Position mapping using BLE mesh networks

Jonathan Persson
jo0462pe-s@student.lu.se
Ernst Padron
er6200pa-s@student.lu.se

Department of Electrical and Information Technology
Lund University

Supervisor: Kaan Bür

Examiner: Maria Kihl

June 12, 2026

Printed in Sweden
Tryckeriet i E-huset, Lund

Abstract

Device location is a practical problem in indoor wireless deployments. Without knowing where devices are positioned in a building, technicians struggle to locate failing equipment, understand coverage issues, and perform maintenance on devices. To facilitate technicians, they often get elevated access to the intranetwork, which not only gives them access to everything, but they still fail to know where devices are located physically in the building. To fully solve this problem, an indoor positioning system is helpful, as it not only displays the devices' locations in a building but also displays the statuses of the devices. This thesis explores the viability of creating such a system using Bluetooth 6.0 with mesh technology and utilizes three different distance measuring techniques called Phase-Based Ranging (PBR), Round-Trip Time (RTT), and Inverse Fast Fourier Transform (IFFT).

For the intended room-level installer support use case, meter-level accuracy is sufficient to identify the correct area rather than exact sub-meter coordinates. Across both Line of Sight (LOS) and Non-Line of Sight (NLOS) conditions, Inverse Fast Fourier Transform (IFFT) performs best, with a mean error of 1.40 m in LOS and 2.40 m in NLOS. In contrast, Phase-Based Ranging (PBR) performs worst with a mean error above 5.00 m for both conditions, while Round-Trip Time (RTT) often performs worse than 1.70 m. The degradation from LOS to NLOS is approximately 40% for PBR, 71% for IFFT, and 102% for RTT.

Based on these ranging results, IFFT is chosen for the indoor positioning system. The position map is then reconstructed using two algorithms,Multidimensional Scaling (MDS) and Semidefinite Programming (SDP). The error from IFFT is low enough for reconstruction attempts. Although the IFFT ranging accuracy is higher, reconstruction introduces additional errors in our experiments, SDP performs better with a mean topology error of 3.37 m, compared to 4.07 m for MDS. This increase in error relative to the IFFT input precision is influenced by multiple factors, especially the network topology and the number of devices. Likewise, the performance of each reconstruction algorithm depends heavily on the deployment environment, the number of devices, and node connectivity. It is therefore difficult to draw a general conclusion about which method is superior across all scenarios.

Popular Science Summary

Finding devices in a large office environment can be a difficult task without a map. Using the new Bluetooth 6.0 distance measurement capabilities and Bluetooth Low Energy (BLE) mesh technology, we can create a map of the devices in the office and their relative positions to each other.

In this thesis, we tested three ways of estimating distance between devices: phase-based ranging (PBR), round-trip time (RTT), and inverse fast Fourier transform (IFFT). The tests were performed both in line-of-sight and in non-line-of-sight office environments where walls, glass, and furniture block or distort signals. This matters because devices in reality aren't placed in ideal conditions.

After comparing the methods, IFFT gave the most reliable distance estimates overall, especially when measurements were filtered over multiple samples. We then used these distances to build a map of where devices were located relative to each other. Two reconstruction methods were evaluated: Multidimensional Scaling (MDS) and Semidefinite Programming (SDP). In our multi-device test, SDP produced the best overall position accuracy, while MDS gave somewhat better consistency for individual pairwise links.

Using Bluetooth Mesh to coordinate the devices for creating indoor positioning to support installers during installation and maintenance, without relying on external positioning infrastructure is a promising approach. For example, an installer could use the system to locate defective devices, verify if a speaker is likely mounted in the wrong corridor or grouped correctly.

The results show clear difficulties in accuracy in challenging non-line-of-sight situations, and reconstruction quality depends on network geometry and connectivity between nodes. Overall the approach is a promising and cost-effective solution with interesting future work, but further development is needed before its full potential can be realized.

Acknowledgements

We would like to express our gratitude to:

- Kaan Bür, academic supervisor, Department of Electrical and Information Technology, Lund University.
- Maria Kihl, examiner, Department of Electrical and Information Technology, Lund University.

Declaration

We hereby affirm that this Master thesis was composed by ourselves, that the work herein is our own except where explicitly stated otherwise in the text. This work has not been submitted for any other degree or professional qualification except as specified; nor has it been published. Where other people's work has been used (either from a printed source, internet, or any other source), this has been carefully acknowledged and referenced. During the preparation of this thesis, we have used Claude to assist us in the writing process to improve language, readability, and correct spelling errors. After using this tool/service, we have reviewed and edited the content as needed and we take full responsibility for the content of the whole thesis.

Table of Contents

List of Figures

List of Tables

Acronyms

API Application Programming Interface
BLE Bluetooth Low Energy
CDF Cumulative Distribution Function
DSP Digital Signal Processing
IFFT Inverse Fast Fourier Transform
IPS Indoor Positioning System
LOS Line of Sight
MAE Mean Absolute Error
MDS Multidimensional Scaling
NLOS Non-Line of Sight
NRF Nordic Semiconductor
PA Public Announcement
PBR Phase-Based Ranging
PoE Power over Ethernet
RMSE Root Mean Square Error
RSSI Received Signal Strength Indicator
RTOS Real Time Operating System
RTT Round-Trip Time
SDP Semidefinite Programming
SMACOF Scaling by MAjorizing a COmplicated Function
SoC System on Chip
ToF Time of Flight
TTL Time To Live
UWB Ultra Wide Band

Chapter 1
Introduction

Bluetooth has long been the standard for peer-to-peer connections between low-power devices, particularly with the introduction of BLE. Its main advantages over other wireless communication technologies include low power consumption, widespread adoption, cost-effectiveness, and strong interoperability across devices from different manufacturers.

One of the limitations of conventional BLE, however, is its inability to create fully distributed networks. Without mesh support, communication is typically limited to direct peer-to-peer connections or centralized topologies, which restrict device-to-device communication across larger networks.

BLE Mesh addresses this challenge by enabling many-to-many communication through a managed flooding mechanism. In a mesh topology, nodes can relay messages for one another, extending network coverage and improving robustness through path redundancy. This architecture enhances scalability, fault tolerance, and flexibility, making it well suited for distributed sensing and control applications.

Beyond communication, mesh networking also enables data exchange among distributed devices. By sharing metrics of spatial location, timestamps, or other proximity-related measurements, nodes can estimate relative distances and infer spatial relationships. This capability provides the foundation for indoor localization systems.

The goal of a 2D Indoor localization system is to provide an alternate way of understanding the topology of a network as well as provide ease of installation when connecting new devices. BLE-based localization offers a low-power and cost-effective alternative to other Indoor localization systems by leveraging signal characteristics to estimate device positions.

For the application studied in this thesis, the relevant requirement is room-level localization rather than exact sub-meter coordinates. Many indoor location-aware services only require determining the correct room or region, and that practical Indoor Positioning System (IPS) solutions often provide average coordinate accuracies in the range of 1-5 m [1]. The office environment used in this thesis has smaller rooms of approximately 12 m^2, which means that the upper end of that error range can be large enough to confuse adjacent rooms or nearby corridor space. Room-level accuracy is therefore considered achieved only when the system is accurate enough to distinguish the correct room reliably, making the lower part of the meter-level error range substantially more useful for the intended installer support scenario.

Traditionally, BLE-based localization has relied on Received Signal Strength Indicator (RSSI) to estimate distances. RSSI is available on every BLE chipset and has been extensively studied, but it is fundamentally limited in accuracy: signal strength is highly sensitive to multipath propagation, absorption, and environmental variation, making reliable ranging difficult to achieve in practice [2], [3]. With the introduction of Bluetooth 6.0, a new capability called Channel Sounding became available, enabling three new ranging techniques: PBR, RTT [4] and IFFT. These methods exploit the physical properties of the radio channel more directly than signal strength alone, offering the potential for improved ranging accuracy. Because Channel Sounding is newly introduced, its performance and limitations in practical mesh deployment scenarios have not yet been broadly characterized. This makes it a compelling subject for investigation.

In this thesis, we study BLE Mesh technology and investigate these Channel Sounding-based ranging methods for constructing a 2D indoor positioning map. The goal is to design and evaluate a decentralized localization approach capable of generating spatial representations of devices in indoor environments without relying on external positioning infrastructure.

1.1 Network Speakers

Network speakers serve as Public Announcement (PA) systems in a wide range of environments, including schools, offices, and public spaces. As a complement to security cameras, they enable real-time delivery of alerts, notifications, and announcements triggered by security events such as intrusion detection, fire alarms, or system malfunctions.

Network-enabled speakers offer audio-based security solutions in environments where cameras are insufficient. These speakers are typically connected to the network via Power over Ethernet (PoE) cables, which supply both power and data over a single connection. Depending on the deployment, multiple speakers may be distributed across different wings or sections of a building to ensure adequate coverage throughout the space.

Each speaker must be individually configured with parameters such as IP address, physical location, audio zone assignment, volume level, sound type, and event timers. This configuration process demands technical expertise and, in many cases, specialized tools or software, making it both time-consuming and error-prone. A particular challenge arises for installers who are not always familiar with the physical layout of a building: without knowing where each speaker is positioned relative to others, verifying correct placement and coverage becomes cumbersome. Additionally, connecting to an internal ethernet network solely for testing and commissioning purposes adds friction to the installation workflow.

With these challenges in mind, this work explores a wireless alternative that can streamline installation and configuration process. We therefore investigate the possibility of using a BLE Mesh network to assist installers in mapping and configuring network speakers without requiring a wired network connection.

1.2 Thesis Objectives and Contributions

The goal for this master's thesis was to evaluate the potential of BLE Mesh technology for indoor positioning applications, specifically in the context of network speaker installation and configuration. The thesis aims to contribute to the understanding of how BLE Mesh can be utilized to create an IPS that assists installers in mapping and configuring network speakers without requiring a wired network connection.

The aim of the project is to explore and evaluate different distance measuring techniques and understand how they can dynamically create a virtual map of devices physical position. We evaluate the accuracy of measuring algorithms and identify both the opportunities and limitations associated with BLE mesh in practical applications. The thesis problem is formulated with the following research question:

How can different distance measuring techniques be used to create a 2D indoor positioning system based on BLE Mesh, and what are the accuracy, opportunities, and limitations of such a system in practical applications?

The rest of this thesis is organized as follows. In chapter 2 we review previous work on distance measurement methods for BLE Mesh and indoor positioning systems. Chapter 3 describes the method and implementation, detailing the experimental setup, and evaluation metrics. Chapter 4 presents the results from line-of-sight and non-line-of-sight distance measurements using Nordic Semiconductor (NRF) devices, comparing the performance of PBR, RTT, and IFFT methods with and without filtering. Chapter 5 discusses the evaluation criteria and comparative analysis of the ranging techniques across different environmental conditions. Finally, chapter 6 concludes the thesis by identifying IFFT as the most suitable method for the room-level indoor positioning objective studied here and suggests future work on network scaling and multi-floor reconstructions.

Chapter 2

Previous Work & Background

BLE Mesh was standardized in 2017 and has since become a significant research platform for IPS, particularly in multi-hop network scenarios. Distance measurement and indoor positioning using BLE has evolved through several approaches, each with distinct trade-offs in accuracy, hardware requirements, and energy consumption. The following review covers distance measurement methods, system architectures, and practical considerations relevant to distributed BLE Mesh positioning.

2.1 Distance Measurement Methods

Accurate distance measurement is important for indoor positioning systems. Several techniques exist which could be used for BLE networks, each offering different trade-offs. The discussion below examines the primary distance measurement methods available, ranging from signal strength-based approaches to time and phase-based techniques.

RSSI is a measure of the power level that a device receives from a signal transmitted by another device. It works by measuring the strength of the signal as it arrives at the receiver. The stronger the signal, the closer the transmitting device is likely to be. However, RSSI can be affected by various factors such as obstacles, interference, and multipath effects, which can lead to inaccuracies in distance estimation. A good signal has a signal strength close to 0 dBm, while a weak signal has a value close to -100 dBm.

RSSI-based methods form the foundation of most practical BLE positioning systems. Several papers like [2] and [3] present RSSI as the main method to distance estimation. The received signal strength is converted into distance estimates using path-loss models. The results from these papers show that while RSSI can provide rough distance estimates, it is highly susceptible to noise and environmental factors, leading to significant errors in many cases. The accuracy of RSSI-based positioning can vary widely depending on the specific environment and the presence of obstacles or interference.

The widespread adoption of RSSI comes from its availability on all BLE chipsets without requiring specialized hardware. Research by [5], [6] and [3] look at challenges including multipath interference, environment specific calibration requirements, and the effect of multi-hop signal chaining in mesh topologies. These studies demonstrate that while RSSI can achieve room-level or meter-scale accuracy, performance is very dependent on environmental conditions. Performance indoors is also significantly worse than outdoors due to increased multipath and signal attenuation.

[2] and [6] explore various filtering techniques to mitigate RSSI noise, such as Kalman filters and average filters. The best results are often achieved using simple averaging techniques. [2] explores how to create a BLE mesh networks and defines its own application layer protocol to send information to all nodes. The network is then used to explore how an indoor positioning system can be created mainly using RSSI, fingerprinting and trilateration. The results show that it is possible to create an IPS with unfiltered RSSI and trilateration and achieve 4.3 meter accuracy. Fingerprinting achieved better results at around 2.68 meters.

2.1.1 Round-Trip Time

RTT is a method used to measure the time it takes for a signal to travel from a source device to a destination device and back again, as illustrated in figure 2.1 [7]. In the context of BLE mesh networks, RTT can be utilized to estimate the distance between two devices by calculating the time delay associated with the round-trip of a signal. The performance of RTT can be affected by many different factors.

Processing delays within devices introduce errors into RTT measurements. The time required for a device to detect an incoming signal, generate a response packet and transmit an acknowledgment can vary depending on hardware and software implementation. These delays then affect distance estimation errors as the measured time includes both signal propagation time and device processing time [7].

Multipath propagation and reflections present major obstacles in indoor environments. Radio signals reflect off walls, furniture, and other surfaces, creating multiple propagation paths with different delays. Even though reflected paths are longer than the direct line-of-sight path, the superposition of multiple components can bias timing estimation in simplified channel models, which may cause both underestimation and overestimation of the true distance between devices [8].

Physical distance and signal attenuation also constrain RTT performance. Like all wireless systems, RTT distance estimates are fundamentally limited by the speed of light and cannot overcome the constraint that signals must physically propagate between devices. Additionally, signal attenuation with distance makes it harder to detect the response signal reliably, which can introduce detection delays that further affect measurement accuracy [7].

In a BLE mesh network, RTT is implemented by having one device send a signal to another device, which then immediately responds with an acknowledgment signal. The source device measures the total time taken for the signal to travel to the destination and back. By knowing the speed of radio waves the device can calculate the distance based on the measured time delay. However, RTT measurements can be influenced by processing delays within the devices, so it is essential to account for these factors to improve accuracy [4].

Time-based estimates have also been considered for a distance application, the following paper [9] shows how it is feasible to use Time of Flight (ToF) between Bluetooth 6 devices to get accurate distance estimates. The position estimations showed good accuracy in indoor environments when LOS conditions were available.

Recent developments with BLE 6 introduce RTT as a standardized time-based ranging mechanism, showing promising results for accurate ranging in challenging environments over distances from 1 to 40 meters [10]. The paper shows that outdoors, RTT achieved a mean distance error of 0.35 m with a standard deviation of 0.26 m. One advantage noted is that RTT only needs a single channel, making it less demanding than other algorithms. Indoors the RTT measurement degraded significantly with overall higher mean errors.

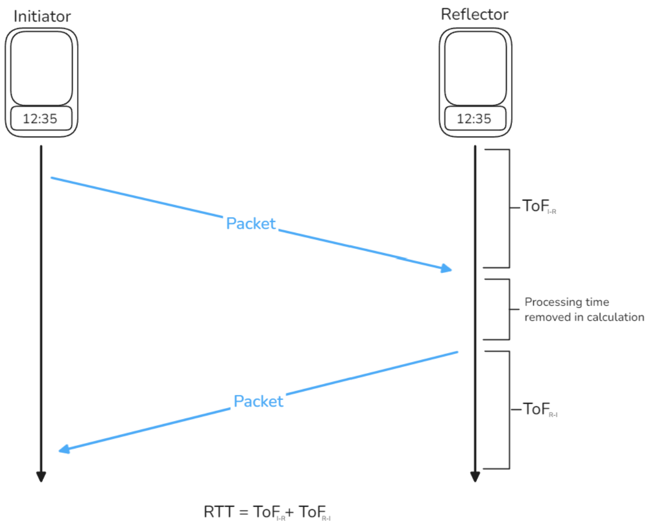

Figure 2.1: An illustration of the Round-Trip Time measurement process between two BLE devices.

2.1.2 Channel Sounding

Channel Sounding represents an emerging technique in BLE 6.0 that allows two compatible devices to measure distance with greater accuracy than traditional RSSI methods [11]. This technology has yet to be widely adopted and is seldom discussed in IPS literature despite its potential advantages. Channel Sounding works by having devices send specially designed reference signals and analyzing the response to extract distance information. The fundamental principle is that radio signals experience delays proportional to physical distance, allowing distance inference from propagation time or phase characteristics [10].

Channel Sounding in BLE 6.0 supports two ranging algorithms: RTT and PBR. RTT, discussed in the previous section, measures round-trip propagation time by sending a packet and measuring response arrival time. PBR estimates distance by measuring the phase shift of a signal transmitted across multiple frequencies [4]. Since the phase of a received signal is proportional to the propagation delay, comparing phase offsets across frequencies allows time-of-flight and distance to be inferred. Sweeping multiple frequencies is essential to resolve the inherent 2π phase ambiguity present when using a single frequency, enabling unambiguous distance estimates over longer ranges. This method can be seen in figure 2.2. Under favorable LOS conditions this method can achieve centimeter-level accuracy, making it a strong candidate for high-precision indoor positioning [12].

The two algorithms have complementary strengths and limitations. RTT is more robust to certain channel conditions but sensitive to processing delays and multipath effects, while PBR offers superior accuracy in line-of-sight scenarios but can be affected by environmental factors impacting phase measurements across the frequency bandwidth. The choice between RTT and PBR depends on specific deployment requirements [10].

Environmental factors significantly impact Channel Sounding performance. In open line-of-sight environments, both algorithms achieve best accuracy. However, indoor multipath propagation, where signals reflect off walls and furniture creating multiple propagation paths with different delays, can cause measurements to deviate from true distance. Signal shadowing from obstacles attenuates received signal strength, potentially causing detection failures. Non-line-of-sight scenarios typically result in significantly degraded performance for both algorithms [10].

Channel Sounding hardware requirements are more stringent than traditional RSSI. Devices must implement precise timing circuits for RTT or high-resolution phase measurement for PBR. The Nordic Semiconductor nRF54 series includes integrated Channel Sounding support. However, not all commercial BLE devices support Channel Sounding, limiting deployment flexibility. The emerging nature of this technology means practical experience remains limited compared to established RSSI-based methods.

2.1.3 Inverse Fast Fourier Transform

Another way of measuring distance is by using an Inverse Fast Fourier Transform on the IQ values from PBR. The IFFT algorithm works by looking at the same IQ measurements from the channel sounding algorithm but extracting information in a different way [13]. An IQ measurement is a complex-valued sample consisting of an in-phase (I) component and a quadrature (Q) component. Together they represent the amplitude and phase of a signal at a given frequency [14]. When the channel sounding procedure collects these measurements across a range of subcarriers the result is a frequency-domain representation of the channel.

The time-domain impulse response produced by the IFFT transformation directly represents the channel's response to an impulse signal across time. Each peak in the impulse response corresponds to a distinct signal propagation path, with the peak's position revealing the propagation delay and thus the distance to that path [15]. In a line-of-sight scenario, the first and strongest peak corresponds to the direct path between transmitter and receiver, while subsequent smaller peaks represent reflections from walls, furniture, and other obstacles. The index of the primary peak is then converted into a distance estimate. The IFFT algorithm uses 75 different frequencies to give multiple IQ values. The frequencies

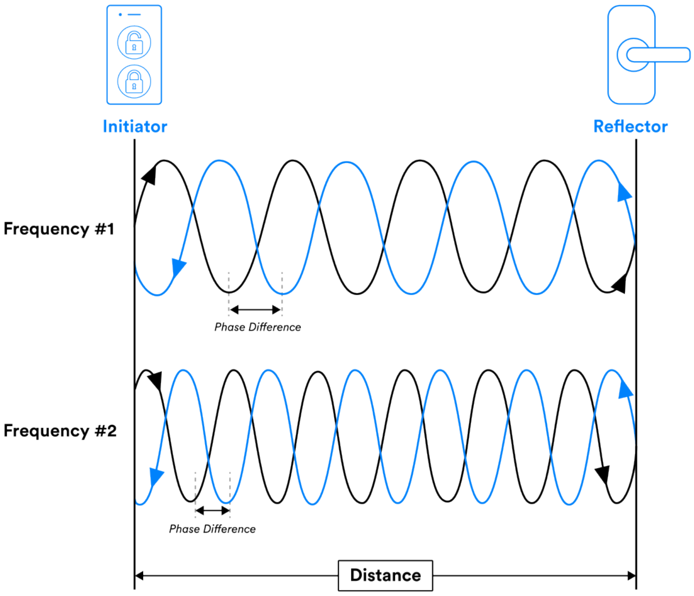

Figure 2.2: An illustration of the Phase-Based Ranging process between two BLE devices. Credit: Bluetooth.com.

are spaced aparts with 1 MHz with all being around 2.4 GHz [16].

[10] evaluates four ranging algorithms, RTT, Slope, IFFT, and a proposed enhancement IFFT2, across both open-field and underground utility basement environments, finding substantial performance differences between them and between the two settings.

In line-of-sight conditions, the phase-based methods significantly outperformed RTT. IFFT2 achieved a mean error of just 13 cm with a standard deviation of approximately 1 cm, while the standard IFFT reached 21 cm mean error. RTT managed 35 cm, and the Slope method performed poorly at nearly 1 m mean error.

Indoor performance degraded substantially for all methods. In a circa-1910 building's basement with concrete ceilings and metal-doored corridors the multipath propagation and signal shadowing drove errors into the multi-metre range, with mean standard deviations of 1.6 m (IFFT2) to 2.1 m (IFFT). Several measurement positions were entirely unreachable [10].

2.1.4 Ultra-Wideband

Ultra Wide Band (UWB) is a wireless technology that offers superior ranging accuracy compared to conventional BLE-based methods. UWB works by transmitting extremely short duration pulses across a wide frequency spectrum, resulting in precise time-of-arrival measurements with accuracy around 50 centimeters [17]. The fundamental mechanism underlying UWB's precision is its use of time-of-flight measurement, where the propagation delay of the transmitted pulses between devices is converted into distance estimates. By using a broad frequency bandwidth around 3.1 to 10.6 GHz, UWB achieves nanosecond-level resolution that enables sub-50 centimeter accuracy. This is more precise than signal strength-based methods such as RSSI, which usually achieve meter-scale accuracy. The wide frequency spectrum and low transmission power density of UWB signals result in minimal interference with narrowband technologies such as Wi-Fi and Bluetooth operating in the same or similar frequency bands, making UWB a good choice for coexistence with other wireless systems in shared indoor environments [17].

However, despite its ranging performance, UWB faces significant practical barriers to widespread deployment in cost-sensitive applications. UWB hardware is more expensive than BLE alternatives, with transceiver modules and positioning system deployments often costing several times more per device. The energy consumption is also not on par with BLE, as BLE remains more energy efficient [18], [19].

A study reviewing UWB in context of an IPS examined multiple wireless technologies for indoor localization, including Wi-Fi, BLE, ZigBee, and RFID, alongside UWB [20]. This study highlighted that while UWB offers better ranging accuracy, other challenges remain, particularly in non-line-of-sight conditions where multipath effects and signal obstruction degrade performance. To address these problems, the research looked at machine learning approaches. This approach aims to reduce ranging error by learning nonlinear relationships between signal characteristics and actual distance. The results achieved positioning errors below 10 meters in indoor environments [20].

The thesis does not look at UWB because of the drawbacks compared to BLE. The evaluation presented in subsequent chapters focuses on BLE-based distance measurement methods, where cost and ecosystem maturity are more favorable for the intended installer support application.

2.2 System Architectures

The architecture of an indoor positioning system fundamentally shapes its performance characteristics, scalability, and operational requirements. Two primary architectural models exist, centralized server-based systems and distributed mesh-based systems. Each approach offers different advantages and trade offs in terms of computational complexity, infrastructure costs, latency, and suitability for different deployment scenarios. The comparison below examines both architectures and discusses how they influence the design of practical BLE positioning systems[21].

2.2.1 Centralized Server-Based Systems

Many existing IPS implementations use a centralized architecture where BLE beacons transmit signals that are received by multiple access points. A central server then collects measurements and computes device positions [5]. In practice, fixed beacons are deployed at known coordinates throughout the environment and continuously broadcast signals. A mobile scanner, or a set of fixed access points, captures the RSSI values from surrounding beacons, packages the measurements, and forwards them to a remote server for processing [22]. The server stores the incoming data and runs positioning algorithms, such as trilateration combined with Kalman filtering, to estimate the device's location from the aggregated signal strengths. The topology of the network can be observed in figure 2.3.

This approach simplifies position computation and allows sophisticated algorithms to run on powerful centralized hardware. The primary advantages of centralized systems include simplified synchronization between different devices and straightforward position computation logic. However, centralized architectures introduce significant limitations. All BLE nodes must maintain a direct connection to the central server, creating a dependency where the entire positioning system fails if the central server becomes unavailable due to hardware failure, software issues, or maintenance. The centralization of all positioning data on a single server presents both a scalability bottleneck and a security concern, as compromising the central server grants access to all collected measurements and device location information. Additionally, centralized systems incur higher infrastructure costs due to the need for multiple access points and robust network connectivity to ensure the central server can reliably communicate with all nodes in the deployment area [5].

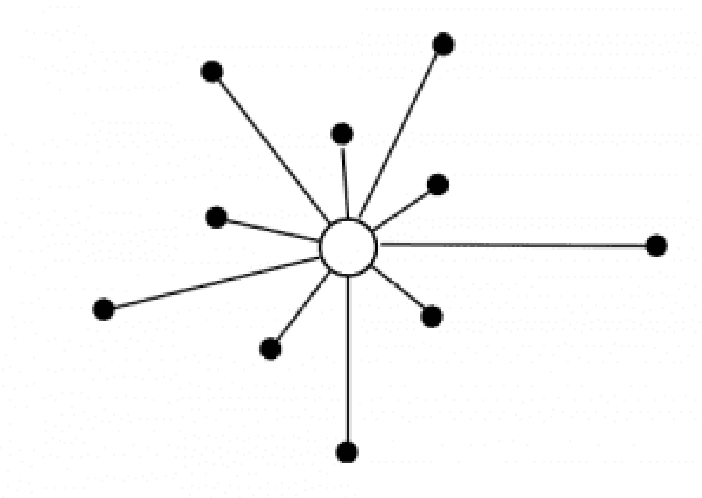

Figure 2.3: An illustration of a centralized network.

2.2.2 Distributed Mesh Architectures

To combat the disadvantages of a centralized system, distributed architectures have been proposed where BLE Mesh nodes perform local position computation based on measurements from neighboring nodes. This approach can reduce latency and infrastructure costs, as it does not require a central server or extensive network connectivity [5], [23].

In a distributed mesh network, positioning functions and computational tasks are distributed evenly among many equal nodes across the deployment area. Rather than relying on a centralized server or a cluster of servers, all measurement processing, computation, and network coordination functions are shared by nodes distributed spatially throughout the indoor environment. This distributed architecture allows nodes to share processing power and work collaboratively, with each node contributing to the overall positioning capability [5], [23]. The topology of the network can be observed in figure 2.4.

Each node in a distributed mesh architecture is responsible for collecting measurements from its neighbors and performing local computations to estimate its own position and that of neighboring devices. The BLE Mesh standard's managed flooding and Publish-Subscribe mechanisms enable nodes to automatically coordinate with each other for position sharing and network updates without requiring centralized orchestration [5], [23].

Distributed mesh architectures offer significant advantages in terms of fault tolerance and scalability. The system is fault-tolerant because any single node can fail independently without impacting the rest of the network. The architecture is scalable because new nodes can be added anywhere in the deployment area at any time and the mesh automatically adds them into the network. Distributed positioning also experiences lower latency than centralized approaches because measurement processing and position computation occur locally at each node rather than requiring communication with a distant central server [5], [23].

However, distributed mesh architectures introduce additional complexity compared to centralized systems. The distributed nature of the system means that synchronization between nodes can be more complex, as nodes must coordinate position estimates across multiple independent local computations. The accuracy of position estimates may be affected by the quality of measurements and the algorithms used for local computation at each node. Additionally, implementing robust distributed positioning algorithms on resource-constrained mesh nodes requires careful algorithm selection and optimization. In an office space this distributed approach is highly desired where the number of nodes can be in the hundreds, and a centralized system would require a large number of access points to ensure coverage [5], [23].

2.3 Reconstruction of Spatial Maps

Reconstructing spatial maps from relative distance measurements is a well-studied problem in the field of sensor networks. The goal is to create a spatial representation of the environment based on pairwise distance estimates between devices. The methods studied in this paper include MDS and SDP and trilateration/multilateration. Each method has its own advantages and limitations in terms of accuracy and computational complexity [23], [24], [25].

2.3.1 Multidimensional Scaling

MDS is a classical approach to sensor network localization. In [26], the authors exploit the low-rank structure of the squared distance matrix to recover node positions from noisy and incomplete measurements. A key theoretical result is that,

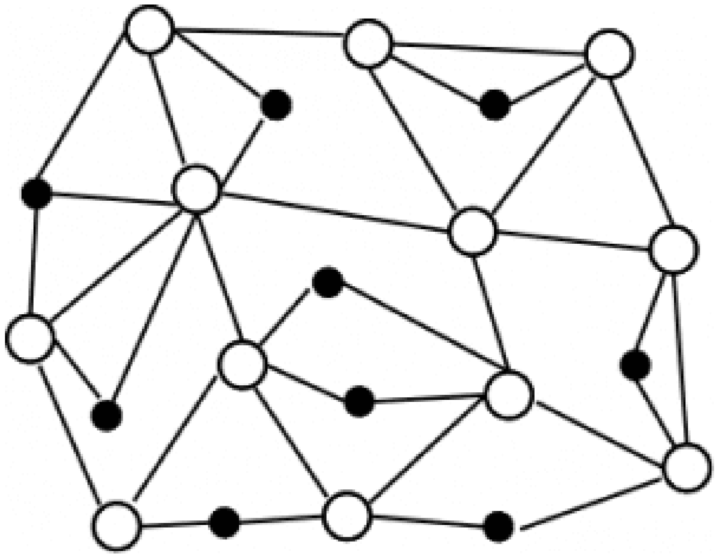

Figure 2.4: An illustration of a distributed network.

in two-dimensional space, the squared distance matrix D has rank at most 4. This means that although D contains n^2 entries, the effective degrees of freedom scale only linearly with n. Their SVD-reconstruct method uses this structure to build an estimator S of the full distance matrix from partial measurements and then computes a best rank-4 approximation S_4 via Singular Value Decomposition. The analysis shows that the average squared reconstruction error per entry decreases proportionally to $1/\sqrt{n}$, providing a formal scaling guarantee that is often missing from purely heuristic localization approaches [26].

An important implication of [26] is that reconstruction quality is not governed only by measurement noise magnitude. The fraction and pattern of missing distance measurements can dominate performance, especially in sparse graphs. This distinction is directly relevant to indoor BLE deployments, where missing links are often structured by walls, occlusions, and connectivity limitations rather than random dropouts.

An alternative MDS formulation that handles missing measurements directly is the Scaling by MAjorizing a COmplicated Function (SMACOF) algorithm [27]. Rather than first reconstructing the full distance matrix, SMACOF minimizes a weighted stress function over the available pairwise distances:

$$\sigma(X) = \sum_{i,j=1}^{n} \omega_{ij}(\hat{d}_{ij} - d_{ij})^2 \tag{2.1}$$

where $\omega_{ij} = 1$ when a measurement between nodes i and j is available, and $\omega_{ij} = 0$ otherwise. Missing measurements are therefore excluded naturally, without a separate completion stage. The algorithm iteratively applies a Guttman-transform update until stress converges [27]. As shown in [27], robustness is strongly affected by three factors: measurement noise, outliers, and the structure of missing data.

In particular, structured missing data can degrade reconstruction more severely than random missing data, and may require different initialization strategies.

Regarding noise, [27] reports that localization error grows approximately linearly with the noise standard deviation σ, with MDS-based methods remaining well behaved at moderate noise levels. Regarding outliers, where some measurements are not merely noisy but entirely wrong, the SMACOF-based approach remains robust up to approximately 40% corrupted measurements, after which performance degrades significantly [27]. Regarding missing data, random missing patterns and range-based missing patterns, where distances beyond a communication range are absent, affect reconstruction differently and require different initialization choices [27]. In particular, range-based missing patterns require Laplacian Eigenmap initialization to avoid fold-over solutions, while classical MDS initialization is suitable for random missing patterns.

Since real indoor BLE deployments naturally produce range-based missing data due to walls, occlusions, and connectivity limits, this initialization distinction is directly relevant for this thesis. This formulation is the one adopted in this work. It should also be noted that the number of nodes used in our experiments is much smaller than in the larger-scale settings typically analyzed in [26]. Consequently, the theoretical $1/\sqrt{n}$ scaling effect is expected to be less pronounced in our test configurations.

2.3.2 Semidefinite Programming

An alternative approach to sensor network localization is SDP [24]. In this framework, localization is written as a graph-realization problem and relaxed to a convex optimization problem with semidefinite constraints. Specifically, the non-convex relationship $Y = X^T X$ is relaxed to $Y \succeq X^T X$, yielding a tractable convex program. In [24], this formulation is analyzed through upper and lower bounds on the objective, which provides a direct way to reason about expected localization error under noisy measurements.

A key result is that the expected objective value increases with the measurement-noise level τ, and empirical results show that positioning error grows approximately with $\sqrt{\tau}$ over relevant noise regimes [24]. This gives a practical noise-to-error interpretation that is useful when setting realistic performance expectations in indoor BLE environments.

A central practical challenge is the high-rank behavior of SDP relaxations under noisy distance constraints. When constraints become inconsistent, the optimizer can move the solution into a higher-dimensional embedding to reduce objective value, and projecting back to 2D then introduces geometric distortion, often visible as crowding toward the center of the configuration [24]. This behavior is affected by geometric conditioning, including anchor placement in anchor-based formulations.

The authors introduce a regularization strategy that discourages folding, followed by a gradient-descent refinement stage initialized from the regularized SDP solution. Their experiments report approximately 30% improvement in positioning accuracy over unregularized SDP, and show that the refinement stage is substantially more reliable when warm-started from the SDP solution than from random initialization [24].

From a noise-tolerance perspective, convex constraints together with regularization can limit how strongly local distance errors propagate into global layout distortion [24]. In practice, this means that SDP often trades tighter pairwise-distance fitting for improved global geometric consistency. This tradeoff is directly relevant to this thesis, where reconstruction quality is evaluated both by pairwise-distance fidelity and by final positional accuracy.

The main drawback is computational cost. The SDP formulation introduces a large number of constraints that grow as $\mathcal{O}(n^2)$ with network size, which can limit applicability in resource-constrained BLE mesh deployments [24]. For the node counts used in this thesis this cost is manageable, but it remains a scalability limitation for larger deployments that the paper [23] addresses.

2.3.3 Trilateration and Multilateration

A common approach to localization in wireless sensor networks is trilateration, which is a geometric method that estimates the position of a device by finding the intersection of circles centred on anchor nodes, where the radius of each circle corresponds to the measured distance between the device and that anchor [28]. This requires a minimum of three non-collinear anchor nodes with known coordinates. However, real-world distance measurements are rarely perfect, and in practice more than three anchor nodes are needed to achieve reliable localization.

To address the sensitivity of trilateration to noisy measurements, multilateration extends the approach by incorporating distance measurements from a larger number of neighboring nodes, depicted in figure 2.5, and applying Maximum Likelihood estimation to minimize the difference between the measured and estimated distances [28]. Comparing trilateration and multilateration based algorithms in a wireless sensor network setting, the paper [28] found that the multilateration based approach achieved lower mean location error, particularly as the number of available anchor nodes increased.

Both methods are well-established and effective in scenarios where anchor nodes with known positions are available. In this work, however, we operate under an anchor-free assumption, as our ad-hoc BLE mesh network does not rely on devices with pre-determined coordinates [28]. Anchor-based approaches such as trilateration and multilateration therefore fall outside the scope of this work but represent a relevant direction for future investigation.

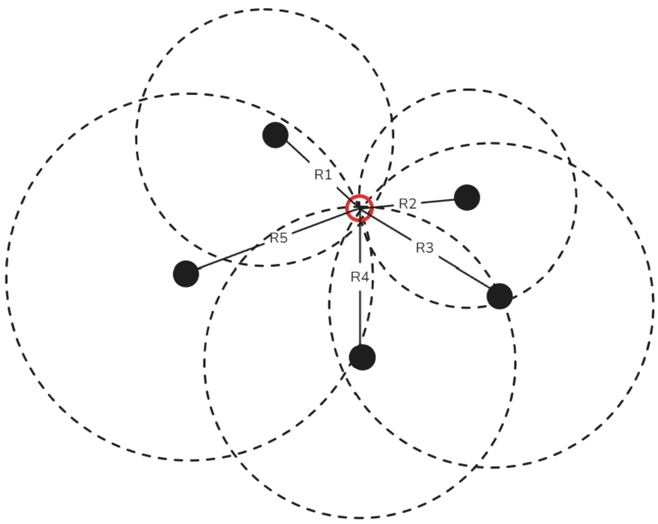

Figure 2.5: An illustration of a multilateration.

Chapter 3

Method & Implementation

This chapter describes the methodology and system implementation developed to build a BLE mesh-based IPS to address the research question.

How can different distance measuring techniques be used to create a 2D indoor positioning system based on BLE Mesh, and what are the accuracy, opportunities, and limitations of such a system in practical applications?

To answer the question, the chapter begins by motivating the methodological approach, followed by the mathematical background and algorithms for signal processing and localization. The system implementation is then described, covering the mesh network protocol, hardware platform, and firmware design. The chapter concludes with the performance criteria used to evaluate the system and the test setups designed to validate it.

3.1 Methodological Approach

The methodology follows a staged process. First, pairwise distance measurements are collected between devices using various ranging techniques. Second, digital signal processing is applied to reduce measurement variance. Third, localization algorithms reconstruct node positions from the pairwise distance matrix. Finally, predefined performance metrics are used to evaluate distance estimation quality and reconstruction quality.

The steps are tested through three experimental setups: line-of-sight measurements, non-line-of-sight measurements for pairwise distance measurements, and a multi-device deployment for localization. This enables evaluation of baseline accuracy, robustness to indoor obstruction, and reconstruction performance.

3.2 Algorithms and Mathematical Background

The following material establishes the theoretical foundation used to transform raw channel sounding observations into reliable position estimates. It first presents the distance estimation methods considered in this work (RTT, PBR, and IFFT), then describes the signal-processing approach used to reduce measurement noise. Finally, it outlines the localization formulations used to reconstruct a 2D device map from pairwise distances, establishing the mathematical basis for the experimental evaluation in later chapters.

3.2.1 Distance Estimation Methods

The distance estimation methods were selected to align with the thesis objective of evaluating newer Bluetooth 6.0 Channel Sounding capabilities and their practical potential for improved indoor accuracy. Although RSSI-based ranging is well established and widely available, it has already been extensively studied and is known to be highly sensitive to environment-dependent effects such as multipath and attenuation. This thesis focuses on RTT, PBR, and IFFT, where current evidence is promising but practical mesh-deployment behavior is still less explored.

3.2.1.1 Round-Trip Time (RTT)

RTT estimates distance by measuring the elapsed time for a signal exchange between two devices and converting that delay to range using radio-wave propagation speed. In this thesis, one device initiates a channel sounding exchange and the peer responds, yielding a round-trip timing measurement that is mapped to distance [4]. The main advantage of RTT is that it can operate with relatively low frequency-domain complexity compared with phase-sweep approaches. However, practical accuracy is limited by processing delays, timestamp resolution, and multipath propagation, especially in indoor NLOS conditions where reflected paths can bias delay estimates [7], [8], [10].

3.2.1.2 Phase-Based Ranging (PBR)

PBR estimates distance from phase differences measured across multiple carrier frequencies during channel sounding. Because phase shift is proportional to propagation delay, combining measurements over a frequency sweep enables the delay, and therefore distance, to be inferred more precisely than single-channel strength-based methods [4]. Using multiple frequencies is important to resolve phase ambiguity and extend the unambiguous ranging interval. Under favorable LOS conditions, PBR can provide high accuracy, but its performance degrades when multipath and obstruction distort phase consistency across the swept bandwidth [10], [12].

3.2.1.3 Inverse Fast Fourier Transform (IFFT)

The IFFT method uses the same complex channel sounding measurements as phase-based ranging, but transforms them from the frequency domain to an impulse-response representation in the delay domain. Distance is then estimated from the dominant delay peak, which corresponds to the most likely direct propagation path in LOS conditions [13], [16]. In this work, IFFT is treated as a separate ranging algorithm because it extracts timing information through peak detection rather than direct phase-slope estimation. This can improve robustness in some scenarios, but indoor multipath still introduces spurious peaks and broader delay spread, which increases ranging error in NLOS environments [10], [16].

3.2.2 Digital Signal Processing

To address the precision requirements of the ranging methods, signal processing techniques were applied to the raw distance estimates. Since the devices are assumed to be stationary throughout all measurements, the Digital Signal Processing (DSP) requirements are significantly simplified. The measured distance can be modeled as:

$$d(t) = d_0 + \eta(t), \quad \eta(t) \sim \mathcal{N}(0, \sigma^2) \tag{3.1}$$

where:

- $d(t)$ is the observed distance at time t,
- d_0 is the true (constant) distance,
- $\eta(t)$ is the additive noise, modeled as a Gaussian random variable with zero mean and variance σ^2.

This model explains the measured distance as the sum of the true physical distance and measurement noise. Several filtering approaches can be applied to raw distance estimates. The moving average algorithm offers simplicity and effectiveness at smoothing high-frequency noise by computing the mean of samples within a window. This approach leverages all available measurements to produce a smoothed estimate [29].

The approach adopted in this thesis uses mean filtering, which provides noise reduction through averaging. A data structure accumulates distance measurements from multiple channel sounding procedures, storing estimates from each available ranging technique independently. The sample window was chosen to contain 20 elements from each distance measuring method. For each method, all measurements in the window are averaged to produce the filtered distance estimate. This approach utilizes all available measurements equally [29].

The algorithm operates as follows:

1. **Window Collection:** Accumulate N distance samples into a buffer: $\{d(1), d(2), \ldots, d(N)\}$. Each sample represents a distance estimate from a completed channel sounding procedure.
2. **Filtering by Technique:** Separate samples by their source technique (IFFT, phase-based ranging, RTT). Only finite values (valid measurements) are retained for processing.
3. **Mean Computation:** For each technique, compute the arithmetic mean of the N samples as the filtered distance estimate for that technique.

The resulting mean estimator is:

$$\tilde{d} = \frac{1}{N} \sum_{i=1}^{N} d(i) \tag{3.2}$$

where $d(i)$ denotes the i-th sample in the window.

3.2.3 Localization Algorithms

After the distance measurements are obtained and filtered, the next step is to estimate the positions of the devices in the environment. The methods MDS and SDP are implemented and compared. Both methods operate in an anchor-free setting, requiring only pairwise distance estimates as input and produce a 2D reconstruction of the environment where each device is represented as a point in the plane. The two localization methods were selected to provide a controlled comparison between complementary optimization paradigms that are both established in the literature. For MDS, this thesis uses the SMACOF variant, which minimizes a weighted stress objective and is therefore well suited to sparse or incomplete pairwise-distance matrices while remaining computationally efficient for iterative reconstruction [27]. For SDP, we use a centered Gram-matrix relaxation in which squared-distance constraints are enforced through a positive-semidefinite formulation, followed by rank-2 projection to recover coordinates [24]. This pairing enables a direct evaluation of the practical trade-off between computational cost and reconstruction fidelity in BLE mesh-based indoor positioning.

3.2.3.1 Multidimensional Scaling

The MDS-based localization is implemented with the SMACOF procedure following [27]. Let $\hat{D} = [\hat{d}_{ij}]$ denote the measured distance matrix and $X \in \mathbb{R}^{n\times 2}$ the unknown node coordinates. SMACOF minimizes the weighted stress criterion

$$\sigma(X) = \sum_{i=1}^{n}\sum_{j=1}^{n} \omega_{ij}\left(\hat{d}_{ij} - d_{ij}(X)\right)^2, \tag{3.3}$$

where $d_{ij}(X) = \|x_i - x_j\|_2$ is the Euclidean distance in the current embedding and ω_{ij} is a binary mask ($\omega_{ij} = 1$ for available measurements, $\omega_{ij} = 0$ otherwise). This weighting allows direct handling of missing distances by excluding unavailable pairs from the objective. Before optimization, the distance matrix is symmetrized by averaging reciprocal entries when both are present. Starting from an initial configuration, SMACOF applies iterative Guttman-transform updates until convergence, defined by a stress decrease below a fixed tolerance or by a maximum iteration limit. The tolerance was set to 10^{-6} and the maximum number of iterations was set to 400.

3.2.3.2 Semidefinite Programming

The SDP formulation follows [24]. Instead of estimating coordinates directly, it estimates a Gram matrix $G \in \mathbb{R}^{n\times n}$ constrained to be positive semidefinite ($G \succeq 0$). Squared distances are represented linearly in G:

$$\hat{d}_{ij}^2 \approx G_{ii} + G_{jj} - 2G_{ij}. \tag{3.4}$$

The optimization minimizes the residual between measured and modeled squared distances over all known pairs:

$$\min_{G\succeq 0} \sum_{(i,j)\in\mathcal{E}} \left(\hat{d}_{ij}^2 - (G_{ii} + G_{jj} - 2G_{ij})\right)^2, \tag{3.5}$$

subject to a centering constraint,

$$G\mathbf{1} = \mathbf{0}, \tag{3.6}$$

which removes translational ambiguity. After solving the relaxed problem, node coordinates are recovered by rank-2 projection of G. If $G = U\Lambda U^\top$ is the eigendecomposition, the 2D embedding is obtained as

$$X = U_2\Lambda_2^{1/2}, \tag{3.7}$$

where U_2 contains the two dominant eigenvectors and Λ_2 the corresponding eigenvalues.

3.3 System Implementation

The proposed IPS is implemented using a BLE mesh platform. The implementation first presents the mesh network and application-layer packet design used to coordinate measurements and data exchange, followed by the orchestrator logic used to schedule pairwise ranging sessions across devices. It then details the hardware and firmware implementation on the nRF54L15-DK, including the software state-machine behavior that enables switching between mesh communication and channel sounding operations.

3.3.1 Mesh Network Structure

The NRF SDK provides built-in support for BLE mesh connectivity, abstracting the underlying protocol stack and exposing a high-level Application Programming Interface (API) for sending and receiving messages. However, since the SDK operates at the transport layer, a custom application-layer protocol was developed to structure and interpret the data exchanged between devices. The devices were part of a fully distributed BLE mesh network with no central device to control how the network operated. To design the application level protocol, the first step was to quantify the abilities each node needed to have; we identified the following as the most crucial:

- **Distance measurement** — each node must be able to determine its distance to neighboring nodes.
- **Information relaying** — nodes must forward data across the network so that measurements propagate beyond direct neighbours.
- **Diagnostics** — nodes must support self-diagnostic tasks such as relaying debug information.
- **Control configuration** — nodes must be able to send configuration information such as which audio to play and when.

With this in place we designed an application protocol for relaying the information, as seen in table 3.1. In a distributed system like this, the nodes do not know which other nodes are part of the network. It is therefore necessary to include information in the packets such as sender id and the id of the intended receiver. To perform distance measuring using channel sounding two devices have to assume two different roles to begin the distance exchange. One node has the role of initiator and the other as reflector. In order for the nodes to know which role they have to assume they read from the assume_role field in the packet.

Field	**Description**	**Size**
`sender_id`	Identifier of the sender device	1 byte
`receiver_id`	Identifier of the recipient device	1 byte
`packet_type`	Type of packet (control or data)	1 byte
`payload`	Packet content depending on type	Variable

Table 3.1: Main packet structure

Field	**Description**	**Size**
`assume_role`	Role assignment (initiator or reflector)	1 byte
`sound_index`	Index of acoustic sound to emit	1 byte
`start_time`	Time to start transmission	2 bytes

Table 3.2: Control packet payload

Field	**Description**	**Size**
`rssi`	Received signal strength indicator	2 bytes
`pbr`	Phase-based ranging distance estimate	2 bytes
`rtt`	Round-trip time distance estimate	2 bytes
`ifft`	IFFT-based ranging distance estimate	2 bytes
`distance_to_id`	Identifier of target device	1 byte

Table 3.3: Data packet payload

In a BLE mesh network, devices may assume different roles in the system. For overall system simplicity, each node can assume the role of Proxy and Relay. This gave the devices the abilities of both roles. This was deemed essential since all devices had to be able to relay and send messages further into the network as well as communicate with a non-mesh device like a phone.

To support these roles and enable efficient distance measurement coordination, the application protocol defines two primary packet types: control packets and data packets. Control packets are used to synchronize measurement sessions and assign initiator/reflector roles, while data packets carry the measured distances computed using different techniques.

Control Packets

Control packets are used to synchronize measurement sessions and assign initiator/reflector roles. These packets are always sent out before a ranging test in order to configure the nodes in the network. As seen in table 3.2, the structure of control packets includes the following fields:

`assume_role`

The assume_role field assigns each node a role (initiator or reflector) for the measurement session. This field is essential for coordinating distance measurements between device pairs.

`sound_index`

The sound_index field carries acoustic sound configuration parameters, specifying which sound to emit during the measurement. This allows the orchestrator to control which signal is used for distance estimation.

`start_time`

The start_time field specifies when transmission should begin, enabling synchronized measurement sessions across the network. Additionally, control packets include session and transaction identifiers for tracking purposes and message-type information to support various control operations.

Data Packets

Data packets convey the results of completed measurements. Once a ranging exchange concludes, the data packet contains multiple distance estimates computed using different techniques. All distances are expressed in centimeters. As seen in table 3.3, the structure of data packets includes the following fields:

`rssi`

The rssi field contains the Received Signal Strength Indicator, a basic measure of signal strength used in distance estimation.

`pbr`

The pbr field contains the Phase-Based Ranging distance estimate, representing one of the ranging techniques used to calculate distance between devices.

`rtt`

The rtt field contains the Round-Trip Time distance estimate, another ranging technique that measures the time for a signal to travel to a device and back.

`ifft`

The ifft field contains the Inverse Fast Fourier Transform-based ranging distance estimate, a signal processing-based technique for distance measurement.

`distance_to_id`

The distance_to_id field contains the identifier of the target peer, allowing the receiver to determine which device the measurements correspond to. The data packet further includes session and transaction identifiers for measurement tracking, status information, and diagnostic reason codes.

When sending messages through the network there are a number of mechanisms making sure the network doesn't become congested. One feature is the Time To Live (TTL) header in each packet. For our experiments default TTL of 7 was used. This essentially means that each packet may only travel through 7 other nodes in the network. This mechanism is essential to stop packets which may be relayed multiple times through the network. When a packet travels through the network it may take multiple paths through the network as there may exist multiple ways to travel from node A to node B. This means that when a node gets a packet it may ignore multiple other duplicates which may have traveled through other paths.

In order to create the mesh network, each device had to be provisioned on first boot. This process added the device to the network and transferred the necessary keys: the network key, used to authenticate devices within the network, and the application key, used to encrypt application-layer communication between nodes. Once provisioned, a device retained these keys in non-volatile storage and rejoined the network automatically on subsequent reboots. After this step the devices were ready to begin communication. To provision the devices and network we used the nRF Mesh mobile application [30].

3.3.2 Orchestrator

In order to facilitate a full network with synchronization between a number of combinations we developed our own algorithm to support this mechanism. Many approaches were discussed from a distributed approach to a centralized one. We opted for a centralized approach since it was simpler to implement and we had a limited number of devices. The algorithm is based on the idea of an orchestrator which is responsible for coordinating the measurement sessions across the network, ensuring that nodes assume their assigned roles and that measurements are conducted in a synchronized manner.

It starts by the orchestrator sending a discovery packet broadcast on the network to find out what devices are part of the network. This is a consequence of working on a pub-sub network where the devices do not know which other devices are part of the network. Once the orchestrator has a list of all devices in the network it can start sending control packets to coordinate the measurement sessions. The orchestrator sends control packets to pairs of devices, assigning one as the initiator and the other as the reflector for each session.

The orchestrator iterates through all possible pairs of devices in the network, sending control packets to each pair to initiate measurements. This process continues until all pairs have been covered. The orchestrator also handles the collection of data packets from the devices, aggregating the distance estimates for analysis. This flow is visualized in figure 3.1.

This is quadratic time complexity that is not scalable for large networks but for our experiments it was sufficient and results in the largest amount of data gathered.

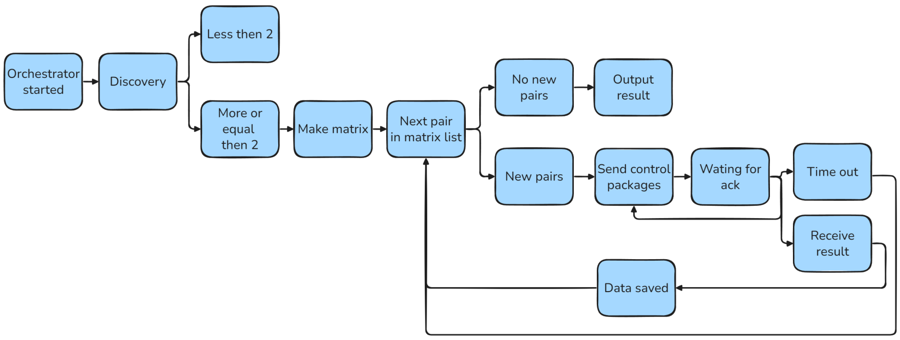

Figure 3.1: Orchestrator Algorithm Overview

3.3.3 Hardware Platform: nRF54L15-DK

The nRF54L15-DK [31] is a development kit designed for prototyping and testing applications using the nRF54L15 System on Chip (SoC) from Nordic Semiconductor. The nRF54L15 is a highly integrated SoC that supports Bluetooth 6.0 and channel sounding, making it suitable for developing BLE mesh networks. Figure 3.2 depicts the development kit used.

To use the nRF54L15-DK, we employed the Zephyr Real Time Operating System (RTOS) to implement our code. The application software was inspired by examples found at [32] and [33]. Zephyr RTOS was chosen because it was recommended by Nordic Semiconductor and already included development tools and support for the nRF SDK required for sound channeling and BLE Mesh development. The nRF54L15-DK has one pcb antenna and for the tests its configured with maximum antenna power set to +8 dBm.

3.3.4 Software state machine

The custom firmware developed for the nRF54L15-DK development kit has two main functions. The first function is to maintain a connection to the mesh network, this includes sending, receiving and relaying messages. The second function is to

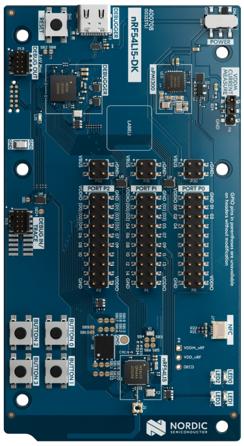

Figure 3.2: nRF54L15-DK Credit: www.nordicsemi.com

perform distance measurements to nearby nodes. Since both of these technologies are built on BLE the program has to switch between these functionalities as they cannot be enabled at the same time. The logic sequence of the program can be seen in figure 3.3.

The states of the nRF devices start by starting up and provisioning on the mesh network. Once the device is provisioned it waits in an IDLE mesh state for packages. It can receive four types of packages: discovery, control, play and configure. Play triggers a sound on the speaker and configure changes the sound configuration. The most important one is the control packet which triggers a measurement session with another device. The device reads the control packet to determine which role it has to assume, either initiator or reflector, and then perform the necessary steps to complete the measurement session. Once the session is completed the device returns to the IDLE state waiting for new packages.

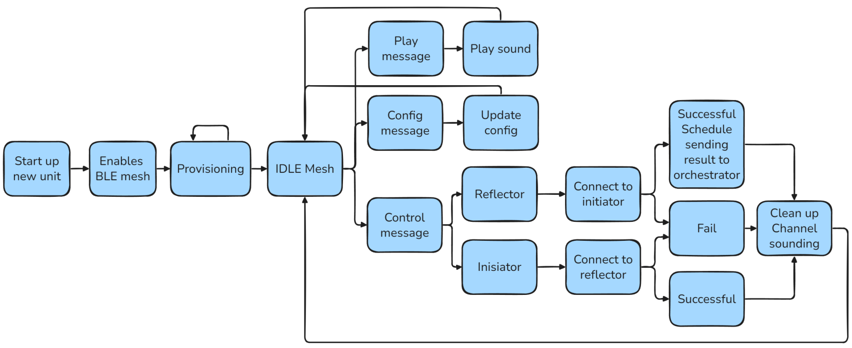

Figure 3.3: Software State Machine Overview

3.4 Evaluation Metrics

In order to compare different distance measuring methods, we define a common set of evaluation metrics. The main metrics considered were the following:

Metric	Description	Quantification (Unit)
Accuracy	Estimated distance matches actual physical distance (ground truth) with minimal error.	Mean Absolute Error (MAE) / RMSE (m)
Precision	Measurement consistency and repeatability across multiple samples in the same conditions.	error reduction (%)
Robustness	Performs reliably across different indoor environments with varying signal obstruction.	RMSE degradation (%)

Accuracy is a fundamental metric for an IPS, as can be seen in multiple studies [2], [10], [34], [35]. High accuracy ensures that the final reconstruction closely resembles the real environment. Robustness is equally important since indoor environments are often non-ideal, with signal reflections and varying device orientations degrading distance estimates. By measuring robustness degradation between line-of-sight and obstructed conditions, we can characterize how methods handle realistic deployment scenarios [2]. Finally, precision ensures measurement consistency and repeatability, which is fundamental for obtaining stable distance estimates that can reliably support the localization process [35], [36].

3.4.1 Localization Evaluation Metrics

Two distinct error types are used to evaluate reconstruction quality, and it is important to distinguish between them.

Known-distance error measures how well the reconstructed coordinates reproduce the observed pairwise distances. For every pair (i, j) with a valid measurement $\hat{d}_{ij}$, the residual is $r_{ij} = \|x_i - x_j\|_2 - \hat{d}_{ij}$. This yields:

$$\mathrm{MAE}_{\mathrm{known}} = \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} |r_{ij}|, \qquad \mathrm{RMSE}_{\mathrm{known}} = \sqrt{\frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} r_{ij}^2}, \tag{3.8}$$

where $\mathcal{E}$ is the set of known pairs. This is a fit-to-input metric, not a ground-truth accuracy measure [24], [27].

Localization error measures positional accuracy relative to ground truth. Because both algorithms are anchor-free, raw reconstructions have an arbitrary pose (rotation, translation, and reflection). Before computing positional errors, each estimated map is therefore aligned to the reference using a rigid 2D [37] transformation comprising rotation, translation, and optional reflection, but no scale change that minimizes point-wise discrepancies. Preserving scale is critical because any scale correction would mask systematic distance errors. Visualization overlays in this thesis additionally use a similarity alignment (which also optimizes scale) for cleaner shape comparison, but all quantitative metrics are computed from the rigid no-scale alignment only [10], [36] .

After rigid alignment, the per-node Euclidean error is $e_i = \|\tilde{x}_i - x_i^\star\|_2$, where $\tilde{x}_i$ is the aligned estimate and $x_i^\star$ the ground-truth position. MAE captures the average positional offset, RMSE penalizes large individual errors more heavily, and MAX identifies the worst-case node placement.

To characterize reconstruction uniformity across nodes, the standard deviation of per-node positional errors is additionally reported:

$$\sigma_e = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(e_i - \bar{e})^2}, \qquad \bar{e} = \frac{1}{n}\sum_{i=1}^{n} e_i. \tag{3.9}$$

A lower σ_e indicates a more uniformly distributed reconstruction error across devices.

3.5 Test Setups

Three test setups were designed to systematically evaluate the metrics established earlier: Accuracy, Robustness and Scalability. All tests were conducted in an office environment within corridors. Ground truth measurements were obtained using a *Bosch GLM 40* laser distance measurer.

Figure 3.4: Bosch PRO GLM 40 Credit: bosch-professional.com

3.5.1 Test Locations

The first two tests were conducted in two different corridors within a typical office environment. The first corridor provided unobstructed line-of-sight propagation up to approximately 40 meters, as shown in Figure 3.5. The second corridor had similar physical characteristics but included a meeting room with glass and drywall partitions that obstructed the direct signal path between measurement points. The second environment can be seen in Figure 3.6.

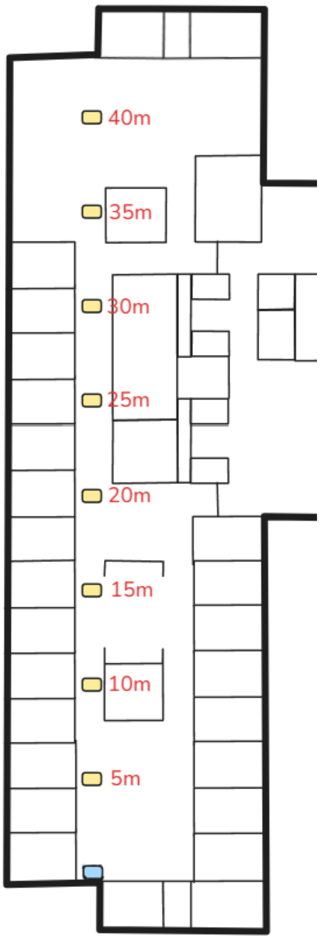

Figure 3.5: Line-of-sight test corridor with an unobstructed propagation path up to approximately 40 meters.

Test 1: Accuracy Baseline (Line-of-Sight)

The first test assessed measurement accuracy under ideal LOS conditions. Two mesh-enabled devices measured distances at 8 fixed distances in direct LOS: 5, 10, 15, 20, 25, 30, 35 and 40 meters. For each distance, multiple measurements were collected from all available ranging techniques (PBR, RTT, and IFFT). Accuracy was verified using absolute error plots and Cumulative Distribution Func-

Figure 3.6: Non-line-of-sight test corridor where glass and drywall partitions obstruct the direct signal path.

tion (CDF) analysis to characterize the distribution of measurement errors. It is illustrated in figure 3.5.

Test 2: Robustness (Non-Line-of-Sight)

The second test evaluated robustness by repeating measurements at the same 6 first distances under NLOS conditions, where signal propagation was obstructed by walls and other indoor obstacles. Robustness was assessed by comparing the absolute error distribution and CDF plots between NLOS and LOS conditions across all distances and techniques, quantifying the percent degradation in measurement performance. It is illustrated in figure 3.6.

Test 3: Multi-Device 2D Reconstruction

The third test evaluated whether the ranging measurements obtained from the previous tests can be used to reconstruct the 2D spatial layout of a multi-device deployment. Eight mesh-enabled devices were placed at known positions across the office environment in a topology designed to include both LOS and NLOS node pairs. Each device attempted to measure its distance to all reachable peers using the orchestrator described in section 3.3.2, yielding up to 28 unique pairwise distance estimates.

Chapter 4

Results and Evaluation

This chapter presents and evaluates the outcomes from the three test setups described in chapter 3. It first shows results from test setups under section 3.5.1, that compares distance estimation performance for PBR, RTT and IFFT under both LOS and NLOS conditions. Comparisons are made using error statistics and cumulative distributions and then analyzes method-specific behavior and robustness. It then presents the multi-device network experiment, where IFFT-based pairwise distances are used for 2D reconstruction with MDS and SDP. Evaluation of pairwise fit, positional accuracy, reconstruction uniformity and scaling implications. All measurements were conducted using Nordic Semiconductor nRF54L15 Development Kits 3.3.3.

4.1 Distance Estimation

4.1.1 Line-of-Sight

Results from test setup 3.5.1 are reported here, where two mesh-enabled devices measured distances in direct line of sight at various distances. Measurements were conducted in an office corridor environment at distances of 5, 10, 15, 20, 25, 30, 35, and 40 meters. The estimates are compared against ground truth measurements obtained using a Bosch GLM 40 laser distance measurer. The devices were configured with maximum antenna power set to +8 dBm for all tests to ensure consistent transmit conditions across measurements.

4.1.1.1 Distance Measurements

Figure 4.1 presents the measured distances for each ranging technique plotted against the ground truth distances obtained from the laser distance measurer. The plot displays measurements across all seven test distances for the three ranging techniques: PBR, RTT, and IFFT. Each point represents a measurement, with the ideal case being all points aligned along the diagonal (measured distance = ground truth distance). Deviations from this diagonal indicate measurement error. Deviations from the true distance can be more easily analyzed by plotting said deviation, this can be seen in Figure 4.2.

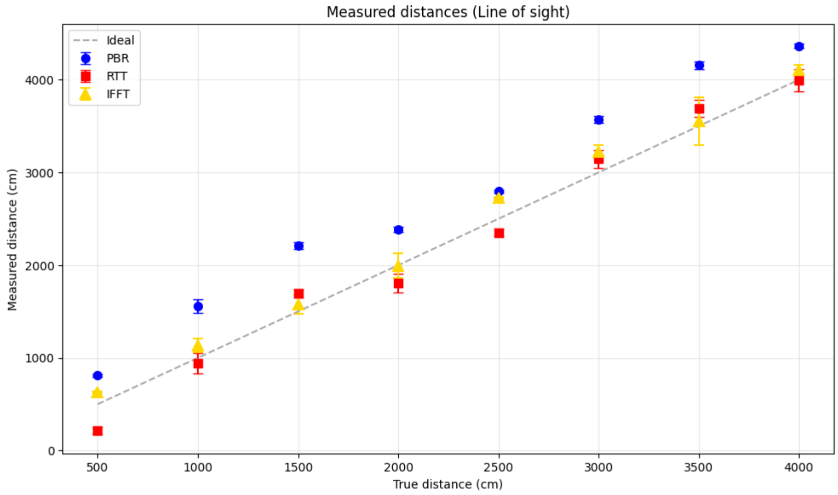

Figure 4.1: Measured distances (Line of sight)

4.1.1.2 Accuracy Analysis

Figure 4.2 shows the error for each ranging technique across all tested distances in line-of-sight conditions. The error is calculated as the difference between the measured distance and the ground truth distance. Mean error values from this dataset are 4.82 m for PBR, −0.19 m for RTT and 1.19 m for IFFT. To complement the distance-dependent error view, the cumulative distribution function is presented in Figure 4.3.

Figure 4.3 presents the cumulative distribution function of errors for each ranging technique under line-of-sight conditions. This representation shows the proportion of measurements that fall within a given error threshold for each method.

Table 4.1 summarizes minimum, maximum, mean, and standard deviation of the error distributions for all three methods in the line-of-sight scenario.

Table 4.1: Results for Line-of-Sight Tests (m)

Metric	PBR Error	RTT Error	IFFT Error
Min	2.96	−2.85	−0.07
Max	7.13	1.98	2.26
Mean	4.82	−0.19	1.19
Std Dev	1.63	1.84	0.80

4.1.2 Non-Line-of-Sight

The second test setup showcase in 3.5.1 reports measurements collected under non-line-of-sight conditions at the same distances as the line-of-sight test. This enables evaluation of the robustness of the positioning methods under obstructed

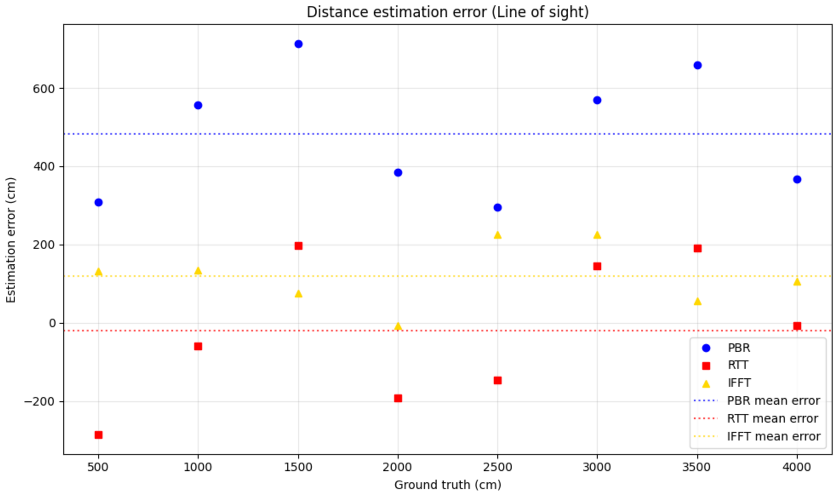

Figure 4.2: Distance estimation error (Line of sight)

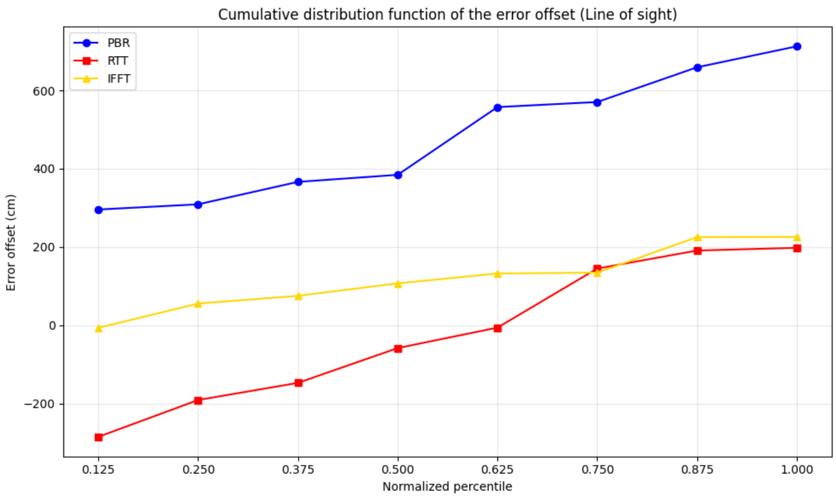

Figure 4.3: Cumulative distribution function (Line of sight)

signal conditions.

4.1.2.1 Distance Measurements

Figure 4.4 displays the measured distances for each ranging technique under non-line-of-sight conditions plotted against ground truth. The presence of signal obstruction (glass and drywall partitions) introduces additional spread relative to the ideal diagonal line.

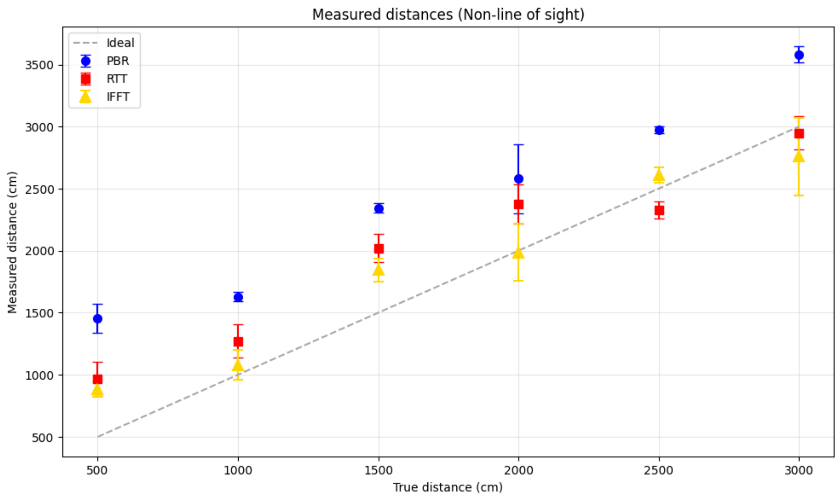

Figure 4.4: Measured distances (Non-line of sight)

4.1.2.2 Accuracy Analysis

These tests were conducted under non-line-of-sight conditions up to 30 meters (5, 10, 15, 20, 25, 30 meters) as the line-of-sight tests. Signal obstruction was provided by glass and drywall partitions within the office corridor environment. Figure 4.5 presents the error for each ranging technique across all tested distances in non-line-of-sight conditions. The error is calculated as the difference between the measured distance and the ground truth distance obtained using the laser distance measurer.

PBR produced a mean error of 6.77 m. RTT produced a mean error of 2.36 m. IFFT produced a mean error of 1.11 m. Detailed error statistics including minimum, maximum, and standard deviation values are presented in Table 4.2.

Figure 4.6 presents the cumulative distribution function of errors for each ranging technique under non-line-of-sight conditions. This representation shows the proportion of measurements that fall within a given error threshold.

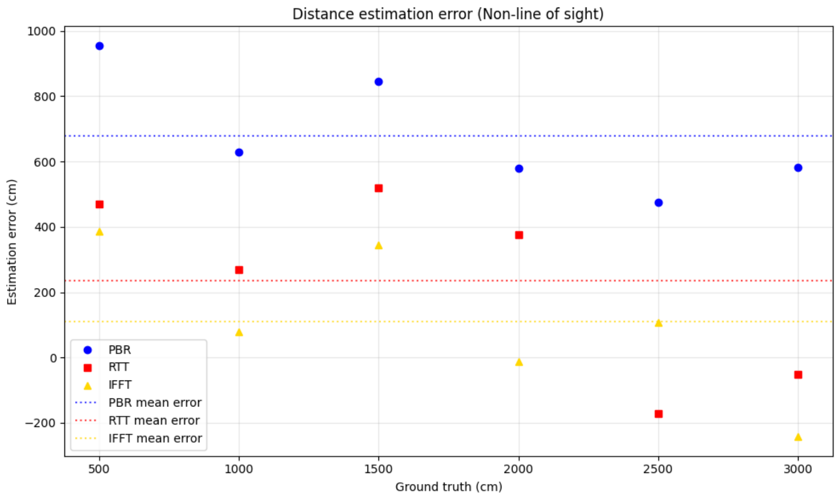

Figure 4.5: Distance estimation error (Non-line of sight)

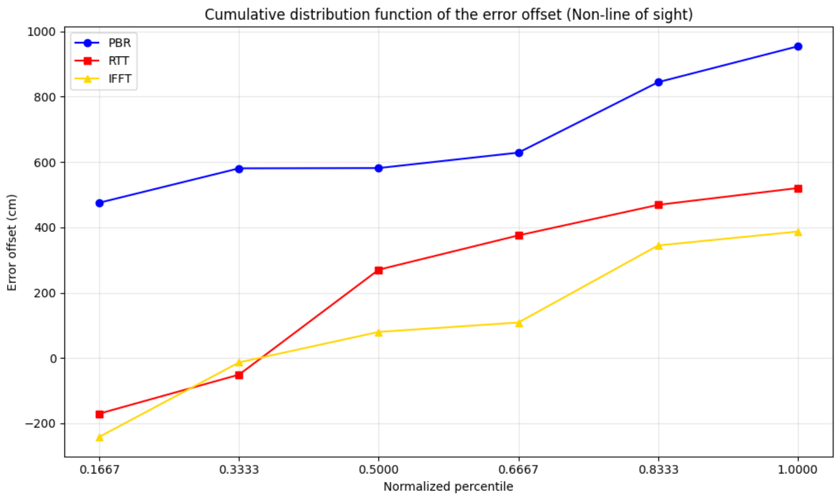

Figure 4.6: Cumulative distribution function (Non-line of sight)

Table 4.2: Results for Non-line of Sight Tests (m)

Metric	PBR Error	RTT Error	IFFT Error
Min	4.75	−1.70	−2.41
Max	9.54	5.20	3.87
Mean	6.77	2.36	1.11
Std Dev	1.82	2.84	2.33

4.1.3 Comparative Method Evaluation

To evaluate overall method performance beyond per-distance mean errors, mean absolute error (MAE) and root mean squared error (RMSE) are calculated across all measurements. While RTT exhibits a mean error very close to zero under LOS conditions, this metric alone does not explain the measurement's reliability. High variance in measurements can occur despite a near-zero mean, resulting in unpredictable errors across individual measurements. To account for this limitation and better represent overall method performance, both metrics are considered, as RMSE penalizes larger deviations more heavily and thus better reflects measurement consistency and reliability. Table 4.3 presents these metrics for the LOS scenario.

Table 4.3: Line-of-Sight Ranging Method Performance Comparison (m)

Method	MAE (m)	RMSE (m)
PBR	4.82	5.05
RTT	−0.19	1.73
IFFT	1.19	1.40

The RMSE values reveal that despite RTT's near-zero mean error, the IFFT method achieves the lowest RMSE of 1.40 m, followed by RTT with 1.73 m. This indicates that IFFT produces more consistent and reliable distance estimates with lower overall measurement variance. The substantially higher RMSE for PBR confirms its poor performance across all evaluation metrics.

The same evaluation metrics are applied to the non-line-of-sight measurements. Table 4.4 presents the MAE and RMSE values for each method in the NLOS environment.

Table 4.4: Non-Line-of-Sight Ranging Method Performance Comparison (m)

Method	MAE (m)	RMSE (m)
PBR	6.77	6.98
RTT	2.35	3.50
IFFT	1.11	2.40

The NLOS results maintain the same ranking as observed in LOS conditions, with IFFT achieving the lowest error followed by RTT and PBR having the highest error. Comparing LOS and NLOS performance reveals significant performance degradation across all methods. IFFT shows a 71.1% increase in RMSE, RTT increases by 102.4%, and PBR increases by 38.1% 4.5. Looking at these results it is clear that IFFT remains the most robust method under obstructed signal conditions, suggesting superior resilience to multipath propagation and signal obstruction.

4.1.3.1 Robustness Analysis

To quantify robustness as defined in the evaluation criteria, the performance degradation from line-of-sight to non-line-of-sight conditions is calculated for each method. Table 4.5 presents the percentage increase in RMSE when transitioning from LOS to NLOS environments.

Table 4.5: Method Robustness: RMSE Degradation from LOS to NLOS

Method	RMSE Degradation (%)
PBR	38.1%
IFFT	71.1%
RTT	102.4%

The analysis reveals that all methods experience increased error in NLOS conditions but the relative impact differs significantly between methods. PBR demonstrates the lowest absolute degradation but it maintains poor absolute accuracy in both conditions. IFFT shows moderate degradation but retains the best absolute performance in NLOS conditions. RTT experiences the most severe degradation, more than doubling its error when transitioning to obstructed environments. This analysis indicates that while IFFT may have higher relative degradation than PBR, its superior absolute accuracy in both conditions and more moderate degradation than RTT make it the most practical choice for indoor positioning systems where NLOS scenarios are common.

These consolidated views enable direct assessment of the accuracy-precision-robustness tradeoffs for each method across all evaluation criteria.

4.1.3.2 Method-Specific Performance Characteristics

Each ranging method exhibits distinct performance characteristics rooted in its susceptibility to environmental factors. Understanding these characteristics provides insight into the observed measurement behaviors.

Phase-Based Ranging (PBR)

Throughout the tests, PBR performs significantly worse compared to the other two methods, exhibiting the highest distance estimation errors in both LOS and

NLOS conditions. The root cause of this poor performance lies in the fundamental approach used by PBR. The method sends out multiple 2.4 GHz waves at different frequencies to estimate the phase shift between transmission and reception. However, in indoor environments with complex multipath propagation, this phase-based approach produces substantially less accurate results than frequency-domain methods.

The degradation of PBR performance can be attributed to phase ambiguity and multipath effects. When signals reflect off surfaces in an indoor environment, the received signal becomes a superposition of direct and reflected paths, each with its own phase shift. This multipath interference causes ambiguity in phase measurement, making it difficult to accurately determine the true distance. Additionally, the relatively low bandwidth of the 2.4 GHz signal limits the ability to distinguish between multiple reflection paths, further degrading accuracy.

Figures 4.7, 4.8, 4.11, 4.9, 4.10, and 4.12 present IQ constellation plots and magnitude spectra from PBR measurements in various indoor and outdoor conditions. These visualizations illustrate the signal characteristics that underlie the performance differences observed.

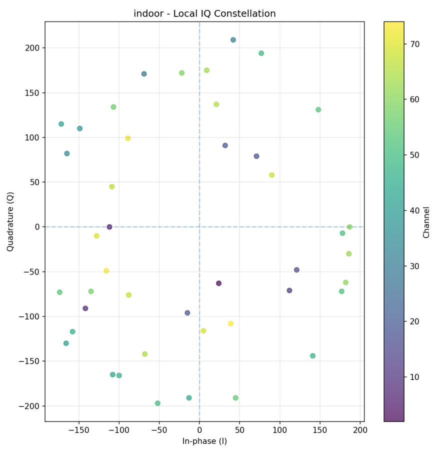

Figure 4.7: IQ Constellation: Indoor Local

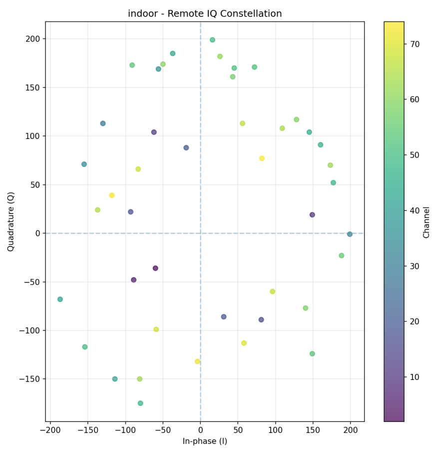

Figure 4.8: IQ Constellation: Indoor Remote

The IQ plots reveal the relationship between in-phase and quadrature components of the received signal. In an outdoor environment with minimal multipath, the constellation points form tight clusters which indicates consistent phase relationships. However, in an indoor environment where multipath interference is much stronger the points become more scattered. Plotting the magnitude also shows how signals change with frequency, this shows that indoor environments show broader and more diffuse characteristics compared to outdoor environments.

This demonstrates that PBR has inherently more problems in estimating distance indoors since the noise present creates phase ambiguity and measurement scatter. The IQ plots in indoor conditions demonstrate that the phase relationships become unreliable, which makes it difficult to extract accurate distance estimates.

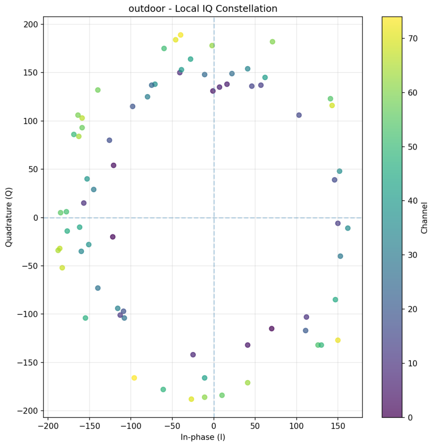

Figure 4.9: IQ Constellation: Outdoor Local

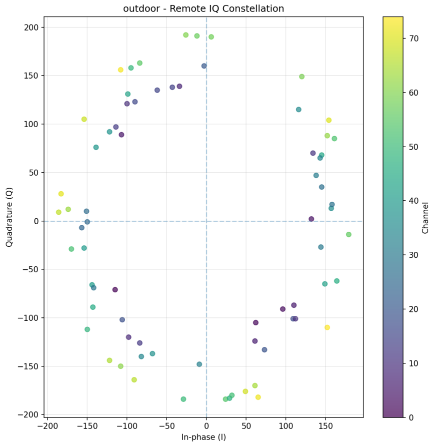

Figure 4.10: IQ Constellation: Outdoor Remote

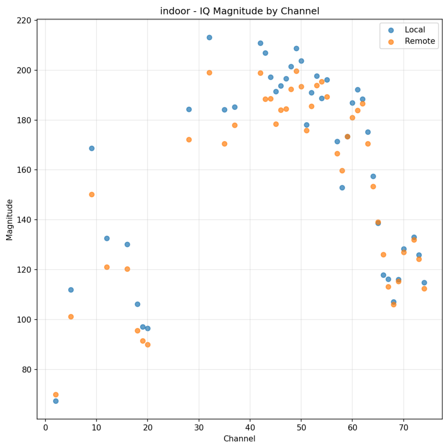

Figure 4.11: Magnitude Spectrum: Indoor

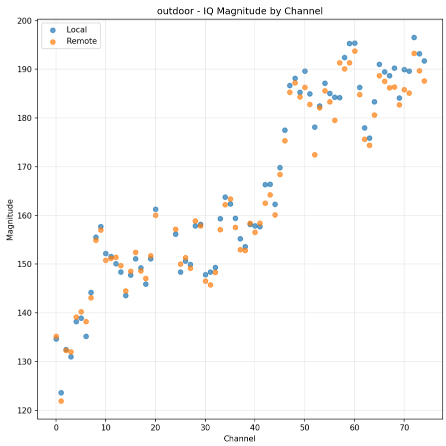

Figure 4.12: Magnitude Spectrum: Outdoor

Round-Trip Time (RTT)

In an indoor environment with a lot of noise and reflections RTT exhibits an interesting but problematic characteristic: a near-zero mean error combined with exceptionally high variance. The primary source of RTT's high variance is sensitivity to signal processing delays and synchronization errors within the devices. RTT depends on precise time measurements between transmission and reception. Any variability in processing time, hardware clock drift, or synchronization between devices directly translates into ranging error. In multipath environments, reflections arriving at slightly different times can confuse the measurement, causing some samples to be triggered by a reflected signal rather than the direct path, introducing substantial error for those specific measurements. This explains why RTT shows the most severe degradation in NLOS conditions, with a 102.4% increase in RMSE compared to LOS performance.

Inverse Fast Fourier Transform (IFFT)

The IFFT method achieves better performance by transforming the channel sounding measurements from the frequency domain into the time domain impulse response. This frequency-domain approach provides advantages in multipath indoor environments. By analyzing the magnitude spectrum the method identifies distinct power peaks which correspond to the signal components with the shortest propagation delay which represents the direct line-of-sight path. Noise and multipath reflections show up as low magnitude peaks that can be ignored during the distance estimation process. The robustness of the method can be seen in Figure 4.13 where outdoor and indoor measurement characteristics are displayed. It becomes evident that the peaks identified are not affected by the surrounding noise and can therefore provide better estimates of the distance.

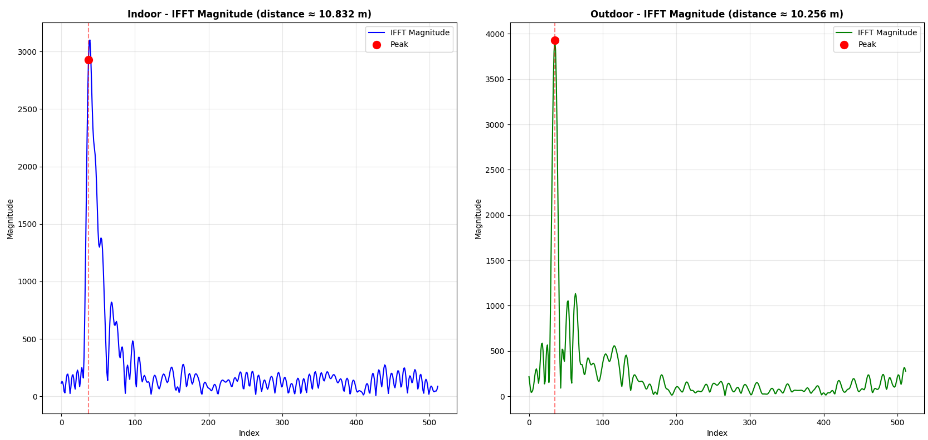

Figure 4.13: IFFT of IQ values

4.2 Multi-Device Network

The third test 3.5.1 evaluates a multi-device deployment in which nodes were placed at different locations throughout the office. Each device measures distance to other reachable nodes in the network, providing a multi-node dataset for reconstruction. Based on the single-pair tests, IFFT measurements were used as input to the multi-device reconstruction with MDS and SDP. The resulting outputs are reported in terms of pairwise distance fit and positional reconstruction error, followed by an evaluation of how reconstruction quality scales with network size.

4.2.1 Network Topology

Figure 4.14: Network topology for multi-device test, showing the placement of devices and their distances.

The devices have been placed in a topology as shown in Figure 4.14, creating a mix of LOS and NLOS conditions where some nodes have direct visibility to each other while others are separated by walls or obstacles. Each node attempts to measure distances to all reachable peers, and the resulting pairwise measurements are used as input to the 2D spatial reconstruction algorithms.

The closest neighboring devices are separated by approximately 9 meters in this deployment. Given that the best-performing ranging method (IFFT) produces a mean error of approximately 1.18 m under LOS conditions and up to approximately 2.40 m under NLOS conditions, this separation corresponds to a worst-case ranging error of roughly 13–27% of the true distance. This places the deployment within the noise tolerance limits reported for both reconstruction methods evaluated in this thesis [24], [27]

4.2.2 Reconstruction Results

Figure 4.15 and 4.16 overlays the reconstruction results on the building floor plan. Blue circles indicate the reference (ground truth) positions of the devices, red triangles show the MDS-reconstructed positions, and green squares show the SDP-reconstructed positions, both rigidly aligned to the reference layout. Dotted lines connect device pairs for which a distance measurement was successfully obtained, the absence of a line between two nodes indicates that the ranging exchange was unreachable during the test.

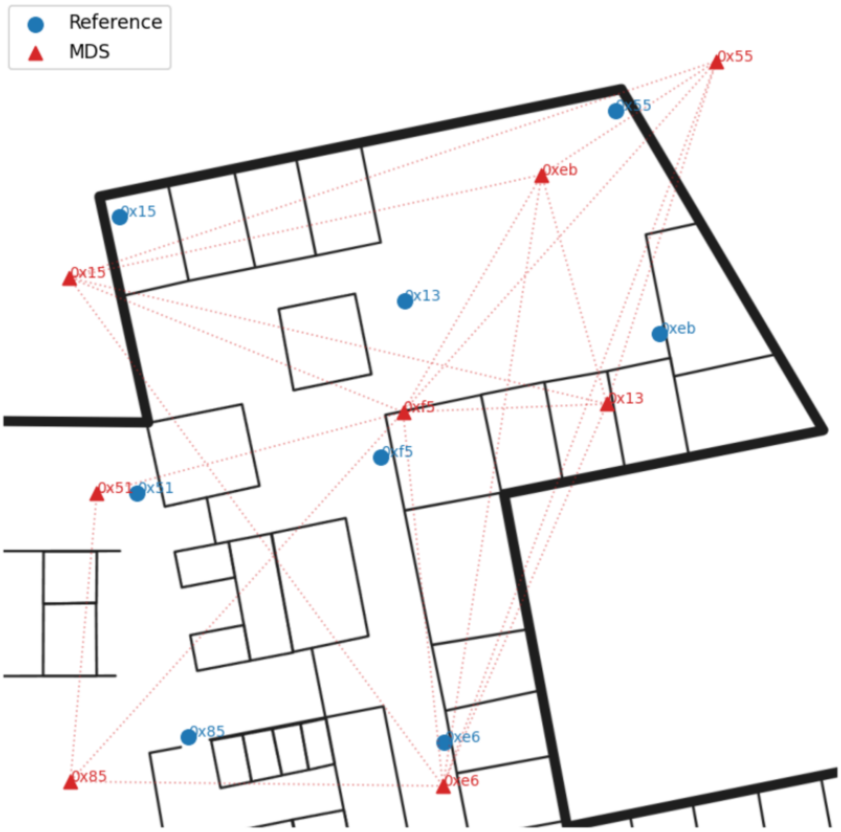

Figure 4.15: MDS reconstruction overlaid on the building floor plan. Dotted lines indicate device pairs for which distance measurements were successfully obtained. Missing lines denote unreachable links.

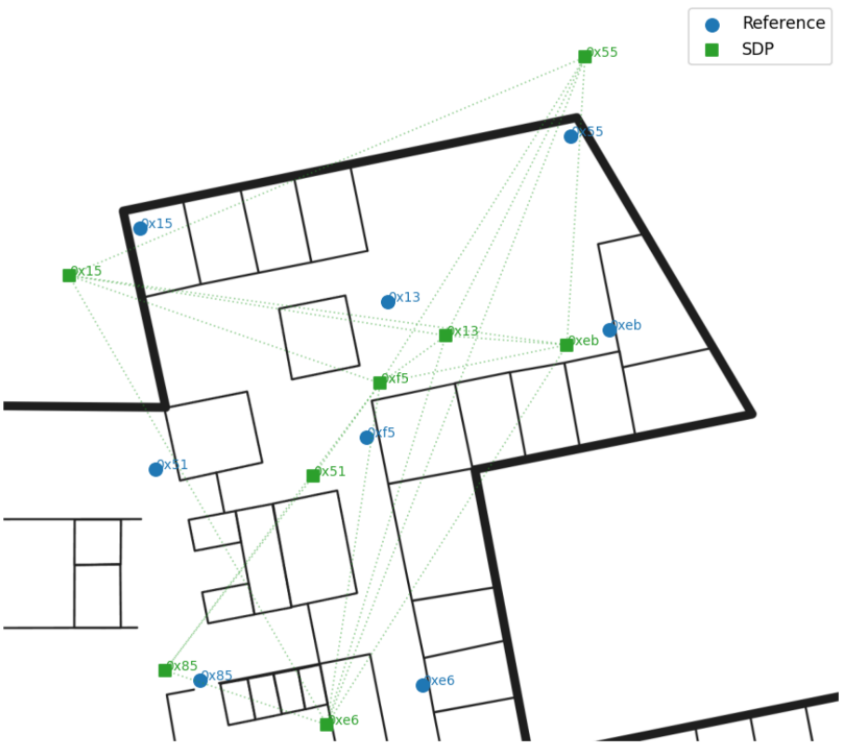

Figure 4.16: SDP reconstruction overlaid on the building floor plan. Dotted lines indicate device pairs for which distance measurements were successfully obtained. Missing lines denote unreachable links.

Table 4.6 summarizes the overall reconstruction accuracy for both methods, while Table 4.7 provides the per-node position errors after alignment. MDS achieves a known-distance MAE of 1.93 m and RMSE of 2.45 m, with a positional MAE of 4.07 m and RMSE of 4.76 m. SDP produces a known-distance MAE of 2.71 m and RMSE of 3.55 m, with a positional MAE of 3.37 m and RMSE of 3.71 m.

Table 4.6: Multi-Device MDS and SDP Reconstruction Error Summary (m)

Metric	MDS (m)	SDP (m)
Known-distance MAE	1.93	2.71
Known-distance RMSE	2.45	3.55
Position mean average error (MAE)	4.07	3.37
Position RMSE	4.76	3.71
Maximum position error	8.43	6.75

Table 4.7: Node-wise Position Error in MDS and SDP Reconstruction (m)

Node	MDS Error (m)	SDP Error (m)
0x13	8.43	2.88
0x15	2.92	3.63
0x51	1.49	6.75
0x55	4.18	3.43
0x85	4.69	1.56
0xe6	1.64	4.41
0xeb	7.31	1.91
0xf5	1.87	2.37

4.2.3 Scaling with Network Size

[26] provides a theoretical foundation for understanding how reconstruction accuracy scales with network size. The authors show that for distance-matrix completion via Singular Value Decomposition, the average squared reconstruction error per entry decreases proportionally to $1/\sqrt{n}$, where n is the number of nodes. This guarantee implies that denser networks are inherently more self-correcting: each additional node introduces new pairwise constraints that help the reconstruction algorithm resolve ambiguities introduced by noisy or missing measurements. With only eight nodes in the multi-device test, the network operates well below the regime where this scaling benefit is meaningful, which is consistent with the elevated positional errors observed.

This has a direct practical implication: increasing the number of BLE channel sounding devices in the environment is expected to improve reconstruction accuracy, not merely provide more coverage. The adoption of BLE 6.0 with native channel sounding support lowers the barrier to deploying larger networks, as more devices become capable of performing IFFT-based ranging without specialized hardware. However, scaling the network introduces a collection bottleneck. The centralized orchestrator developed in this work iterates over all node pairs sequentially, resulting in $O(n^2)$ measurement sessions. For a network of n devices this means the total collection time grows quadratically, making the current approach impractical beyond a moderate number of nodes.

Addressing this bottleneck would require a shift from centralized sequential coordination to a distributed or parallelized measurement strategy, where non-overlapping node pairs conduct measurements concurrently. Such an approach would keep collection time approximately linear in n while still providing the full pairwise distance matrix needed for reconstruction. Designing and evaluating such a protocol is a natural direction for future work 5.2, and would be a prerequisite for deploying the system at the scale where the theoretical accuracy improvements become significant.

4.2.4 Reconstruction Evaluation

4.2.4.1 Reconstruction Evaluation Criteria

The reconstruction analysis follows the predefined metrics from Section 3.4.1: known-distance MAE/RMSE, positional MAE/RMSE/MAX after rigid alignment, and reconstruction uniformity quantified by the standard deviation of per-node positional errors. This section focuses on interpreting the results using those criteria rather than redefining them.

4.2.4.2 Pairwise Distance Accuracy

The pairwise distance inputs used for reconstruction were obtained from the IFFT ranging method. As established in the distance methods analysis, IFFT measurements carry a systematic positive bias of approximately 1.18 m under LOS conditions and 2.40 m under NLOS conditions. This bias is present in all distance inputs supplied to both reconstruction algorithms, and its effect on the reconstruction quality is analyzed in the following sections.

MDS achieves lower known-distance error on both metrics: its MAE of 1.93 m is 41.0% lower than SDP's 2.71 m, and its RMSE of 2.45 m is 44.9% lower than SDP's 3.55 m. This indicates that MDS produces a solution that more faithfully reproduces the input pairwise distances. However, as the positional accuracy analysis shows, this advantage in distance reproduction does not carry over to reconstruction quality.

4.2.4.3 Positional Reconstruction Accuracy

Despite its superior pairwise distance accuracy, MDS produces worse positional reconstruction outcomes across all three positional metrics. SDP achieves a mean position error of 3.37 m, which is 20.8% lower than MDS's 4.07 m. The position RMSE follows the same pattern: SDP achieves 3.71 m compared to 4.76 m for MDS, a 28.3% improvement. The maximum per-node error is also lower for

SDP (6.75 m versus 8.43 m) 4.7, representing a 25.0% reduction in the worst-case positional displacement.

This disconnect between distance accuracy and positional accuracy is the central observation of the multi-device test. The ability to faithfully reproduce individual pairwise distances does not guarantee a high-quality global layout reconstruction, and each algorithm's handling of systematic input bias and sparse graph structure differs fundamentally.

To quantify the uniformity of reconstruction across nodes, the standard deviation of per-node position errors is computed from table 4.7. Seen in table 4.8, MDS yields a per-node error standard deviation of approximately 2.47 m, while SDP achieves a substantially lower value of approximately 1.55 m. This 37.2% reduction in spread indicates that SDP not only achieves a lower mean error but distributes the residual error more evenly across the network, which is a significant advantage in positioning applications where per-device accuracy matters.

Table 4.8: Positional Reconstruction Uniformity

Metric	MDS	SDP
Per-node error Std Dev (m)	2.47	1.55

4.2.4.4 Per-Node Error Analysis

Table 4.7 presents the per-node position errors for both methods after rigid alignment to the reference layout. Examining this table reveals a notably mixed pattern: rather than one method uniformly outperforming the other across all nodes, each algorithm produces low errors for certain nodes and high errors for others.

Table 4.9: Node-wise Position Error in MDS and SDP Reconstruction (m) with Favored Method

Node	MDS Error (m)	SDP Error (m)	Favored Method
0x13	8.43	2.88	SDP
0x15	2.92	3.63	MDS
0x51	1.49	6.75	MDS
0x55	4.18	3.43	SDP
0x85	4.69	1.56	SDP
0xe6	1.64	4.41	MDS
0xeb	7.31	1.91	SDP
0xf5	1.87	2.37	MDS (slight)

Nodes can be grouped into three categories based on which algorithm handles them better. SDP-favored nodes (0x13, 0x55, 0x85, 0xeb) are those for which SDP achieves substantially lower error, these tend to be peripheral or boundary nodes in the measurement graph, where the convex geometric constraints enforced by SDP prevent the large global displacements that MDS can produce for nodes with sparse connectivity. MDS-favored nodes (0x15, 0x51, 0xe6) are those where the SMACOF stress minimization yields a locally consistent fit to the available pairwise distances, these nodes typically have measurement relationships that are internally consistent and well constrained by observed links. Node 0xf5 represents a near-neutral case, performing reasonably well under both methods (1.87 m and 2.37 m), likely because it benefits from a higher number of redundant and reliable range measurements that both algorithms can exploit.

Node 0x13 provides the most striking illustration of the divergence: it records the highest error under MDS (8.43 m) yet is among the better-performing nodes under SDP (2.88 m). Conversely, node 0x51 achieves the lowest MDS error (1.49 m) while recording the highest SDP error (6.75 m). This opposing behavior confirms that the two algorithms are sensitive to structurally different properties of the measurement graph, and that a node's position in the network topology determines which algorithm will handle it well.

4.2.4.5 Method-Specific Reconstruction Characteristics

Each reconstruction algorithm exhibits distinct behavior rooted in how it formulates the localization problem and how it responds to the systematic biases and sparse connectivity present in the eight-device test.

Multidimensional Scaling (MDS)

The SMACOF formulation (Section 3.2.3.1) minimizes a weighted stress function over the available pairwise distances, which directly encourages the reconstructed layout to match the measured links. This is consistent with the lower known-distance MAE and RMSE observed for MDS in this experiment.

The weakness of this approach becomes apparent when the input distances carry systematic bias or when the measurement graph is sparse. The systematic positive bias of approximately 1.18 m in the IFFT inputs means that the stress-minimizing solution is encouraged to preserve distances that are already too large, which can produce an expanded or distorted layout. For nodes with few measured links, this effect has larger consequences because those nodes have less constraint information anchoring their position in the embedding. This is consistent with the high per-node error variance observed in the MDS results: good agreement with the measured pairwise distances can coexist with poor positional accuracy for individual nodes 4.9.

Semidefinite Programming (SDP)

The centered Gram-matrix relaxation used in SDP (Section 3.2.3.2) imposes stronger global geometric structure on the recovered solution than MDS, although it does not guarantee exact recovery of the true 2D configuration under noisy measurements.

This distinction is important because noisy or biased constraints can produce a higher-rank relaxed solution, and the subsequent projection back to 2D can itself introduce distortion. Even so, in this experiment SDP achieved lower positional RMSE and lower per-node error variance than MDS despite its higher known-distance error. This indicates that, for the evaluated eight-node network, the SDP relaxation traded tighter pairwise-distance fitting for a more geometrically useful global layout.

The trade-off is that SDP accepted larger per-link distance residuals than MDS, which is consistent with its higher known-distance MAE and RMSE. If the primary objective is to reproduce the measured pairwise distances as closely as possible, MDS is preferable in this dataset. If the goal is to recover a more accurate and more uniform spatial layout, the results here favor SDP.

4.2.4.6 Distance Reconstruction Quality Tradeoff

The results reveal a fundamental tradeoff between achieving accurate pairwise distance reproduction and achieving accurate positional reconstruction. Table 4.10 presents this tradeoff explicitly for both methods, including the per-node error standard deviation as the reconstruction uniformity metric.

Table 4.10: Distance Reconstruction Quality Tradeoff Summary

Method	**Known-dist MAE (m)**	**Position MAE (m)**	**Per-node Std Dev (m)**
MDS	1.93	4.07	2.47
SDP	2.71	3.37	1.55

MDS sacrifices 20.8% of positional accuracy relative to SDP in order to achieve a 41.0% advantage in known-distance MAE. SDP accepts this worse distance reproduction in exchange for a more geometrically consistent and uniformly distributed positional reconstruction. The selection of algorithm should therefore be driven by the application's primary objective: if accurate pairwise distance reporting is required (for example, to feed downstream range-based protocols), MDS is the preferred choice. If the goal is to reconstruct a spatial map of device positions with the lowest average and worst-case positioning error, SDP is the more suitable method.

4.2.4.7 Comprehensive Reconstruction Summary

To provide a consolidated view of reconstruction performance, Tables 4.8 and 4.6 present all evaluation criteria.

Across all positional accuracy and uniformity criteria, SDP outperforms MDS, while MDS leads on the distance-reproduction criteria. The non-uniform per-node error distributions observed in both methods are a direct consequence of the small and sparse eight-node network topology, where limited measurement redundancy reduces each algorithm's ability to self-correct for locally poor distance estimates. Table 4.9 shows that none of the two models outshines the other in per-node performance in out topology and environment. In addition, as established in section 4.2.3, increasing the number of devices would introduce additional pairwise constraints into both reconstruction algorithms, which is expected to reduce both the mean positional error and the per-node error spread, making the choice between MDS and SDP more predictable and the overall reconstruction quality substantially higher.

Chapter 5

Conclusion & Future Work

This thesis investigated the viability of using Bluetooth 6.0 Channel Sounding with mesh technology to construct an anchor-free IPS. To address this, three distance estimation techniques PBR, RTT, and IFFT were implemented and evaluated under both LOS and NLOS conditions. The best-performing ranging method was then used as input to two position-reconstruction algorithms, MDS and SDP, which together produce a topology map of the network without requiring any pre-deployed anchor infrastructure.

Among the three distance estimation methods evaluated, IFFT emerges as the best choice for indoor positioning applications. Despite having high performance degradation in NLOS conditions IFFT maintains the lowest absolute error in both LOS and NLOS environments, achieving 1.40 m and 2.40 m respectively. This superior absolute accuracy, combined with its moderate robustness profile, makes it the most practical method for real-world indoor deployments where NLOS scenarios are unavoidable.

PBR, by contrast, demonstrates fundamental limitations in indoor environments. Multipath effects and environmental noise, which are prevalent in typical office settings, introduce phase ambiguities that severely degrade the method's accuracy. These phase measurement inconsistencies result in the highest mean error among all evaluated methods. Additionally, interference from multiple signals operating in the same 2.4 GHz frequency band as Bluetooth further compounds the phase estimation problem, reducing overall measurement accuracy.

RTT shows varying performance throughout all tests. The method exhibits high variance with relatively low mean error compared to the other algorithms. However, this high variance makes it unsuitable for indoor positioning systems where constant performance is required. The performance degradation is also the highest out of all methods making it non-ideal for non-line-of-sight measurements.

With IFFT being the most performant method, it is used to create the indoor position map. Looking at the results presented in section 4.2.4, SDP achieves better overall positional accuracy than MDS in the evaluated eight-node network. However, the relative performance of each algorithm depends heavily on the deployment environment, the number of devices in the network, and the connectivity of each node. It is therefore difficult to draw a general conclusion about which method is superior across all scenarios.

An important practical consideration is how computational cost scales with network size. MDS is less computationally complex than SDP [27]. As the network grows larger and denser, positional accuracy is expected to improve for both methods, as additional pairwise constraints reduce per-node error and improve reconstruction accuracy.

If a network scales and the MDS errors are reduced to a reasonable amount, the impact of the computational cost would likely make MDS more preferable. Future work on how to handle this is discussed in section 5.2.

With these results we have shown how different distance measuring algorithms perform and can be used to create an IPS. BLE Mesh is used to create an IPS and its performance, accuracy and limitations have been presented in regards to practical applications.

5.1 Limitations

The reported results should be interpreted under the study assumptions and scope limitations. Within the scope of this thesis, all devices are assumed to be static, and therefore the methods and setup were designed accordingly. The implementation relies on Bluetooth 6.0 Channel Sounding support in the selected hardware and firmware platform. In addition, the evaluation was performed in a single-floor office environment with an eight-node deployment, which means that performance in larger, dynamic, or materially different environments may differ from the results presented here.

5.2 Future Work

As discussed in section 4.2.3, a primary future-work direction is handling network scaling. The centralized orchestrator 3.3.2 architecture suffers from an $O(n^2)$ measurement collection time, which becomes impractical as the number of devices grows. To address this bottleneck, a distributed measurement strategy could be adopted. A gossip-based protocol, for example, would allow non-overlapping node pairs to conduct measurements concurrently, reducing collection time to approximately $O(n)$ while still providing the full pairwise distance matrix required for reconstruction. Such an approach would be compatible with the mesh networking capabilities of BLE, which supports deployments of up to 32,767 devices theoretically. With an increased network scale, investigating performance becomes a natural and promising direction for future investigation.

A second direction is extending the reconstruction from a single-floor setting to multi-floor and fully three-dimensional deployments. The current evaluation focuses on a single-floor office environment, but future work could explore using MDS and SDP in three dimensions rather than restricting them to a plane. An alternative approach could be to use two-dimensional reconstructions from each floor as input to a higher-level algorithm that estimates the relative vertical positions of floors.

A third direction concerns robustness improvements through controlled structural information and measurement-quality weighting. In this work, we explore IPS without anchor devices to keep setup time low and maintain a truly ad-hoc system. However, introducing a small number of anchors could improve overall accuracy [24] and resolve rotational ambiguity without significantly increasing setup complexity. In the same vein, MDS and SDP were applied without explicitly weighting measurements by quality, so incorporating measurement-quality-aware weighting into the reconstruction pipeline is promising. The bias-reduction framework in [38] shows that improved weighting models and explicit treatment of higher-order noise effects can reduce localization bias and approach the theoretical bound, while [39] indicates that RSSI-based quality indicators can improve reconstruction quality. Combining these ideas, for example by using RSSI or related link-quality metrics to construct adaptive weights for MDS/SDP, could improve robustness in heterogeneous LOS/NLOS conditions. Evaluating the resulting accuracy-complexity trade-off, as well as exploring additional quality factors beyond RSSI, is therefore a relevant and practical future research direction.

A fourth direction is to investigate a quality-adaptive hybrid reconstruction strategy that combines MDS and SDP in a single weighted objective. The motivation comes directly from the trade-off observed in this thesis where MDS better reproduces measured pairwise distances while SDP yields lower overall positional error and more uniform node-wise reconstruction. A hybrid strategy could assign stronger distance fitting influence to high-quality and well connected measurements, where MDS is most reliable, while applying stronger geometric regularization to sparse or poorly connected measurements, where the global constraints of SDP are most beneficial. A key challenge is that MDS and SDP operate on fundamentally different optimization variables, so a unified formulation would need to be carefully derived. The approach could be evaluated across different LOS/NLOS ratios, connectivity patterns, and network sizes using both pairwise-fit and positional metrics to determine whether adaptive weighting consistently outperforms either method alone.

Bibliography

[1] T. Kim, K.-S. Kim, and K.-J. Li, "Improving room-level location for indoor trajectory tracking with low ips accuracy," *ISPRS International Journal of Geo-Information*, vol. 10, no. 9, 2021, ISSN: 2220-9964. DOI: 10.3390/ijgi10090620 [Online]. Available: https://www.mdpi.com/2220-9964/10/9/620

[2] O. Silver, "An indoor localization system based on ble mesh network," Master's thesis in Communication Systems, Linköping University, Department of Electrical Engineering, Linköping, Sweden, 2016. [Online]. Available: https://liu.diva-portal.org/smash/get/diva2:935745/FULLTEXT01.pdf

[3] Y. Cao, H. Kandula, and X. Li, "Measurement and analysis of rss using bluetooth mesh network for localization applications," *Network*, vol. 1, no. 3, pp. 315–334, 2021, ISSN: 2673-8732. DOI: 10.3390/network1030018 [Online]. Available: https://www.mdpi.com/2673-8732/1/3/18

[4] I. Bluetooth SIG, *Bluetooth® channel sounding*, https://www.bluetooth.com/learn-about-bluetooth/feature-enhancements/channel-sounding/, Accessed: 2026-02-03, 2026.

[5] R. Ramirez, C.-Y. Huang, C.-A. Liao, P.-T. Lin, H.-W. Lin, and S.-H. Liang, "A practice of ble rssi measurement for indoor positioning," *Sensors*, vol. 21, no. 15, 2021, ISSN: 1424-8220. DOI: 10.3390/s21155181 [Online]. Available: https://www.mdpi.com/1424-8220/21/15/5181

[6] B. Black, J. Rafferty, J. A. Santos, and A. Ennis, "Bluetooth rssi signal chaining for mesh network analysis," in *2021 IEEE SmartWorld (SCALCOM/UIC/ATC/IOP/SCI)*, 2021, pp. 248–253. DOI: 10.1109/SWC50871.2021.00042

[7] Cloudflare, Inc. "What is round-trip time (rtt)? | rtt definition," Accessed: May 4, 2026. [Online]. Available: https://www.cloudflare.com/learning/cdn/glossary/round-trip-time-rtt/

[8] M. Kochláň and J. Miček, "Indoor propagation of 2.4ghz radio signal propagation models and experimental results," in *The 10th International Conference on Digital Technologies 2014*, 2014, pp. 125–129. DOI: 10.1109/DT.2014.6868703

[9] A. Comuniello, A. De Angelis, A. Moschitta, and P. Carbone, "Using bluetooth low energy technology to perform tof-based positioning," *Electronics*, vol. 11, no. 1, 2022, ISSN: 2079-9292. DOI: 10.3390/electronics11010111 [Online]. Available: https://www.mdpi.com/2079-9292/11/1/111

[10] T. Ulsamer, C. Mazur, and S. Knauth, "Study of ble6 ranging performance in a utility basement," in *Proceedings of the IPIN-WCAL 2025: Workshop for Computing & Advanced Localization at the Fifteenth International Conference on Indoor Positioning and Indoor Navigation*, Corresponding author: Stefan Knauth, Tampere, Finland: CEUR-WS.org, 2025. [Online]. Available: https://ceur-ws.org/Vol-4047/short26.pdf

[11] I. Bluetooth SIG, *High-accuracy low-power secure ranging using bluetooth® channel sounding*, https://www.bluetooth.com/bluetooth-resources/high-accuracy-low-power-secure-ranging-using-bluetooth-channel-sounding/, Accessed: 2026-02-13, 2024-04-17.

[12] J. Wieme et al., "Performance evaluation of bluetooth channel sounding on commercial hardware," *IEEE Access*, 2023. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=11207662

[13] Metirionic GmbH. "Bluetooth channel sounding: Unlocking precision in distance measurement," Accessed: May 4, 2026. [Online]. Available: https://metirionic.com/metirionic-technology/how-it-works/bluetooth-channel-sounding-unlocking-precision-in-distance-measurement/

[14] All About Circuits. "Understanding i/q signals and quadrature modulation," Accessed: May 4, 2026. [Online]. Available: https://www.allaboutcircuits.com/textbook/radio-frequency-analysis-design/radio-frequency-demodulation/understanding-i-q-signals-and-quadrature-modulation/

[15] S. Boyd, *The fourier transform*, EE 102: Introduction to Signals and Systems, Handout #23, Lecture 11, Stanford University, 2002. Accessed: May 4, 2026. [Online]. Available: https://web.stanford.edu/class/ee102/lectures/fourtran

[16] Nordic Semiconductor ASA, *nRF Connect SDK: cs_de.c and cs_de.h (Bluetooth Channel Sounding Distance Estimation)*, https://github.com/nrfconnect/sdk-nrf/tree/main/subsys/bluetooth/cs_de, Includes both implementation and header files, 2026.

[17] Inpixon. "Ultra-wideband (uwb)." Accessed: 2026-05-11, Inpixon, Accessed: May 11, 2026. [Online]. Available: https://www.inpixon.com/technology/standards/ultra-wideband

[18] RTLS Alliance Team. "Technology comparison: Ble vs uwb for rtls use cases, Ble vs uwb: Cost, accuracy & deployment comparison." Accessed: 2026-05-11, RTLS Alliance, Accessed: May 11, 2026. [Online]. Available: https://www.rtlsalliance.com/resources/ble-vs-uwb-cost-accuracy-deployment-comparison

[19] M. R. Rashidi and S. Chinara, "A comprehensive study and performance evaluation of indoor positioning methods in ble and uwb technologies," *IEEE Internet of Things Journal*, vol. 12, no. 22, pp. 46 475–46 499, 2025. DOI: 10.1109/JIOT.2025.3611794

[20] F. Che, Q. Z. Ahmed, P. I. Lazaridis, P. Sureephong, and T. Alade, "Indoor positioning system (ips) using ultra-wide bandwidth (uwb)—for industrial internet of things (iiot)," *Sensors*, vol. 23, no. 12, 2023, ISSN: 1424-8220. DOI: 10.3390/s23125710 [Online]. Available: https://www.mdpi.com/1424-8220/23/12/5710

[21] ZPE Systems. "Centralized vs. distributed network management: Which one to choose?" Accessed: May 4, 2026. [Online]. Available: https://zpesystems.com/centralized-vs-distributed-network-management-zs/

[22] M. Hultgren and D. Papathanopoulos, "Evaluating the usage of bluetooth low energy beacons and smartphones for indoor positioning systems," M.S. thesis, University of Gothenburg, Chalmers University of Technology, Department of Computer Science, and Engineering, Gothenburg, Sweden, Jun. 2014. Accessed: May 4, 2026. [Online]. Available: https://gupea.ub.gu.se/server/api/core/bitstreams/8555a56d-d20d-4652-9843-8cb40793f701/content

[23] A. Tewfik, "Distributed sensor network localization using socp relaxation," *IEEE Transactions on Wireless Communications*, 2008. DOI: 10.1109/T-WC.2008.070241

[24] P. Biswas, T.-C. Liang, K.-C. Toh, Y. Ye, and T.-C. Wang, "Semidefinite programming approaches for sensor network localization with noisy distance measurements," *IEEE Transactions on Automation Science and Engineering*, vol. 3, no. 4, pp. 360–371, 2006. DOI: 10.1109/TASE.2006.877401

[25] P. A. Regalia and J. Wang, "On distance reconstruction for sensor network localization," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2010, pp. 2866–2869. DOI: 10.1109/ICASSP.2010.5496180

[26] P. Drineas, M. Magdon-Ismail, G. Pandurangan, R. Virrankoski, and A. Savvides, "Distance matrix reconstruction from incomplete distance information for sensor network localization," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2005, pp. 1077–1080.

[27] M. Patane, B. Rossi, P. Fragneto, and A. Fusiello, "Wireless sensor networks localization with outliers and structured missing data," in *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, 2017, pp. 1–7. DOI: 10.1109/PIMRC.2017.8292346

[28] A. Labyad, H. K. Aroussi, and A. Mouloudi, "Localization algorithms research in wireless sensor networks based on multilateration and trilateration techniques," in *Proceedings of the IEEE Conference*, IEEE, 2014, pp. 415–419, ISBN: 978-1-4799-5979-2.

[29] J. Gonzalez-Barajas and D. Montenegro, "Average filtering: Theory, design and implementation," in Apr. 2016, ISBN: 978-1-63485-210-4.

[30] Nordic Semiconductor, *Nrf mesh mobile app*, Accessed: 2026-03-19, 2026. [Online]. Available: https://www.nordicsemi.com/Products/Development-tools/nRF-Mesh

[31] Nordic Semiconductor. "Nrf54l15-dk multiprotocol development kit." Product page and technical overview, Mouser Electronics, Accessed: Mar. 8, 2025. [Online]. Available: https://www.mouser.se/ProductDetail/Nordic-Semiconductor/nRF54L15-DK

[32] I. Bluetooth SIG, *Developer study guide: An introduction to bluetooth mesh networking*, https://www.bluetooth.com/bluetooth-resources/bluetooth-mesh-developer-study-guide/?utm_campaign=mesh&utm_source=internal&utm_medium=blog&utm_content=bluetooth-mesh-developer-study-guide-v2.0, Accessed: 2026-02-02, 2020.

[33] zephyrproject rtos, *Zephyr*, `https://github.com/zephyrproject-rtos/zephyr/tree/main/samples/bluetooth`, Accessed: 2026-02-03, 2026.

[34] H. Cho, J. Ji, Z. Chen, H. Park, and W. Lee, "Measuring a distance between things with improved accuracy," vol. 52, Jun. 2015. DOI: `10.1016/j.procs.2015.05.119`

[35] M. Hägglund Eriksson. "Accuracy and precision analysis of total station measurements," Accessed: May 4, 2026. [Online]. Available: `https://kth.diva-portal.org/smash/get/diva2%3A727374/FULLTEXT01.pdf`

[36] P. Li, W. Wu, Z. Zhao, and G. Q. Huang, "Indoor positioning systems in industry 4.0 applications: Current status, opportunities, and future trends," *Digital Engineering*, vol. 3, p. 100 020, 2024, ISSN: 2950-550X. DOI: `https://doi.org/10.1016/j.dte.2024.100020` [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S2950550X24000207`

[37] J. Aspnes et al., "A theory of network localization," *IEEE Transactions on Mobile Computing*, vol. 5, no. 12, pp. 1663–1678, 2006. DOI: `10.1109/TMC.2006.174`

[38] Z. Zhang, D. Wang, B. Yang, and J. Yin, "Weighted multidimensional scaling localization method with bias reduction based on toa," *IEEE Sensors Journal*, vol. 23, no. 17, pp. 19 803–19 814, 2023. DOI: `10.1109/JSEN.2023.3296986`

[39] X. Zhang, L. Chen, M. Feng, and T. Jiang, "Toward reliable non-line-of-sight localization using multipath reflections," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 1, 2022, ISSN: 2474-9567. DOI: `10.1145/3517244` [Online]. Available: `https://dl.acm.org/doi/10.1145/3517244`