernels applied to the same image: The results were obtained from erosion using a (5×5) kernel and 10 iterations applied to (a).

The kernel is a crucial element as it is required in any type of transformation function. The two basic transformation functions are erosion and dilation. The erosion follows the same principle as in nature, eroding the boundaries of the foreground object. The defined kernel slides through the image. A pixel in the original image (either 1 or 0) will be considered 1 (white) only if all the pixels under the kernel are 1, otherwise, it is eroded (set to 0, black), Figure 3.5.

(a) (b)

Figure 3.5: Eroded object with a (5×5) rectangular kernel [57]: (a) Test image; (b) Eroded image.

In dilation, the complete opposite happens. A pixel is 1 (white) if at least one pixel below the kernel is 1, therefore increasing the foreground object (white area), Figure 3.6.

(a) (b)

Figure 3.6: Dilated object with a (5×5) rectangular kernel [57]: (a) Test image; (b) Dilated image.

There are more complex transformations, based on those base transformations mentioned previously. For instance, a closing is a dilation followed by an erosion. It is useful for closing small holes inside the foreground objects, or small black points on the object, like in Figure 3.7a. On the other hand, an opening is an erosion followed by a dilation. Its main use is for reducing noise in the background, as shown in Figure 3.7b.

(a) (b)

Figure 3.7: Complex morphological transformations [57]: (a) Closing; (b) Opening.

To fine-tune the impact of the morphological transformations applied, it is possible to vary the number of iterations. The number of iterations refers to the repetition of a specific operation on an image. Increasing the number of iterations can lead to a more pronounced effect, such as further thinning or expansion of image features. This flexibility is essential for customising morphological operations to specific image processing tasks and adapting to different characteristics within images. Figure 3.8 shows the result that an increase in iterations has on a closing transformation with a (2,2) rectangular kernel.

(a) (b) (c)

Figure 3.8: Iterations effect on a Closing transformation with a (2,2) rectangular kernel: (a) Original image (binary) [54]; (b) Closing with 2 iterations; (c) Closing with 10 iterations.

3.3 Contours

Contours are a fundamental concept in image processing and computer vision. They represent the boundaries of objects within an image and are essential for various tasks, including object detection, shape analysis, and image segmentation. Contours provide a way to identify and extract the structural information of objects, making them a valuable tool for analysing and manipulating images.

In the context of this work and its research objectives, the emphasis lies on image segmentation, primarily to distinguish between indoor and outdoor areas, however, it could also be used for rooms, and obstacle recognition among others. Establishing a clear demarcation between these environments is critical, as it enables the utilisation of the location method discussed earlier in Subsection 2.3.2.

In OpenCV, the hierarchy of contours is a mechanism used to establish relationships between contours in a given image. The hierarchy describes the nesting and connectivity of contours, providing information about how one contour is positioned in relation to others. This hierarchical information is organised as a tree structure, where each contour is represented as a node, and relationships are defined through parent-child connections.

The hierarchy information for each contour is typically represented as a four-element array, which includes the indices of the next, previous, first child, and parent contours. These elements are used to navigate the contour hierarchy.

- “Next” refers to the index of the next contour on the same hierarchical level. In Figure 3.9, contours 0, 1, and 2 are at the same hierarchy level. Contour 0 is linked to contour 1 (“Next” = 1), and contour 1 is linked to contour 2 (“Next” = 2). Since no other contours are on this level, contour 2 is noted as “Next” = -1.

- “Previous” indicates the index of the preceding contour at the same level.

- "First child" points to the first child contour nested within the current contour. In Figure 3.9, contour 2a is a child of contour 2, denoted as "First child" = 2a for contour 2. Similarly, for contour 3a with two children, it selects the first child, 4, with "First child" = 4.
- "Parent" specifies the index of the parent contour that encloses the current contour.

By examining these hierarchy relationships, it becomes possible to understand how contours are organised in the image. This is particularly useful in tasks like identifying holes within objects or distinguishing between external and internal contours.

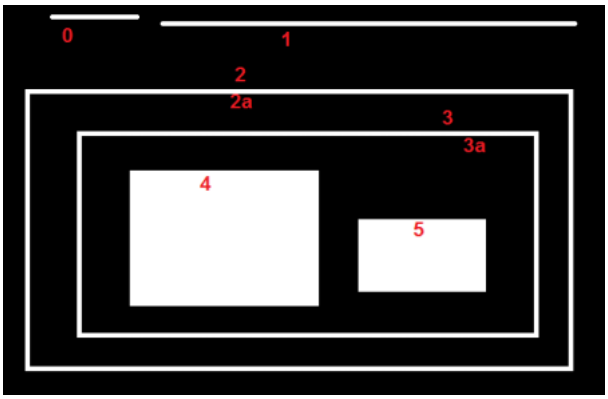

Figure 3.9: Hierarchy explanation [54]: The numbers marked in red are the indexes attributed to the contours.

There are four retrieval modes for extracting contours from binary images. The "list" mode retrieves all contours in a flat, unstructured list. "External" mode identifies only the outer contours, useful when you need primary object outlines, Figure 3.10. "Ccomp" mode categorises contours into external (level 1) and internal (level 2) ones, creating a two-level hierarchy. In Figure 3.9 contours 0, 1, 2, 3, 4 and 5 are level 1 type, while 2a and 3a are level 2a. "Tree" mode provides a detailed hierarchical structure with parent-child relationships, ideal for complex tasks like object tracking and region segmentation. Figure 3.11 shows the classification of contours using the "Tree" mode.

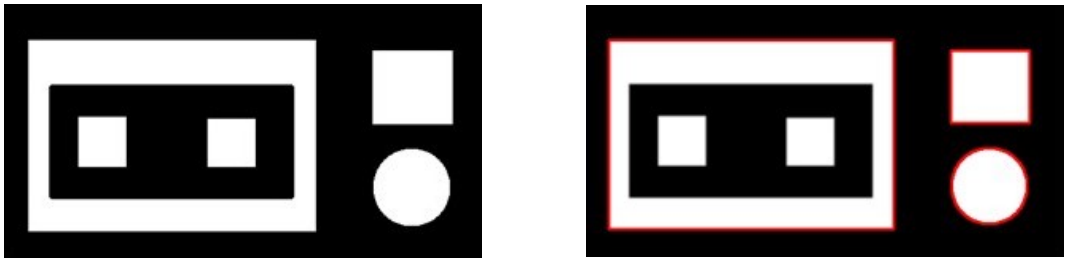

(a) (b)

Figure 3.10: Obtaining the external contour from an object [58]: (a) Threshold image; (b) Threshold image with external contours in red.

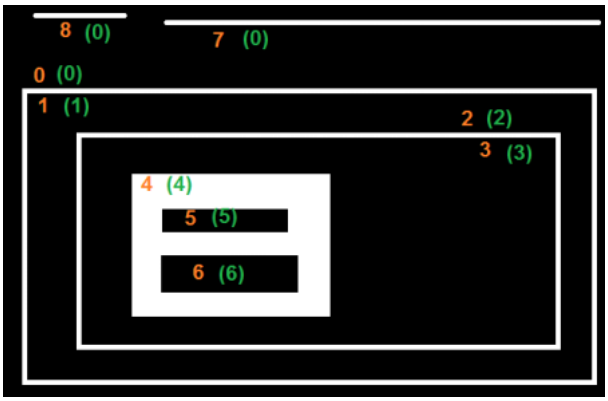

Figure 3.11: Tree Hierarchy [54]: The numbers in orange are the indexes attributed to the contours and in green it is shown its hierarchy level.

Contours are essentially a collection of connected points that encircle the object's edges, providing a structured representation of its shape and spatial characteristics. When retrieving contours it is possible to use approximation methods to simplify complex contours by reducing the number of vertices while retaining essential shape characteristics. This simplification enhances memory efficiency, reduces computational requirements, and streamlines subsequent image-processing tasks, making them crucial for various applications in image analysis and computer vision. Figure 3.12 shows the difference between a contour obtained with and without approximation.

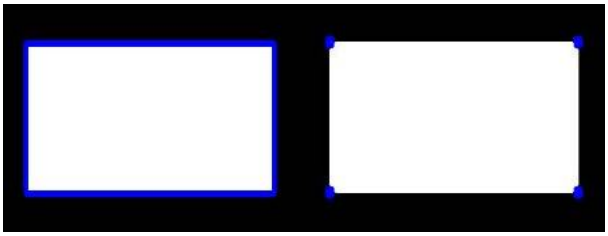

Figure 3.12: No approximation vs. simple approximation [54].

Chapter 4

Methodology

In this chapter, the goal is to present and explain the work developed, starting with a summarised overview of the system and then detailing its specifics along Sections 4.1 and 4.2. Furthermore, various routes that were not included in the final solution, but that played a pivotal role in shaping the outcome, are also described.

The purpose of this work is to develop an algorithm which can output a layout for the BSs to perform multilateration in a floor schematic. This layout will be tested against the previous layout, attained through basic rules for beacon placement provided by experts [59]. The main requirement is for the layout to optimise efficiency by employing the minimum number of BSs necessary.

To achieve this, the project was divided into two separate phases: the initial setup and the algorithm itself. Although the algorithm is the focal point of this project it is important to clarify this initial phase, since it incorporates important procedures and decisions that are relevant later.

The system input starts from a raw floor plan with walls and free-space information about the building. During the initial setup, Section 4.1, the aim is not only to work on this schematic, by separating the different materials and calculating the image scale, but also to acquire data concerning the hardware used for the BSs, such as antenna gains, radio transmitted and received power among others. Section 4.2, which is the beacon placement algorithm, focuses on using this information and finding which layout best suits the input schematic. By the end of its execution, the algorithm should be able to display the resulting BS layout and respective coverage

in the form of a heat map. Figure 4.1, shows a summarised pipeline of the project and an output table containing the coordinates for the BS positions in the layout.

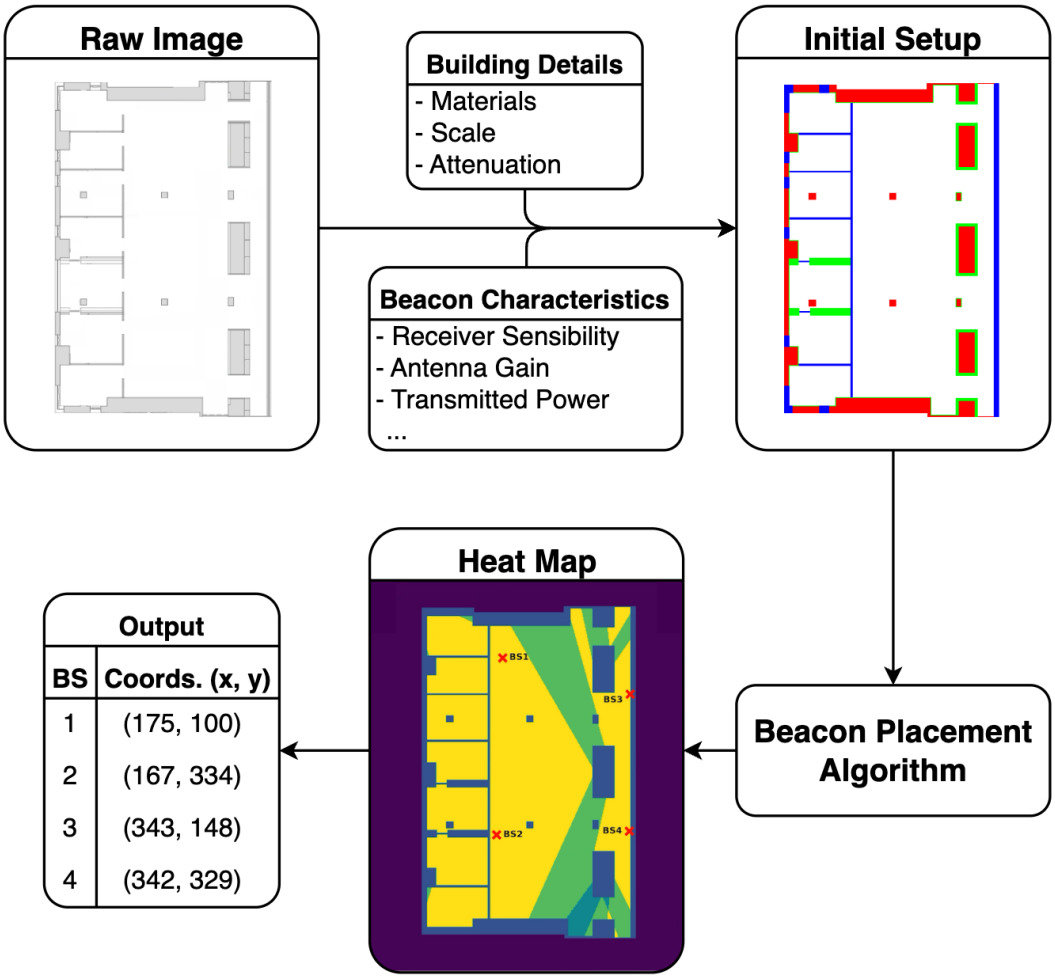

Figure 4.1: Pipeline of the project: Raw Image - Original Schematic from the building with information about the walls and free-space; Initial Setup - Modified Floor plan with building details and beacon characteristics; Beacon Placement Algorithm - Calculates the best layout for the schematic; Heat Map - Output of the Beacon Placement Algorithm, coverage heat map from the resulting layout with the BSs plotted in red. The output table shows the coordinates of each BS in the image.

4.1 Initial Setup

When setting up an indoor location system, it is crucial to consider material and infrastructure attenuation, algorithm schematic scale, and the acquisition of hardware characteristics for Bluetooth beacons. These factors are vital for the system's accuracy and performance. In this section, the goal is to explain how these variables were acquired and present their use case in the algorithm.

4.1.1 Materials

Materials play a pivotal role in this algorithm, as their different characteristics can create different constraints. In [50], the materials found in the floor plan serve as limiting factors for the range of a BS. A similar strategy is adopted in this work, as by using the pixel's colour the algorithm will be able to distinguish between different materials.

To apply this concept it is first necessary to identify the main materials that make up the building, as well as their location. In an attempt to simplify this initial step, only three materials were considered - concrete (red), glass (blue), and drywall (green), as presented in Figure 4.2b. As mentioned in the final stages of subsection 2.3.1, the output result will always be a simplified version of the environment, since it is extremely difficult to replicate the real-world environment. Figure 4.2 shows both the raw schematic and the floor plan with the materials. It is worth mentioning that the closer the schematic is to the real world, the better the output will be.

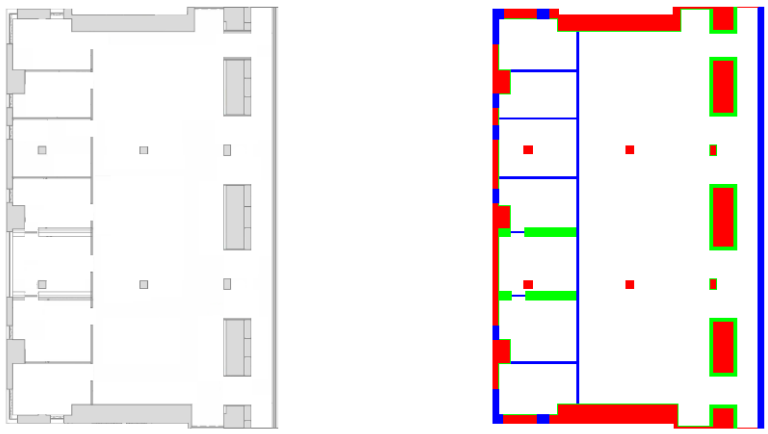

(a) (b)

Figure 4.2: Floor plan with materials: (a) Raw floor plan; (b) Floor plan with materials.

Notice that, in Figure 4.2b, the zones where doors (wall gaps) used to be (Figure 4.2a) are now closed as in the real world, having a door closed or open can influence signal propagation. Thus the goal of this approach is to mitigate this problem and compute the result considering the worst-case scenario. This way, it is expected that the final result will be able to cope with both situations.

4.1.2 Scale

Another important feature to take into consideration is the image scale. To find the actual scale it was necessary to measure the distance between two points in the real

environment and also count the number of pixels between those two points in the schematic. Then, by using Equation 4.1 it is possible to calculate the actual scale in pixels.

$$\frac{1}{Distance(px)} = \frac{x}{Distance(m)} \tag{4.1}$$

Due to modifications to the building infrastructure after its construction, the floor plan in disposal was not updated. As a result, some objects' representation in the schematic did not correspond with reality. For example, some pillars were larger in reality due to the added material from the restoration. To avoid miscalculating the scale, it was necessary to obtain a measurement in an unaltered area. The two central pillars of the schematic were not altered during restoration, so this was the place selected to obtain the real-world measurement, Figure 4.3.

This was achieved by positioning the laser, Figure 4.3a, on the side of one of the pillars and pointing it directly to the other pillar. The returned distance was 7.67 m, and the homologous measurement in the image resulted in 131 pixels, Figure 4.3b. Substituting these values in the above Equation 4.1, results in a scale of 0.05856 meters per pixel.

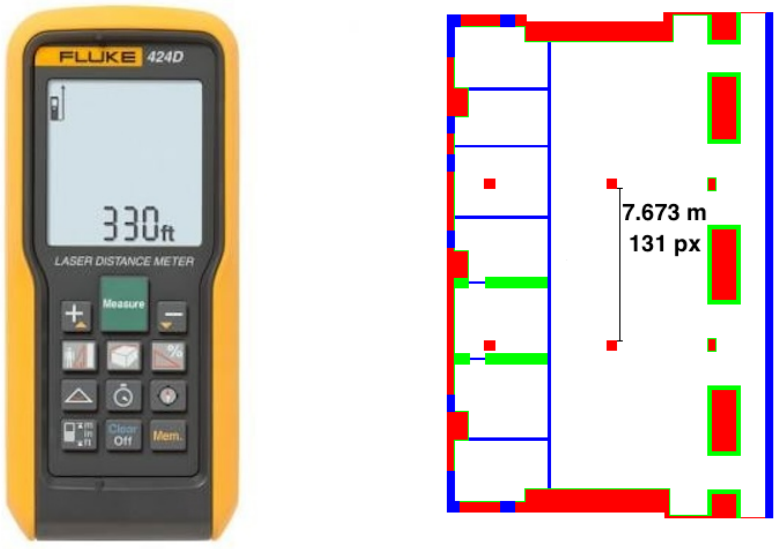

(a) (b)

Figure 4.3: Laser and Measurements: (a) Laser used for measurements [60][61]; (b) Calculated values in meters and pixels.

4.1.3 Hardware Characteristics

The purpose of the colours in the floor plan is to consider the attenuation of the signal through an object and verify if the asset can be reached by a BS. To achieve this, it is necessary to understand the physics behind this attenuation. However, such a

topic requires quite extensive research, which is far from the scope of this project. Therefore, the method used follows a rather simpler approach called link budget. The link budget is the accounting of signal gains and losses travelling through a medium from a transmitter to a receiver [62]. The link budget Equation 4.2, contains five variables: Transmitted Power (P_{TX}) (dBm), Receiver Sensibility (P_{RX}) (dBm), Transmitter Antenna Gain (G_{TX}) (dB), Receiver Antenna Gain (G_{RX}) (dB) and Losses (L) (dB), [63]. The L aggregate both coupling losses between different parts of the system and losses due to absorption/attenuation of the transmission medium. The goal is for the P_{RX} to be higher than the minimum receiver sensibility (-96 dBm).

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L \tag{4.2}$$

This equation is the mechanism used by the system to decide whether a BS can reach an asset or not. Apart from the L, all the values are found in the datasheet of the hardware manufacturer. In this case, Nordic nRF52833 DK (Devkit) was used (Figure 4.4a) for the radio and a standard monopole 2.4 GH_z antenna (Figure 4.4b). According to the radio datasheet [38], the P_{TX} and P_{RX} values are 8 dBm and -96 dBm, respectively. For G_{TX} and G_{RX}, it was considered -4.4 dBi as this is the average value advertised by the datasheet [64].

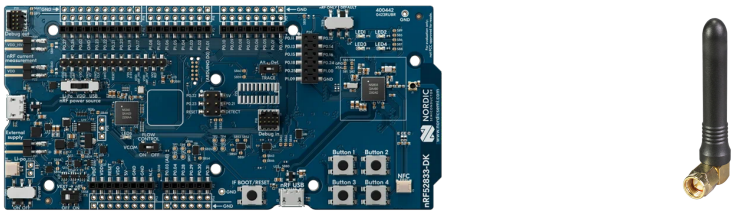

(a) (b)

Figure 4.4: Nordic Radio and Monopole Antenna: (a) nRF52833 [38]; (b) 2.4 GH_z Monopole Antenna [64].

In Equation 4.2, L corresponds to the sum of all the propagation losses in a defined path. The values used for the material losses can be acquired through new measurements or repurposed from another project. Although the first option would offer a more representative scenario from the schematic used, this process can be time-consuming and complex to obtain. Therefore it was opted to use information from previous works. The studies described in [65], [66], and [67] obtained propagation value losses in the 2.4 GH_z band for different materials found in buildings. It is important to highlight that attenuation is not a value that increases linearly according to the thickness of the material. As a result, it is important to guarantee that the values used for this application are represented in dB/m. Both in [65] and [66] the measured values are not represented in dB/m; instead, they present a path

loss value for a referred thickness of the material. Due to this, the values taken into consideration for this application were the ones reported in [67], as they were represented in dB/m. The values considered for each material are shown in Table 4.1.

Table 4.1: Attenuation values for the materials at 2.4 GH_z [67].

Material	**Attenuation (dB/m)**
Concrete	50
Drywall	15
Glass	8

Notice that these values can easily vary from one building to another and since they were acquired under such specific conditions, using them under different circumstances may also cause slight errors. Another key factor to consider is the substantial variation in the properties assigned to different types of drywall, resulting in significant differences in attenuation values.

4.2 Beacon Placement

The algorithm uses an image, obtained from the initial setup, as its input. After loading the initial information stored in the configuration files, the algorithm proceeds to acquire the rest of the data. During the algorithm's execution, different layouts are tested and the ones with greater scores are saved. The final step consists of plotting the best layout on top of the schematic image alongside the coverage heat map. Figure 4.5 shows a summarised pipeline from the algorithm.

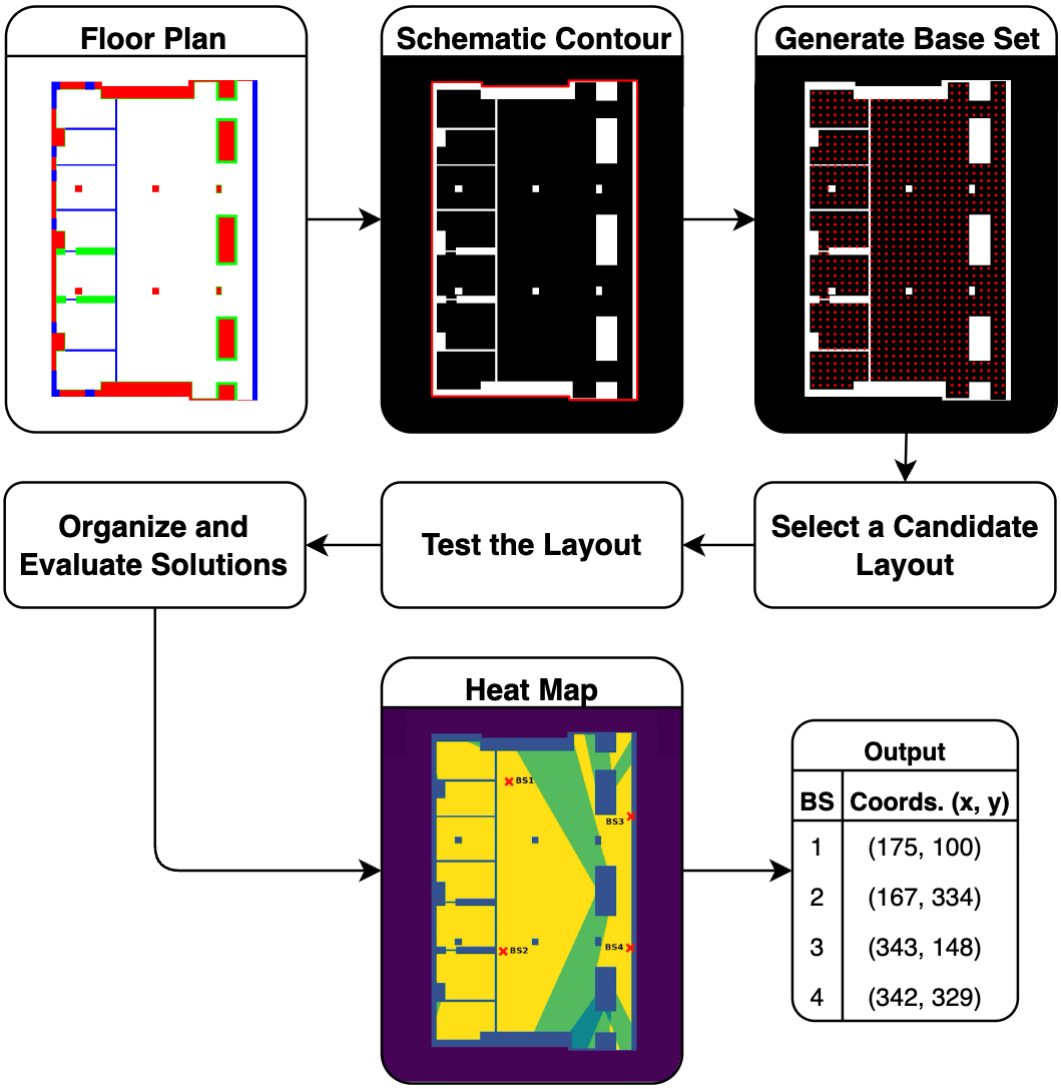

Figure 4.5: Pipeline of the algorithm.

Since the algorithm is rather extensive and more complex than what is shown in the previous figure, this section breaks it down into smaller parts. It is important to highlight that the development of an algorithm is not a straightforward process. As a result, the final solution does not reflect the entire work developed. Due to that, some subsections contain an explanation of preliminary approaches that help to justify the final decision.

4.2.1 Finding Boundaries of the Schematic

The developed algorithm makes use of a building schematic as its primary input source. In the context of this system, the floor plan acts as a reference for determining whether a point is inside or outside the building. To establish such a limit it is important to analyse the image and trace this edge for usage in later stages.

Since the input schematic is represented by an image, this initial step is primarily composed of image-processing tasks. As it was alluded to in Section 3.1, the image processing typically starts with converting the original image into a binary one.

The input image used is the one presented in Figure 4.2b, applying the default binary threshold from OpenCV to it would result in the white areas remaining white and the coloured ones turning black. In Chapter 3, it is noticeable that the binary images are normally represented with a white foreground over a black background, this is also a recommendation seen in OpenCV documentation [54]. To follow this recommendation, the inverted binary threshold was used instead, since it allows the foreground (walls and pillars) to be represented in white while leaving the background black. An Otsu's threshold was employed for improved automation.

Another notable aspect is the lack of any morphological transformation. The image needed editing to acquire materials' information, as outlined in Subsection 4.1.1. This preliminary step addressed imperfections or potential noise in the image, eliminating the necessity for employing morphological transformations afterwards.

Before deciding which method from Section 3.3 is going to be used to acquire the image contours, it is important to understand the goal of the contour in the algorithm. From Figure 4.6, the algorithm needs to be able to consider the black area enclosed by the walls as the inside area (or working area) and the rest as the outside (white spaces and background).

To achieve such a thing the approach was to obtain the external contour from the image and, as shown by Figure 4.6, the retrieved contour, traced in red. This contour marks the exact boundary between the background and the area enclosed by the building walls. This first condition discards all the points that are part of the background. The second condition checks if the point is black, and therefore ensures it is part of the desired area. As shown by the flowchart from Figure 4.7, if the point checks both conditions, it is considered inside or outside otherwise.

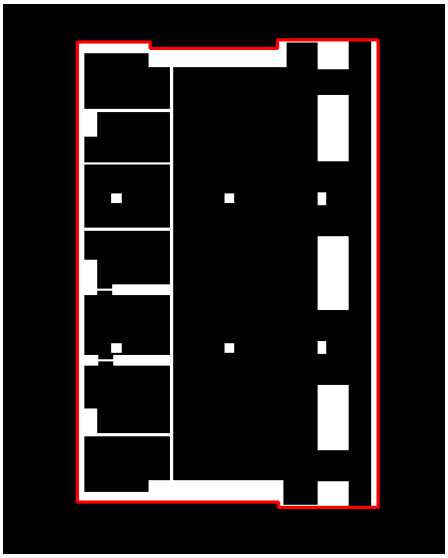

Figure 4.6: Reference contour represented in red used to discard the points from the background.

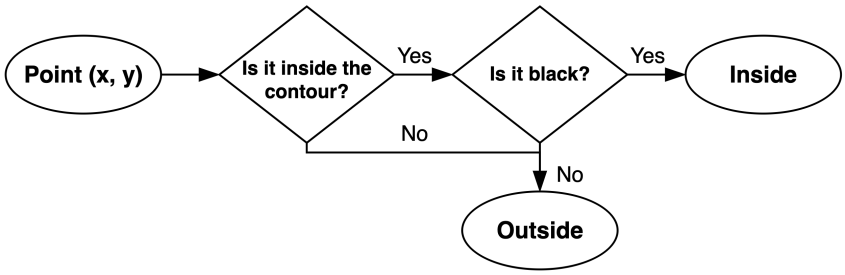

Figure 4.7: Flowchart to identify a point inside: Point (x,y) represents the point being tested, passing both verifications assures the point is inside the schematic and not part of an obstacle.

4.2.2 Generate a Beacon Set

Recalling what was shown in Subsection 2.3.4, the methodology behind the algorithm was to have a random set of generated points that acts as the base set of BSs. Then the algorithm starts to select two points from the set, creating a candidate layout. After testing this layout, it can either enlarge with more points from the base set or stop in case the candidate set provides complete area coverage.

To apply the same principle to this application, the initial approach was to randomly generate each point until 1000 points were obtained. This value was chosen after trial with lower values not attaining a good representation of the floor plan given. An example of the output result can be seen in Figure 4.8a. Due to this randomness, the initial set varies significantly from run to run.

To avoid such variations the final solution followed a different approach. As demonstrated by the flowchart displayed in Figure 4.9, instead of randomly generating each point, the approach is to create a grid space where the points are evenly spaced across the schematic, as shown in Figure 4.8b. In this case, the value of the distance between points was 10 pixels. As this distance is applied in both axes the area left in between is going to be 100 pixels squared ($10 \times 10 = 100$). Due to this, the resulting base set is going to be approximately 100 times smaller than the total number of points in the working area. Notice that depending on the scale the distance between points may be different.

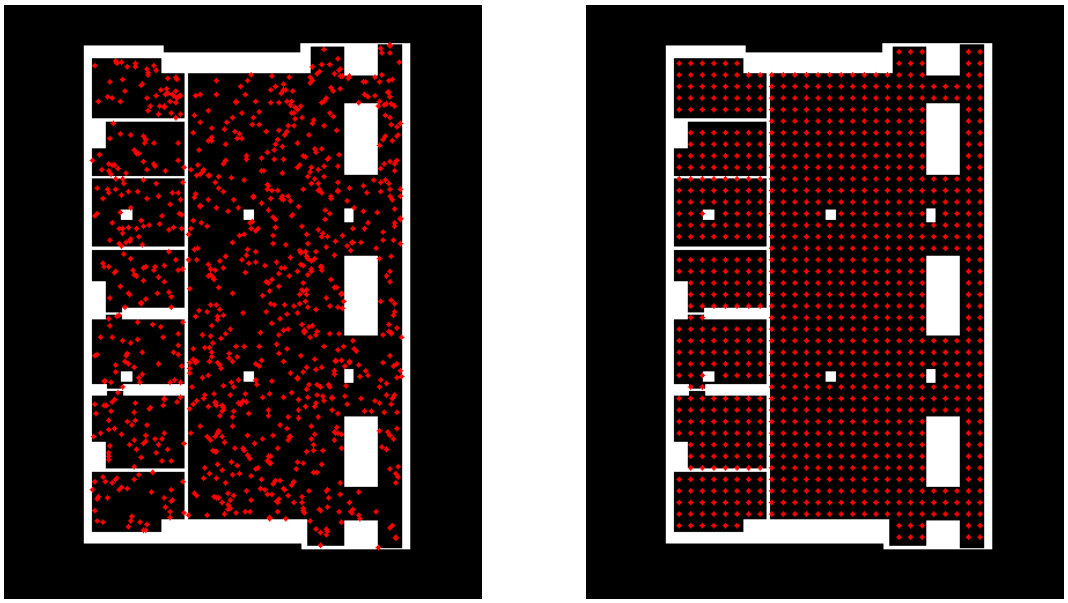

(a) (b)

Figure 4.8: Floor plan with points: (a) Set of 1000 points obtained from a random generation; (b) Set of 983 points obtained from a 96968 point sample by fixing a distance of 10 pixels between the sample.

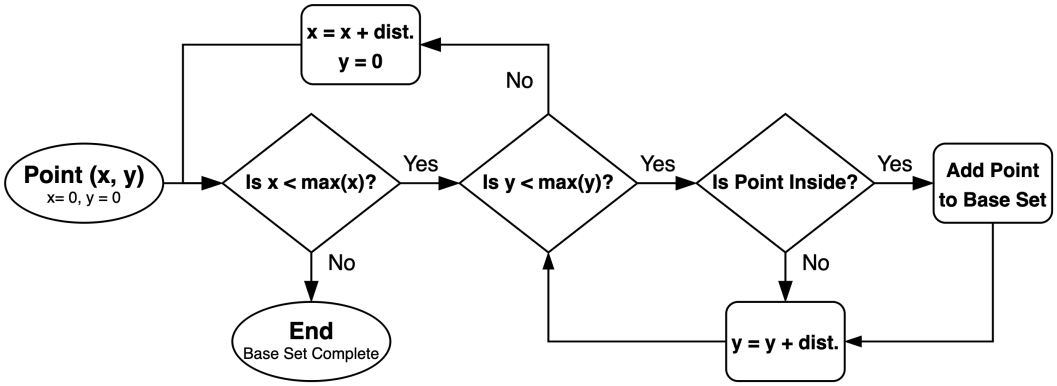

Figure 4.9: Flowchart to generate the base set of BS: Point (x,y) represents the point being tested, with both x and y starting from 0, both max(x) and max(y) are the maximum values that x and y can have in the image, dist. is the distance left in pixels between points, in this work was 10.

4.2.3 Selecting a Candidate Layout

The goal of the previously generated points is to serve as a base for the selection process. As in [49], the approach is to start by selecting two points and assuming them as a candidate layout. The rest of the base set is then used if the candidate layout needs to be enlarged to expand its coverage.

In [14, 51], it is explained that trilateration does not perform correctly if the BSs used for the measurement are near each other. To make sure this issue is avoided, it is necessary to implement a minimum distance between the points selected from the base set. Since this minimum distance represents the distance in between BSs, the lower the number of BSs in a candidate layout, the higher the minimum distance is going to be and vice-versa. Thus, the minimum distance value is expected to be higher at the start of each run, as the number of BSs per layout is at its lowest value.

As shown by the flowchart in Figure 4.10, the candidate layout is complete when there are no additional points in the base set that could satisfy the minimum distance requirement. Thus, defining the right value for the minimum distance is essential, as otherwise, the candidate layout could have more or less BSs than required.

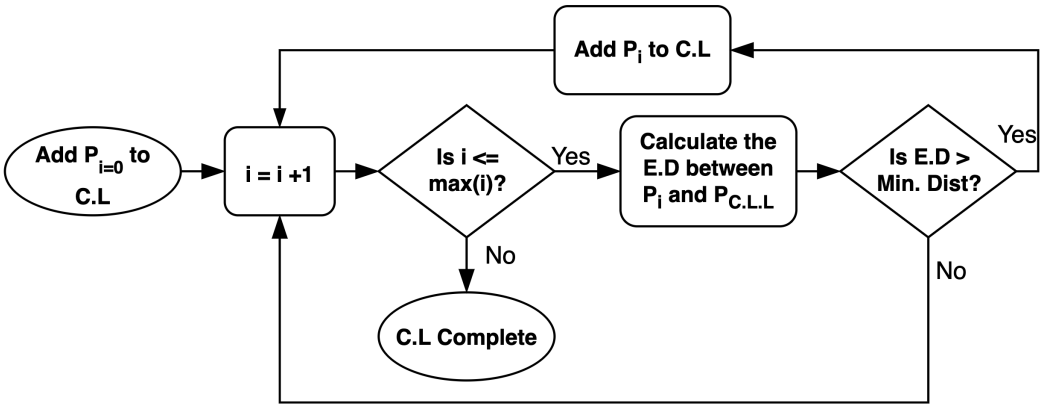

Figure 4.10: Flowchart to Select a Candidate Set: $P_i, i = 0, ..., max(i)$ represents the points from the Base Set, C.L means Candidate Layout, $P_{C.L.L}$ represents the latest point added to the C.L, E.D means Euclidean Distance and Min. Dist means minimum distance; After adding the first point from the Base Set to the C.L, the next point is only added in case its distance to the latest addition is higher than the minimum distance.

This methodology imposes the problem of finding the right minimum distance to guarantee an initial layout of two BSs from the base set.

The approach was to divide the schematic area (red, Figure 4.11) into smaller squares, the number of squares is determined by the number of BSs we are aiming for, in this case, 2 (green, Figure 4.11). Then calculating the distance between the centre (violet, Figure 4.11) of both squares should return a close estimation of the minimum value to be used. Since the squares are supposed to be side by side with no gap in between the distance between centres is the same as the sides of the squares, which is easily achieved with the square root from the area of one of the squares (cyan, Figure 4.11).

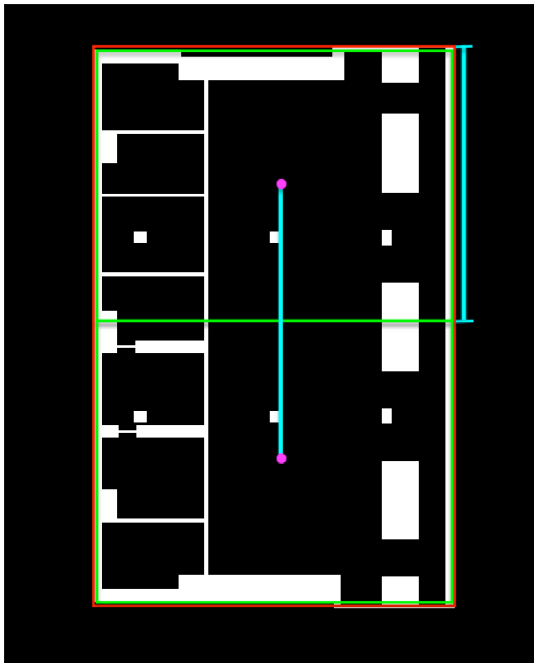

Figure 4.11: Minimum Distance Calculation: The red contour marks the limit of the Working Area; The green are the squares resulting from the division of the schematic in two; The violet are the square centres; The Cyan is the distance between centres which is equal to the square side and represents the initial estimation for the minimum distance.

It is important to acknowledge that the returned value is not more than an approximate estimation. However, this method proves to be very close to the target value in the majority of the cases and, more importantly, it is less resource-intensive than continuously increasing the value of the minimum distance until the candidate layout returns the desired result. Also, before the testing of a new candidate layout, the algorithm always checks if it is composed of two BSs only. If it fails, the minimum distance value is modified and the candidate layout is recalculated until the criteria are met.

If a candidate layout fails to meet the criteria imposed by the test, the algorithm responds by decreasing the minimum distance value between the candidate BSs. This adjustment results in a larger candidate layout. Importantly, the decrease is calibrated in such a way that the new candidate layout is only one BS larger than the previous one, as illustrated in Figure 4.12 featuring the flowchart of this segment.

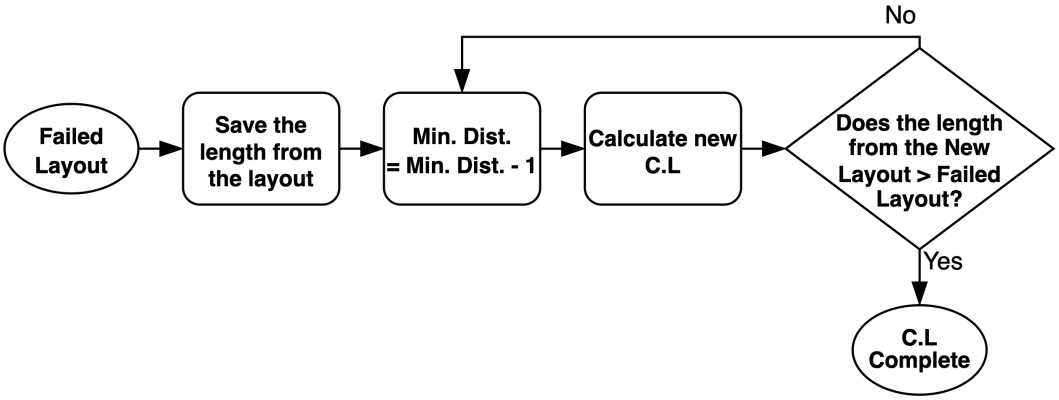

Figure 4.12: Flowchart Enlarge Candidate Layout: C.L means candidate layout and Min. Dist means minimum distance; The goal is to save size from the current C.L, then decrease the minimum distance by one to enlarge the C.L.

Recalling Figure 4.10, it is important to point out that by default the algorithm selects the first point from the base set to become part of the candidate set. If the initial conditions do not change, the output is expected to remain unchanged, which was confirmed in later experiments. This is not ideal, as the algorithm could be permanently stuck into a local maximum. As a result, the solution involved shuffling the points from the base set, so the initial point selected for the candidate layout would change for each run. Figure 4.13 shows the flowchart explaining the implemented approach, which consists of shuffling the base set after finding a feasible solution.

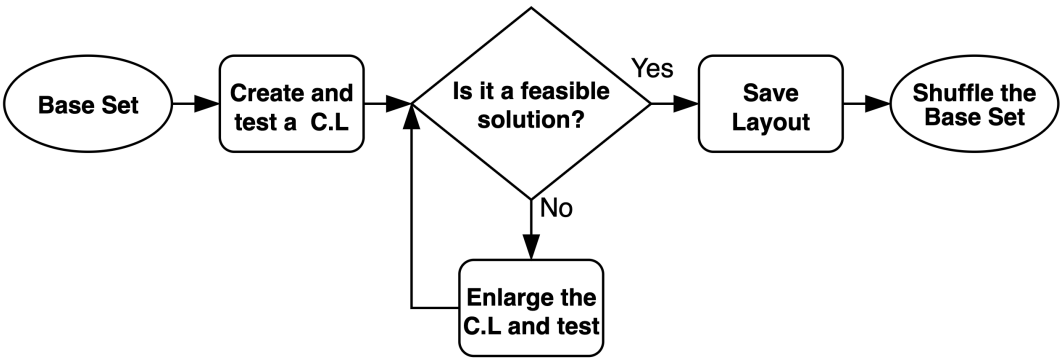

Figure 4.13: Flowchart for Shuffling the Base Set: C.L means Candidate Layout; After achieving the first solution the Base Set is shuffled to change the starting point from the Base Set and therefore generate a different C.L.

4.2.4 Testing the layout

After defining the candidate layout, the algorithm proceeds to test it. In summary, the test consists of ensuring that any point from the working area can be localised using trilateration.

Similar to what is explained in Subsection 2.3.2 and seen in the flowchart from Figure 4.14, for every point in the working area the algorithm has to check which BSs from the candidate layout cover it. Then, based on the number of BSs a point can either be classified as locatable or unlocatable. As defined in 2.3.4 and on the flowchart from Figure 4.15, one BS corresponds to a point being unlocatable, three or more to locatable. The classification for two BSs depends on other factors, such as the number of points resulting from the intersection, schematic verification, etc. A candidate layout is considered a feasible solution for the algorithm when every point from the working area is locatable.

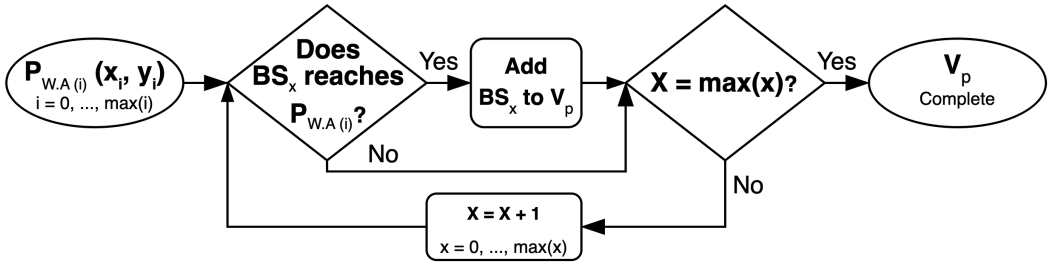

Figure 4.14: Flowchart Test Point Visibility: $P_{W.A(i)}, i = 0, ..., max(i)$ represents every point from the Working Area, $BS_x, x = 0, ..., max(x)$ is every BS from the Candidate Layout and V_p is the point visibility defined in Definition 4. Each point from the working area is tested to find its visibility, the visible BSs are saved in V_p.

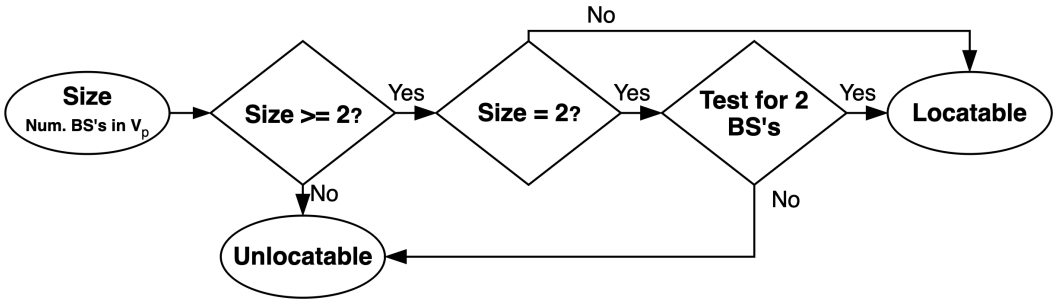

Figure 4.15: Flowchart Test Classification: The Size variable represents the number of BSs in V_p. If it is not higher or equal to two then it is unlocatable, since it means it has one or no BSs visible to the point. If the Size is not two in the second verification it means it has three or more visible BSs, therefore locatable. If the Size is equal to two then it is necessary to apply the methodology explained in 2.3.4.

In Subsection 4.1.3 the link budget was introduced as the selected method for checking BS coverage. The Equation 4.2 expects to receive a value for L, which is the theoretical attenuation between a BS and a point, which is called a pair. To calculate this value the algorithm uses the binary image, to trace a straight segment in between the pair, and checks which pixels from the segment are White (1), as

shown in Figure 4.16. Recalling that the white pixels in the binary image represent the obstacles, that possess an attenuation value.

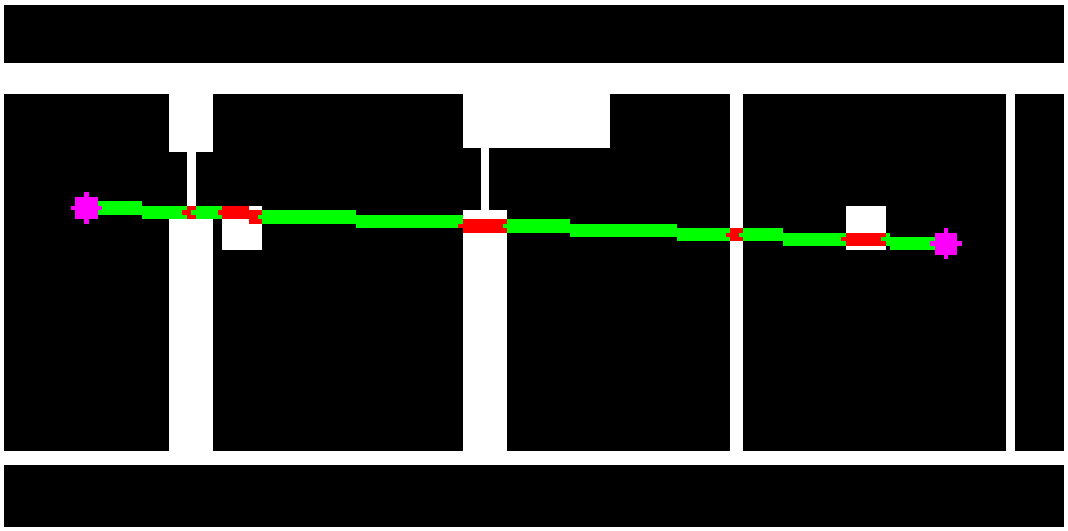

Figure 4.16: Line Segment between BS and Point: The magenta pixels represent both the BS and a point from the Working Area, the green pixels the point from the segment with no theoretical attenuation association (air), the red pixels are the ones that cross obstacles and have an attenuation associated.

After obtaining information about the obstacle pixels - red pixels from Figure 4.16 - the algorithm then checks which colour they are represented by in the image containing the materials, Figure 4.2b. Based on the retrieved RGB code, the algorithm can look for the corresponding attenuation value in the data collected from the initial setup, 4.1.3. Adding all the attenuation values from each obstacle pixel results in the total attenuation, which is then adjusted to its domain when multiplied by the scale.

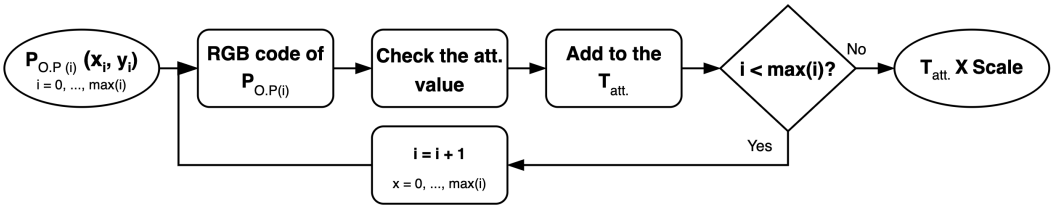

Figure 4.17: Flowchart to Calculate the Theoretical Attenuation: $P_{O.B(i)}, i = 0, ..., max(i)$ represents the obstacle pixels, att. represents attenuation and $T_{att.}$ total attenuation. The returned RGB code from the pixel is then associated with its respective theoretical attenuation value, after summing all the attenuation values the result is multiplied by the scale to find the value of L.

The retrieved value allows the algorithm to calculate the link budget and interpret the result. If the link budget is higher than the P_{RX} value (-96 dbm), then the

BS can reach the point. Otherwise, it cannot.

Recalling the pipeline presented in Figure 4.1, the goal is to create a heat map based on the layout coverage. This coverage variable is the number of BSs visible to each point from the working area (the same one used in the classification process). This is also represented in Figure 4.15 as Size.

Another variable used by the algorithm to classify its solutions is the LOS. The purpose of this variable is explained in a later Subsection 4.2.5, here the goal is only to explain how it was calculated. The LOS condition is met every time a point is locatable and has no obstacles in between it and its visible BSs. However, it is important to note that in case a point has visibility to more than three BSs, it only needs to guarantee LOS with three of them, as this is the minimum requirement for trilateration. The concept of coverage and LOS is explained and illustrated in Figure 4.18.

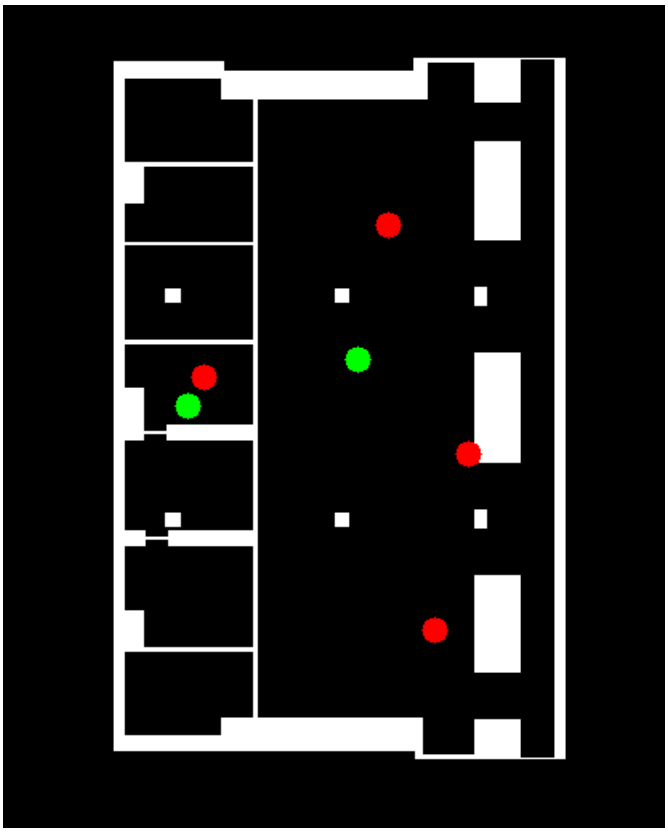

Figure 4.18: LOS Example. Assuming the pixels in red are BSs that can cover both points from the working area represented by the green pixels. In this scenario, both points would have coverage of four BSs, however, only the point in the middle would be considered to have LOS since it has no obstruction with at least three BSs.

4.2.5 Organising and Evaluating Solutions

Recalling the process of selection of the candidate layout explained in Subsection 4.2.3, it was mentioned that shuffling the order of the base set resulted in a different starting point, resulting in a more diversified set of solutions. Although it is important to have some degree of diversification in the output layouts, it is also important to ensure coherence across multiple tests. The main issue with the shuffling approach was that the obtained solutions were not coherent, varying both in BSs per layout and in their positioning in the floor plan.

To avoid such variances, a local maximum for the layout size was set. The goal is for this variable to store the size value from the shortest solution generated, and then use it as a condition to evaluate future solutions. As a result, newer solutions are only kept in case they are equal or lower in size than the current best. In case a layout with fewer BSs is found, all the previous results are deleted, as their purpose is no longer relevant since better alternatives are available.

This method solves the problem of having different numbers of BSs per layout. However, there is still a problem related to the variable positioning of the BSs. To solve this, a large data set of solutions was generated and then passed through a clustering method like the K-Means. The goal was to find the area where the point accumulation would be higher and therefore indicate the most probable solution.

As Figure 4.19 demonstrates, when plotting the results, no clear patterns were found, as the points were randomly spread across the floor plan. As a result, the generated clusters had a large dispersion. Note that the number of clusters is obtained from the number of BSs per layout. For example, if the saved layouts have four BSs, the number of clusters is four.

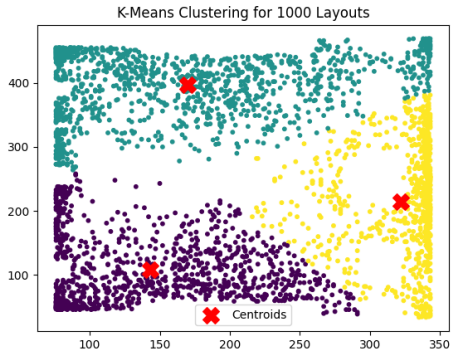

(a) Run 1

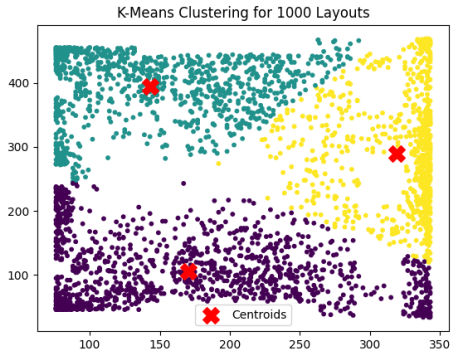

(b) Run 2

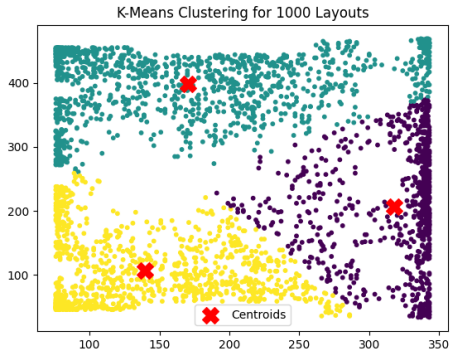

(c) Run 3

Figure 4.19: Clustering method using K-Means: The three plots presented were obtained from three distinct runs that gathered 1000 solutions each with 3 BSs per layout.

Since the clustering method did not produce the results expected, the approach shifted towards a different solution which is the one featured in the final result. In Subsection 2.1.6, it was noted that the Bluetooth technology would achieve more accurate measurements when in LOS condition with the asset. This theory was also confirmed during the measurements taken in Section 5.3. Based on this, the goal was to use the LOS indication mentioned in Subsection 4.2.4 to classify the output layouts and organise them in a descendent order, from the one with the highest percentage of points with LOS to the one with the lowest percentage.

4.2.6 Heat Map

After ordering the obtained solutions based on the LOS percentage, the algorithm returns a plot of the solution. The output result is the one shown in Figure 4.20. As mentioned in Subsection 4.2.4, the heat map is built upon the coverage parameter obtained during the testing of the layout, with the plotted result featuring the best layout obtained based on the LOS percentage.

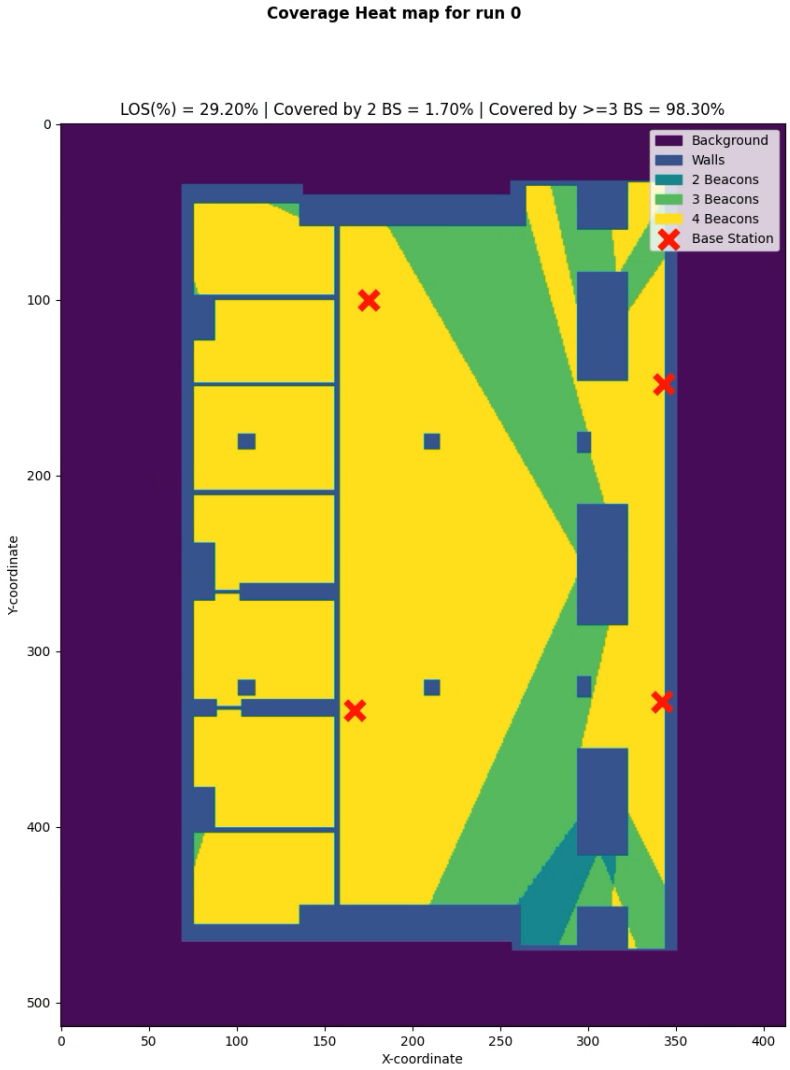

Figure 4.20: Coverage Heat Map: This heat map features the best layout calculated by the algorithm after experimenting with 100 different layouts. The obtained solution has LOS with 29.20% of the working area, with 98.20% of that area being covered by three or more BSs and only 1.70% being covered by two BSs.

Chapter 5

Experiments and Results

5.1 Testing Infrastructure

The experiments described in this chapter were carried out in a pre-existing testing infrastructure. This infrastructure is composed of the BS, the asset, the database and the Asset Tracking Engine (ATE). Each element is explained in the following Subsections. It is important to mention that the elements which compose this infrastructure were developed to perform in a preceding study [31]. Thus as it is mentioned in the introductory chapter of this work, Chapter 1, the goal is to validate the developed indoor location algorithm with the pre-existing software and hardware.

5.1.1 Base Station

The primary role of the BS is to receive BLE signals and to calculate the distances from the assets that are being tracked (Figure 5.1). It comprises a Nordic nRF5340 DevKit (Figure 4.4a) equipped with specialised distance measurement software leveraging novel Bluetooth 5.1 distancing capabilities. This DevKit is linked to a Raspberry Pi running a Python script, responsible for collecting the data sent to the serial port from the DevKit and processing the data before storing it in a database.

The performance of the system significantly depends on the properties of the antenna used. Factors like its directionality, gain, and sensitivity influence the coverage

area and can potentially introduce errors in the estimated positions. To address this, an external monopole antenna was utilised in place of the DevKit's built-in printed antenna to ensure consistent signal coverage throughout the area surrounding the base station. To power all of this it is used a generic power bank with an output rated for 5 W. All base stations are supported by a 3D-printed support and were attached to a wall at an approximate height of 2.15 meters.

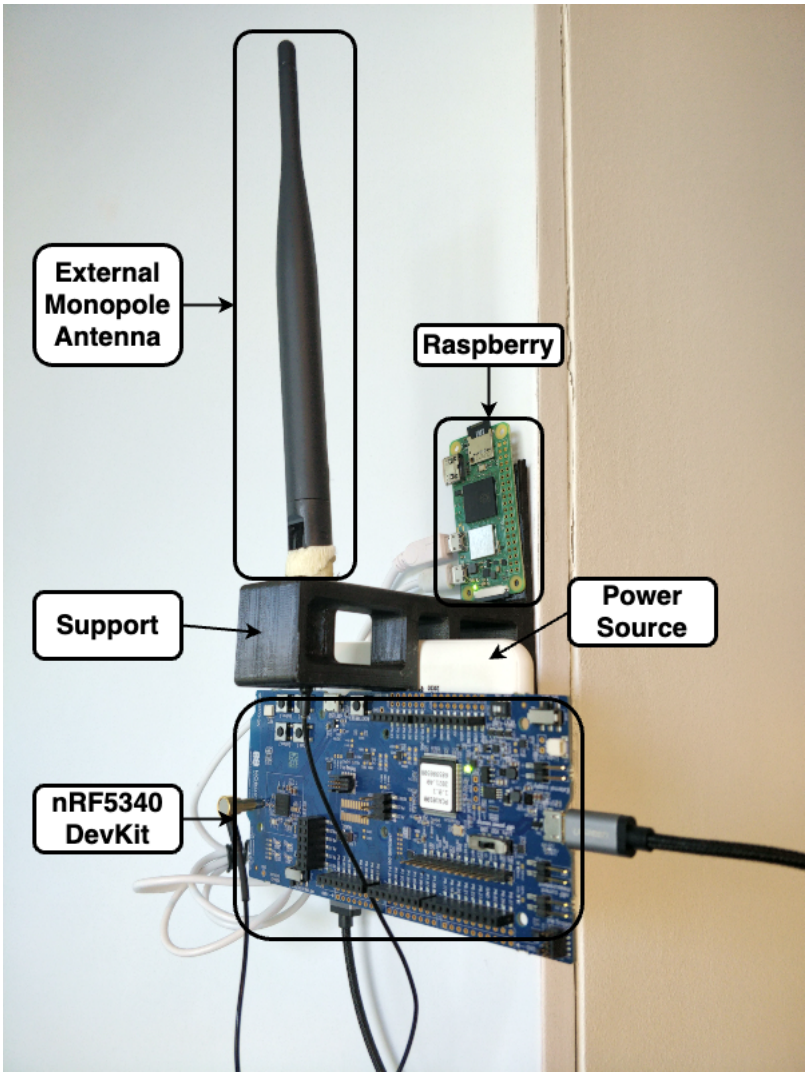

Figure 5.1: Base Stations Assembly

5.1.2 Asset

The Asset is the object the system is trying to locate. For that, it sends BLE Signals for the BSs, Figure 5.2b. It consists of a Nordic nRF5340 DevKit (Figure 4.4a), the same one used for the BS, and for the power source it was used a regular 5 W transformer. The experiments were done with the asset mounted on the back of an

office chair which translates to a height of approximately 1.20 meters from the floor. The asset is attached to a 3D-printed support.

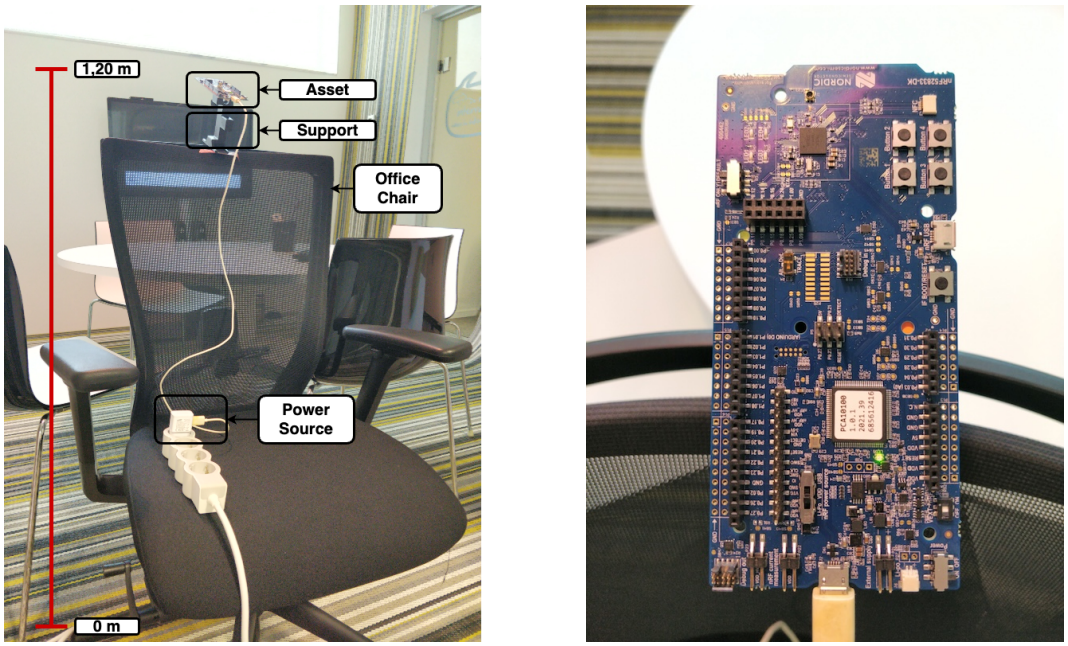

(a) Asset Assembly (b) Asset Radio Module

Figure 5.2: Asset.

5.1.3 Database

The Database is where all the measurements are stored after being acquired and parsed from the BS. It is a relational database deployed with *MySQL*. It consists of two tables, one for the BSs and the other for the asset tracking. The first one stores the id and MAC address of each BS in the testing infrastructure. The second table keeps the recordings from the measurements sent by the BSs. Each line from the table has the id and MAC address from the BS that took the measurement, the MAC address from the asset, a time stamp from when the measurement was taken and five more columns for the values measured from the Devkit.

Those five values incorporate a quality grading from the signal measured and four different ways to estimate the distance. The quality grading has two possible values “ok” or “poor”. These values are attributed by the Devkit and are helpful for the ATE to identify which measurements should be used. The four methods used to obtain the distance are the phase slope estimation, the Inverse Fast Fourier Transform (IFFT) of the spectrum, the Friis path loss (RSSI) and the High Precision (HPrecision), which is the best performing.

Nordic Semiconductor© maintains its library in a pre-compiled state, making it challenging to make deterministic mathematical analysis assumptions about the

HPrecision method. The library employs the IFFT spectrum method, augmented with additional computational steps to enhance precision [31].

5.1.4 Asset Tracking Engine (ATE)

The ATE is the application layer used to compute the localisation of the assets. To use this engine it is necessary to define a scenario. These scenarios serve as comprehensive repositories of crucial information essential for conducting measurements and accurately determining the location of a targeted asset. The contents of a scenario encompass:

- Layout Information: List of the BSs used in the experiment. Each BS needs to be declared with its MAC address for identification and the coordinates of its position;
- Timestamps: Temporal details indicating the commencement and conclusion of the experimental phase, enabling a temporal context for the collected data;
- Position: Holds Ground truth location from the asset location and MAC address for identification. This ground truth enables the ATE to calculate the performance parameters relative to the computed value. A scenario can hold multiple asset positions, and each position has to encompass the commencement timestamp from the time the asset was left in that position.

This information is then used by the ATE to query the database and fetch the respective recorded measurements for later being computed and return the asset location. Moreover, the ATE also displays some performance metrics such as the average error, minimum error, maximum error and standard deviation that are used in the following Sections to compare the obtained results.

It is important to mention that the ATE is still in an early stage of development and therefore does not take into consideration some assumptions considered in this work. One of them is the fact that the engine is not ready to calculate an asset positioning using two BSs only. As a result, the layouts used in the experiments cannot be inferior to three BSs. Another limitation is not taking into consideration which BSs should be used to perform trilateration and therefore compute the asset location using all the BSs in the layout. This raises the possibility of computing inaccurate positions due to using measured values from a BS with no NLOS conditions to the asset and thick obstacles in the way.

5.2 Experiments

The experiments were conducted in a real-world environment: an open-space office, on a normal working day, with desks and WiFi and Bluetooth-connected electronic

devices. All these aspects are expected to influence the accuracy of distance estimation. The positions of both the BSs and assets were recorded and then compared to the actual, ground truths.

The experiment consists of testing two different layouts for the same three different positions, Figure 5.3a. Layout 1 consisted of three BSs and was positioned based on a simplistic approach that aimed to guarantee the best coverage possible. To achieve that no BS was positioned in a concave corner that might limit the radio range. Another concern was to position the BSs towards the middle of the schematic to ensure an even coverage of the signal.

Layout 2 was obtained by using the Beacon Placement algorithm developed. As described in the work, this approach takes into consideration more technical factors when compared with Layout 1. Some of them are the link budget to check a BS range, the coverage verification for each point on the schematic, and the LOS effect among other considerations. As a result, the output layout considered four BSs.

Note the fact that Layout 2 has more BSs than Layout 1. This possibility was mentioned in Section 2.3 and pointed out as a downside of the human-made layouts as the increased complexity of a schematic could lead to an overestimation of BSs in search to find full coverage. On the other hand, the contrary may also be applied. Since no in-depth coverage mechanism was used to verify all the locations in the schematic, the result may be a layout with an underestimation of necessary BSs.

The determined positions for the asset are shown in Figure 5.3b. The first position chosen is in the open space (P1) of the schematic, the reason being the fact that this should represent the best scenario where the asset has the best coverage for both layouts, with LOS for three BSs at least. The second position is in the corridor (P2), the goal with this position is to evaluate a scenario on an extreme edge of the schematic where some BSs may have LOS but it is not guaranteed. The third position is inside an office room (P3), this position ensures that no BS has LOS with the asset and therefore evaluates the impact of the walls in the measured values.

Those positions were defined before the second layout was generated, therefore those assumptions made above in terms of LOS and coverage were made based on Layout 1.

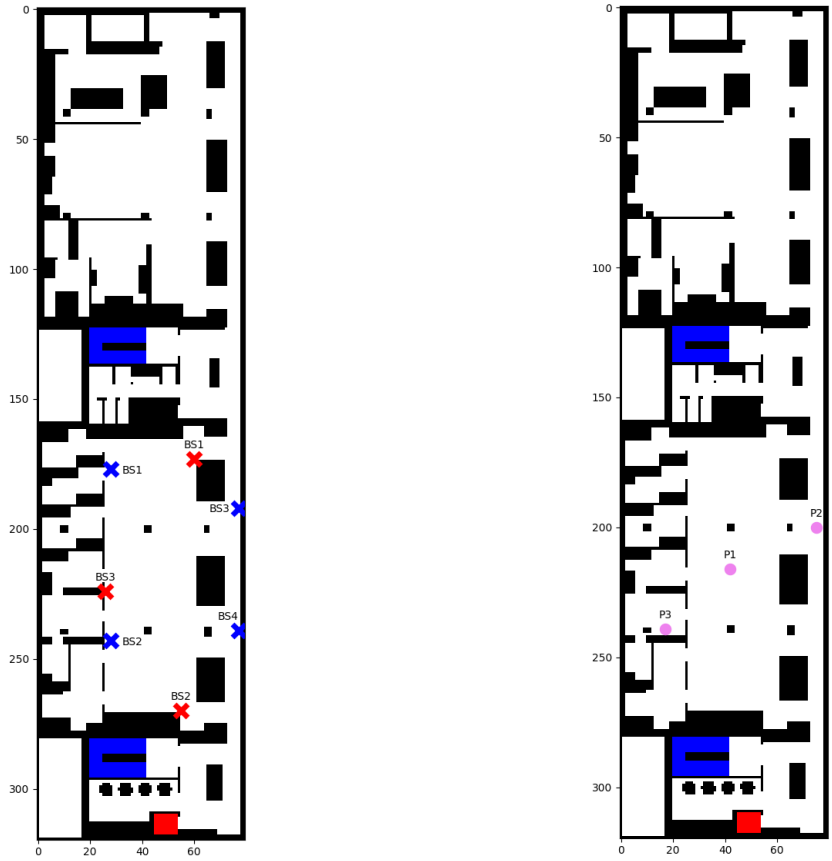

(a) Layouts Tested (b) Defined Positions

Figure 5.3: Layouts Tested and Positions Defined: (a) Layouts Tested - The red crosses are the BSs positions used for Layout 1, the blue crosses are the BSs positions used for Layout 2, each BS has its identification near it; (b) Defined Positions - The violet points are the defined positions for the measurements, each position has its identification near it.

The recording of the measurements followed the same methodology for both layouts. The asset was positioned in the conditions described in Subsection 5.1.2 and left there for an hour. Despite having registered one hour of measurements in each position, the scenarios used were reduced to twenty minutes only. This shortening in the dataset of measurements is due to the movement of the asset from one position to the other causing the initial values to be in a transitory period and therefore not reporting the most accurate results. Another limiting factor is the environment as, as mentioned at the beginning of this Section 5.2, the testing environment is an office space with many elements that can influence the acquisition of the measurements. As a result, the twenty-minute dataset extracted from the main dataset (one hour) had to reflect a period where the measurements were more or less stable.

5.3 Results and Discussion

This Section is divided into four Subsections. The first three are dedicated to the analysis of the results obtained for both layouts defined previously. In the last Subsection, the goal is to lay the main conclusions from the experiments.

5.3.1 Open Space (P1)

The reason for choosing this position was due to its LOS condition with the three BSs from Layout 1. The same LOS conditions are also verified for Layout 2, which in theory should result in similar computed positions or better results. However, the results shown in Figure 5.4 do not confirm this theory.

To understand this, it is worth remembering how the ATE works. As previously mentioned, the ATE, as well as the remaining infrastructure, was developed to work in a previous work [31]. In that work, the layout used for the experiments is the one identified in this work as Layout 1. Since the ATE uses trilateration the resulting intersection is expected to represent the asset position. Moreover, Layout 1 only has three BS, meaning that no matter the position of the asset the same BSs were used.

As the ATE does not possess any verification method to evaluate which BSs are in the best conditions to calculate the asset position, it uses all the available ones. Layout 2 has four BSs, multilateration with four BSs can result in more than one intersection, which is not a result expected by the ATE.

In Figure 5.4b, it is visible that BS_3 does not have LOS with the asset and thus the measurements reported from this BS are inaccurate. Performing multilateration with this BS results in a second intersection, and consequently in a computed position assumed by the ATE to be in between the two intersections.

Figure 5.4b also shows that in case the ATE had selected the best three BSs, BS_1, BS_2 and BS_4 which have LOS with the asset, the trilateration reported result would have been close to the ground truth defined.

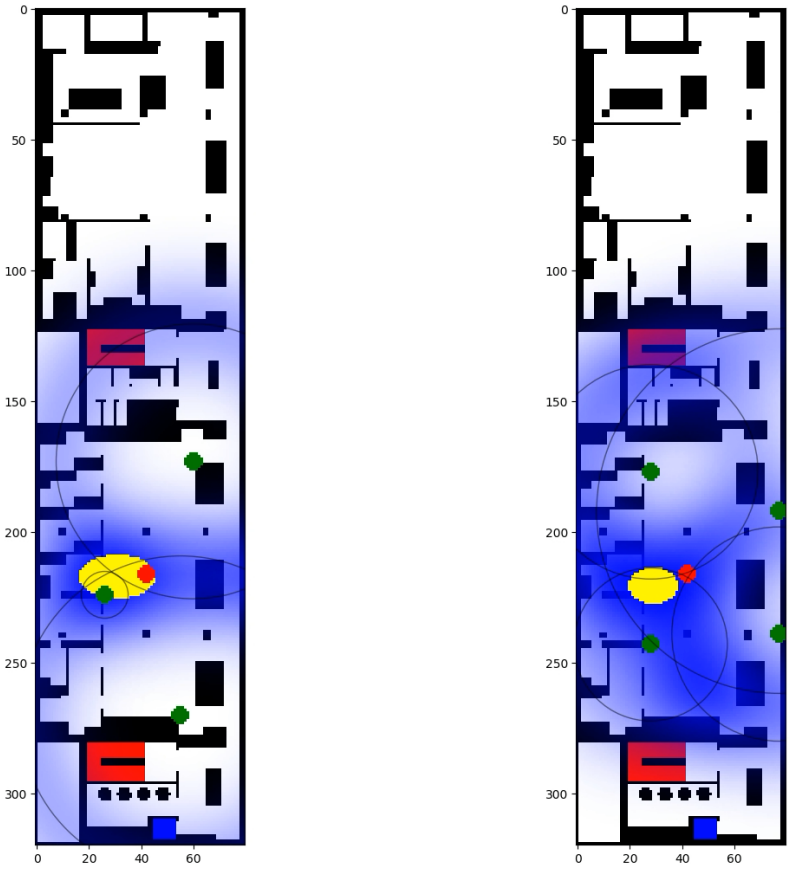

(a) Layout 1 - Prediction Rings for P1

(b) Layout 2 - Prediction Rings for P1

Figure 5.4: Prediction Rings for P1 Position: The green points are the BSs, the dark circular traces are the range of the respective BS with a radius equal to the measured distance, the red point is the ground truth for the asset and the yellow spot represents the predicted position by the engine for the asset location.

From this analysis of the open space position, it is possible to conclude that the accuracy from the computed position was limited by the ATE rather than the layout itself. It was proved that Layout 2 could improve the point localisation as it was shown in Figure 5.4b.

5.3.2 Corridor (P2)

The second position represents the middle scenario, in the case of Layout 1 only BS_3 has LOS with this position. On the other hand, Layout 2 has all the BSs with LOS conditions with the asset. As a result, it is expected that the computed position for Layout 2 to be better than Layout 1. Figure 5.5 shows the comparison of the 3D traces respectively calculated by the ATE for both layouts.

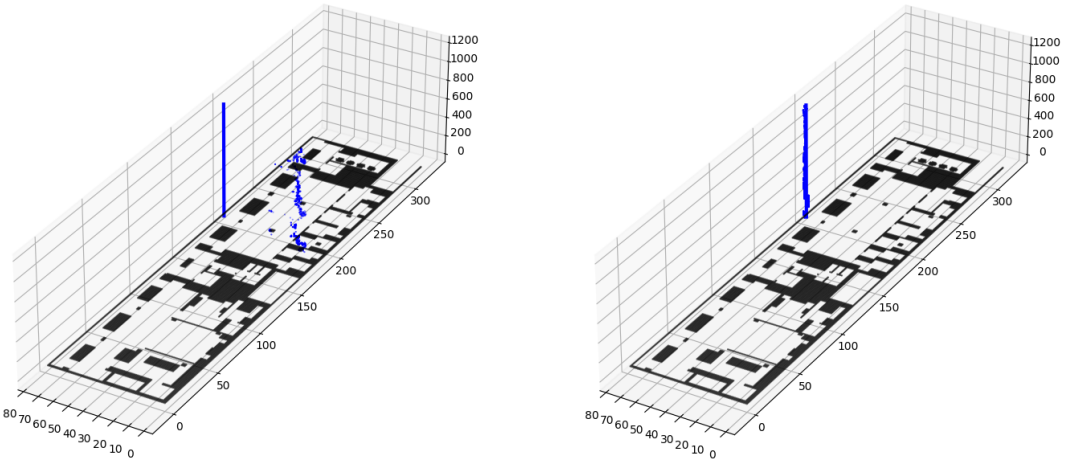

(a) Layout 1 - 3D Trace

(b) Layout 2 - 3D Trace

Figure 5.5: Comparing 3D Traces for P2 Position: The height axis represents the time step, the computed positions are plotted at the height of its corresponding time step, and the straight bar is the ground truth of the asset.

As shown visually by the traces (Figure 5.5) and the performance Table 5.1, the computed position for Layout 2 has improved the results obtained by a significant margin, the average error decreased 97.22%. Moreover, the success of Layout 2 was such that sub-meter performance was achieved, therefore leading to an extremely accurate computed position. These values contrast drastically with what is observed with Layout 1 and enhance the importance of having either LOS or few obstacles in between with the asset aimed to locate.

Table 5.1: Performance results for the P2 position using the two layouts.

Layout	**Layout 1**	**Layout 2**
Avg. Error[1](m)	13.68 ± 1.39	0.38 ± 0.31
Max. Error	14.75	1.21
Min. Error	8.59	0.11

[1] Excluding 60 s before and after asset position change;

This scenario also proves that positioning BS near the centre as in Layout 1 may limit the performance for scenarios like this. It is worth mentioning that due to all BSs having LOS with the asset, no inaccurate values are reported and thus engine computation performs without issues, Figure 5.6b. The following Figure 5.6a shows how the obstacles in between the BS and the asset can affect the measurements acquired.

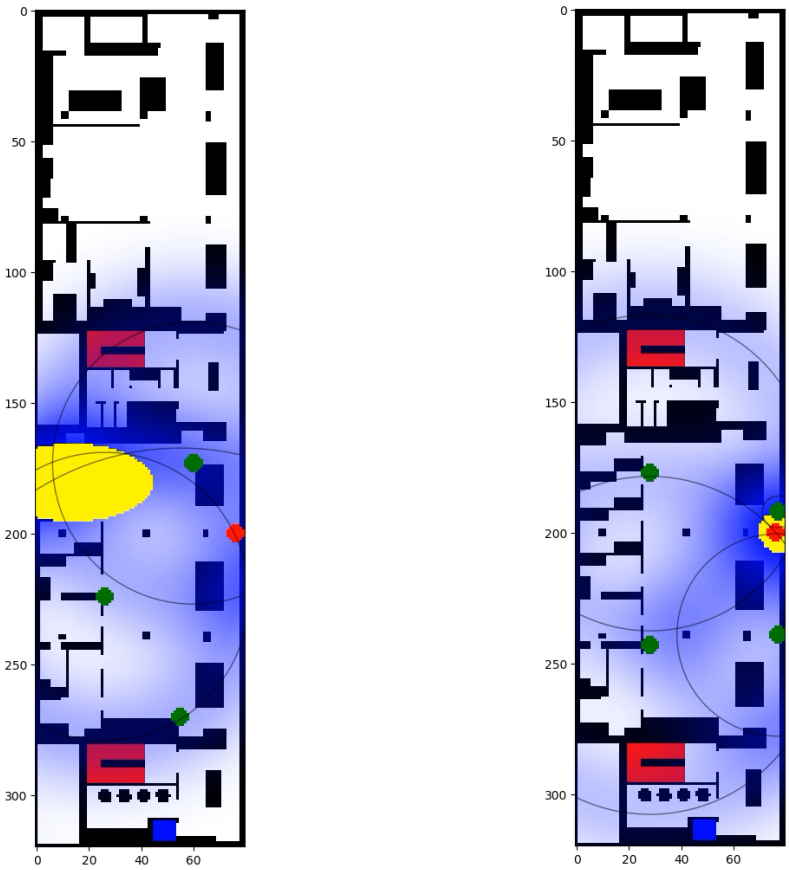

(a) Layout 1 - Prediction Rings for P2

(b) Layout 2 - Prediction Rings for P2

Figure 5.6: Prediction Rings for P2 Position: The green points are the BSs, the dark circular traces are the range of the respective BS with a radius equal to the measured distance, the red point is the ground truth for the asset and the yellow spot represents the predicted position by the engine for the asset location.

5.3.3 Office Room (P3)

The third recorded position is inside the office room, this scenario should represent the worst-case situation where no BS has LOS with the asset. Figure 5.7 shows the comparison of the 3D traces respectively calculated by the ATE for both layouts.

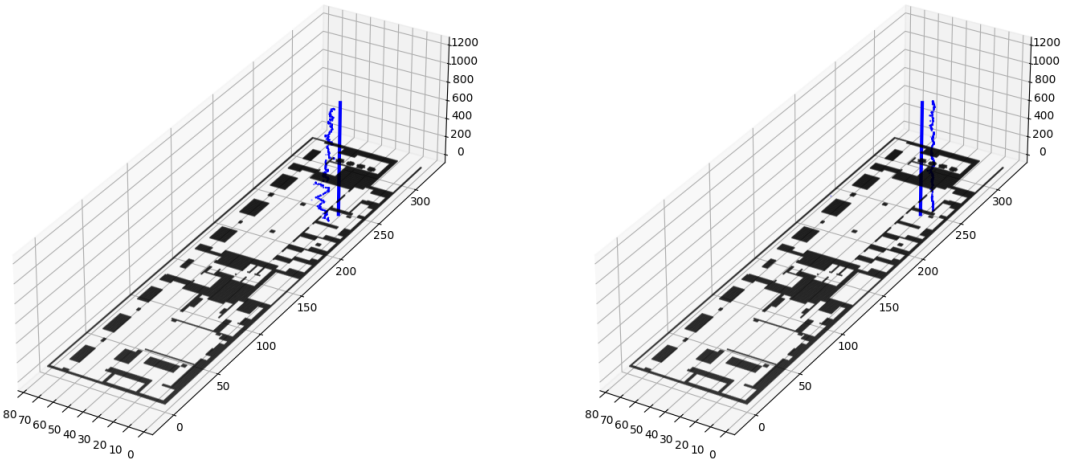

(a) Layout 1 - 3D Trace

(b) Layout 2 - 3D Trace

Figure 5.7: Comparing 3D Traces for P3 Position: The height axis represents the time step, the computed positions are plotted at the height of its corresponding time step, and the straight bar is the ground truth of the asset.

Figure 5.7 shows a similar error for both layouts, which was expected since neither of the layouts had at least three BSs with LOS conditions to the asset. In terms of the performance metrics reported in Table 5.2, most of the indexes are better for Layout 2 with the average error being 34.15% lower. Despite the lower value, this is still not comparable to sub-meter values achieved for the corridor, Subsection 5.3.2. It is also important to address the maximum error of 11.91 m for Layout 2, this value can have multiple natures from the interference of electronic devices to a person standing in between the asset and the BS creating further attenuation and reflections. These factors are a result of the real-world environment in which these experiments were conducted.

Table 5.2: Performance results for the P3 position using the two layouts.

Layout	**Layout 1**	**Layout 2**
Avg. Error[1](m)	2.46 ± 0.40	1.62 ± 0.32
Max. Error	3.80	11.91
Min. Error	1.61	1.19

[1] Excluding 60 s before and after asset position change;

Analysing the prediction rings from Figure 5.8, it is possible to see how the obstacles can influence the acquired measurements, just like in Figure 5.6a. Note that although there is a BS closeby in both layouts, having an obstruction in between (in this case, a glass panel) can alter the measurement by a significant margin.

Despite the schematic trying to replicate the best representation of the real-world scenario, this is not always possible, so it is worth highlighting that the BS two from Layout 2 had a slight disadvantage when compared with BS three from Layout 1, due to a wooden shelf causing more obstruction when acquiring the measurements.

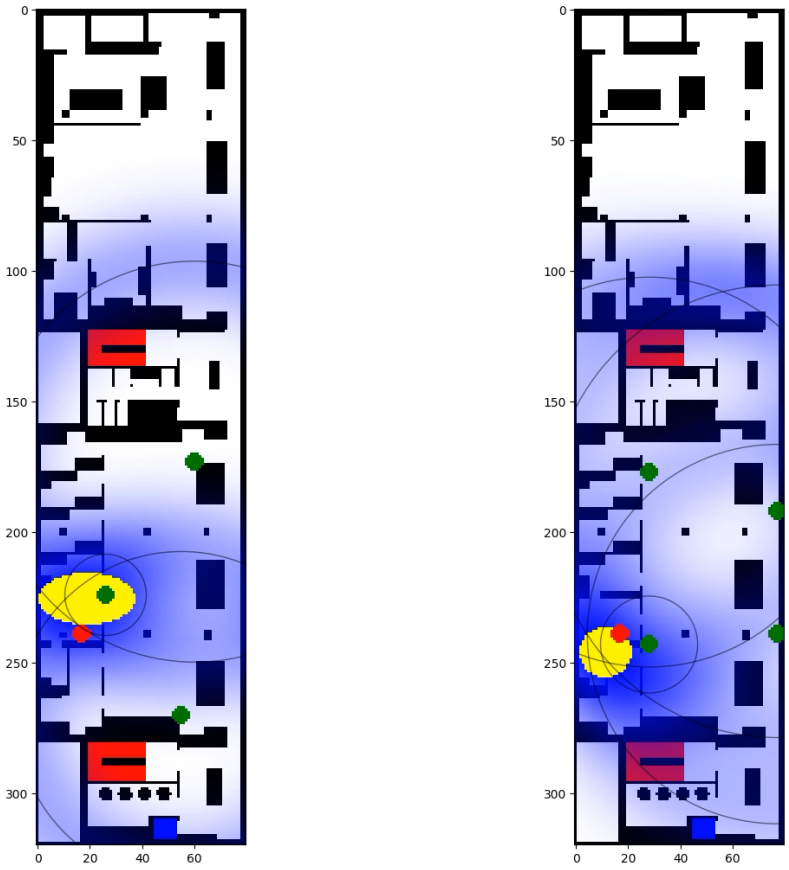

(a) Layout 1 - Prediction Rings for P3

(b) Layout 2 - Prediction Rings for P3

Figure 5.8: Prediction Rings for P3 Position: The green points are the BSs, the dark circular traces are the range of the respective BS with a radius equal to the measured distance, the red point is the ground truth for the asset and the yellow spot represents the predicted position by the engine for the asset location.

5.3.4 Main Conclusions

Analysing the results as a whole, one characteristic is evident: if a system using BLE signals wants to have sub-meter accuracy, then LOS is required for at least three BSs or two depending on the ATE and asset location. Nonetheless, analysing all of these scenarios based on performance metrics alone, especially P3 (Subsection 5.3.3), would not be fair without contextualisation. When working with ILS, precision is not always the key. Taking for example the last position analysed P3, the goal was to represent an extreme scenario and in those conditions, the average error was kept below 2 m, which for indoor room location may be enough if the goal is to identify

which room the asset is in or close by. It depends on the level of precision the application requires.

Another conclusion taken from the experiments is that having a more in-depth analysis of the layout coverage is particularly important than guessing its coverage. The results for P2 showed that Layout 1 due to having all its BSs towards the centre made the layout perform poorly when trying to locate a point in an extreme edge. The way the Beacon Placement algorithm positioned both BSs three and four taking advantage of those side openings of the right side of the schematic allowed the measurements to be as accurate for P2 (extreme scenarios) as they are for P1.

As a result, optimising the Beacon Placement algorithm to work towards achieving full coverage of the space and then selecting the best result based on the LOS concept made a significant difference in the final result.

Chapter 6

Conclusion

The research presented in this work has provided valuable insights into the critical aspect of Bluetooth beacon placement in indoor location applications. The primary goal of this study was to develop and evaluate a beacon placement algorithm aimed at optimising the layout of Bluetooth beacons. The significance of this work lies in understanding the profound impact that beacon layout has on the overall effectiveness and accuracy of indoor positioning systems.

6.1 Current Achievements

The proposed Beacon Placement algorithm has met its intended objectives defined at the beginning of this work. The algorithm can leverage schematic images, effectively translating them into optimised layouts for indoor location-based systems. By harnessing the power of visual data, this algorithm has streamlined the process of beacon deployment, significantly reducing the complexities associated with manual placement while ensuring accuracy and efficiency.

Moreover, the algorithm's thorough examination of BS ranges and its dedication to attaining complete area coverage has played a crucial role in its success. These qualities distinguish the algorithm as a thorough and reliable solution for indoor positioning, marking a notable advancement in the field of beacon placement. This showcases the practicality of employing advanced methods to improve indoor location services.

In addition, it's important to note that the Beacon Placement algorithm stands out in enhancing localisation accuracy, especially in the peripheral areas of the schematic, as demonstrated by the results obtained in Subsection 5.3.2. Through strategic beacon placement in these regions, the algorithm expands its reach, delivering reliable and precise location data. This extended coverage underscores the algorithm's versatility and adaptability, contrary to what was observed when setting up the first layout.

Additionally, the algorithm's effective operation in situations with an obstructed LOS demonstrates its resilience. It can handle challenging conditions, ensuring accurate location data in complex indoor layouts, where obstacles may impede conventional positioning methods. These achievements highlight the algorithm's practical value in real-world applications, solidifying its role as a versatile and dependable tool for optimising beacon placement and improving indoor location services.

6.2 Limitations

Throughout the development of this work, several limitations emerged, encompassing various aspects of the work, from the technology selection to the algorithm design and culminating in the ATE. Firstly, it became evident that the choice of technology posed constraints, as it was mentioned in Subsection 2.1.6 Bluetooth has some difficulties when dealing with NLOS measurements. This was later confirmed by the obtained results in Section 5.3.

The beacon placement algorithm, despite its promising potential, exhibits several limitations that warrant attention. One critical challenge lies in the algorithm's susceptibility to getting trapped at local maxima during the optimisation process. This issue can compromise the final beacon layout, as the algorithm may not always find the global optimum. While local optimisations are often necessary for efficiency, finding a balance that reduces the risk of settling at a sub-optimal solution is essential.

Another significant limitation relates to the extensive run-times required for analysing complex floor plans. The algorithm's computational demands can become burdensome, particularly in large and intricate indoor environments. The need for a substantial amount of time for analysis can hinder real-time beacon placement, which is often crucial in dynamic settings.

Furthermore, the algorithm's reliance on a limited beacon range test, primarily using the link budget, can introduce constraints. A more comprehensive evaluation of the beacon range, including its interactions with obstacles and the physical environment, is essential for a thorough understanding of the system's performance. Relying solely on link budget measurements may not capture the nuances of real-world beacon propagation.

In addition to these concerns, the limited testing with various materials is a notable drawback. Different materials, such as walls, furniture, and architectural elements, can have a significant impact on signal propagation and, consequently, on beacon placement. Ignoring the diversity of materials and their effects can lead to sub-optimal beacon layouts, as the algorithm may not adequately account for these variations in its placement decisions.

Addressing these limitations is crucial for refining the beacon placement algorithm and ensuring its effectiveness across a broader spectrum of indoor environments and use cases. Future developments should focus on mitigating the risk of local maxima, optimising computational efficiency, expanding beacon range testing, and incorporating a more comprehensive understanding of materials to enhance the algorithm's robustness and practical utility.

The ATE also presents constraints that require attention. First and foremost, the engine faces a constraint in its ability to compute with only two BSs. This limitation restricts its applicability in scenarios where a more extensive network of BSs might not be feasible or necessary. In such cases, the engine's inability to effectively function with only two BSs can hinder its practical utility and may necessitate the incorporation of additional infrastructure.

Another critical limitation pertains to the engine's approach to localisation. The engine relies on utilising data from all available BSs rather than selecting the best ones to compute the positioning. without attributing different weights to them based on their signal quality or reliability. This methodology can lead to sub-optimal location estimates, particularly when certain BSs may be obstructed or experiencing interference.

6.3 Future Work

Improvements to the beacon placement algorithm are essential to enhance its effectiveness and efficiency. Firstly, an optimal starting position for the first BS should be identified to initiate the layout calculation from an advantageous point. This ensures that the algorithm's optimisation process begins with a strong foundation and reduces the likelihood of getting stuck in local maxima.

Secondly, to address the issue of extensive run-times, it is advisable to downscale the input image. This downsizing reduces the processing overhead, expediting the analysis of the floor plan. By working with a smaller image scale, the algorithm can maintain computational efficiency without compromising the quality of the layout optimisation. The same methodology was mentioned in Subsection 2.3.1, and the conclusions were that it enhanced the algorithm's run-time.

To advance the algorithm's accuracy, more sophisticated methods for calculating beacon range should be implemented, such as utilising path loss equations. These

equations provide a more precise and realistic estimation of signal propagation, allowing for more informed beacon placement decisions.

Lastly, to mitigate the limitations associated with materials, it is imperative to conduct direct testing to measure the attenuation characteristics of the materials used in the specific indoor environment. Relying on values from other studies may not account for unique material properties, and conducting on-site material attenuation tests ensures the algorithm's accuracy in the real-world setting.

Improving the ATE is crucial for its enhanced functionality and reliability. Firstly, the engine should be upgraded to accommodate schematic data, allowing it to compute effectively with just two BSs, similar to the beacon placement algorithm. This enhancement would broaden the engine's applicability, making it more versatile in scenarios where a more extensive BS network is not available or feasible.

Secondly, an important improvement would involve attributing weights to the available BSs when computing the asset's position. By assigning varying weights based on factors such as signal quality, reliability, and distance, the engine can make informed decisions about which BSs to prioritise for asset tracking. This approach ensures that the engine considers the best-performing BSs, leading to more accurate and robust asset positioning, even in challenging environments with signal obstructions or interference.

References

[1] S. M. Asaad and H. S. Maghdid, "A comprehensive review of indoor/outdoor localization solutions in iot era: Research challenges and future perspectives," *Computer Networks*, vol. 212, p. 109041, 2022. [Cited on page 5]

[2] G. Deak, K. Curran, and J. Condell, "A survey of active and passive indoor localisation systems," *Computer Communications*, vol. 35, no. 16, pp. 1939–1954, 2012. [Cited on page 5]

[3] Ubisense, "Transforming physical space into smartspace." `https://ubisense.com/`, 2023. Accessed: 2023-11-03. [Cited on pages 5 and 12]

[4] Centrak, "Indoor locating for real-time visibility and better healthcare outcomes." `https://centrak.com/`, 2023. Accessed: 2023-11-03. [Cited on pages 5 and 12]

[5] wavecom, "wavecom." `https://wavecom.pt/`, 2020. Accessed: 2023-10-29. [Cited on pages 5 and 12]

[6] J.-S. Lee, Y.-W. Su, and C.-C. Shen, "A comparative study of wireless protocols: Bluetooth, uwb, zigbee, and wi-fi," in *IECON 2007 - 33rd Annual Conference of the IEEE Industrial Electronics Society*, pp. 46–51, 2007. [Cited on pages ix, 5, and 13]

[7] Z. Farid, R. Nordin, and M. Ismail, "Recent Advances in Wireless Indoor Localization Techniques and System," *Journal of Computer Networks and Communications*, vol. 2013, p. e185138, Sept. 2013. Publisher: Hindawi. [Cited on pages vii, ix, 6, 7, 8, 13, and 18]

[8] A. R. J. Jiménez Ruiz, F. Seco Granja, J. Carlos Prieto Honorato, and J. I. Guevara Rosas, "Pedestrian indoor navigation by aiding a foot-mounted imu with rfid signal strength measurements," in *2010 International Conference on Indoor Positioning and Indoor Navigation*, pp. 1–7, 2010. [Cited on page 6]

[9] Y. Gu, A. Lo, and I. Niemegeers, "A survey of indoor positioning systems for wireless personal networks," *IEEE Communications Surveys & Tutorials*, vol. 11, no. 1, pp. 13–32, 2009. [Cited on pages 7 and 8]

[10] H. Zheng, Z. Xu, C. Yu, and M. Gurusamy, "A 3-d high accuracy positioning system based on visible light communication with novel positioning algorithm," *Optics Communications*, vol. 396, pp. 160–168, 2017. [Cited on page 7]

[11] W. A. Cahyadi, Y. H. Chung, and T. Adiono, "Infrared indoor positioning using invisible beacon," in *2019 Eleventh International Conference on Ubiquitous and Future Networks (ICUFN)*, pp. 341–345, 2019. [Cited on page 7]

[12] J. Wang, M. Zhang, Z. Wang, S. Sun, Y. Ning, X. Yang, and W. Pang, "An ultra-low power, small size and high precision indoor localization system based on mems ultrasonic transducer chips," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 69, no. 4, pp. 1469–1477, 2022. [Cited on page 7]

[13] C. Benavente-Peces, V. M. Moracho-Oliva, A. Domínguez-García, and M. Lugilde-Rodríguez, "Global system for location and guidance of disabled people: Indoor and outdoor technologies integration," in *2009 Fifth International Conference on Networking and Services*, pp. 370–375, 2009. [Cited on pages 7 and 8]

[14] F. Zafari, A. Gkelias, and K. K. Leung, "A survey of indoor localization systems and technologies," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2568–2599, 2019. [Cited on pages vii, 8, 9, 16, and 45]

[15] DecaWave, "Real time location systems: An introduction (application note aps003)," tech. rep., DecaWave, Lda, 2014. [Cited on page 8]

[16] A. Correa, M. Barcelo, A. Morell, and J. L. Vicario, "A review of pedestrian indoor positioning systems for mass market applications," *Sensors*, vol. 17, no. 8, 2017. [Cited on page 8]

[17] "What is Ultra-Wide Band (UWB) Technology? - everything RF." [Cited on pages vii and 8]

[18] D. Vasisht, S. Kumar, and D. Katabi, "{Decimeter-Level} localization with a single {WiFi} access point," in *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pp. 165–178, 2016. [Cited on page 9]

[19] S. Kumar, S. Gil, D. Katabi, and D. Rus, "Accurate indoor localization with zero start-up cost," in *Proceedings of the 20th annual international conference on Mobile computing and networking*, pp. 483–494, 2014. [Cited on page 9]

[20] "IEEE SA - The Evolution of Wi-Fi Technology and Standards." [Cited on pages 9 and 10]

[21] H. Jiang, B. Liu, and C. W. Chen, "Performance analysis for zigbee under wifi interference in smart home," in *2017 IEEE International Conference on Communications (ICC)*, pp. 1–6, 2017. [Cited on page 9]

[22] R. Natarajan, P. Zand, and M. Nabi, "Analysis of coexistence between ieee 802.15.4, ble and ieee 802.11 in the 2.4 ghz ism band," in *IECON 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society*, pp. 6025–6032, 2016. [Cited on pages 9 and 11]

[23] N. A. M. Maung and W. Zaw, "Comparative study of rss-based indoor positioning techniques on two different wi-fi frequency bands," in *2020 17th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology (ECTI-CON)*, pp. 185–188, 2020. [Cited on page 9]

[24] Q. D. La, D. Nguyen-Nam, M. v. Ngo, H. Hoang, and T. Q. Quek, "Dense deployment of ble-based body area networks: A coexistence study," *IEEE Transactions on Green Communications and Networking*, vol. PP, pp. 1–1, 07 2018. [Cited on pages vii, 10, and 11]

[25] B. Kim, M. Kwak, J. Lee, and T. T. Kwon, "A multi-pronged approach for indoor positioning with wifi, magnetic and cellular signals," in *2014 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 723–726, 2014. [Cited on page 10]

[26] T. Kim Geok, K. Zar Aung, M. Sandar Aung, M. Thu Soe, A. Abdaziz, C. Pao Liew, F. Hossain, C. P. Tso, and W. H. Yong, "Review of indoor positioning: Radio wave technology," *Applied Sciences*, vol. 11, no. 1, 2021. [Cited on page 10]

[27] S. J. Hayward, K. v. Lopik, C. Hinde, and A. A. West, "A Survey of Indoor Location Technologies, Techniques and Applications in Industry," *Internet of Things*, vol. 20, p. 100608, 2022. [Cited on pages ix, 10, and 13]

[28] M. Cäsar, T. Pawelke, J. Steffan, and G. Terhorst, "A survey on bluetooth low energy security and privacy," *Computer Networks*, vol. 205, p. 108712, 2022. [Cited on pages 10 and 11]

[29] A. Nikoukar, M. Abboud, B. Samadi, M. Güneş, and B. Dezfouli, "Empirical analysis and modeling of bluetooth low-energy (ble) advertisement channels," in *2018 17th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*, pp. 1–6, 2018. [Cited on pages 10 and 11]

[30] A. A. Eltholth, "Improved spectrum coexistence in 2.4 ghz ism band using optimized chaotic frequency hopping for wi-fi and bluetooth signals," *Sensors*, vol. 23, no. 11, 2023. [Cited on page 11]

[31] M. Roque, T. Rocha, D. Pereira, V. Guimarães, and R. Santos, "Novel bluetooth 5.1 location services for indoor asset tracking using multilateration," California State Polytechnic University, IEEE Communications Society, 2023. [Cited on pages vii, 11, 12, 55, 58, and 61]

[32] S. Subedi and J.-Y. Pyun, "A survey of smartphone-based indoor positioning system using rf-based wireless technologies," *Sensors*, vol. 20, no. 24, 2020. [Cited on pages 11 and 17]

[33] "RF-RANGE-ESTIMATOR Calculation tool | TI.com." [Cited on page 11]

[34] "nRF9160: Asset Tracker v2 — nRF Connect SDK 2.3.99 documentation." [Cited on page 11]

[35] L. Klinkusoom, T. Sanpechuda, K. Maneerat, K. Chinda, and L.-O. Kovavisaruch, "Accuracy comparison of bluetooth low energy indoor positioning system based on measurement techniques," in *2022 IEEE International Conference on Consumer Electronics-Asia (ICCE-Asia)*, pp. 1–5, 2022. [Cited on pages 11 and 17]

[36] Nordic, "Bluetooth 5.1 aoa/aod direction finding solution brings high-precision to location services." `https://www.nordicsemi.com/`, 2021. Accessed: 2023-11-11. [Cited on page 12]

[37] E. Hernández-Orallo, A. Armero-Martínez, C. T. Calafate, and P. Manzoni, "Analysis and Evaluation of Tracker Tag Efficiency," *Wireless Communications and Mobile Computing*, vol. 2023, p. e2880229, Jan. 2023. Publisher: Hindawi. [Cited on page 12]

[38] Nordic Semiconductor, *nRF52833*. Nordic Semiconductor, Trondheim, Noruega, 2022. [Cited on pages viii, 12, and 39]

[39] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of wireless indoor positioning techniques and systems," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 6, pp. 1067–1080, 2007. [Cited on pages vii, 13, and 15]

[40] A. Yassin, Y. Nasser, M. Awad, A. Al-Dubai, R. Liu, C. Yuen, R. Raulefs, and E. Aboutanios, "Recent advances in indoor localization: A survey on theoretical approaches and applications," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 1327–1346, 2017. [Cited on pages vii, 14, 15, 17, and 18]

[41] X. Li, Z. D. Deng, L. T. Rauchenstein, and T. J. Carlson, "Contributed review: Source-localization algorithms and applications using time of arrival and time difference of arrival measurements," *Review of Scientific Instruments*, vol. 87, no. 4, 2016. [Cited on pages 14 and 15]

[42] A. Zvikhachevskaya, V. Gourov, A. Awang Md Isa, L. Mihaylova, and G. Markarian, "Enhanced positioning techniques for hybrid wireless networks," in *INFORMATIK 2011 – Informatik schafft Communities* (H.-U. Heiß, P. Pepper, H. Schlingloff, and J. Schneider, eds.), (Bonn), pp. 483–483, Gesellschaft für Informatik e.V., 2011. [Cited on page 14]

[43] J. He, Y. Yu, and Q. Wang, "Rss assisted toa-based indoor geolocation," *International journal of wireless information networks*, vol. 20, pp. 157–165, 2013. [Cited on page 14]

[44] M. Vossiek, L. Wiebking, P. Gulden, J. Weighardt, and C. Hoffmann, "Wireless local positioning - concepts, solutions, applications," in *Radio and Wireless Conference, 2003. RAWCON '03. Proceedings*, pp. 219–224, 2003. [Cited on page 16]

[45] Nordic Semiconductor, *Direction Finding*. Nordic Semiconductor, Trondheim, Noruega, 2020. [Cited on pages vii, 16, and 17]

[46] D. Lu, *Multipath mitigation in TOA estimation based on AOA*. PhD thesis, Ph. D. dissertation, University of Calgary, Alberta, Canada, 2007. [Cited on page 17]

[47] Y. Wen, X. Tian, X. Wang, and S. Lu, "Fundamental limits of rss fingerprinting based indoor localization," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, pp. 2479–2487, 2015. [Cited on page 17]

[48] M. Pušnik, M. Galun, and B. Šumak, "Improved bluetooth low energy sensor detection for indoor localization services," *Sensors*, vol. 20, no. 8, 2020. [Cited on page 18]

[49] H. Wang, N. Rajagopal, A. Rowe, B. Sinopoli, and J. Gao, "Efficient beacon placement algorithms for time-of-flight indoor localization," in *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPATIAL '19, (New York, NY, USA), p. 119–128, Association for Computing Machinery, 2019. [Cited on pages vii, 19, 22, 23, 24, 25, 26, and 45]

[50] S. Löffler, F. Kroll, I. Becker, and P. Hofstedt, "Optimal beacon placement for indoor positioning using constraint programming," in *2022 IEEE/ACS 19th International Conference on Computer Systems and Applications (AICCSA)*, pp. 1–8, 2022. [Cited on pages vii, 19, 20, 21, 27, and 37]

[51] N. Rajagopal, S. Chayapathy, B. Sinopoli, and A. Rowe, "Beacon placement for range-based indoor localization," in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, 2016. [Cited on pages 22, 23, 24, and 45]

[52] S. Curtis, "The classification of greedy algorithms," *Science of Computer Programming*, vol. 49, no. 1, pp. 125–157, 2003. [Cited on page 24]

[53] Brilliant.org, "Greedy algorithms." `https://brilliant.org/wiki/greedy-algorithm/`, 2016. Accessed: 2023-04-26. [Cited on pages vii and 24]

[54] OpenCV, "Opencv." `https://docs.opencv.org/4.x/index.html`, 2023. Accessed: 2023-08-25. [Cited on pages vii, viii, 27, 28, 32, 33, 34, and 42]

[55] Shivanee Jaiswal, "Contrast enhancement of grayscale images using morphological operators." `https://towardsdatascience.com/`, 2022. Accessed: 2023-11-01. [Cited on pages vii and 29]

[56] Jiwon Jeong, "Computer vision for beginners: Part 2." `https://towardsdatascience.com/`, 2022. Accessed: 2023-11-01. [Cited on page 30]

[57] OpenCV, "Opencv-python tutorials." `https://opencv24-python-tutorials.readthedocs.io/en/latest/index.html`, 2016. Accessed: 2023-08-25. [Cited on pages vii, 30, and 31]

[58] AJLearn.com, "Find contours in image using opencv." `https://ajlearn.net/opencv_contours`, 2023. Accessed: 2023-09-18. [Cited on pages vii and 34]

[59] Steerpath, *Beacon placement guide*. Steerpath, Alberga Business Park, Bertel Jungin aukio 3, 02600 Espoo, Finland, 2018. [Cited on page 35]

[60] Flike, "Fluke 424d laser distance meter." `https://www.fluke.com/`, 2023. Accessed: 2023-09-18. [Cited on page 38]

[61] Fluke, *Fluke 424D Laser Distance Meter*. Fluke, Everett, Washington, USA, 2023. [Cited on page 38]

[62] N. Ya'acob, J. Johari, M. Zolkapli, A. L. Yusof, S. S. Sarnin, and N. F. Naim, "Link budget calculator system for satellite communication," in *2017 International Conference on Electrical, Electronics and System Engineering (ICEESE)*, pp. 115–119, 2017. [Cited on page 39]

[63] Glover, Daniel R., *Satellite Radio Communications Fundamentals and Link Budgets*, pp. 293–324. New York, NY: Springer New York, 2013. [Cited on page 39]

[64] Linx Technologies, *CW-RCS Series WiFi 6/6E/7 Right-Angle Whip Antenna.* Linx Technologies, Ort Ln, Merlin, Oregon, 97532, United States, 2018. [Cited on pages viii and 39]

[65] R. Wilson, "Propagation losses through common building materials 2.4 ghz vs 5 ghz," tech. rep., University of Southern California, Aug. 2002. [Cited on page 39]

[66] A. Magiera and J. Solecka, "Radiofrequency electromagnetic radiation from wi-fi and its effects on human health, in particular children and adolescents. review," *Roczniki Panstwowego Zakladu Higieny*, vol. 71, pp. 251–259, 10 2020. [Cited on page 39]

[67] R. Rudd, K. Craig, M. Ganley, and R. Hartless, "Building materials and propagation," *Final Report, Ofcom*, vol. 2604, 2014. [Cited on pages ix, 39, and 40]