Master Thesis

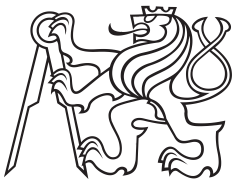

Czech Technical University in Prague

F3

Faculty of Electrical Engineering
Department of Cybernetics

Robot Following Human with Limited Visual Contact

Bc. Bohumil Brož

Supervisor: prof. Ing. Tomáš Svoboda, Ph.D.
Field of study: Cybernetics and Robotics
May 2024

MASTER‘S THESIS ASSIGNMENT

I. Personal and study details

Student's name: **Brož Bohumil** Personal ID number: **492348**

Faculty / Institute: **Faculty of Electrical Engineering**

Department / Institute: **Department of Cybernetics**

Study program: **Cybernetics and Robotics**

II. Master’s thesis details

Master’s thesis title in English:

Robot Following Human with Limited Visual Contact

Master’s thesis title in Czech:

Robot následující osobu v prostředí se sníženou viditelností

Guidelines:

Field rescuers use robots mostly in a fully teleoperated mode by using some kind of remote control device. This type of deployment requires a skilled operator and suffers from a loose connection between the robot and other field rescuers. This loose connection results in inefficiency and, most importantly, a lack of trust and low acceptance. The overall goal of the project is to develop and test a device and algorithm for navigating a robot close to a leading human in a way that is perceived as safe and understandable. Especially focus on cases when visual contact is difficult or even impossible (e.g., in a smoky environment).
The plan for the diploma project is the following:
• Suggest suitable sensory configurations and mount positions on robots and humans.
• Design and implement a robust estimation of human-to-robot transformation (mutual pose).
• Design and implement a method for robots to follow the human pose.
• Design and implement a human path following method and discuss possible limitations compared to the simpler pose(only) following method.
• Define a simple (proxy) test scenario for evaluating perceived safety and understandability.

Bibliography / sources:

[1] P. Teja Singamaneni, A. Favier, and R. Alami, Human-Aware Navigation Planner for Diverse Human-Robot Interaction Contexts, 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).
[2] Thibault Kruse, Amit Kumar Pandey, Rachid Alami, Alexandra Kirsch, Human-aware robot navigation: A survey. Robotics and Autonomous Systems, Volume 61, Issue 12, 2013, Pages 1726-1743, ISSN 0921-8890.
[3] P. Šimek, Precision of Position Determination Radio Systems. Bachelor’s Thesis, CTU Faculty of Electrical Engineering, 2022.

Name and workplace of master’s thesis supervisor:

prof. Ing. Tomáš Svoboda, Ph.D. Vision for Robotics and Autonomous Systems FEE

Name and workplace of second master’s thesis supervisor or consultant:

Date of master’s thesis assignment: **11.01.2024** Deadline for master's thesis submission: **24.05.2024**

Assignment valid until: **21.09.2025**

prof. Ing. Tomáš Svoboda, Ph.D. Supervisor’s signature	prof. Dr. Ing. Jan Kybic Head of department’s signature	prof. Mgr. Petr Páta, Ph.D. Dean’s signature

III. Assignment receipt

The student acknowledges that the master's thesis is an individual work. The student must produce his thesis without the assistance of others, with the exception of provided consultations. Within the master's thesis, the author must state the names of consultants and include a list of references.

.	
Date of assignment receipt	Student's signature

Acknowledgements

First of all I would like to thank my supervisor prof. Ing. Tomáš Svoboda, Ph.D. for his guidance, valuable feedback, and advice provided while working on the thesis. Furthermore, I would like to thank Ing. Bedřich Himmel for advice regarding the hardware and for help with organising the experiments.

I would also like to thank Ing. Václav Navrátil, Ph.D. and Ing. Josef Krška from the Department of Radioelectronics for consultations regarding the radio-based localisation systems and for lending part of the hardware used in this work.

Finally, I would like to thank Ing. Miloš Prágr from the Department of Computer Science for giving me the opportunity to perform the experiment with the six-wheeled autonomous vehicle.

Declaration

I declare that the presented work was developed independently and that I have listed all sources of information used within it in accordance with the methodical instructions for observing the ethical principles in the preparation of university theses.

Prague, date 24. May 2024

.....
Signature

Abstract

In some situations, it may be advantageous for a robot to autonomously follow a human moving in its vicinity. However, direct visibility is not always guaranteed, for example, due to smoke or dust. This work deals with the relative localisation of a human to a robot without the use of cameras and LiDAR. Radio-based localisation systems — Ultra-wideband and Bluetooth are used to obtain the relative position of the human. A method is proposed to estimate the relative position of the human from inaccurate measurements provided by radio receivers, and the resulting algorithm is implemented as a ROS package. The thesis also presents the effect of robot design on localisation accuracy and provides general guidelines for placing sensors on the robot to obtain the best results. The experiments performed show that the system is able to adapt to different robotic platforms and that the robot is able to track humans under different conditions.

Keywords: Relative Localisation, Ultra-Wideband, Two-Way Ranging, Time Difference of Arrival, Bluetooth, Angle of Arrival

Supervisor: prof. Ing. Tomáš Svoboda, Ph.D.

Abstrakt

V některých situacích může být výhodné, aby robot autonomně sledoval člověka, který se pohybuje v jeho blízkosti. Ne vždy je ale zaručena přímá viditelnost, například kvůli kouři nebo prachu. Tato práce se zabývá relativní lokalizací člověka vůči robotovi bez použití kamer a LiDARu. K získání relativní pozice člověka jsou využity radiové lokalizační systémy — Ultra-wideband a Bluetooth. Je navržena metoda pro odhad relativní pozice člověka z nepřesných měření poskytnutých radiovými přijímači a výsledný algoritmus je implementován jako ROS package. Práce dále prezentuje vliv konstrukce robota na přesnost lokalizace a poskytuje obecný návod na umístění senzorů na robota za účelem dosažení co nejlepších výsledků. Provedené experimenty ukazují, že systém je schopný se přizpůsobit různým robotickým platformám a že robot je schopný sledovat člověka v různých podmínkách.

Klíčová slova: Relativní lokalizace, Ultra-Wideband, Two-Way Ranging, Time Difference of Arrival, Bluetooth, Angle of Arrival

Překlad názvu: Robot následující osobu v prostředí se sníženou viditelností

Contents

Chapter 1

Introduction

The ability of a robot to autonomously follow a human can be beneficial in many situations because no operator is required for the robot. For example, in a rescue scenario, the robot could carry the necessary equipment while rescuers perform their tasks. There are already commercially available robots with the "follow me" mode. These robots are often intended for military use [1, 2], and therefore not much information is publicly available. An example of such a robot is Rheinmetall Mission Master [3] with the Path autonomy kit [4], the PLATON kit [5] by Diehl Defence, PROBOT [6] by Elbit Systems or THeMIS [7] by Milrem Robotics equipped with MIFIK [8].

The goal of this thesis is to develop a system that would allow a robot to autonomously follow a human without needing any visual information such as an image from cameras or a scan from LiDAR. The system should allow the robot to obtain the relative position of the human with respect to the robot and follow the path that the human has taken. In case the robot is not capable of localising itself within the environment, the robot will drive straight towards the human (it is impossible to store the path of the human for following, since no static coordinate frame exists). Figure 1.1 shows the desired functioning of the system.

The proposed solution should be safe but also acceptable to the user, meaning that it is clear that the robot is still following the human. To achieve this, the robot has to never be too far behind the human. The human has to never feel threatened by the robot; as a result, the robot has to maintain a large enough gap to the human and stop immediately if the human is too close. The system should also report its state in a way that is easy to understand, and controlling the system should not require any deeper understanding of the internal functioning.

Ultra-wideband and Bluetooth angle of arrival were selected as two methods for localising the human without the need for any visual information. The system uses speech synthesis to give status messages that are easy to understand, and it is possible to enable or disable the following by pressing a single button.

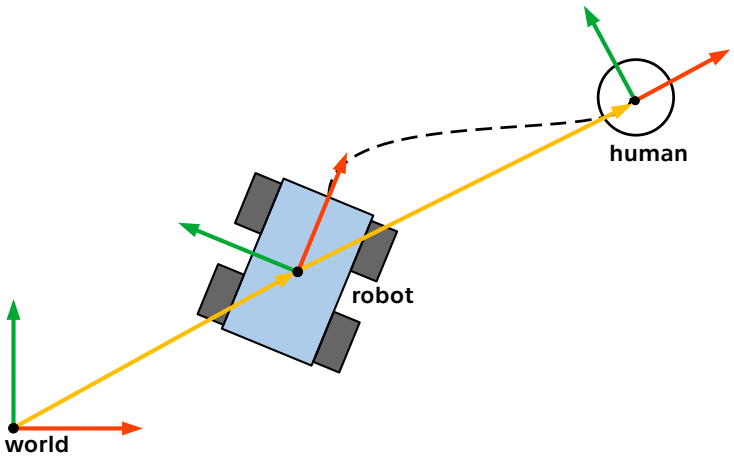

Figure 1.1: The desired functionality of the system in a 2D scheme. The robot estimates the relative position of the human (yellow arrow connecting the robot and human reference frames). This allows the robot to follow the human path (dashed black line). The robot might also be capable of localising itself in relation to some fixed world coordinate frame, this capability is shown with a yellow arrow connecting the world coordinate frame to the robot coordinate frame.

Chapter 2

Related Works

Autonomously following a human is a task relatively similar to autonomous driving in a convoy of vehicles. The agent has to follow the path that its leader has taken. The main task that has to be solved to allow a robot to autonomously follow its leader is the relative localisation of the leader with respect to the robot.

There are multiple approaches to the relative localisation of the leader. The main two categories are localisation with markers and marker-less localisation. The use of markers could result in higher accuracy and reliability of the localisation. The markers could be active (LEDs, radio transmitters) or passive (e.g. AprilTag [9] — the algorithm searches for a special pattern similar to a QR code). Some use cases limit the usage of markers; for example, an active marker is not suitable for military usage because the opponent could detect the marker.

The localisation without markers relies on distinguishing the tracked object from the environment. This requires prior knowledge of the appearance of the leader. There are two possible data sources, RGB cameras [10, 11, 12] and a sensor that provides depth scans, for example, LiDAR [13, 14, 15, 16]. Both of these data sources share the same limitations — they require direct line of sight visibility of the leader. If the environment is full of obstacles that could occlude the target or if the environment denies usage of visual information (due to smoke or dust particles in the air, or a completely dark environment where cameras do not work), the system loses the ability to localise the leader.

The use of markers enables a more robust localisation of the leader. By adding visual markers (fiducial marker), such as in [17], localisation from the camera image is possible even without knowing the appearance of the leader. Active visual markers enable localisation even in environments that degrade visual information (smoke, dust, darkness); such an active marker can be a set of LEDs with a predefined shape, [18] uses UV light, to make localisation robust to lighting conditions.

Markers do not have to be only visual; the leader could carry a radio transmitter that the robot tries to localise. Radio-based localisation systems are often used for indoor localisation, a fixed infrastructure is mounted around the area, and the target carries a transmitter. One of the technologies used

for indoor localisation is ultra-wideband. Localisation with ultra-wideband uses time of flight of the signal to obtain the position of one or several nodes of the UWB network. Ultra-wideband can be used on its own, as in [19, 20, 21, 22, 23], or combined with additional data sources [24, 25, 26].

Another radio-based system for localisation uses Bluetooth. It is possible to estimate the distance between the moving transmitter and the receiver with a known position based on the strength of the received signal, as used in [27, 28, 29]. The position of the target is then obtained via triangulation. Another option is to measure the angle of arrival/departure, such as in [30, 31, 32]. When the heading from multiple fixed locations is known, the position of the moving target can be obtained as an intersection of several half-lines.

There are additional challenges that a human following system could try to solve. One of them is to make the robot behave in a way that is acceptable to other humans present in the environment. This task is related to human-aware path planning. Human-aware path planning solves the task of planning a path in an environment where humans are present. The main objective is to get the robot to its goal while making the robot's behaviour as comfortable for the humans as possible. The works [33, 34] present some of the requirements on the system: do not enter an intimate or personal zone of humans, signal intentions of the robot in advance, slow down while close to humans, approach humans from the front, respect social rules. Another requirement, solved in [35], is to avoid crossing the paths of humans and not disturb humans interacting with objects in the environment. Human-aware planning requires the knowledge of position of non-cooperating humans, therefore visual information (cameras, LiDAR) is required to detect them. This work focusses on scenarios with limited or no visual information; therefore, human-aware planning is not considered.

Chapter 3

Localisation with active markers

Two radio-based localisation systems that are often used for the indoor localisation of moving objects were selected — ultra-wideband and Bluetooth. In both of the systems, a single sensor does not measure the full position of the target, instead only a ditance or direction towards the target is measured. To obtain the full position of the target, multiple measurements from sensors mounted in multiple known positions are needed. The measurements provided by the fixed infrastructure are then used to calculate the position of the target. The task of human localisation has different scales of fixed infrastructure compared to indoor localisation; the modules mounted on the robot can usually be at most one metre apart from each other as opposed to a distance of several metres in indoor localisation. Both systems require the human to carry an active transmitter.

3.1 Ultra-wideband

Ultra-wideband (UWB) is a radio technology that offers short-range communication and allows locating some of the nodes in the UWB network. It is based on the standards IEEE 802.15.4a [36] and IEEE 802.15.4z [37]. There are two types of device in the UWB localisation system, static anchors with known position in the world coordinate frame and moving tags that are being localised. Two main principles are used to localise the tags; both are based on measuring Time of Flight of the signal.

3.1.1 Two-Way Ranging

Two-way Ranging localisation is based on measuring the distance between the moving device (tag) and each node of the static UWB network (anchors). The range is not measured directly, the time of flight of the signal is measured instead, and the range is calculated from the known speed of the signal in the air.

Bidirectional communication is used to measure the time of flight of the signal between the tag and the anchor. The tag sends a poll; it is received by the anchors. Each anchor waits for a predefined time after receiving the poll message. After this time passes, a response is sent. The communication

and the measured times are shown in Figure 3.1. The time between the transmission of the poll message and reception of the response is:

$$t_{roundtrip} = TOF + t_{reply} + TOF \tag{3.1}$$

Time of flight can be therefore computed as:

$$TOF = \frac{t_2 - t_1 - t_{reply}}{2} \tag{3.2}$$

where t_1 is the timestamp of the transmission of the poll message and t_2 is the timestamp of the reception of the anchor's response. All measurements are performed by the tag, meaning that the tag can directly compute its own position. The position can be estimated if multiple range measurements are available. The anchors are fixed; their position is therefore known. The tag lies on the intersection of multiple spheres with centres at the anchors and radius given by the measured range.

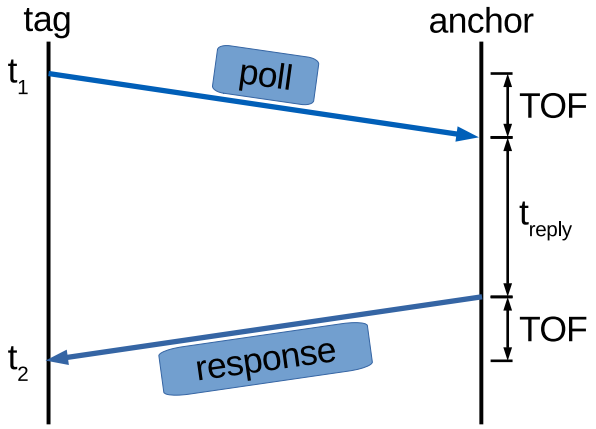

Figure 3.1: Scheme showing the principle of UWB Two-way ranging.

This work uses Qorvo-made DWM1001C modules (three modules can be seen in Figure 3.2). All of the modules in the network have identical hardware; it is possible to switch the mode of each module (anchor or tag). The modules use serial communication over USB to publish the range measurements. As was already stated earlier, the range measurements are available at the side of the tag, this means that the anchors are passive and do not provide any measurements. Having the range measurements immediately available on the side of the robot is beneficial for the safety of the system; therefore, all of the modules mounted on the robot are operating in the tag mode and the human carries an anchor. This means that the system operates in a configuration multiple tags to one anchor, and indoor localisation uses the opposite configuration — one tag to multiple anchors.

The UWB modules do not directly measure the 2D or 3D position of the human, but the accuracy of the range measurements is closely related to the accuracy of localisation. To evaluate the accuracy of the range measurements,

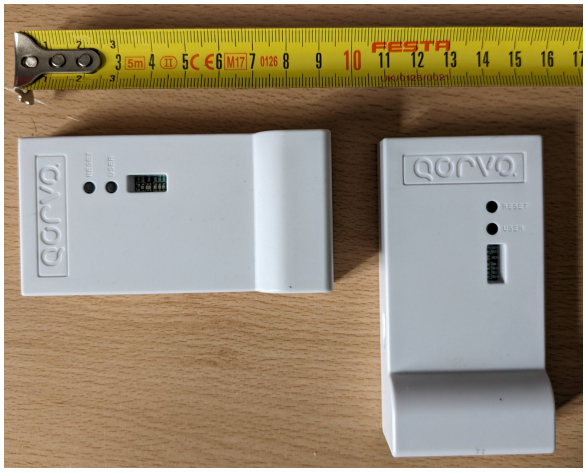

Figure 3.2: The Qorvo DWM1001C modules.

the reported ranges for various distances between the UWB tag and the anchor were recorded for one minute. As can be seen in Figure 3.3, the measurements follow an approximately normal distribution. Therefore, the mean and standard deviation of the recorded measurements can be computed.

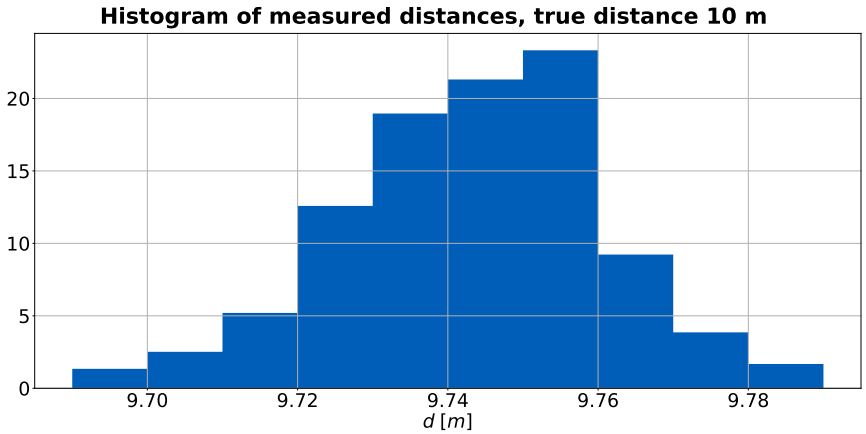

Figure 3.3: The histogram of UWB range measurements for the true distance of 10 metres.

As can be seen in Figure 3.4, the precision of the measurements does not depend on the distance between the modules, the standard deviation fluctuates between the individual measurements, but remains lower than 4 centimetres even for a distance of 12 metres.

The measurements are also sufficiently accurate for the given task. The reported distance is consistently lower than the true one (as can be seen in Table 3.1), but this does not pose a problem for the "follow me" mode; the robot will maintain a gap that is slightly higher than the desired one. Therefore, the safety is not compromised by the bias of the range measurements.

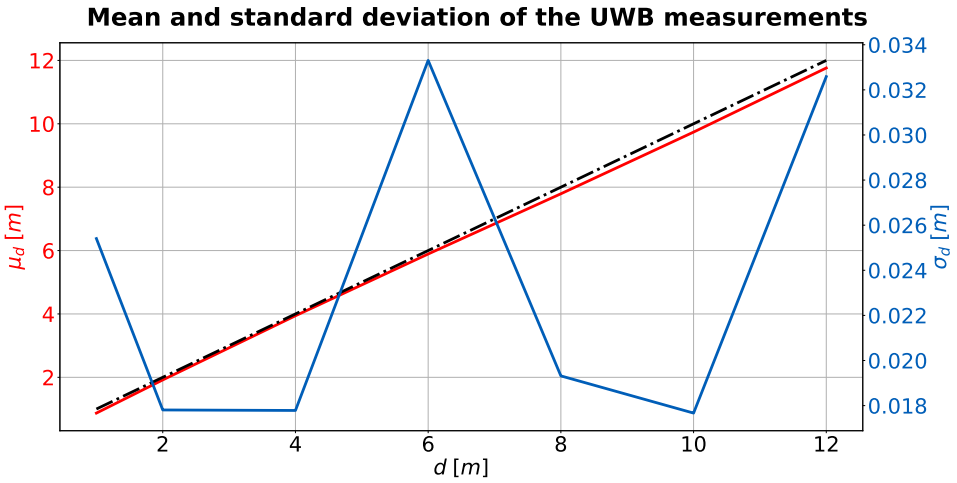

Figure 3.4: Comparison of the mean of the measurements (red) and the true range values (black). The blue line shows the dependency of the standard deviation on the distance.

distance	μ [m]	σ [m]
1	0.87	0.025
2	1.92	0.018
4	3.94	0.018
6	5.89	0.033
8	7.79	0.019
10	9.74	0.018
12	11.76	0.033

Table 3.1: The mean and standard deviation of the range measurement

3.1.2 Time Difference of Arrival

Time Difference of Arrival also uses a fixed infrastructure, but this time the moving device (tag) does not have to communicate with each node of the static network (anchors). The anchors are time-synchronised, they are either connected with a cable or they can utilise the UWB communication. The tag sends a blink message; each anchor receives it with a different timestamp. The time of transmission of the blink is unknown; therefore, it is not possible to compute the distance between the anchors and the tag. The time differences between the timestamps of pairs of anchors can be used to compute the position of the tag. Each pair of anchors defines a pair of fixed points; the time difference can be transformed into a difference of distances from the two fixed points. A single time difference measurement therefore defines a hyperbola (in 2D) or a hyperboloid (in 3D); in case the difference of distances is equal to the distance between the anchors, the hyperbola is reduced to a half-line. The relation between the position of the anchors, the position of

the tag, and the measured difference of distances can be formulated as:

$$\|\mathbf{a}_1 - \mathbf{t}\| - \|\mathbf{a}_2 - \mathbf{t}\| = d \tag{3.3}$$

If multiple time-difference measurements are available, the tag is located on the intersection of the hyperbolas/hyperboloids. The results of the localisation are available on the side of the static infrastructure, meaning that if the moving target requires to know its position within the environment, the results have to be transmitted to the tag node.

This work uses TDoA anchors developed at the Department of Radioelectronics of CTU FEE, details about the hardware and the algorithms used can be found in [38, 39]. These anchors use wireless communication to synchronise the clocks. One of the anchors is a master; the other anchors synchronise their clocks with the master. The tag is again the DWM1001 module from Qorvo. There must be ongoing communication between the TWR modules, which means that the robot must be equipped with at least one TWR module.

Each pair of anchors defines one time difference. For n anchors, the n-1 time differences can be used to localise the tag. Each time difference is a hyperbola where the fixed points are the positions of the anchors. The tag lies at the intersection of the hyperbolas. Figure 3.5 shows the hyperbolas for all combinations of available anchors. TDoA is used mainly in cases where the tag is inside the area delimited by the anchors. This is because even a small error in the time differences results in a large shift of the outer parts of the hyperbolas. If the time difference has an error of one nanosecond, the error of distance difference is already 30 centimetres. This error is comparable to the mutual distance of the anchors in this use case, the anchors in the experiments shown in Figure 3.5 were 50 centimetres apart from each other.

The error increases as the tag moves further away from the robot. Therefore, it is better not to use TDoA to obtain a precise position of the tag; instead a direction towards the tag can be estimated. Even after this change, the accuracy of the estimated direction is not high and the measurements are affected strongly affected by the body of the robot. The anchors synchronise their clocks wirelessly, but this means that the time of flight of the synchronisation signal has to be known. If the direct path between the tag and one of the anchors is blocked, this obstacle affects the whole pair of anchors, making the measured time difference inaccurate. For these reasons, using TDoA for the “follow me“ task is complicated.

3.2 Bluetooth Angle of Arrival

Bluetooth can be used to obtain range measurements, but the estimate is based purely on the strength of the received signal. The reception is affected by the presence of obstacles between the transmitter and the receiver, and the accuracy is much lower than in the case of UWB. Bluetooth version 5.1 added a new method for localisation of Bluetooth devices – Angle of Arrival and Angle of Departure. It adds a special new type of Bluetooth packet that

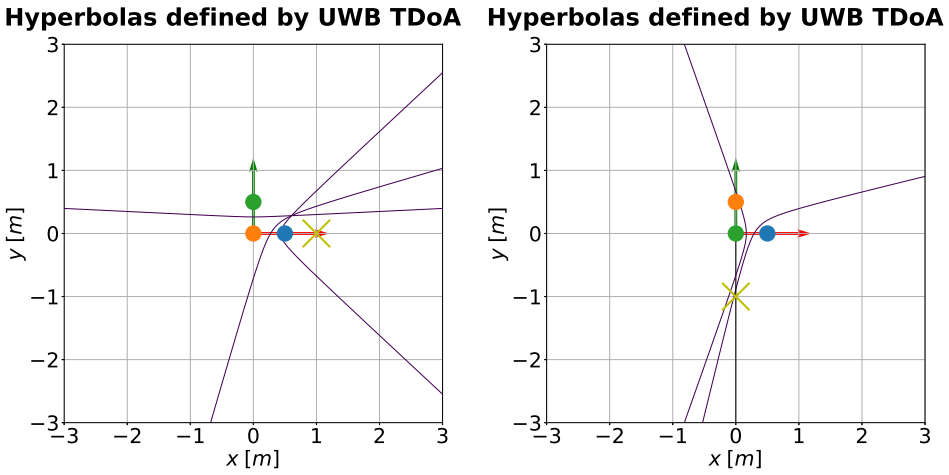

(a) (b)

Figure 3.5: TDoA hyperbolas for two positions of the tag. The position of the tag is shown with a yellow cross and the position of the anchors is shown with coloured dots. The results were obtained by averaging the time differences recorded for one minute. In the figure on the left, the time differences were incorrectly measured, leading to incorrect localisation of the tag.

enables the receiving device to estimate the direction towards the sender. The receiver has to have a whole antenna array to enable Angle of Arrival measurements, Angle of Departure requires the transmitter to use several antennas simultaneously. The principle of operation of the two methods is shown in Figure 3.6.

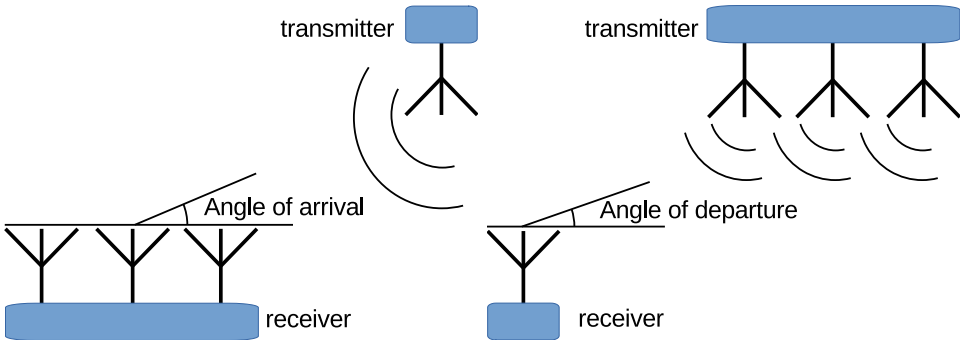

(a) The Angle of Arrival measurement (b) The Angle of Departure measurement

Figure 3.6: The two techniques for Bluetooth localisation

Angle of arrival requires two devices – a moving transmitter (tag) and several static receivers (anchors), each with an array of antennas. The tag sends periodic localisation packets; these are standard Bluetooth packets with Constant Tone Extension (CTE) added to their end. Bluetooth uses a specific frequency for each channel to transmit a symbol representing 0 or 1, CTE is a series of 1's, as a result this part of the packet is a signal

with constant frequency. As the tag moves with respect to the antenna array, the angles (azimuth and elevation) at which the signal arrives at the array change. Each individual antenna of the array receives the signal with a time delay, and this delay is related to the position of the tag. The CTE ensures that all antennas receive signal with a constant frequency. The time delay is therefore transformed into a phase delay. The phase delay measured at each antenna of the array can be used to determine the direction from which the signal is being received. The configuration of the antenna array determines the capabilities of the anchor, antennas arranged in a line allow measuring a single angle, and antennas arranged in a square allow to measure both azimuth and elevation relative to the plane of the antenna array.

To obtain the position of the tag just from the AoA measurements, several antenna arrays have to be used. Each array has a limited "field of view", the areas covered by the arrays have to overlap to enable localisation. Each AoA measurement defines a half-line starting at the reference point of the antenna array. Antenna arrays are static, which means that their position within a common reference frame is known. As a consequence, the position of the tag can be obtained as an intersection of several half-lines.

For this work, two u-blox AoA development kits were used, XPLRAOA-1 and ANT-B10 antenna board. The kits differ by the antenna array configuration — XPLRAOA-1 uses a simpler array comprised of five antennas arranged in an L shape (the antenna array is depicted in Figure 3.7(a)), ANT-B10 has eight antennas arranged in a square (the arrangement can be seen in Figure 3.7(b)). Both kits use UART (XPLRAOA-1 contains a UART to USB converter, ANT-B10 is connected to the UMFT234XF USB to UART development breakout board from FTDI[1]) to report the measured azimuth, elevation, and RSSI (Received Signal Strength Indicator). The documentation states that both kits should be capable of measuring angles in the range of -90 to 90 degrees (the whole half-space in front of the antenna array). Both kits use the same active C209 tag (shown in Figure 3.7(a)), this tag has a configurable advertising frequency and can be powered by a battery or by USB.

The antenna arrays are directional, meaning that the tag can be correctly localised only if it is located in the half-space in front of the antenna array. The array cannot identify the situation when the tag is behind the array, meaning that some measurements (incorrect) are being sent even in this case. The antenna array also measures RSSI of the signal, RSSI is indicator of the strength of the signal used for computing the angles of arrival. RSSI depends on the distance between the tag and the antenna array, the angle of arrival, and the position of the tag (in front of the array or behind it). In case an obstacle is present on the direct path of a signal, a reflected signal could be measured instead, leading to incorrectly reported angles. The antenna array can also be sensitive to occlusion; if part of the array is covered by an obstacle (covering some of the antennas by hand is enough), the reported angles could be wrong. The sensitivity to reflections depends on the transmit

[1] `https://ftdichip.com/products/umft234xf/`

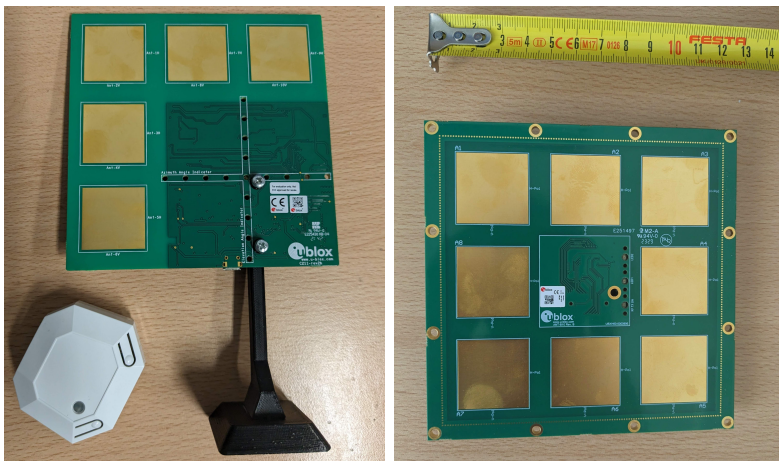

(a) The u-blox XPLRAOA-1 kit and the C209 tag.

(b) The u-blox ANT-B10 board.

Figure 3.7: The two Bluetooth angle of arrival development kits.

power of the tag; if the power is high, the range is increased, but the reflected signal is stronger too. This could mean that even slight occlusion can lead to incorrect measurements, because the direct path signal is attenuated and the reflected one is stronger. The C209 tag has a sensible default transmit power, the range is several metres, but the reflections are limited. It is possible to decrease the power (up to the level that measurements are available only up to around 30 centimetres) or increase it, but leaving it at the default value provides good results.

The accuracy and precision of both kits was evaluated in a static test scenario — the antenna array was mounted on a rotating base, and the tag was placed on a stand at various distances from the antenna array. The measured azimuth and RSSI were recorded for one minute, and their mean and standard deviation were computed. The results are presented in the following two subsections.

3.2.1 Evaluation of Accuracy and Precision of the XPLRAOA-1 Kit

The measured azimuth between the antenna and the tag was recorded for distances of 1, 2 and 4 metres between the antenna and the tag. The antenna was rotated between -80 and 80 degrees. An example of the data recorded for the distance of 4 metres between the antenna and the tag is shown in Figure 3.9. The sensitivity of the measurement to obstacles that were present on the line of sight path between the tag and the antenna array was tested by stepping between the two devices while the data were being recorded. A recording with a human standing approximately in the middle between the tag and the antenna is shown in Figure 3.10.

As can be seen in Figure 3.11, the measured angles roughly follow the

(a) The ANT-B10 antenna board placed on the rotating mount.

(b) The Bluetooth tag placed 2 metres away from the ANT-B10 antenna board.

Figure 3.8: The setup used to measure accuracy and precision of Bluetooth angle of arrival.

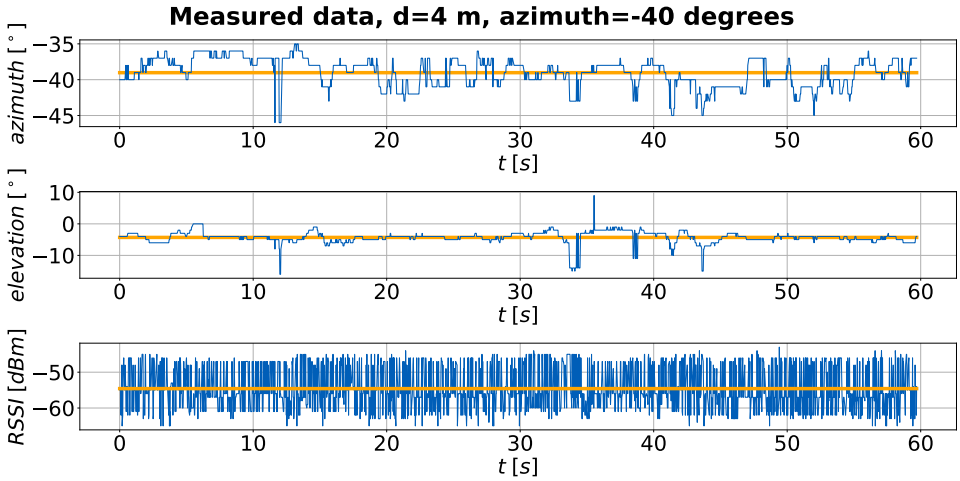

Figure 3.9: The measured azimuth, elevation and RSSI for distance of 4 metres between the antenna and the tag, the azimuth was approximately -40 degrees. The computed mean is shown as an orange line.

normal distribution. The mean and standard deviation can be calculated for each experiment; the results are in Tables 3.2, 3.3 and 3.4.

As can be seen in Figure 3.12, the accuracy of the azimuth measured drops at the edges of the measuring range. The accuracy also decreases with the increase in the distance between the antenna array and the tag. The RSSI also decreases with increasing distance, but it cannot be used as an indicator of the quality of the measurement. For example, for the distance of 4 metres, the RSSI at azimuth -80 degrees is higher than the one for +80 degrees, but the measurements for -80 degrees are not accurate (mean is -6.27 degrees)

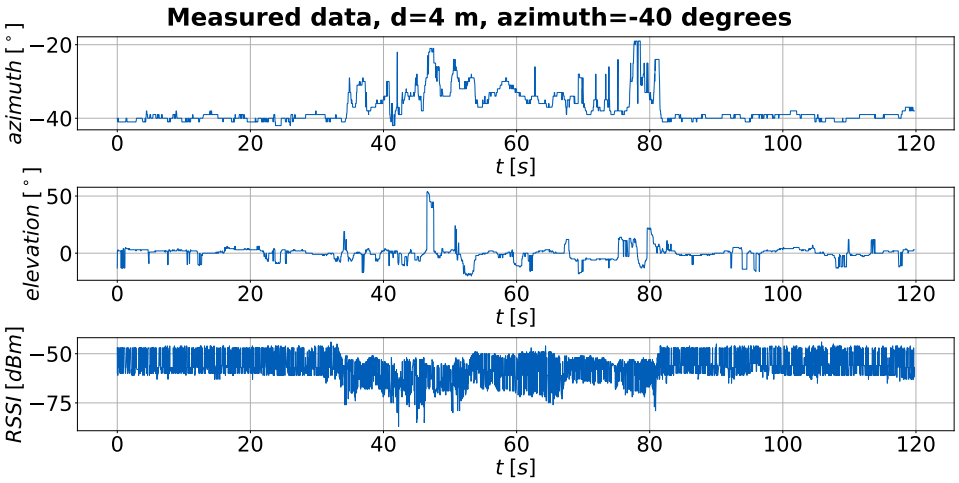

Figure 3.10: The same setup as for figure 3.9, this time a human stepped between the antenna and the tag. The reported RSSI was noticeably dropped and the azimuth measurements were affected by an offset.

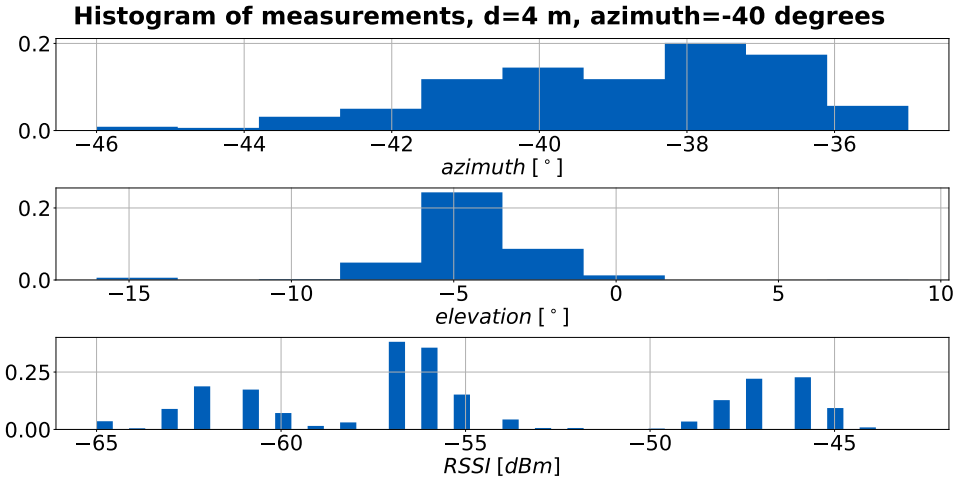

Figure 3.11: The histograms for the measured azimuth, elevation and RSSI shown in the Figure 3.9.

and have low precision (standard deviation 50.77 degrees). Contrary to what the RSSI would indicate, the measurements for +80 degrees are much more accurate (mean is 64.69) and precise (standard deviation is just 4.19 degrees).

The measurements are sufficiently accurate and precise in the middle of the sensor range (from -60 to 60 degrees), as can be seen in Figures 3.12 and 3.13. Measurements from just a single antenna array are sufficient only in the case that the human always stays in front of the robot. In all other cases multiple AoA antennas have to be used or the data from the Bluetooth AoA kit have to be fused with data from other sensors that also provide direction to the target.

azimuth	μ	σ	RSSI μ	RSSI σ
-80	-64.16	1.80	-51.90	3.30
-60	-55.65	2.07	-48.83	2.14
-40	-40.09	0.43	-45.31	1.77
-20	-18.69	1.26	-48.25	1.79
0	3.04	1.94	-47.81	2.29
20	21.04	1.92	-47.37	3.28
40	40.50	2.21	-45.68	3.62
60	60.80	2.03	-46.81	4.07
80	83.24	6.73	-49.31	4.64

Table 3.2: Measurements for distance of 1 meter.

azimuth	μ	σ	RSSI μ	RSSI σ
-80	-28.52	4.24	-60.44	8.91
-60	-55.17	4.36	-54.55	6.72
-40	-41.12	2.14	-52.06	5.27
-20	-19.88	3.51	-53.12	5.06
0	-0.10	5.05	-54.33	5.11
20	21.32	2.29	-53.88	5.41
40	35.48	2.47	-53.76	4.45
60	60.98	7.25	-54.10	5.24
80	66.22	3.01	-57.20	6.24

Table 3.3: Measurements for distance of 2 meters.

azimuth	μ	σ	RSSI μ	RSSI σ
-80	-6.27	50.77	-56.96	6.58
-60	-47.78	2.03	-54.88	6.48
-40	-39.02	1.97	-54.58	5.93
-20	-20.93	0.82	-55.55	5.47
0	-1.54	2.07	-56.34	4.90
20	20.43	3.91	-56.89	5.56
40	39.47	3.27	-54.63	5.91
60	53.82	2.19	-55.01	6.47
80	64.69	4.19	-60.22	8.56

Table 3.4: Measurements for distance of 4 meters.

In addition, RSSI cannot be used to detect an obstacle between the tag and the antenna. When the tag is stationary and an obstacle enters the space between the tag and the antenna, RSSI drops. This drop is noticeable, but when the tag moves, the RSSI changes as the distance and azimuth change. Therefore, any drop caused by an obstacle could also be caused by the change in relative position between the tag and the antenna.

The measured angles are not robust to obstacles present between the

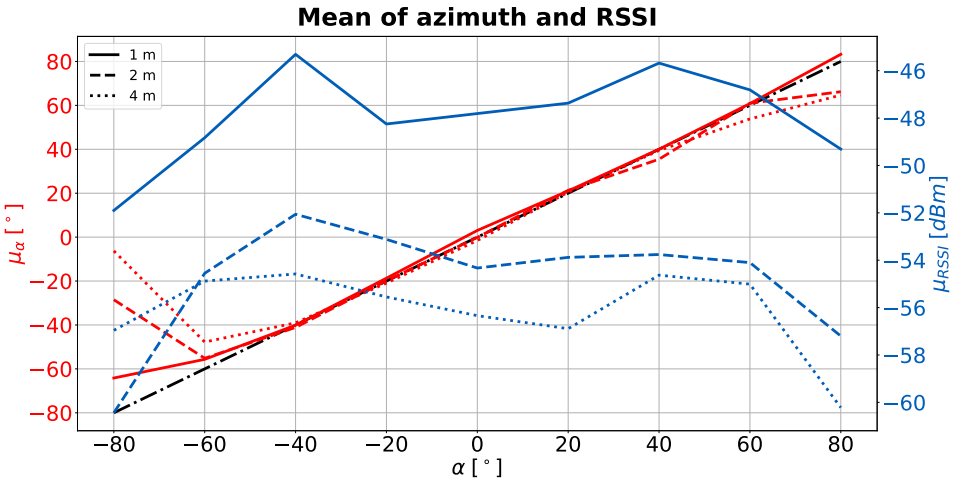

Figure 3.12: The mean of the measured azimuth (red) and RSSI (blue) for distance of 1 (solid line), 2 (dashed line) and 4 metres (dotted line). The correct azimuth is shown as a black dashed line to facilitate the comparison of the actual measurements.

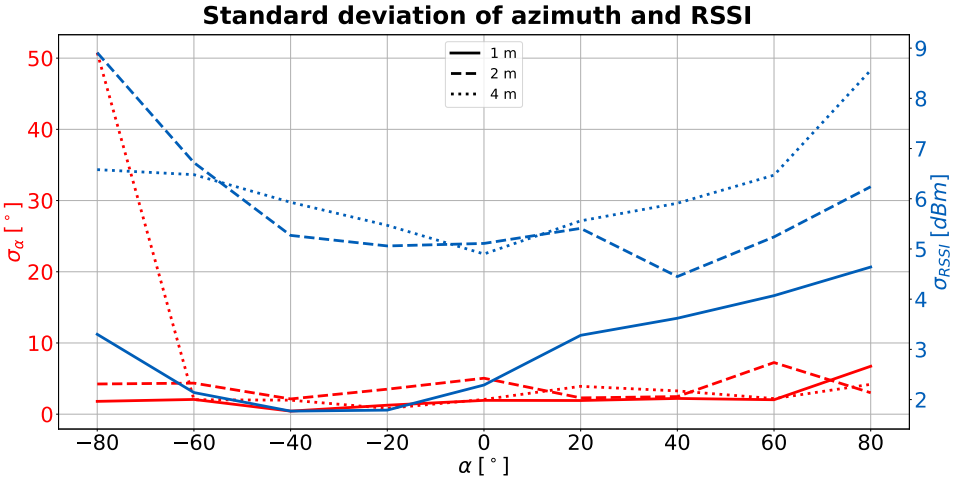

Figure 3.13: The standard deviation of the measured azimuth (red) and RSSI (blue) for distance of 1 (solid line), 2 (dashed line) and 4 metres (dotted line).

antenna and the tag. As can be seen in Figure 3.10, a single obstacle can shift the measured azimuth by a non-negligible amount. An even worse error is present when just part of the antenna is covered (for example, by some part of the robot's body or by a hand).

The low accuracy at the edges of the measuring range is a property of the sensor; the environment does not cause it. The work [40] measured the same development kit in an anechoic chamber and reached a similar conclusion about the usable range of angles. This work also shows the influence of reflections on the accuracy of measurements. The experiment described in this section was carried out in a laboratory that contains many objects;

therefore, there were probably many reflections that caused interference.

3.2.2 Evaluation of Accuracy and Precision of the ANT-B10 Antenna Array

The measurement procedure was identical to that presented for the XPLRAOA-1 kit. The measured data for distances of 1, 2 and 4 metres are shown in Tables 3.5, 3.6 and 3.7.

azimuth	μ	σ	RSSI μ	RSSI σ
-80	-48.34	8.04	-66.22	4.18
-70	-52.92	1.90	-67.08	5.86
-60	-44.66	4.33	-67.14	6.44
-50	-40.82	0.91	-66.35	6.43
-40	-34.63	0.98	-63.71	4.77
-30	-28.52	0.73	-61.90	3.31
-20	-20.95	0.74	-60.36	2.52
-10	-11.90	0.81	-58.82	1.92
0	-3.03	0.62	-58.46	1.79
10	7.15	1.14	-59.18	1.63
20	14.03	1.28	-59.42	1.79
30	24.33	1.91	-59.66	1.44
40	31.04	2.09	-60.76	1.93
50	41.09	3.21	-62.76	2.11
60	56.00	2.33	-68.14	5.14
70	55.92	2.55	-65.45	5.44
80	66.76	3.64	-64.91	3.24

Table 3.5: Measurements for distance of 1 meter.

The antenna array is symmetrical, resulting in symmetric inaccuracy at both edges of the measurement range, the difference is apparent when comparing Figures 3.14 and 3.12.

In case of this antenna array, the RSSI seems to be a good indicator of measurement quality. Figure 3.14 shows that RSSI is noticeably higher when the azimuth is close to 0 degrees (in this case the accuracy is also the best one), as the azimuth approaches the limits of the sensor, RSSI drops. This behaviour corresponds to the increasing inaccuracy of the measurements.

azimuth	μ	σ	RSSI μ	RSSI σ
-80	-39.39	1.35	-66.86	5.95
-70	-42.11	1.96	-64.89	4.16
-60	-43.64	1.86	-63.44	4.79
-50	-38.74	1.40	-63.62	5.62
-40	-33.85	1.17	-63.30	5.26
-30	-26.48	0.84	-61.75	4.01
-20	-18.18	0.91	-61.06	3.72
-10	-10.09	0.61	-61.58	3.99
0	-1.98	0.72	-62.15	4.19
10	5.98	1.12	-65.29	4.47
20	14.67	1.34	-64.72	3.97
30	26.89	2.14	-65.46	4.37
40	37.53	3.12	-65.43	4.90
50	45.05	3.77	-67.12	5.26
60	52.97	3.02	-67.57	3.96
70	59.15	4.60	-69.77	4.15
80	54.44	8.11	-69.53	4.17

Table 3.6: Measurements for distance of 2 meters.

azimuth	μ	σ	RSSI μ	RSSI σ
-80	-23.81	7.39	-68.85	2.83
-70	-28.18	3.67	-68.93	3.29
-60	-36.01	3.03	-69.66	4.37
-50	-31.29	2.87	-70.06	3.91
-40	-28.69	3.50	-69.77	2.07
-30	-23.57	2.06	-69.30	4.10
-20	-18.77	2.23	-69.26	4.08
-10	-13.06	1.64	-69.59	4.35
0	-5.22	1.68	-68.87	4.50
10	3.43	1.59	-69.86	4.80
20	13.10	2.00	-67.06	1.69
30	19.92	1.29	-67.25	2.26
40	25.06	1.72	-68.83	4.26
50	33.80	1.53	-69.34	4.33
60	37.37	1.60	-69.70	4.25
70	37.23	4.76	-70.48	4.48
80	40.97	4.06	-69.55	3.07

Table 3.7: Measurements for distance of 4 meters.

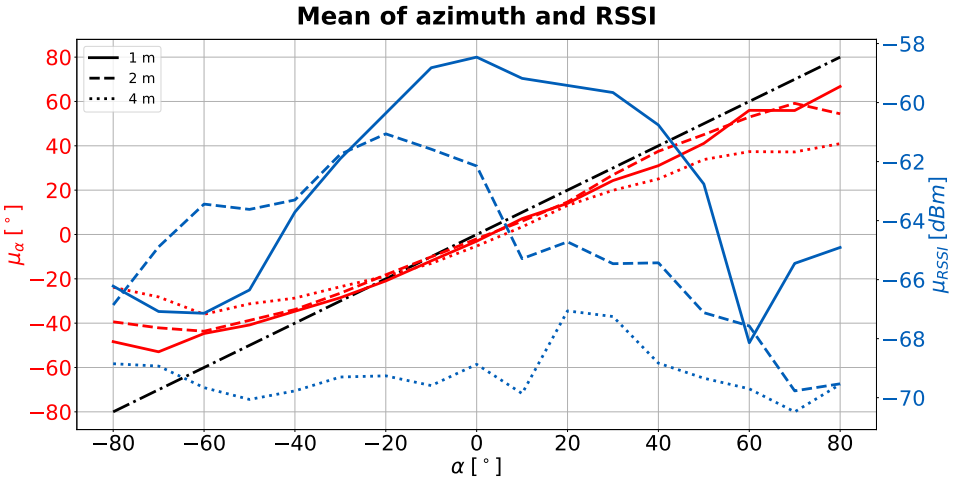

Figure 3.14: The mean of the measured azimuth (red) and RSSI (blue) for distance of 1 (solid line), 2 (dashed line) and 4 meters (dotted line). The ideal characteristic of the sensor is shown with the black dashed line.

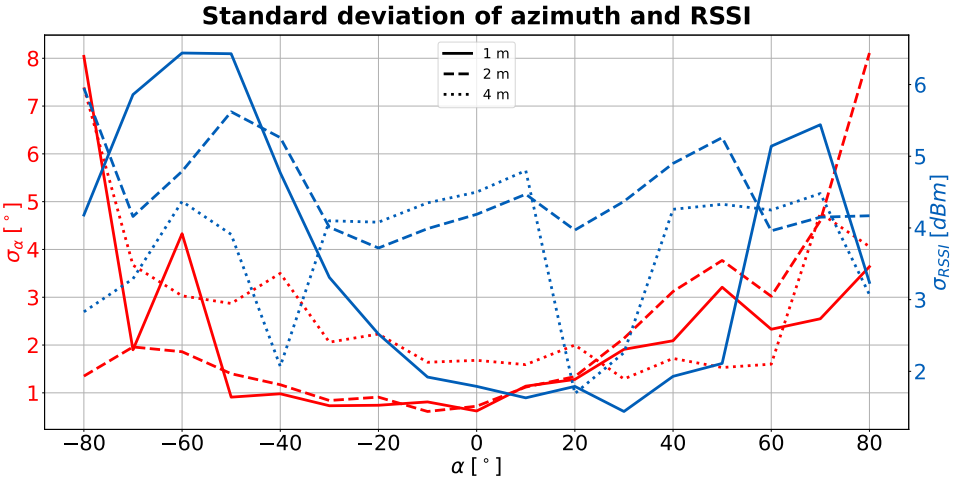

Figure 3.15: The standard deviation of the measured azimuth (red) and RSSI (blue) for distance of 1 (solid line), 2 (dashed line) and 4 meters (dotted line).

3.3 Mounting the Sensors on the Robot

This section describes the requirements for mounting the UWB modules and the Bluetooth antenna boards on a robot. The mounts have to respect the limitations of the used sensors and the positions should be selected in a way that brings useful measurements from each sensor.

3.3.1 Ultra-wideband TWR

Each UWB module receives the signal multiple times due to the reflections from the environment. The first received signal that is strong enough is considered to be the line of sight signal and the range measurement is obtained from it. If there is a large obstacle present on the line of sight trajectory, the signal might fall below the noise threshold and a reflected signal might be used instead, leading to incorrect range measurement. Some level of occlusion is acceptable, but the UWB modules should be mounted in positions that allow direct visibility to all directions.

An important factor that affects the accuracy and precision of the localisation is the configuration in which the modules are mounted on the robot. The closer the modules are to each other, the shorter is the range of localisation. The modules in principle perform triangulation, as the base of the triangle (the distance between the modules) gets shorter relative to the distance to the target, the lines connecting the modules and the target get nearly parallel. As a consequence, any error in the range measurements results in large lateral shift of the estimated position.

The UWB modules should be spread as far apart as possible within the limited footprint of the robot. If more than two modules are available, they should not be mounted in one line. Four and more modules allow 3D localisation of the target if the tags are mounted at different heights. Figures 3.16, 3.17 and 3.18 show a simulated localisation in different configurations of the modules, each of the modules measures range with error ranging from -3 to 3 centimetres, and intersection point for all possible combinations of errors is shown. It is clear that the modules should be spread as far away from each other as possible. The distance to the target will be accurately estimated in any configuration of the modules, but the lateral error greatly depends on the spacing of the modules.

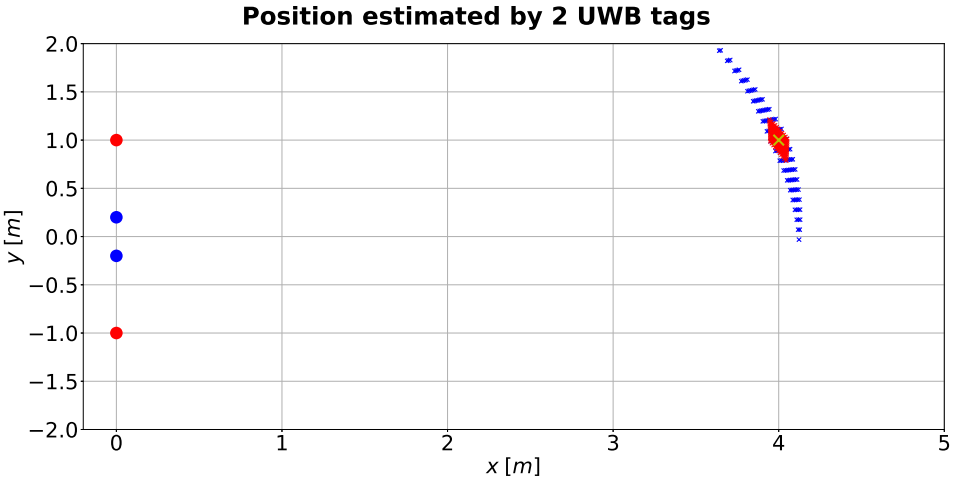

Figure 3.16: Simulated localisation with 2 UWB modules with spacing of 40 centimetres (blue) and 2 metres (red). The true position of the target is shown with a yellow cross.

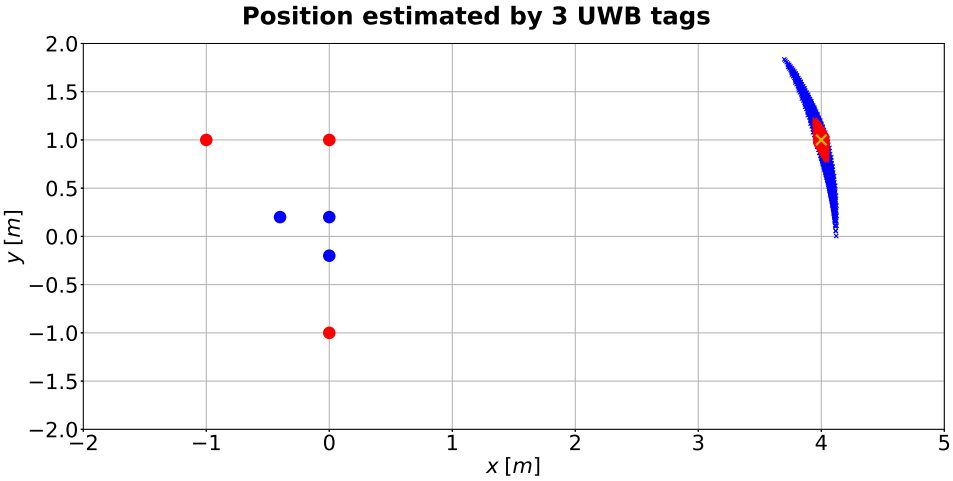

Figure 3.17: Simulated localisation with 3 UWB modules with spacing. In this case, the localisation is more accurate than in Figure 3.16.

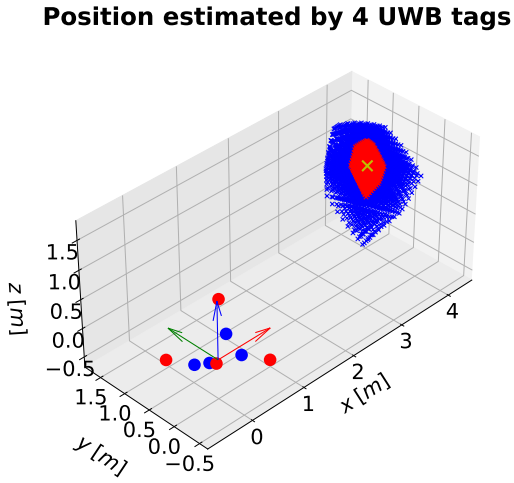

Figure 3.18: Simulated localisation with 4 UWB modules. The modules are not placed on a single plane, allowing for 3D localisation.

3.3.2 Bluetooth AoA

The antenna array requires direct line of sight visibility to the tag without major occlusions. This means that the requirements for mounting the antenna boards are similar to mounting cameras — it has to be mounted on such a place that it has direct visibility of the area of measurement. According to the data sheet for the antenna boards (CITE), the antennas can be under a plastic cover, this cover has to be at least 1 centimetre away from the antenna board. This makes the antenna board with a protective case quite large (12.6 by 12.6 centimetres with thickness of more than 1 centimetre); therefore, they cannot be mounted arbitrarily on the robot.

An example of this is shown in Figure 3.19, the antennas are mounted on a Spot robot, but they should not occlude the cameras and the LiDAR. It is not possible to mount the antennas to the side of the robot body or the payload (frame attached on top of the robot), due to possible collisions with robot's legs. A mounting bracket that allows one to construct a cube from the antenna boards was designed; the bracket is mounted at the bottom and top of the antennas for improved rigidity. The two brackets are identical (the bracket is depicted in Figures 3.20(a) and 3.20(b)), allowing additional equipment to be mounted on top of the cube.

Figure 3.19: The four Bluetooth antenna boards mounted as a single cube on top of Ouster OS-0 LiDAR on Spot robot.

This way of mounting the antenna boards on the robot has the additional benefit of making the deployment much easier — a single measurement is required to determine the position of all four antenna boards (it is necessary to measure the position of the centre of the bottom bracket with respect to some known frame of reference on the body of the robot). A USB hub can be placed inside the cube, meaning that from the point of view of installing the cube, it is one device with a single USB connector.

(a) One of the brackets used to assemble the cube

(b) The assembled cube with ANT-B10 boards

Figure 3.20: The mounting bracket for four u-blox ANT-B10 antenna boards

3.4 Sensor Package for the Human

The proposed solution uses two independent radio-based localisation systems. This means that the human has to carry two transmitters. The DWM1001-dev UWB module and the C209 Bluetooth tag have the option of inserting a battery into the enclosure. The downside is that each device has its own battery that has to be recharged or replaced (the DWM1001-dev boards lack built-in battery charging circuit since 2021) separately and it is not possible to easily power on/off the whole tag consisting of the two transmitters. The DWM1001-dev board has a Raspberry Pi header, the header contains 5 volt power. The Bluetooth tag can be powered through the Micro USB connector; therefore, a cable can be chained to the Raspberry Pi header on the UWB board.

A single enclosure for both modules was designed and 3D printed, the model can be seen in Figure 3.21. This enclosure contains two 18650 cells with a Li-ion cell charging module, this battery powers both transmitters, and the charging module contains a power button that controls both modules. The charging module also contains a USB-C connector to charge the two cells, which means that the tag does not have to be disassembled if the battery is empty. The enclosure leaves the Micro USB port of the UWB module open; this means that its configuration can be changed, and it can even be operated in the tag mode; the measurements can then be read on a computer connected to the USB port. The USB port on the Bluetooth tag is occupied by the power supply, but it is still possible to change its configuration via Bluetooth with the Android nRF Connect app. The assembled tag with its dimensions can be seen in Figure 3.22(a), the internal wiring is shown in Figure 3.22(b).

The enclosure is larger than each of the individual modules, but it can

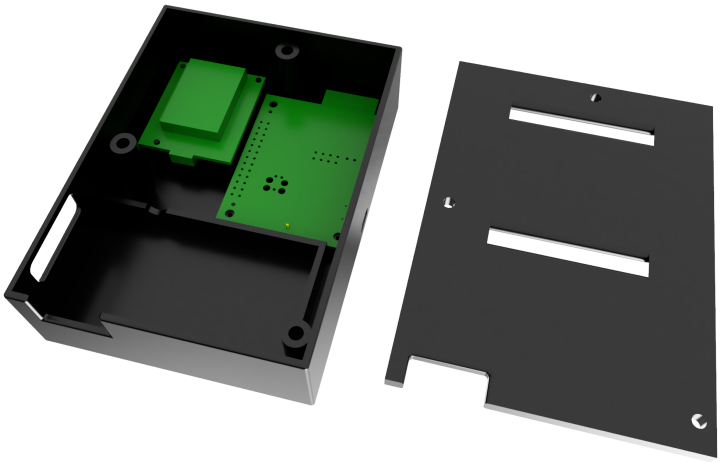

Figure 3.21: The 3D model of the enclosure for both transmitters.

(a) The front of the tag with the holes for the status LEDs

(b) The inside of the tag for human

Figure 3.22: The 3D printed enclosure for UWB module and Bluetooth tag

still fit larger pockets. It also contains loop holes for mounting the box on a strap. On the front side of the box, there are holes that make the status LEDs of both modules visible. It would be possible to make the box smaller by stacking the two transmitters on top of each other, but this placement means that it would not be possible to keep the status LEDs visible for both modules, or the USB port of the UWB module would not be accessible. Stacking the boards could also have a negative impact on the signal strength, one of the boards could obstruct the antenna of the other board, meaning that the strength of the signal received by the robot could be significantly decreased in some orientations of the tag.

Chapter 4

Architecture of the System

The "follow me" system is implemented using Robot Operating System (ROS) [41]. ROS promotes a modular architecture in which software units (nodes) are responsible for small tasks, and ROS provides a communication interface (messages) that nodes use to exchange information. All of the sensors used send the measured data through a serial interface. The sensors are connected to a computer running ROS, a ROS driver was created for each sensor to transform serial messages into ROS messages, details about serial communication are available in the Appendix A and the custom ROS messages used in this work are described in the Appendix B. ROS offers a tool for visualising sensor data — rviz, the nodes built for this system use rviz to visualise the measurements (UWB range measurements, AoA angle measurements).

Figure 4.1 shows the scheme of the system. Each of the sensors has its own driver that publishes the measurements as ROS messages. There is a node that takes measurements from the sensors and uses them to estimate human position. The estimate is then received by the velocity controller; the controller also receives range measurements from UWB TWR modules to stop the robot immediately if the human is too close. The controller outputs desired velocities (linear and angular) that the driver of the robot uses to control the robot's motion.

4.1 Two-Way Ranging Localisation

The Two-Way Ranging (TWR) range measurements define a set of spheres with centres at the mounting points of the TWR modules on the robot. The intersection of these spheres is the position of the moving UWB module. In case of no inaccuracy of the range measurements, the spheres will intersect at the exact position of the UWB module that is being localised (there can be multiple intersection points of the spheres). The range measurements are not perfect; therefore, there will be no point at which all of the spheres intersect (assuming that more than just two measurements are available, in the opposite case an intersection of the two spheres exists in most scenarios); instead, there will be an area where all points are close to all of the spheres. The task of position estimation is formulated as a non-linear least-squares

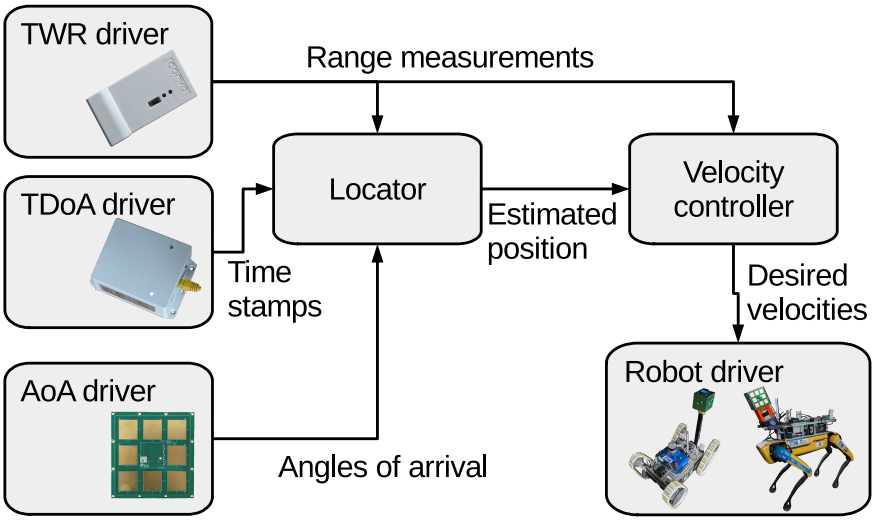

Figure 4.1: The architecture of the "follow me" system. Each box represents a ROS node, and arrows denote the communication between the nodes, e.g. the locator node subscribes messages with angles of arrival published by the AoA driver node.

optimisation problem. The difference between the estimated distances of the human from the UWB tags and the measured ones has to be minimised:

$$\boldsymbol{\Theta}^* = \arg\min_{\boldsymbol{\Theta}} \|\mathbf{f}(\boldsymbol{\Theta}) - \mathbf{m}\|^2 \tag{4.1}$$

where $\boldsymbol{\Theta}$ is the estimated 3D position of the human ($\boldsymbol{\Theta} = \left[\hat{x}, \hat{y}, \hat{z}\right]$) and $\mathbf{m}$ is a vector of measurements ($\mathbf{m} = \left[r_1, \ldots, r_n\right]$). r_1, ...,r_n are the distances reported by the UWB tags. $\mathbf{f}(\boldsymbol{\Theta})$ is a function of the estimated position:

$$\mathbf{f}(\boldsymbol{\Theta}) = \begin{bmatrix} \|\boldsymbol{\Theta} - \mathbf{t}_1\|^2 \\ \vdots \\ \|\boldsymbol{\Theta} - \mathbf{t}_n\|^2 \end{bmatrix} \tag{4.2}$$

where $\mathbf{t}_1$, ...,$\mathbf{t}_n$ are the positions of the UWB tags in the reference frame.

The resulting position is still influenced by the measurement noise. The sequence of the estimates can be filtered with a Kalman filter. Humans are capable of moving in all directions in 2D (forward/back but also sideways), it is expected that the robot moves in such a way that the human stays approximately at the same relative position (in front of the robot at the desired distance). The system can therefore be described as:

$$\begin{bmatrix} x_{k+1} \\ y_{k+1} \\ z_{k+1} \end{bmatrix} = \mathtt{A} \begin{bmatrix} x_k \\ y_k \\ z_k \end{bmatrix} \tag{4.3}$$

$$\begin{bmatrix} x_{k+1} \\ y_{k+1} \\ z_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_k \\ y_k \\ z_k \end{bmatrix} \tag{4.4}$$

The covariance matrices of the LKF were manually set to be:

$$\mathtt{R} = \begin{bmatrix} 50 & 0 & 0 \\ 0 & 50 & 0 \\ 0 & 0 & 50 \end{bmatrix} \tag{4.5}$$

$$\mathtt{Q} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{4.6}$$

The measured states are the position in the x, y, and z axes. The prediction step can, therefore be described as:

$$\hat{\mathbf{x}}_{k+1} = \mathtt{A}\hat{\mathbf{x}}_k \tag{4.7}$$

$$\hat{\mathtt{P}}_{k+1} = \mathtt{A}\hat{\mathtt{P}}_k\mathtt{A}^T + \mathtt{Q} \tag{4.8}$$

The correction step is:

$$\mathtt{H} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{4.9}$$

$$\mathtt{K} = \hat{\mathtt{P}}\mathtt{H}^T(\mathtt{H}\mathtt{P}\mathtt{H}^T + \mathtt{R})^{-1} \tag{4.10}$$

$$\hat{\mathbf{x}} = \hat{\mathbf{x}} + \mathtt{K}(\boldsymbol{\Theta}^* - \mathtt{H}\hat{\mathbf{x}}) \tag{4.11}$$

$$\hat{\mathtt{P}} = (\mathtt{I}_3 - \mathtt{K}\mathtt{H})\hat{\mathtt{P}} \tag{4.12}$$

where $\boldsymbol{\Theta}^*$ is the estimated position of the human obtained in Equation 4.1, and $\mathtt{I}_3$ is the identity matrix of size 3x3.

4.2 Time Difference of Arrival Localisation

The Time Difference of Arrival (TDoA) measurements define several hyperbolas, and the human should be positioned at the intersection of these hyperbolas. The position of the human can again be obtained with the non-linear least squares method. The function that is being minimised has the following form:

$$\boldsymbol{\Theta}^* = \arg\min_{\boldsymbol{\Theta}} \|\mathbf{f}(\boldsymbol{\Theta}) - \mathbf{m}\|^2 \tag{4.13}$$

$\boldsymbol{\Theta}$ is again the estimated 3D position of the human ($\boldsymbol{\Theta} = \begin{bmatrix}\hat{x}, \hat{y}, \hat{z}\end{bmatrix}$) and $\mathbf{m}$ is a vector of measurements ($\mathbf{m} = \begin{bmatrix}d_1, \ldots, d_n\end{bmatrix}$). d_1, ...,d_n are the distance differences of a pair of TDoA anchors. $\mathbf{f}(\boldsymbol{\Theta})$ is a function of the estimated position:

$$\mathbf{f}(\boldsymbol{\Theta}) = \begin{bmatrix} \|\boldsymbol{\Theta} - \mathbf{a}_1\| - \|\boldsymbol{\Theta} - \mathbf{a}_2\| \\ \vdots \\ \|\boldsymbol{\Theta} - \mathbf{a}_{n-1}\| - \|\boldsymbol{\Theta} - \mathbf{a}_n\| \end{bmatrix} \tag{4.14}$$

where $\mathbf{a}_1$, ...,$\mathbf{a}_n$ are the positions of the TDoA anchors in the reference frame.

TDoA is usually used in the way the tag moves inside the polygon defined by the TDoA anchors. In this scenario, the human is always outside of this polygon. This results in a much greater influence of the measurement uncertainty on the final estimate. The outer parts of each hyperbola noticeably change position even if the measurement changes just a little bit. As the tag gets further away from the anchors, this effect becomes stronger. At the scale of the robot, the tag will always be placed at a distance that is several times larger than the mutual distance of the anchors. Therefore, it is not possible to try to obtain a 2D or 3D position of the human, but the heading towards the human can be estimated. This estimate is still affected by a large error caused by the inaccuracy of each measurement, meaning that the estimate obtained is not usable for the task of following a human.

The output of NLS optimisation ($\mathbf{\Theta}^*$) is a point in the 3D space. Its distance from the robot is very inaccurate, but the direction defined by this point corresponds to the true direction towards the UWB tag. Therefore, the azimuth of the vector defined by $\mathbf{\Theta}^*$ is computed and used the same way as the Bluetooth AoA measurement.

4.3 Angle of Arrival Localisation

When Angle of Arrival (AoA) is used for indoor localisation, the areas covered by the antennas overlap, so the position of the tag can be obtained. The "follow me" scenario requires the localisation to be omnidirectional. This would require mounting at least two antenna arrays to each side of the robot, but such a setup is impractical and takes up a lot of space on the body of the robot. If the goal is to obtain just the direction towards the human, a measurement from a single antenna array is sufficient, meaning that four antenna arrays are sufficient to cover the whole area around the robot.

The antenna that provides the correct measurement has to be selected because the antennas provide some incorrect measurements even if the tag is outside of the usable range. If some other system provides the position of the human (for example, UWB TWR), this estimate can be used to select the correct antenna. If no such information is available or the AoA localisation should be completely independent, the signal parameters have to be used to determine the antenna that provides the correct measurements.

The proposed AoA localisation node can be configured to either use external estimate of the human's position or to operate independently. The independent mode uses the RSSI that each antenna provides with the measured angle to select the most relevant measurement. The queue with the last few RSSI values is kept for each antenna, and the mean and standard deviation are computed for each antenna. The antenna with the highest value of $RSSI_{mean} - 2RSSI_{std}$ is selected as the best antenna at the current time. This value is used to select the antenna that really provides the best measurements, as was shown earlier, the mean but also the standard deviation of RSSI are correlated with the quality of the measurement. Two antennas (one with the correct AoA and the other receiving a reflected signal)

could have nearly identical mean RSSI, but the one that provides incorrect measurements has a higher standard deviation of RSSI. The history of the selection is also kept; after the best antenna is selected, the current selection is compared with the history. If the selected antenna is not in the majority of the last few iterations, the selection is changed to match the majority. The pseudocode of the antenna selection algorithm is shown in Listing 4.1.

Listing 4.1: Algorithm for selecting the correct antenna just from reported RSSI

```
best_rssi = -INF
selected_antenna = None
for each antenna:
    rssi = mean(antenna_queue)-2*std(antenna_queue)
    if rssi > best_rssi:
        best_rssi = rssi
        selected_antenna = antenna
history.append(selected_antenna)
selected_antenna = majority(history)
```

After one of the antennas is selected, the measurement is brought into a common coordinate frame and passed through a Kalman filter, as a point on a unit sphere to make the measurements continuous. The Kalman filter is identical to the one used in Section 4.1, the states are again the coordinates in the x, y and z axis. This results in a steadier estimate with no large jumps between the measurements; these jumps are caused by switching from one antenna to another.

4.4 Fusing Sensor Measurements From All Available Sources

The human localisation node is built on the nonlinear least squares (NLS) optimisation problem described in Equation 4.1. TDoA localisation uses NLS optimisation, too; therefore, it is possible to combine these two optimisation problems into one. The angle of arrival measurements obtained from the algorithm described in Section 4.3 can be included as another element of the vector function $\mathbf{f}(\boldsymbol{\Theta})$. The new modified optimisation problem is:

$$\boldsymbol{\Theta}^* = \arg\min_{\boldsymbol{\Theta}} \left\| \tilde{\mathbf{f}}(\boldsymbol{\Theta}) - \mathbf{m} \right\|^2 \tag{4.15}$$

the function $\tilde{\mathbf{f}}(\boldsymbol{\Theta})$ has the following form:

$$\tilde{\mathbf{f}}(\boldsymbol{\Theta}) = \begin{bmatrix} \|\boldsymbol{\Theta} - \mathbf{t}_1\|^2 \\ \vdots \\ \|\boldsymbol{\Theta} - \mathbf{t}_n\|^2 \\ \|\boldsymbol{\Theta} - \mathbf{a}_1\| - \|\boldsymbol{\Theta} - \mathbf{a}_2\| \\ \vdots \\ \|\boldsymbol{\Theta} - \mathbf{a}_{k-1}\| - \|\boldsymbol{\Theta} - \mathbf{a}_k\| \\ \arccos \frac{\boldsymbol{\Theta}^T \mathbf{b}}{\|\boldsymbol{\Theta}\|\|\mathbf{b}\|} \end{bmatrix} \tag{4.16}$$

where $\boldsymbol{\Theta}$ is the estimated position, $\mathbf{t}_1$, ...,$\mathbf{t}_n$ are the positions of the TWR tags in the reference frame, $\mathbf{a}_1$, ...,$\mathbf{a}_k$ are the positions of the TDoA anchors in the reference frame, and $\mathbf{b}$ is the direction towards the transmitter obtained from Bluetooth AoA. The vector $\mathbf{m}$ now contains all the measurements, $\mathbf{m} = \begin{bmatrix} r_1, \ldots, r_n, d_1, \ldots, d_k, 0 \end{bmatrix}$. The zero is present at the end of the vector, because the last element of the function $\tilde{\mathbf{f}}(\boldsymbol{\Theta})$ is the angle between the true direction to the target (represented by the vector $\boldsymbol{\Theta}$) and the measured one ($\mathbf{b}$).

The sequence of estimates can be filtered with the same Kalman filter as was described in Section 4.1. The resulting estimates of the position of the human can then be used to control the motion of the robot.

4.5 Solving the Non-linear Least-squares Problem

Once the nodes described in Sections 4.1, 4.2 and 4.4 receive the necessary problems, the vector function $\mathbf{f}(\boldsymbol{\Theta})$ or $\tilde{\mathbf{f}}(\boldsymbol{\Theta})$ and the vector $\mathbf{m}$ can be constructed. The next step is to solve the NLS optimisation problem. SciPy Python library is used to find the solution with the `least_squares` function. The Trust Region Reflective algorithm [42] is used.

An initial guess has to be provided; in case the localisation node is already running, the estimate obtained in the previous iteration is used. This is possible because the human is expected to be at approximately the same position as in the previous iteration because the velocity of the human is not that high. At the start-up, the position of the human is unknown, a guess is therefore constructed from one of the TWR measurements. Four points in the circle defined by one of the TWR measurements are selected; the circle has a centre at the TWR module and is parallel to the x-y plane of the main reference system of the robot. This leads to up to four unique solutions to the NLS problem. The solution with the lowest value of the criterion function is selected, since this is the solution that best matches the set of equations. This approach should lead to a good initial estimate that has the same quality regardless of the relative position of the human.

The solution of the optimisation problem is later used as was described in the respective sections. The optimisation might not converge to a solution, the `least_squares` function reports this; in such a case no further processing is done and the position of the human will be estimated in the next iteration.

4.6 Velocity Control

This control node sets the velocities of the robot (linear and angular) so that the robot drives straight to the current position of the human. This means that the robot might take a different path than the human (the robot can cut corners), meaning that if the robot has no obstacle avoidance capability, a collision is possible. This control method works best in simple environments without many obstacles.

The linear and angular velocity are obtained from a PID controller, the linear velocity is set to maintain the desired distance to the human, the angular velocity controller tries to keep the robot heading straight towards the human.

4.7 Path Following

This control node uses the same path as the human, which means that there should be no collisions with the environment (in the case that the environment contains only static obstacles). The robot has to be able to localise itself within the environment, odometry is sufficient for this control method. The path of the human is stored in a static reference frame, the robot follows waypoints on the stored trajectory. The linear velocity of the robot is again controlled to maintain the desired distance from the human.

4.8 State System Feedback

All other parts of the system publish a string containing their status reports. There is an ROS node that subscribes to this topic and places the messages in a queue. Messages are played using ROS package sound_play, the node makes sure that just a single message is being played at a time, it also reduces frequency of identical messages.

4.9 System Deployment

The first step in deploying the "follow me" system on a new robot is to select a good placement for each sensor. Some general guidelines on mounting the sensors were already mentioned in Section 3.3, a few examples of possible sensor configurations are provided in the next chapter.

The next step is to measure the position of each sensor as accurately as possible. The position of all sensors does not have to be measured in the same reference frame, but the system expects the positions of the same sensors to be provided in a single reference frame (for example the position of all Bluetooth antenna boards is measured in the same reference frame). The position should be measured towards the correct point on the body of the sensor (or at least towards the same point for all sensors of the same type) — the Bluetooth

antennas measure the angles from the centre of the board, the UWB modules measure the distance from the antenna (the antenna is located in the upper left corner when looking on the front side of the enclosure as shown in Figure 4.2). The measured positions together with the reference frame have to be set in the configuration file. Each system (UWB and Bluetooth) has its own configuration file, and the values are then loaded as ROS parameters. The system takes care of creating the necessary tf frames (the reference frame should already exist and be connected to the rest of the tf tree).

Figure 4.2: The placement of the antenna of the Qorvo TWR module is marked with green cross.

The last step is to configure the USB ports for all devices in the corresponding launch files. All of the sensors are connected through USB, strong vibrations could cause the connection to drop out temporarily. The drivers for both types of sensors are created in a way that allows to restore the connection if the device reconnects to the same port. Therefore, it is recommended to create an UDEV rule for all of the sensors. This means that the ports have to be set up just once (if some other devices are connected through USB, the ports could change at each start of the computer) and the system should be able to recover from a minor failure.

Appendix C shows what changes have to be made when the system is deployed on a new robot.

The system also allows the user to calibrate the UWB modules by entering coefficients of a polynomial. Calibrating the modules could increase the accuracy, but in general it is hard to come up with a procedure that would improve the results in most scenarios. As the sensors are mounted on the robot, the interference caused by the body of the robot changes with respect to the position of the human. This is caused by the signal having to pass through the body of the robot in some configurations (the body is possibly metal, meaning that the direct path could be completely blocked and some signal reflected from the body of the robot is received instead). This leads

to increased performance for some particular direction (for example, higher precision when the human is in front of the robot) but decrease in some other. If it is possible to ensure that the robot body has no influence on the measurement (the sensors are mounted on rods high above the robot, resulting in an unobstructed view of the whole surroundings), calibration will improve the accuracy.

Chapter 5

Experiments

Three types of experiments were carried out to evaluate the performance of the proposed system. The first test is static; the human sensor pack was placed in various positions around the robot. This test aims to verify the accuracy and precision of the localisation. In the second experiment, the robot's movement is disabled and the human walks on a predefined trajectory around the robot. This test verifies that the tracking works continuously in all directions. The ground truth position of the human is obtained through tracking the human in the LiDAR point clouds. The third test scenario is dynamic - the robot autonomously follows the human. The ground truth is again obtained from LiDAR point clouds; the objective is to successfully follow the human while always maintaining a safe distance (the robot has to never be closer than the pre-defined safety margin while it is moving).

The third experiment was also used to evaluate the perceived safety of the system. Users should feel that the system is beneficial to them. The system does not only have to be safe, but the user should also feel safe while using the system. Perceived safety goes beyond a simple binary feature such as no collision/collision occurred. Evaluation of perceived safety is a research topic that combines robotics and psychology. The work [33] lists several approaches to increase the perceived safety of the human. The robot should maintain a proper distance from the human. This distance is larger than the one that would be necessary to avoid collision, the robot should respect proxemics — it should avoid the intimate and personal zones of the human (less than 1.2 metres). Another factor contributing to the safety of the system is smooth motion and low velocity when the robot is close to the human. This means that the robot should gently slow down as it gets closer, the only exception is if the distance is below a safety threshold — in this case, the robot should stop immediately. The user comfort can be increased by making it clear that the robot is still following the human, this means that the gap remains small. It is also beneficial if the human can understand the intentions of the robot.

5.1 Experimental Setups

Multiple robots with different sensor configurations were used to verify that the system offers acceptable performance under a wide range of operating

conditions. Different numbers and placement of sensors were used to show the possibilities of mounting the sensors on a robot. The tests were carried out outdoors and inside a building to show that the system is capable of working even in the presence of obstacles that would reflect the signal. The outdoor environment can be seen in Figures 5.1(a) and 5.1(b) together with two of the robotic platforms used for the tests. All tests on each robotic platform were carried out at once to allow for a direct comparison of the possible sensor configurations.

(a) (b)

Figure 5.1: The two robotic platforms in the outdoor environment used for testing.

5.1.1 Spot Robot with Three TWR Modules

Due to the dimensions of the robot, only three UWB modules can be mounted in a way that makes their mutual distance sufficiently large for human localisation (the payload is only 18 centimetres wide).

The placement of the modules can be seen in Figure 5.2. The front UWB module had to be mounted on a rod sticking to one side of the robot. At the front of the payload (the frame attached on top of the robot) is a mount for one camera and the 3D LiDAR. This mount is a thick 3D printed part that contains a metal plate. The LiDAR itself has metal heat sinks at the bottom and top. All of this material attenuates the signal. If the front UWB module was placed right behind the mount, it would work unreliably when the human is in front of the robot. Offsetting the module to the side partially mitigates this problem; the mount still affects the module, but in a different direction.

The position of the human is obtained with the method described in Section 4.1.

5.1.2 Spot Robot with Three TWR Modules and Four AoA Antennas

As was already discussed earlier, the mutual distance of the TWR modules mounted on the robot has an impact on the accuracy of the measurements. As the distance between the robot and the human increases, the accuracy

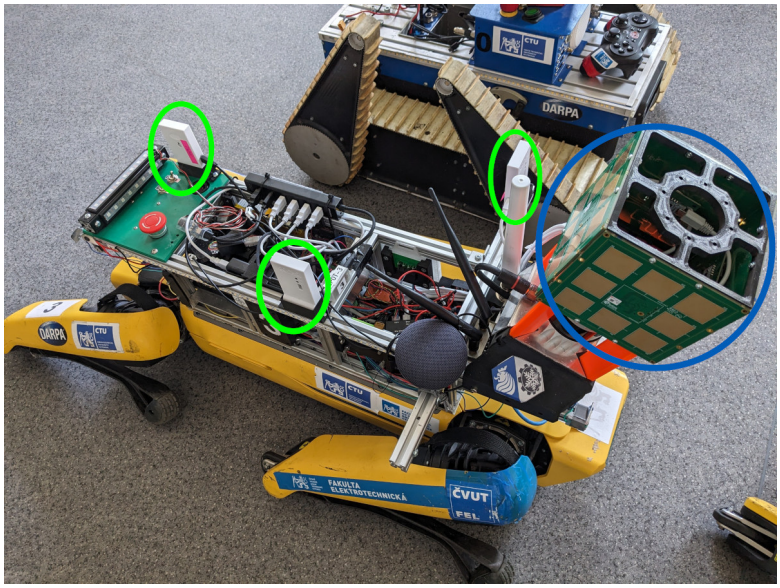

Figure 5.2: The Spot robot equipped with three UWB TWR modules (marked with green colour and four Bluetooth AoA antenna boards (in blue circle).

decreases rapidly. TWR modules on their own can be used to localise human only up to a few metres. The addition of the AoA measurement means that there is another source of information that does not depend on the distance. When the gap between human and robot is large, the localisation in principle reduces to intersection of a circle and a half-line (all of the UWB modules report nearly the same range measurements apart from the noise). The placement of the AoA antenna boards on the robot is depicted in Figure 5.2.

The human position is obtained with the method described in Section 4.4, but since TDoA was not mounted on the robot, its use was disabled (it was omitted from the optimisation problem). A single angle of arrival measurement is obtained with the algorithm described in Section 4.3.

5.1.3 Tracked Robot with Four TWR Modules

Localisation with just range measurements is in principle about finding an intersection point of several spheres. When just three measurements are available, the spheres could have up to two intersection points, meaning that the position of the human is not uniquely determined (the tag can be above or below the plane defined by the three UWB modules). By adding a fourth range measurement, a unique 3D position of the human can be obtained.

Compared to the deployment on the Spot robot, the tracked robot has a bigger footprint, the width is doubled, and the length is approximately 60 centimetres.

The position of the human is obtained with the method described in Section 4.1.

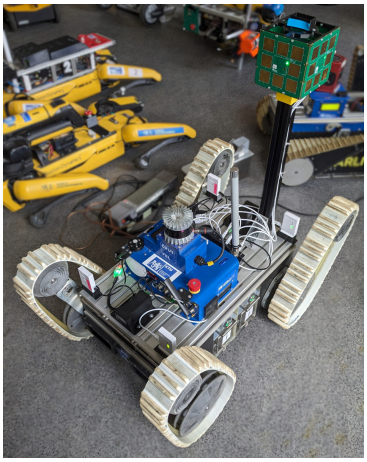

Figure 5.3: The tracked robot with the UWB modules mounted in the corners of the top plate, there is also a cube of AoA antennas mounted on a rod above the body of the robot.

5.1.4 Single Unit Containing Four AoA Antennas and One TWR Module

This sensor configuration could make system deployment really simple. The hardware necessary for localisation can be a single cube (shown in Figure 5.4) that can be easily mounted on the robot and it is just necessary to know the position of the mounting point of the whole cube; the position of each individual sensor with respect to the cube's mounting point is already known.

Figure 5.4: The cube composed of four Bluetooth AoA antennas with a UWB module mounted on top.

This "localisation device" was tested on both robotic platforms, but since the mutual distance of the sensors remains the same, the performance of the localisation should not be affected by the used platform. The only way the robot could influence the accuracy is by attenuation of the signal caused by

large occlusion of the sensors in some particular direction.

The human position is obtained with the method described in Section 4.4, but since TDoA was not mounted on the robot, its use was disabled (it was omitted from the optimisation problem). A single angle of arrival measurement is obtained with the algorithm described in Section 4.3.

5.2 Obtaining Ground Truth Position

All the robots used for the testing of the proposed solution were equipped with 3D LiDAR. In suitable conditions (direct line of sight visibility, no smoke or dust in the air), the LiDAR can provide much more accurate position estimates than all of the other proposed localisation systems.

Detection of the human starts by down-sampling the point cloud and cropping a smaller area around the robot, the full point cloud could have more than 100 000 points with range up to 100 metres (the robots are equipped with Ouster OS0-128 LiDAR[1]). Such a high range is useless for the task of following a human, the robot should stay within a few metres from the human. The point cloud is also unnecessarily dense in the vicinity of the robot, processing a large number of points would slow down the algorithm. The rectangular area centred around the robot is cut out, the size of the bounding box is configurable, and the default side of the cube is 10 metres. The point cloud library is used to down-sample the point cloud with voxel-grid filter. The voxel grid filter splits the space into 3D voxels (cubes) and replaces all of the points inside one voxel with a single representative, their mean. Figure 5.5(a) shows the unaltered point cloud created by the 3D LiDAR, Figure 5.5(b) shows the point cloud after the voxel grid filter is applied and the range is decreased.

The next step is to remove all points that belong to the ground. A 2D grid is assembled from the filtered point cloud; this grid represents the ground. The elevation in each cell is estimated as the minimal z coordinate of all points located inside the cell of the grid. Some cells could contain no points; in this case, the elevation is set to be the same as on the current position of the robot. Some cells contain only points belonging to objects in the environment due to occlusion from the robot's body. These cells would contain an incorrect elevation estimate. It is assumed that the ground is smooth, median filter is applied to the estimated elevation map to remove outlier values, and Gaussian filter is applied to the elevation map to further smooth out the estimated ground. The resulting elevation map is shown in Figure 5.6. All points that lie within some margin of the ground are removed from the point cloud, and the resulting point cloud contains just points belonging to some objects in the environment (humans, walls, obstacles). This point cloud is then used to localise the human (the point cloud is depicted in Figure 5.7).

The detection of the human is purely geometrical; an initial estimate of the humans position is required to initialise the LiDAR localisation. This

[1] `https://ouster.com/products/hardware/os0-lidar-sensor`

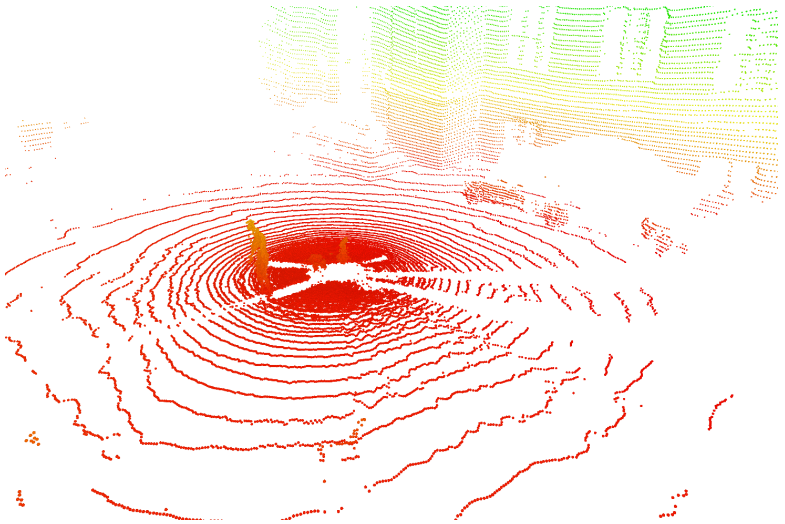

(a) The point cloud publish by the 3D LiDAR

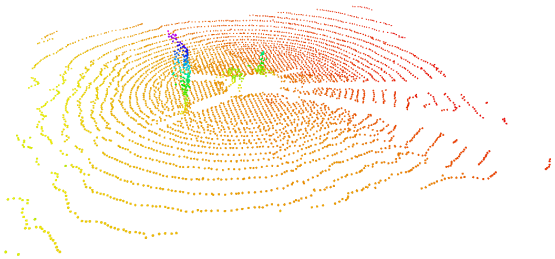

(b) The point cloud after downsampling and applying the bounding box

Figure 5.5: Comparison of the full point cloud and its downsampled version

estimate is provided by the radio-based localisation system (UWB TWR). A smaller area around the last known position of the human (either from TWR at startup or from the previous iteration of the LiDAR tracker) is cut and the DBSCAN clustering algorithm [43] is applied to the remaining points. The human is a sufficiently large cluster that is closest to the previous position of the human. This assumes that the tracking runs at a high enough frequency compared to the velocity of the human. The experiments were conducted outdoors to ensure that the LiDAR has direct line of sight visibility of the human, and therefore the ground truth data is always available. In static experiments, the positions estimated from the LiDAR scans were averaged to obtain a single ground-truth position.

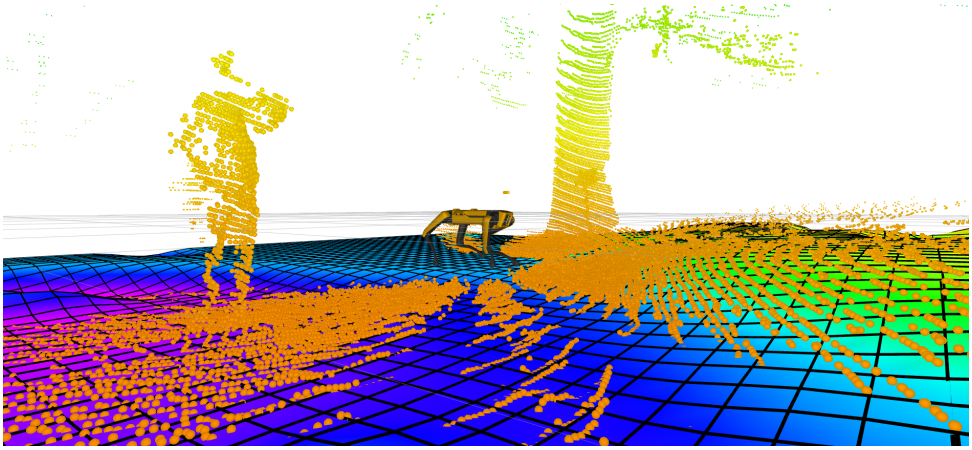

Figure 5.6: The detected ground (the surface plot) and the unfiltered point cloud (points coloured by their z coordinate).

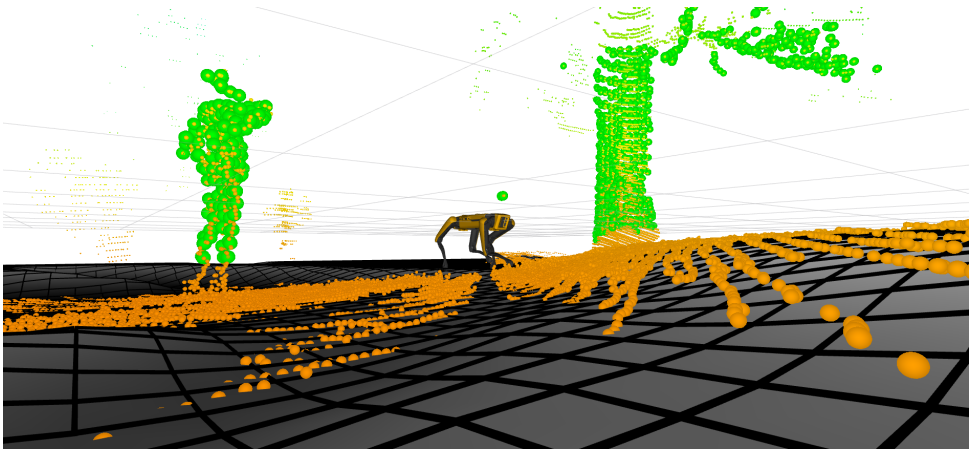

Figure 5.7: The remaining points after ground is removed (green spheres), the input point cloud and the detected ground.

The position estimated from LiDAR point clouds lies inside of the human (it is the centroid of the points seen by LiDAR on the human's body), and therefore this estimate could never match the radio-based estimate exactly, since the tag package is carried on the surface of the body (i.e. in a pocket). The z coordinate will also differ; the LiDAR estimate is shifted towards the most dense area of the human body (probably the chest). The radio estimated is either placed at the location of the active tag or at a fixed height relative to the robot's body (when the radio localisation operates in 2D). Figure 5.8 shows the human position estimated from the LiDAR point cloud. The detected position lies inside the body of the human, the distance between the correct radio-based estimate and the LiDAR-based ground truth will always be nonzero, it is approximately 10-15 centimetres in the xy plane.

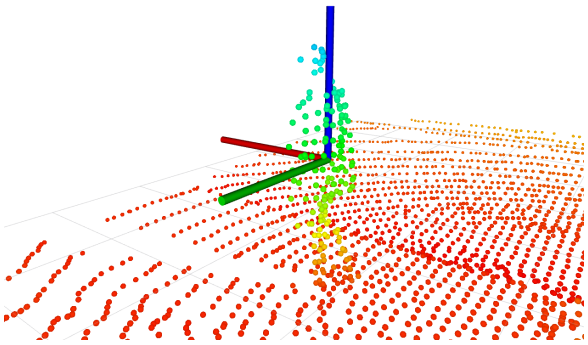

Figure 5.8: The position of the human detected from LiDAR point cloud shown as a coordinate frame and the point cloud used for detection. The points are coloured by their z coordinate

5.3 Experiment 1: Static Test of Accuracy and Precision

In these tests, the robot remains stationary and the human stands in several positions around the robot. The position estimated by the radio-based localisation system is compared to the ground truth positions obtained from LiDAR. All of the sensors were mounted on the robot, and the localisation was performed in multiple modes (just TWR, TWR and AoA) in parallel to make the comparison easier. The measurements were made at the same time, which means that the human was at the same time for all of the presented methods.

Spot Robot

As can be seen in Figure 5.9, the localisation works poorly in a particular direction. In this direction the front UWB module is completely blocked by the LiDAR mount (this can be seen in Figure 5.2, the module that is blocked is the one in the upper right corner of the image, the LiDAR and its mount can be seen right below the blue circle), the mount and the LiDAR itself contain metal, the UWB signal is therefore blocked on the straight path, and the range measurement is incorrect. In all other cases, the localisation performs significantly better, as is evident in Figure 5.10.

This is a general limitation of all radio-based localisation techniques that rely on time of flight. If the direct path signal passes through some object, the velocity of the signal may be different from that expected. In specific conditions, the signal could reflect from objects in the environment, the reflected signal could be even stronger than the direct path one. A really strong attenuation of the direct path signal could bring it under the noise threshold, and a reflected signal is used instead (leading to potentially huge error).

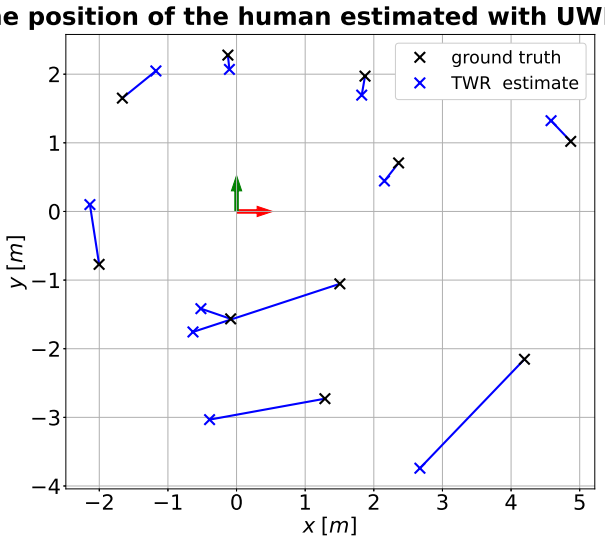

Figure 5.9: The results of the TWR localisation (blue crosses) compared with the true position of the human obtained from LiDAR (black crosses). The position of the human is computed as described in Section 4.1.

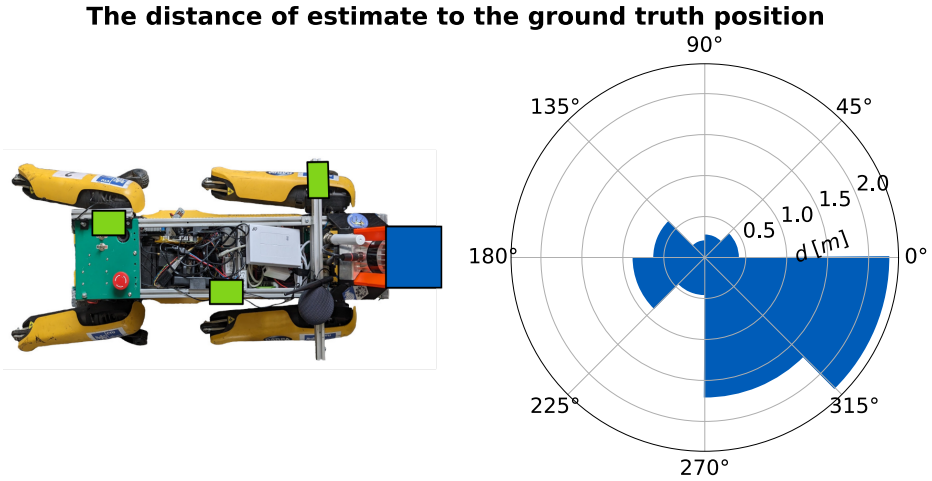

Figure 5.10: A histogram of the distances between the estimated and true positions of the human shown in Figure 5.9. Each bin represents the maximal error in the particular direction. The TWR modules are marked as green rectangles on the left-hand side of the figure. Notice that the error is extremely large when the human is in the front right of the robot. This is caused by the LiDAR mount that can be seen on the front of the robot in the left part of the figure.

The only way this behaviour can be limited is by mounting the modules so that they are never blocked by the robot's body. This might require mounting them on rods high above the base of the robot, but such a mounting system is highly restrictive — the rods might limit other sensors placed on the robot (LiDAR, cameras), they also restrict the access to the top of the robot (in case

the robot serves as a mule). As a result, it is not always possible to achieve similar performance in all directions, and the robotic platform is the limiting factor. The positions of the modules should therefore be selected in a way that maximises performance in the most frequent use case, for example, it is expected that the human stays in front of the robot, therefore the localisation should perform best in the area in front of the robot.

Figure 5.11 shows that the standard deviation of the estimated position remains within 10 centimetres on the x-axis, but the occlusion caused by the LiDAR mount affects even the standard deviation — it is close to one metre. This is caused by the fact that the affected range measurement was jumping between the true value and the one caused by the reflected signal.

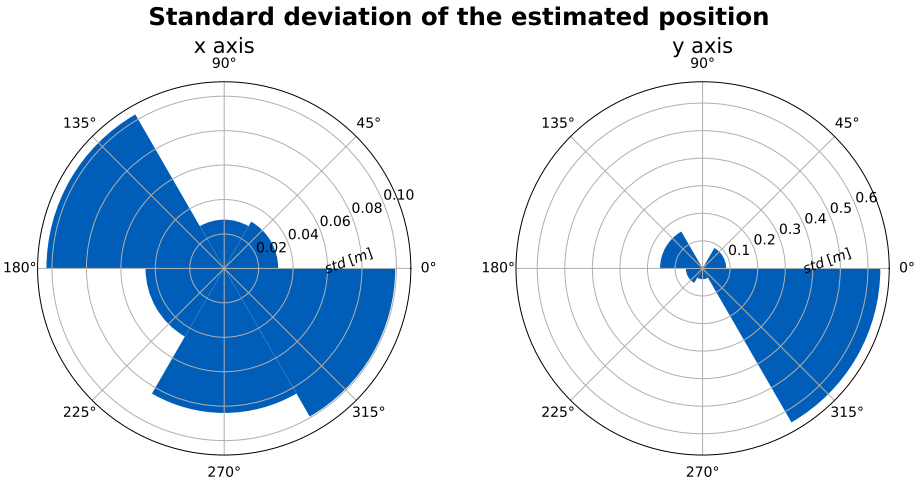

Figure 5.11: The histogram of standard deviations of positions estimated by three TWR modules mounted on the Spot robot. The positions shown in Figure 5.9 are means of estimates recorded for one minute, and each bin of this histogram shows the maximal standard deviation in the corresponding direction.

The addition of AoA measurements mitigates the problem. The estimates are similar in all directions as shown in Figure 5.13. The accuracy of the localisation also does not depend on the distance thanks to the AoA measurement. The same limitation as in the previous case applies here, the modules should not be blocked by any part of the robot. The AoA antennas were mounted high above the robot, meaning that their performance was not degraded in any particular direction.

Angle of Arrival measurements also improved the standard deviation of the estimate. As can be seen in Figure 5.14, the standard deviation remained within 20 centimetres.

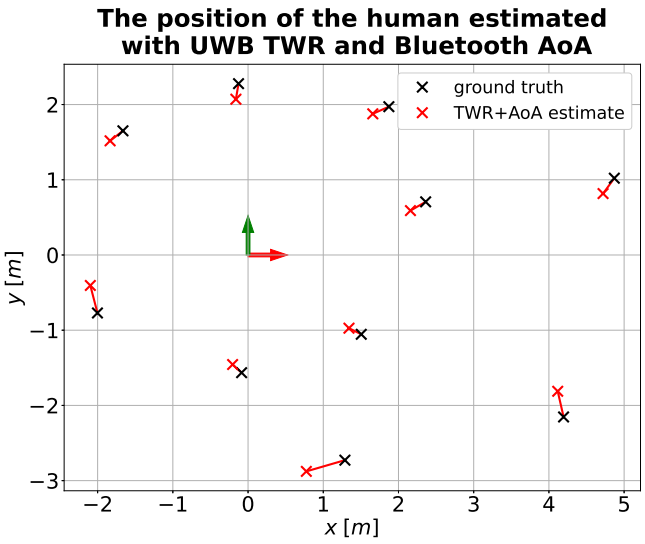

Figure 5.12: The results of combined TWR and AoA localisation (red crosses) compared with the true position of the human obtained from LiDAR.The position of the human is computed as described in Section 4.4.

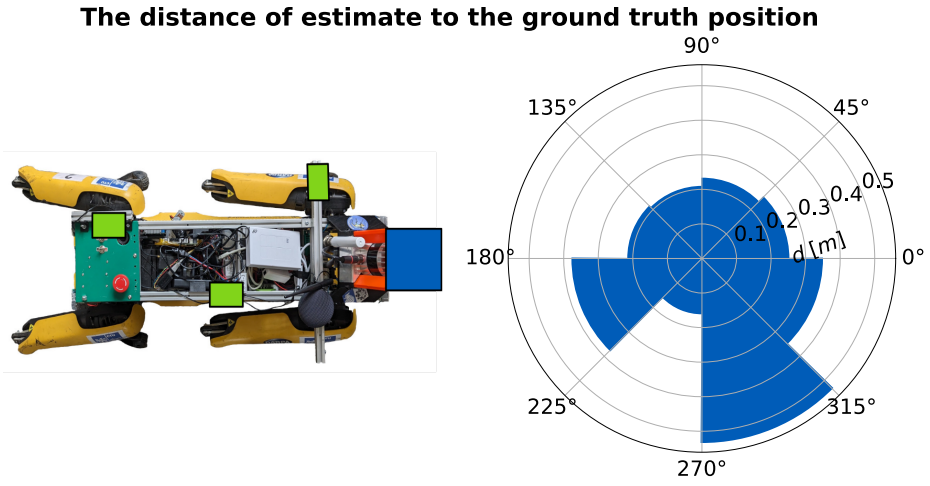

Figure 5.13: A histogram of the distances between the estimated and true positions of the human shown in Figure 5.12. Each bin represents the maximal error in the particular direction. The left part of the figure shows the mounting points of the sensors on the robot, green rectangles represent TWR modules, blue rectangle is the cube composed of AoA antennas.

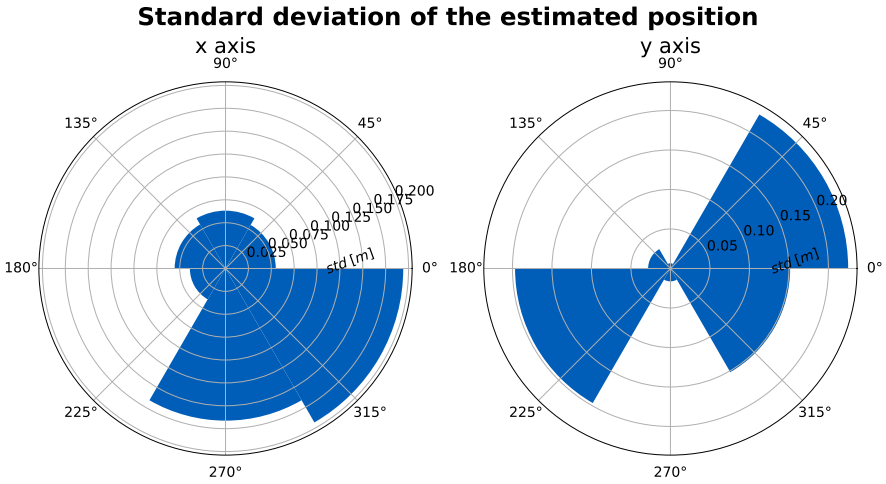

Figure 5.14: The histogram of standard deviations of positions estimated by three TWR modules and four AoA antennas mounted on the Spot robot. The positions shown in Figure 5.12 are means of estimates recorded for one minute, and each bin of this histogram shows the maximal standard deviation in the corresponding direction.

Tracked Robot

Figure 5.15 shows the results of the localisation with four TWR modules mounted on the tracked robot. The increased mutual distance between the UWB modules has led to an increase in precision compared to the experimental setup with the Spot robot (Figure 5.10). There is again a decreased localisation performance in one particular direction, this time it is caused by the rod used to mount the AoA antennas. This rod is mounted to the right of one of the UWB modules, which means that the signal is attenuated if the human is to the right of the robot. This can be seen in Figure 5.16.

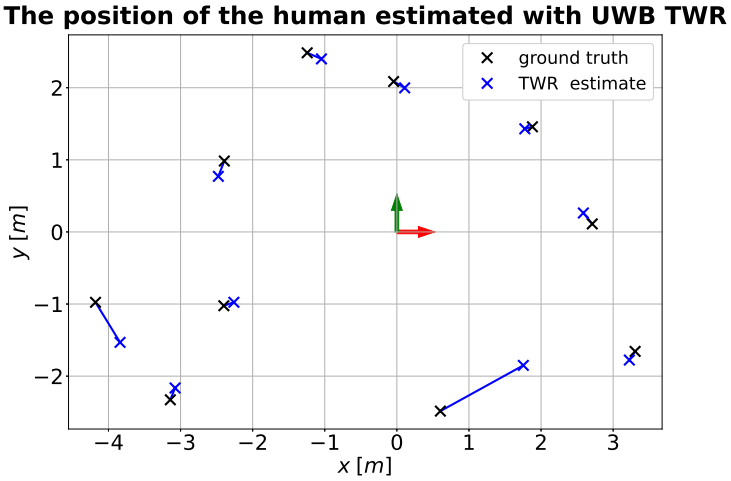

Figure 5.15: The results of the TWR localisation (blue crosses) compared with the true position of the human obtained from LiDAR (black crosses). The position of the human is computed as described in Section 4.1.

The standard deviation of the estimate based purely on TWR measurements is lower than in the case of the Spot robot, this can be seen in Figure 5.17. This is due to the increased mutual distance of the modules and the absence of occlusion by the body of the robot.

The addition of AoA measurements does not have the same benefits as in Section 5.3, since the estimates are already accurate enough. The error with just TWR was in most cases within 20 centimetres, but this error could be caused by the method of obtaining the ground truth position, as was described earlier. Figures 5.18 and 5.19 present the results, it is apparent that an improvement was achieved in the worst-case situations, but in general the error remains the same (around 20 centimetres).

The standard deviation was already low when only TWR was used; nevertheless, the results were improved by including the AoA measurements into the optimisation problem. Figure 5.20 shows that the standard deviation remains under 10 centimetres in all directions.

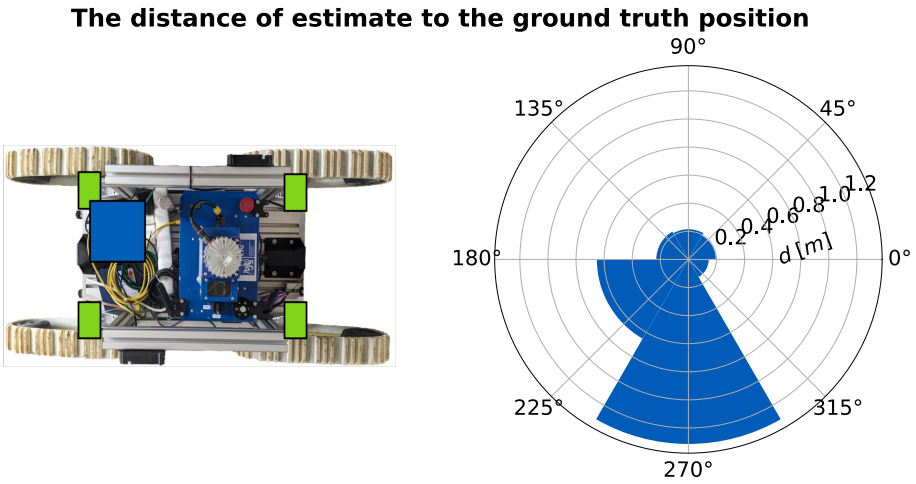

Figure 5.16: A histogram of the distances between the estimated and true positions of the human shown in Figure 5.15. Each bin represents the maximal error in the particular direction. The placement of the TWR modules on the robot can be seen in the left part of the figure; they are marked with green rectangles.

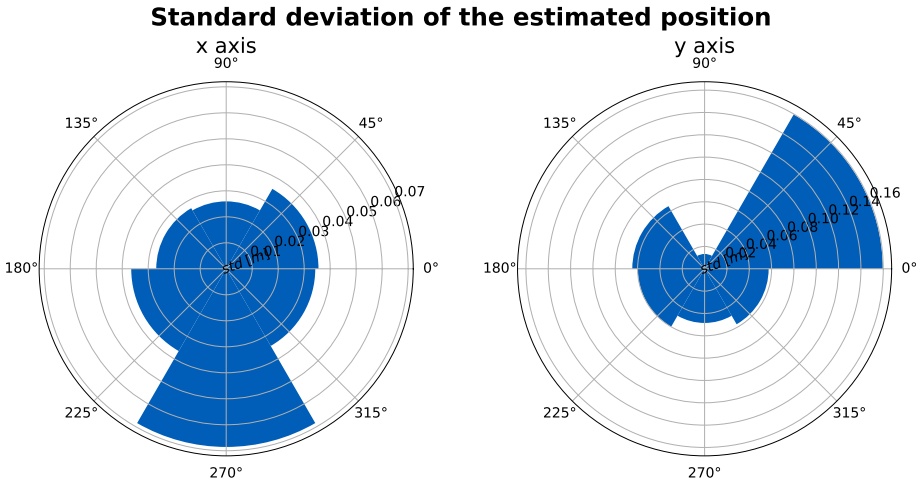

Figure 5.17: The histogram of standard deviations of positions estimated by four TWR modules mounted on the tracked robot. The positions shown in Figure 5.15 are means of estimates recorded for one minute, and each bin of this histogram shows the maximal standard deviation in the corresponding direction.

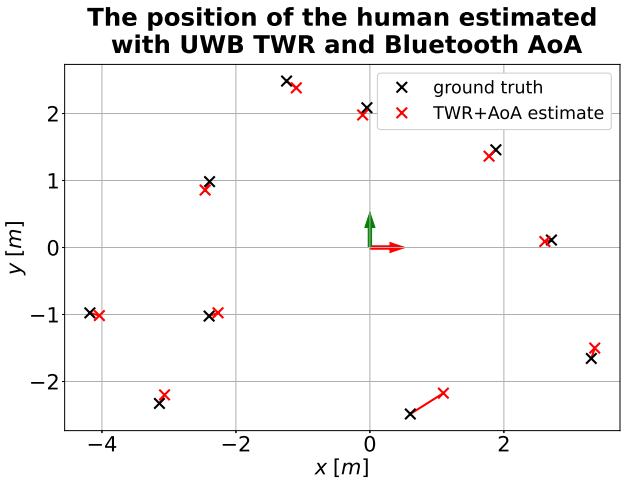

Figure 5.18: The results of combined TWR and AoA localisation (red crosses) compared with the true position of the human obtained from LiDAR.The position of the human is computed as described in Section 4.4.

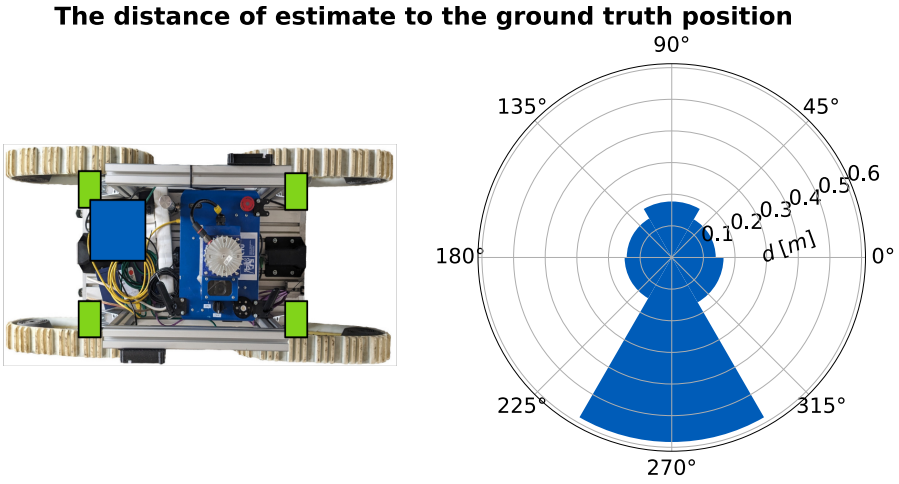

Figure 5.19: A histogram of the distances between the estimated and true positions of the human shown in Figure 5.18. Each bin represents the maximal error in the particular direction. The left part of the figure shows the robot and the placement of the sensors — green rectangles represent the TWR modules, blue rectangle is the cube composed of four AoA antennas.

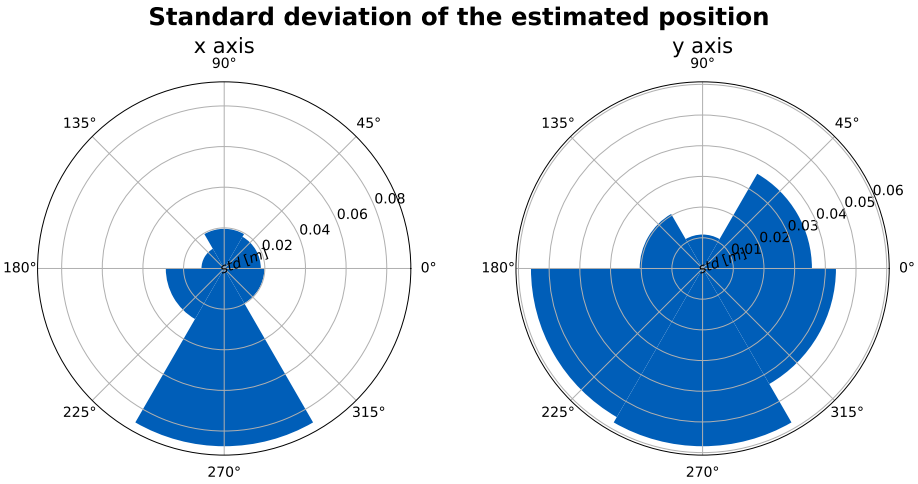

Figure 5.20: The histogram of standard deviations of positions estimated by four TWR modules and four AoA antennas mounted on the tracked robot. The positions shown in Figure 5.18 are means of estimates recorded for one minute, and each bin of this histogram shows the maximal standard deviation in the corresponding direction.

Single Unit Containing Four AoA Antennas and One TWR Module

The AoA cube with the TWR module was mounted on the Spot robot and the tracked robot. On both platforms, the cube was mounted above the robot, guaranteeing an unobstructed view of the whole area around the robot. All of the benefits of mounting the sensors in an unobstructed place seen in Section 5.3 apply here too. The whole sensor package can be mounted as a single unit high above the robot, which means that even the range measurements provided by the single UWB TWR module are not affected by the body of the robot. Figures 5.21 and 5.22 show that the estimated position matches the true one with an error that is mostly within 30 centimetres. As was described earlier, the method for obtaining the ground truth adds an offset that is always present; therefore, the actual error of the localisation is lower than 30 centimetres.

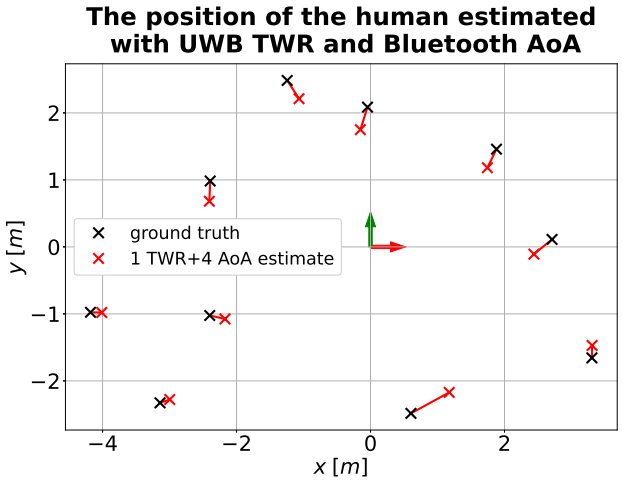

Figure 5.21: The results of combined localisation from single TWR measurement and four AoA measurements (red crosses) compared with the true position of the human obtained from LiDAR. The position of the human is computed as described in Section 4.4.

In case of this setup, the measurements are not redundant. As a result, the noise from each of the measurements has a greater effect on the final estimated position. This is apparent when comparing Figure 5.23 with the results achieved when using four TWR modules and four AoA antennas (Figure 5.20).

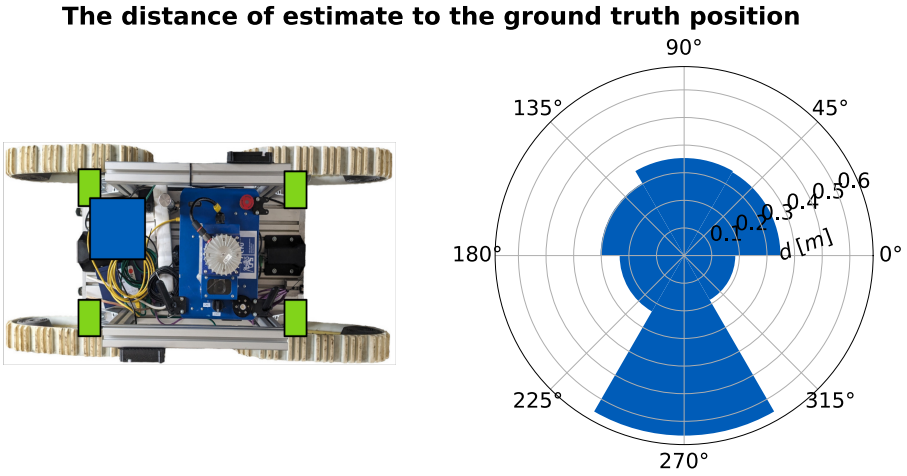

Figure 5.22: A histogram of the distances between the estimated and true positions of the human shown in Figure 5.21. Each bin represents the maximal error in the particular direction. The cube with four AoA antennas and one TWR module is represented by a blue rectangle on the rear of the robot in the left part of the figure.

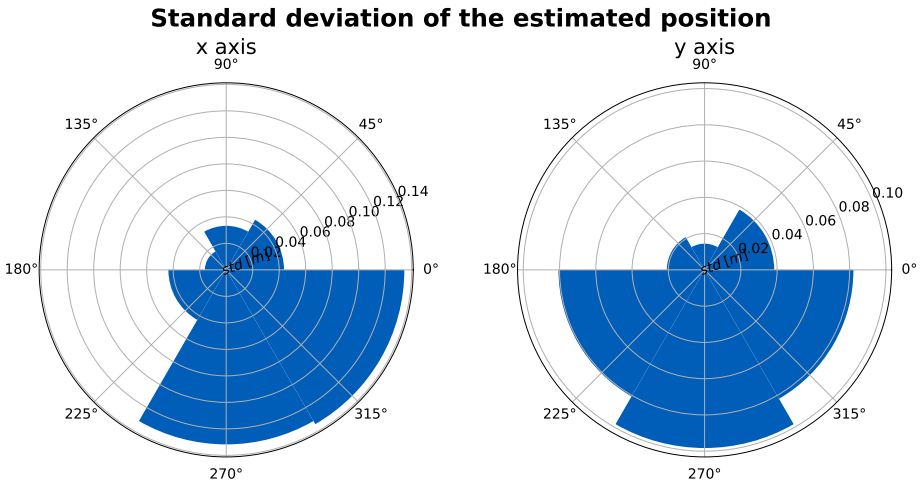

Figure 5.23: The histogram of standard deviations of positions estimated by one TWR module and four AoA antennas mounted on the tracked robot. The positions shown in Figure 5.21 are means of estimates recorded for one minute, and each bin of this histogram shows the maximal standard deviation in the corresponding direction.

5.4 Experiment 2: Human Walking on a Predefined Trajectory

This set of tests aims to verify the ability to track a moving target. The robot again remains stationary,but the human is moving this time. He walks on a path around the robot. The estimated path of the human is recorded and compared with a path obtained from the LiDAR tracker described above.

Spot Robot

The human moved on a trajectory that is approximately half-circle at the front of the robot. Figure 5.24 shows the comparison of all the configurations described earlier. It can be seen that the shape of the trajectory was matched in all of the configurations, but the radius is larger in the ground truth trajectory. This is caused by the method used for obtaining the ground truth, namely that the position estimated from the LiDAR point clouds lies inside the body of the human, but the sensor package is placed in a pocket on the surface of the human's body.

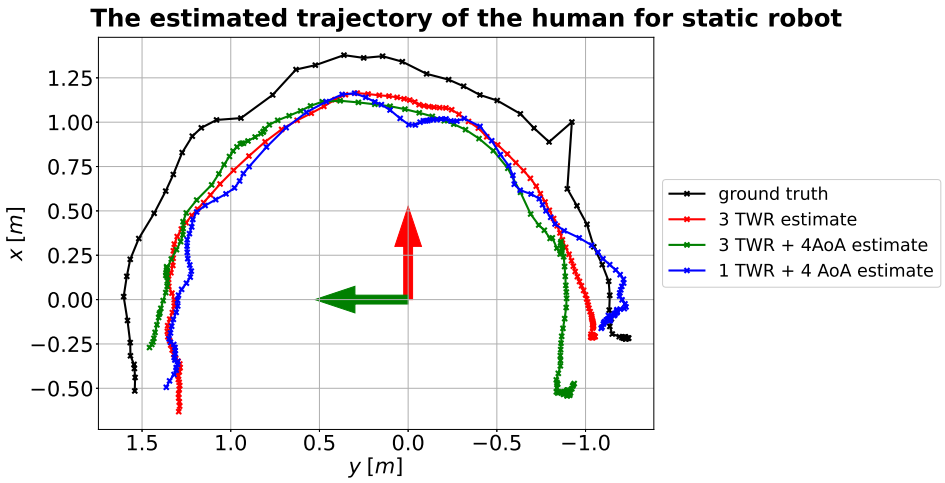

Figure 5.24: The comparison of the possible sensor configurations on Spot robot, the ground truth trajectory was obtained by tracking the human in LiDAR point clouds. Note that the offset of the ground truth trajectory is caused by the method that is used to obtain it, it is not caused by an error of the proposed localisation algorithms.

Figure 5.25 shows that the proposed system is capable of tracking a moving target in multiple sensor configurations. Using just four AoA antennas and a single UWB module is enough to track the moving human, even tough, only the signal strength can be used to select the correct angle measurements.

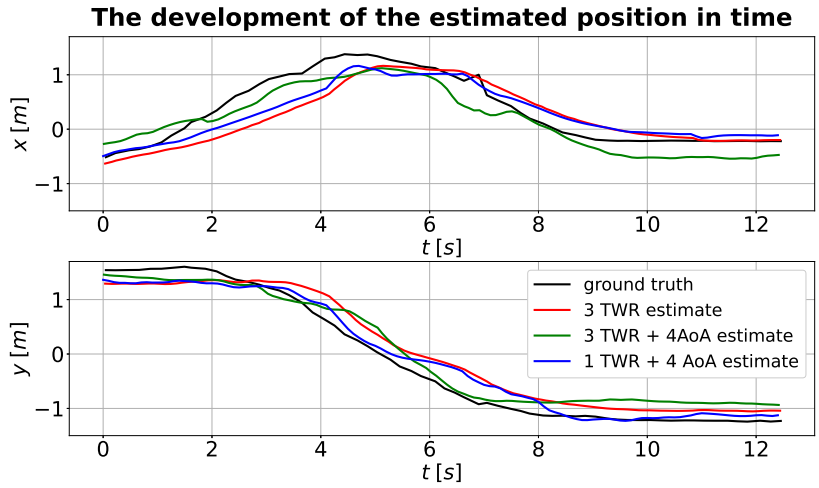

Figure 5.25: Each individual coordinate of the estimated position compared to the ground truth position of the human in time.

Tracked Robot

The human used a trajectory similar to that used in the experiment with the Spot robot. Figure 5.26 again compares the different sensor configurations. All of them manage to approximately track the path of the human

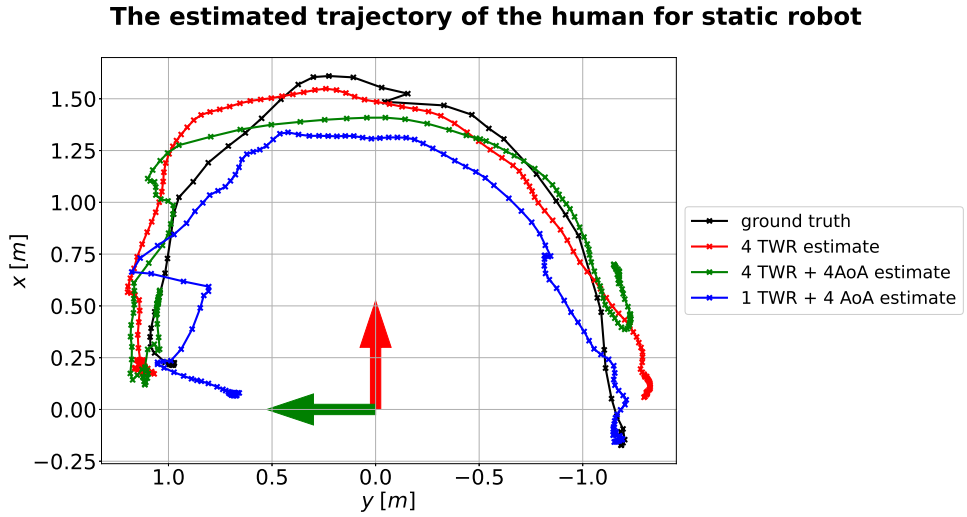

Figure 5.26: The comparison of the possible sensor configurations on tracked robot, the ground truth trajectory was obtained by tracking the human in LiDAR point clouds

Figure 5.27 shows how the x and y coordinates of the position estimate changed over time compared to the ground truth. On this robotic platform, using only TWR (as described in Section 4.1) produces sufficiently accurate estimates due to the larger footprint of the robot and relatively unobstructed

view in all directions.

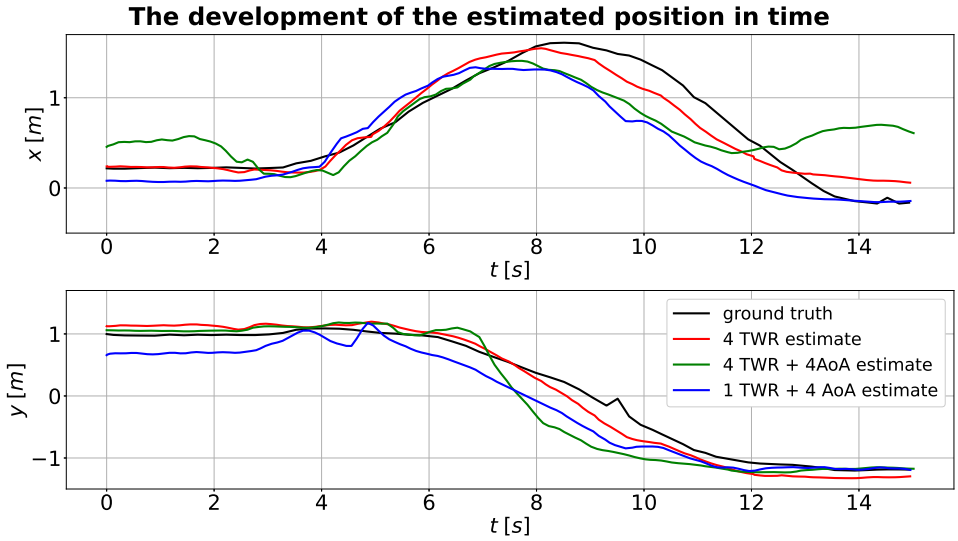

Figure 5.27: Each individual coordinate of the estimated position compared to the ground truth position of the human in time.

5.5 Experiment 3: Autonomous Following of the Human Outdoors

The main goal of these tests is to verify that the robot maintains the prescribed distance from the human and stops when the human approaches, ensuring that the robot does not collide with the human. The robot should also follow the path of the human as closely as possible, meaning that if no visual information is available, the robot should still be capable of navigating through the environment thanks to the collision-free path the human has used. The estimated relative position is again compared to the ground truth obtained from LiDAR. All robots are also capable of localising themselves within the environment (using LiDAR, IMU, and odometry); therefore, the path of the human and the path of the robot in the static frame are compared. The experiment was carried out on the same outdoor testing path with both robots, as can be seen in Figure 5.28.

Figure 5.28: The approximate path of the outdoor experiment shown on an orthophoto map

Spot Robot

Localisation using only TWR measurement had degraded performance when the human was in the front right of the robot. This means that in sharp right turns the robot would have difficulties following the human. Therefore, a localisation combining TWR and AoA measurements was used to control the robot, the human position was obtained from the optimisation problem described in Section 4.4.

The maximum speed of the robot was limited to 1.2 m/s, the robot is in theory capable of reaching higher speeds, but this depends on the current conditions (terrain, load of the robot). The desired distance between the human and the robot was set at 1.6 metres, the robot was set to stop

immediately if the human was less than 1.3 metres away from any of the UWB TWR modules.

Figure 5.29 shows the results of the experiment using TWR and AoA for human localisation. The robot was closely following the estimated path of the human. As a result, the robot was also following the actual path of the human. The only exception is the 270 degree turn in the lower left corner of the figure. The estimated position jumped twice to an incorrect value, leading the robot further away from the correct path of the human. This error was quickly fixed, meaning that the robot took a different path but shared the same

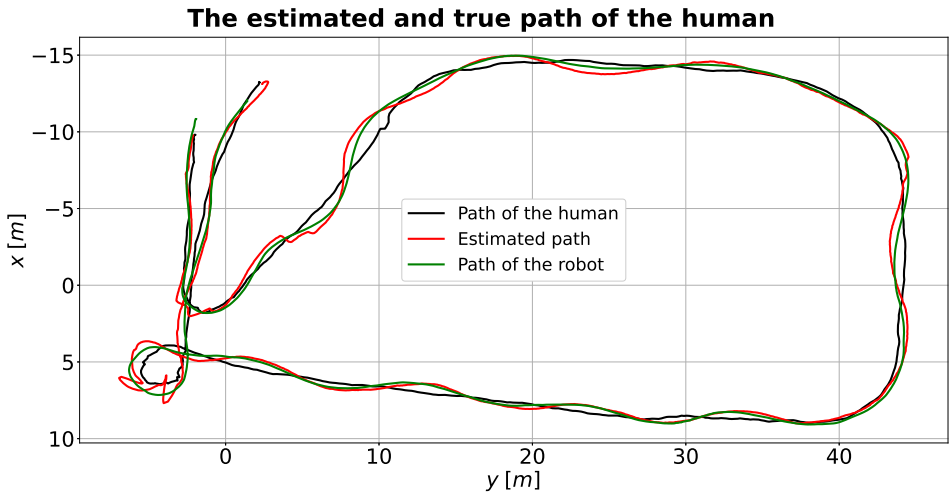

Figure 5.29: The path of the human (black) compared to the estimated path (red) and the actual path of the robot (green).

The distance between the robot and the human and the robot's velocity is shown in Figure 5.30. It can be seen that the human walked too fast during some parts of the experiment. As a result, the gap was growing, but as soon as the human slowed down, the robot closed the gap. The robot never got too close to the human, and an appropriate distance was maintained even during the 270-degree turn when the human walked slowly.

Figure 5.30 shows that the robot was changing its velocity with relatively low acceleration. This contributes to perceived system safety. At the start of the experiment, the speed was gradually increased. At the time of 70 seconds, the robot started to catch up to the human and even though the actual gap was larger than the desired one, the robot started to decrease its velocity, because the human was getting closer. At the end of the experiment the human stopped and started moving against the robot. This caused the distance to rapidly decrease; as it got below the safety threshold, the robot immediately stopped to make the human feel safer.

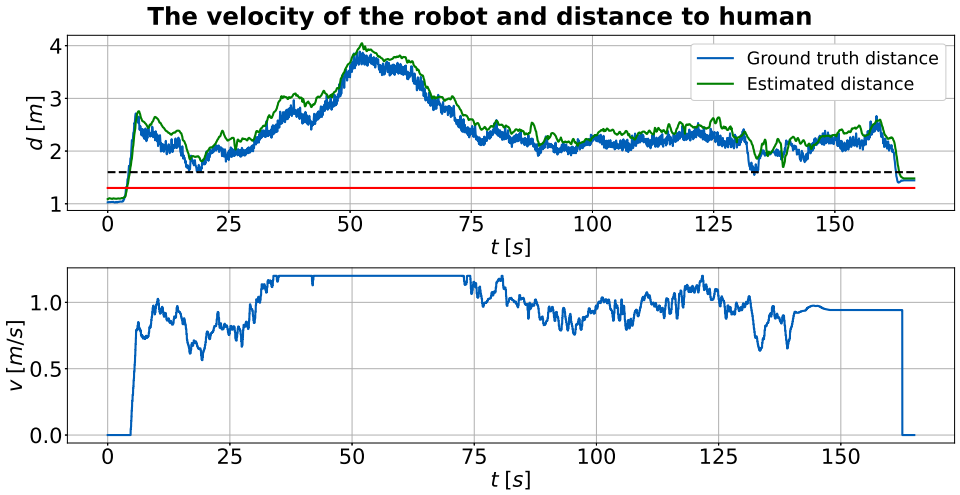

Figure 5.30: The distance between the robot and the ground truth position and between the robot and estimated position are shown in the upper part of the figure. The desired distance is shown with a black line and the distance at which the robot must stop is shown in red. The lower part of the figure shows the velocity commands that are being sent to the robot. Between the times of 30 and 60 seconds, the human was walking faster than the robot's maximal velocity, the gap was therefore increasing. Once the human slowed down, the robot closed the gap to the human.

Tracked Robot

The tracked robot used for the test is not capable of reaching speeds as high as that of the Spot robot (the maximum speed is 1 m / s compared to 1.5 m/s). The velocity in outdoor experiments was further decreased by the paved surface on which the robot moved. The tracks were sometimes slipping, and the robot struggled to turn at higher velocities. As a result, the robot's velocity had to be limited to 0.5 m / s, which required the human to walk really slowly to allow the robot to follow him.

Figure 5.31 shows the estimated path of the human compared to the real one. The estimate was affected by strong noise, but the path of the robot still corresponds to the true path of the human. One of the possible causes of the noise is that the human was farther away than was supposed to (as shown in Figure 5.32), due to the very limited velocity of the robot. As was shown earlier, the accuracy and also the precision of TWR localisation depend on the mutual distance between the modules placed on the robot relative to the distance to the target. The increased distance between human and robot led to a large amount of estimates with high lateral error, the distance was correct but the direction had higher error (this can be again seen in Figure 5.32). If the robot could travel with a higher velocity, it would be easier to maintain the desired distance, and the results of the localisation would be more accurate.

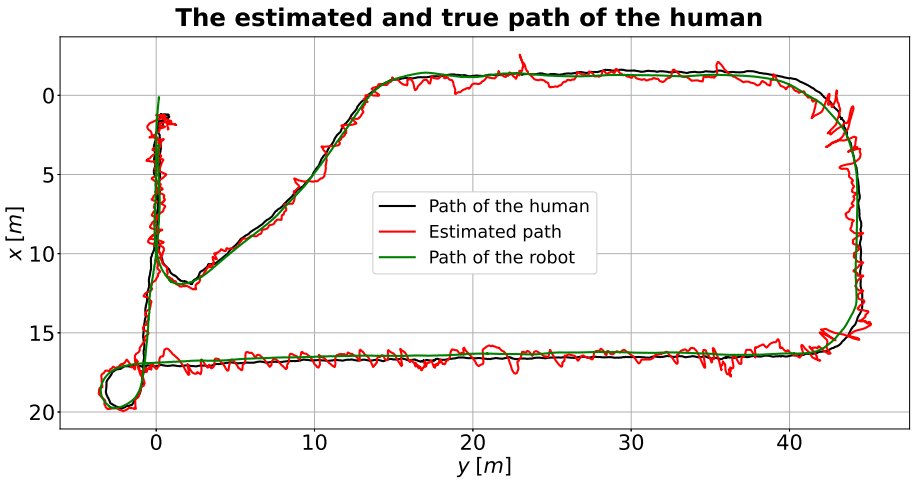

Figure 5.31: The path of the human (black) compared to the estimated path (red) and the actual path of the robot (green).

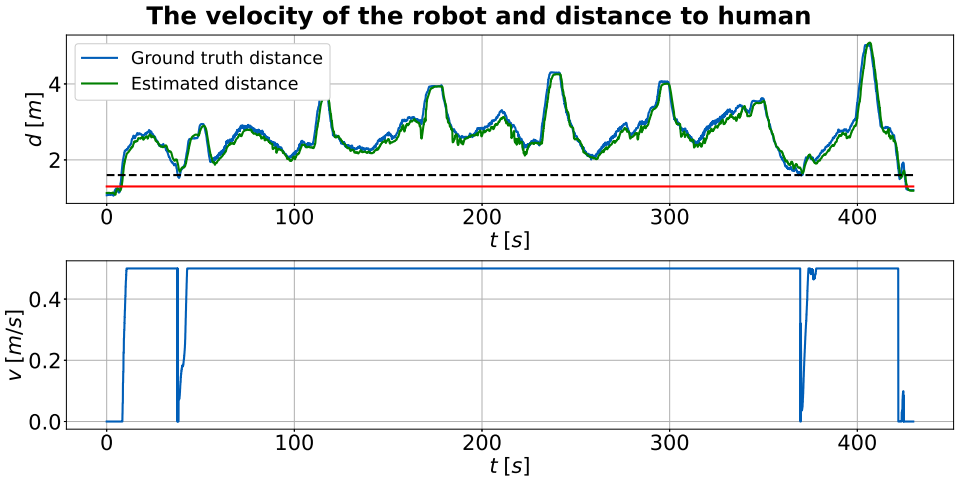

Figure 5.32: The distance between the robot and the ground truth position and between the robot and estimated position are shown in the upper part of the figure. The desired distance is shown with a black line and the distance at which the robot must stop is shown in red. The lower part of the figure shows the velocity commands that are being sent to the robot. On multiple occasions the robot stopped moving even though it was receiving non-zero velocity commands, this caused the spikes in the distance between the robot and the human.

Single Unit Containing Four AoA Antennas and One TWR Module

Following experiment with the cube was carried out with the tracked robot. The path of the human was identical to that used with the Spot robot to make comparison easier. Figure 5.33 shows the trajectories. The robot was able to precisely follow the human path. The only exception is again the 270 degree turn at the lower left corner of the figure. This time the system relies purely on AoA measurements to estimate the heading towards the human. In the other experiments, multiple TWR measurements were available, which meant that the position of the human could be determined even without the AoA measurements. During this sharp turn, the human is at the border of two AoA antenna measurement ranges. As was shown earlier, the error in angle measurements grows as the angles get closer to the maximal value. RSSI is not always a good indicator of measurement quality. As a consequence, the system will be switching between the two AoA antennas, but one of them could be reporting incorrect angles, and there is no way how to distinguish this.

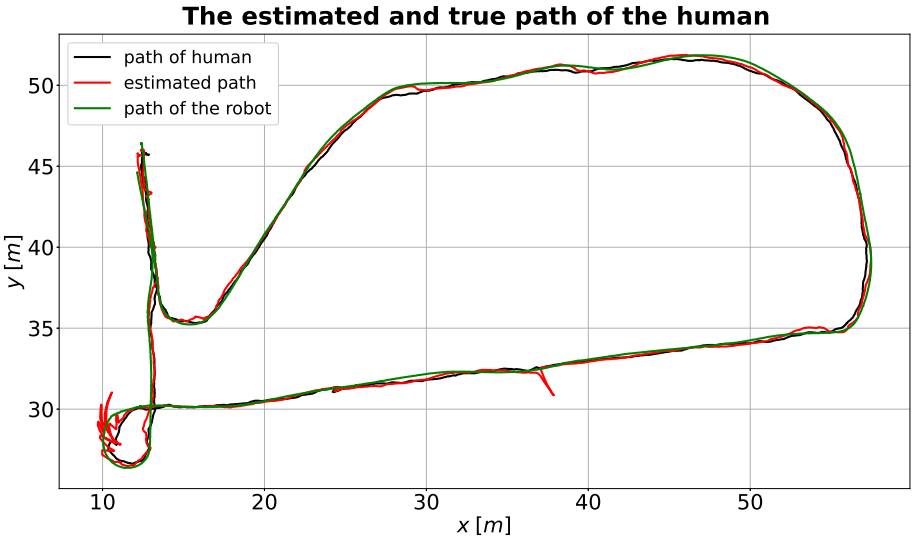

Figure 5.33: The path of the human (black) compared to the estimated path (red) and the actual path of the robot (green).

During the experiment, the robot experienced an internal error that caused it to temporarily stop, this behaviour is visible as three spikes in the distance between the robot and the human in Figure 5.34. Otherwise, the robot stayed close to the desired distance. At around 220 seconds, the human stopped and started moving against the robot. The robot stopped as the human came closer and started moving once the human continued walking.

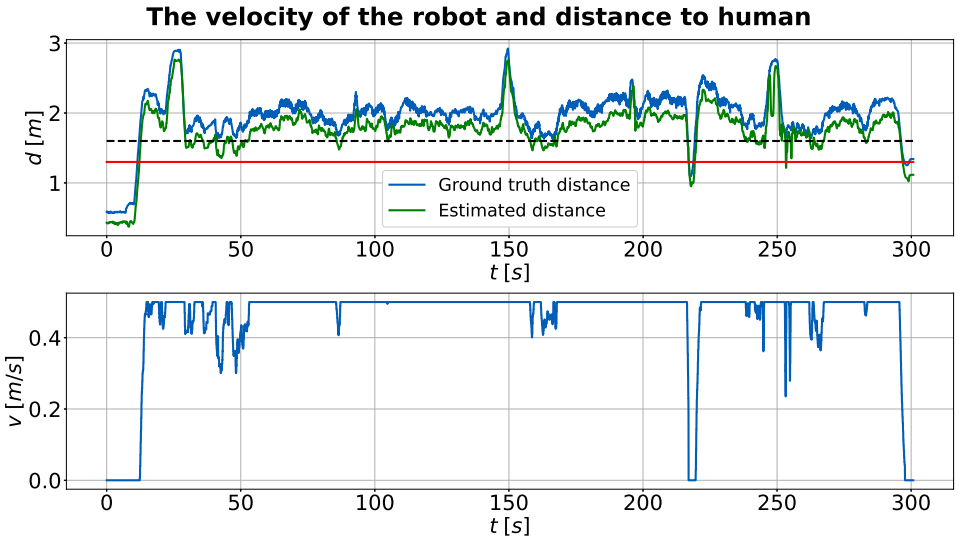

Figure 5.34: The distance between the robot and the ground truth position and between the robot and estimated position are shown in the upper part of the figure. The desired distance is shown with a black line and the distance at which the robot must stop is shown in red. The lower part of the figure shows the velocity of the robot.

5.6 UWB TWR Localisation Deployed on an Autonomous Vehicle

A six-wheeled autonomous vehicle was equipped with four UWB TWR modules. The position of the target was obtained by solving the optimisation problem shown by Equation 4.1. The task of the vehicle was to autonomously follow a car that is driving in front. This test aims to show how the proposed solution scales up; the vehicle used for testing is much larger than any other robot used for testing throughout this work. This means that the modules can be spread further away from each other, leading to an increased range of the localisation.

The UWB modules were mounted on the top of the front part of the robot, they formed a rectangle with size 100 by 65 centimetres. The robot and the placement of the modules is shown in Figures 5.35(a) and 5.35(b).

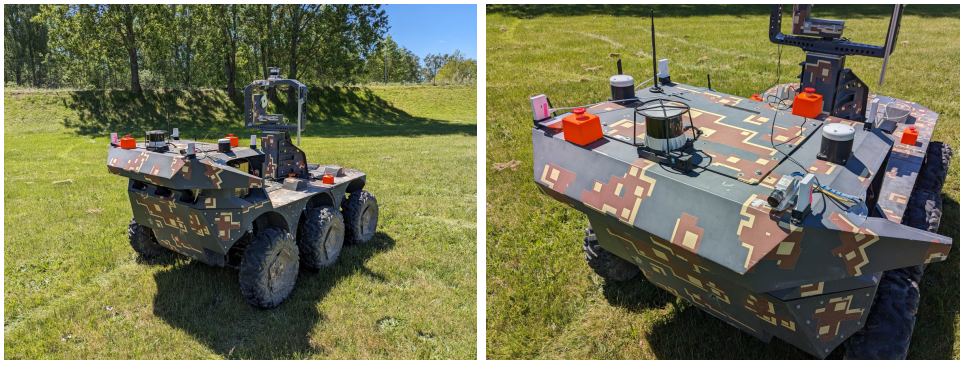

(a) The six-wheeled autonomous vehicle. (b) The placement of the UWB modules on the vehicle.

Figure 5.35: The six-wheeled autonomous vehicle with four UWB TWR modules mounted at the front.

Due to the increased mutual distance of the UWB modules, the localisation had a larger usable range. It was possible to localise a human at a distance greater than 10 metres (shown in Figure 5.36), Figure 5.37 shows the localisation of a car at a distance of 14.3 metres. This shows the importance of the mounting points of the modules on the robot, the modules should be spread as far as possible to increase the range of the localisation. The UWB modules could even be mounted on the back platform of the autonomous vehicle, but this would require it to drive backward to give the modules an unobstructed view (the raised front platform is complete metal, meaning that the signal cannot pass through it).

Figure 5.38 shows a comparison of the position estimated with the UWB TWR compared to the ground truth obtained from the LiDAR point clouds. As can be seen, the localisation works even for greater distances than in the previous experiments, this is again due to the larger mutual distance of the TWR modules. The localisation also has comparable accuracy in all

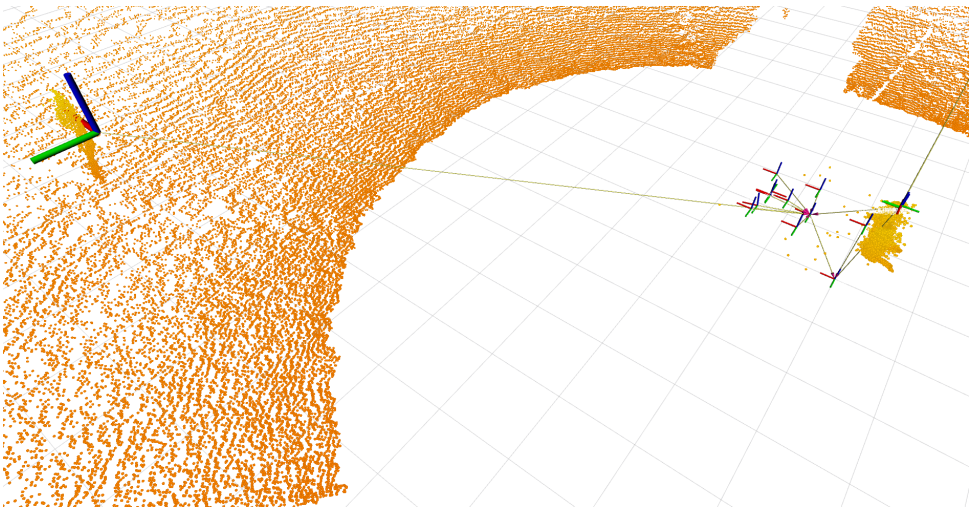

Figure 5.36: A human localised by four TWR modules mounted on the six-wheeled autonomous vehicle. The human is approximately 10.8 metres away from the robot.

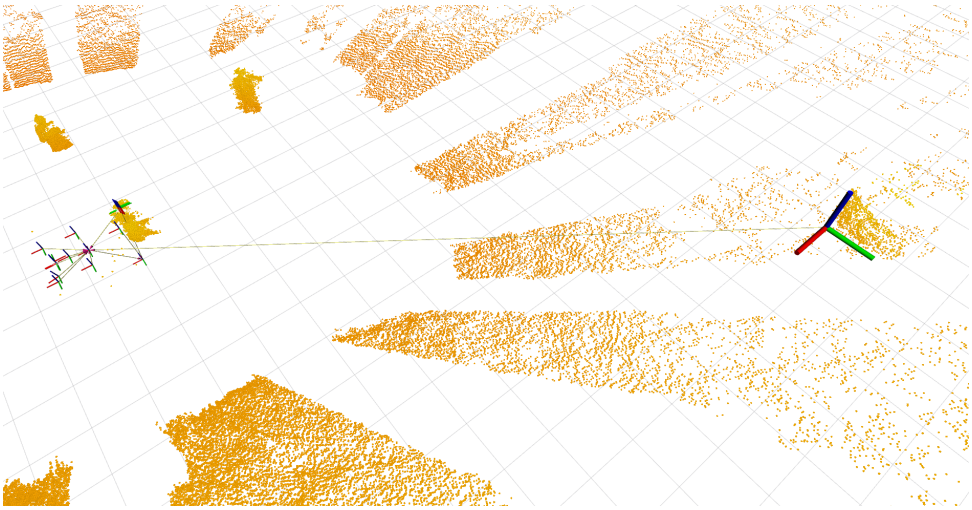

Figure 5.37: A car localised by four TWR modules mounted on the six-wheeled autonomous vehicle, the tag was placed on the rear windshield wiper. The car is approximately 14.3 metres from the robot.

directions; this is because the modules are mounted on top of the robot and they have a mostly unobstructed view of the whole space around the robot.

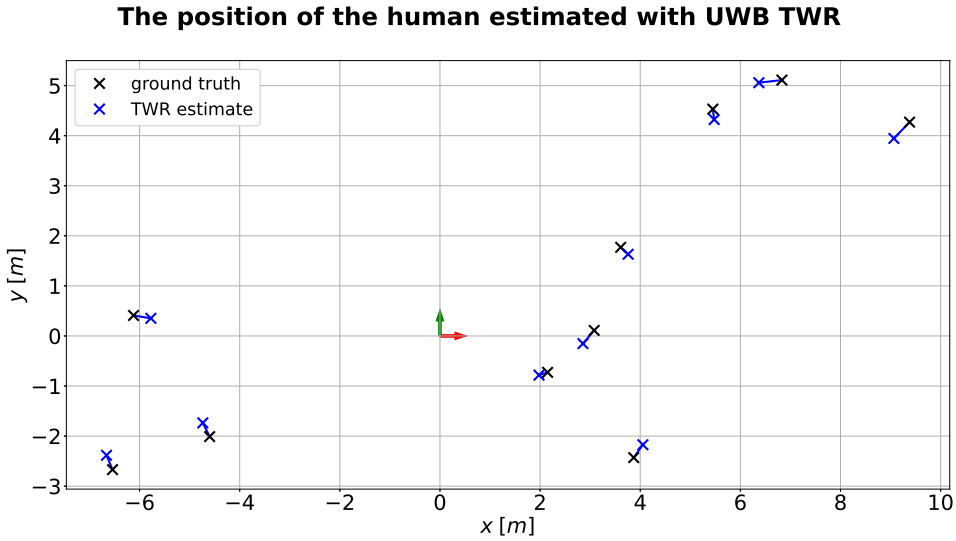

Figure 5.38: Comparison of the TWR estimate (blue crosses) to the true position of the human obtained from LiDAR (black crosses). The corresponding points are connected with a blue line.

5.7 Comparison of the Results Achieved by Different Localisation Methods

This part of the text aims to show the main benefits and limitations of the proposed localisation methods in relation to the robotic platform on which they are used.

The deployment on the Spot robot clearly shows that the main limiting factor of accuracy is the construction of the robotic platform. Spot is narrow, the body is only 18 centimetres wide. Although the body of the robot is relatively long (around 1 metre), the mutual distance of the sensors mounted on the robot is small. This leads to a lower localisation range and a lower overall accuracy. Figure 5.39 shows another factor that has an even greater influence on the functioning of the system — oclusion. In one particular direction, one of the TWR modules was blocked by the mount of the 3D LiDAR, this caused huge error of the estimated position and made it impossible to use just ultra-wideband to localise the human.

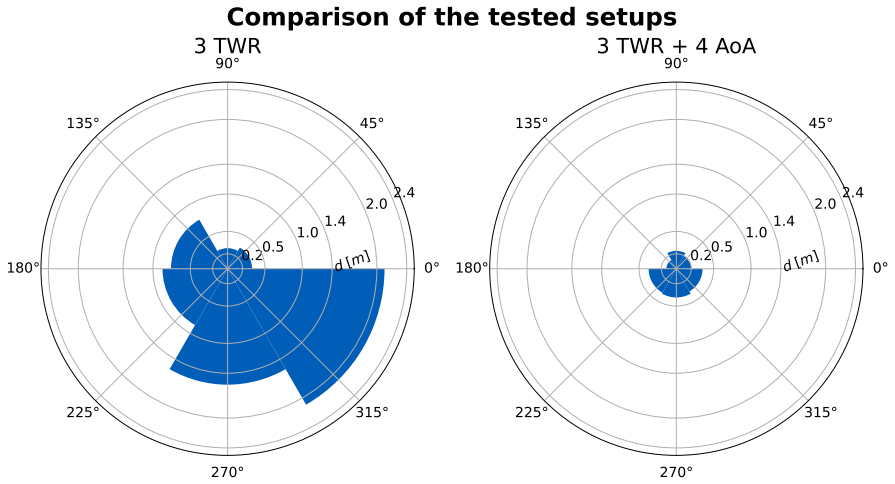

Figure 5.39: Comparison of the sensor configurations used on the Spot robot. Each of the figures shows how the distance between the estimated position and the true one depends on the direction in which the human is located.

It is clear that in this case, adding the Bluetooth Angle of Arrival measurement significantly decreased the error. This is because ultra-wideband is still capable of measuring the range with a small error. Bluetooth AoA adds accurate information about the direction towards the target resulting in an accurate estimate of the human's position.

Deployment on the tracked robot shows that if the UWB modules can be spread further apart, the benefit of Angle of Arrival measurements decreases. This can be seen in Figure 5.40. When only ultra-wideband is used to localise the human, the accuracy is higher than in the case of the Spot robot.

In the case of combined TWR and AoA localisation on the tracked robot, the error was mostly within 20 centimetres. As was stated in Section 5.2,

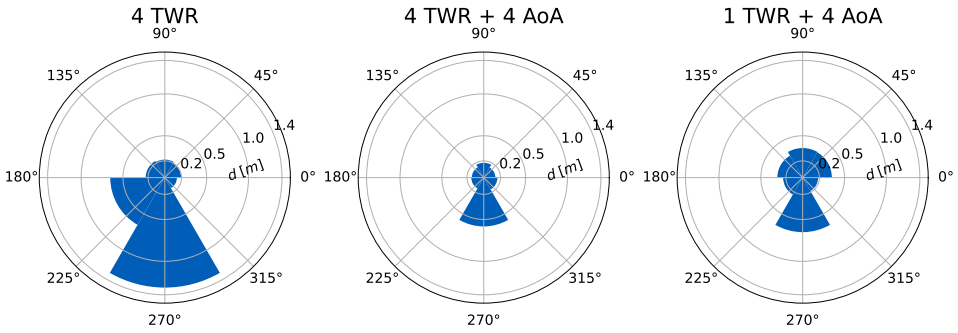

Figure 5.40: Comparison of the sensor configurations used on the tracked robot. Each of the figures shows how the distance between the estimated and true position of the human depends on the direction in which is the human located.

there is an offset between the ground truth position obtained from LiDAR and the true position of the human perceived by the radio-based localisation system. This offset is approximately 10 to 15 centimetres, meaning that in this particular configuration, the estimates were very close to the true position of the human.

The effect of the large lateral error caused by the small error in the range measurements can be seen in Figure 5.41(a). The distance from the human is correctly estimated, but the estimate is less accurate in the direction perpendicular to the motion of the robot. Combining one TWR measurement with four AoA antennas provides significantly better results, as can be seen in Figure 5.41(b). This is because both of the data sources are used in the most accurate way — TWR is used just to obtain the distance from the human and AoA is used to obtain the direction towards the human. The strength of the Bluetooth has to be used to select the most relevant measurement, but this approach cannot be completely reliable, this can be seen as a few jumps of the estimated position in the lower left corner of Figure 5.41(b).

In general, the most important factor that affects accuracy is the direct line-of-sight visibility of the modules. If some of the modules are blocked, the accuracy of the localisation will be low. If a higher operating range is required, the modules should be placed further apart from each other, but the limit is the footprint of the robot. Adding a Bluetooth Angle of Arrival measurement is mainly beneficial if the dimensions of the robot are small. The main drawback of AoA is that the antenna boards are significantly larger than the UWB modules, making it harder to find a suitable position for them.

All of the sensor configurations presented were sufficiently accurate to allow the robot to follow the human. Even the setup with three TWR modules on the Spot robot could be used if it can be guaranteed that the human will stay in the area where the estimates are accurate. Using just a single

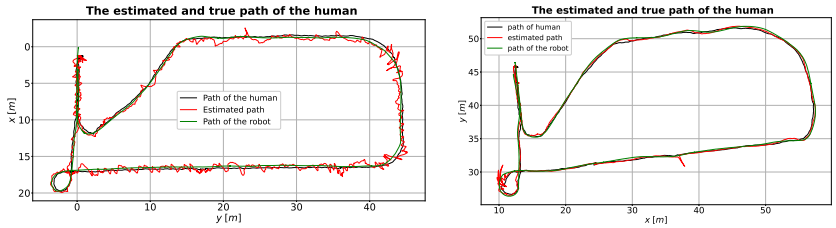

(a) (b)

Figure 5.41: Comparison of localisation with only four TWR modules and with one TWR module and four AoA antennas. Both figures show the estimated path of the human compared to the true one and the actual path of the human. More details can be seen in Figures 5.31 and 5.33.

TWR module and four AoA antennas proved to be sufficient for localising the human, but because only the strength of the Bluetooth signal can be used to determine the direction towards the human, the estimate might not be reliable in some specific situation.

The estimates were also sufficiently precise, multiple of the sensor configurations had the standard deviation of the estimates below 20 centimetres. The estimated position of the human remains therefore mostly within the body of the human or close to it.

Chapter 6
Conclusion

The "follow me" system developed in this work is capable of following the human safely. Thanks to the UWB TWR range measurements, the robot knows at all times at least the distance from the human, meaning that it is possible to stop the motion of the robot if the human gets too close. The accuracy of the localisation strongly depends on the selected robotic platform and the mounting points of the sensors. If the robot has a complex shape that means that the sensors do not have direct visibility in some directions, the localisation accuracy will be degraded. The robot also needs to be sufficiently large to allow the sensors to be mounted at an appropriate mutual distance; the TWR modules should be at least 30 centimetres apart from each other.

One of the proposed configurations, the cube with four AoA antennas and one TWR module, is suitable for easy deployment on new robots. Only the position of a single mounting point has to be measured, and only a single USB cable has to be routed to the sensor package. This means that the system is ready to be used in a short time.

Mounting all of the sensors on a single position might not always be possible; therefore, several other sensor configurations are proposed together with general guidelines for mounting the sensors.

6.1 Future Work

One of the possible extensions of this work is the addition of human-aware planning. The "follow me" system would be used to generate navigation targets for the planner. The planner would be responsible for creating an acceptable trajectory in an environment where other humans are present. A path following controller would then create steering inputs for the robot, the velocity should be controlled in a way that makes sure that the robot stays at a safe distance from the human. The system has to have two operating modes, one that uses the human-aware planner and another that is identical to the solution proposed in this work. The second mode acts as a backup in the event that vision is compromised due to the conditions of the environment. This ensures that the robot is still capable of fulfilling the main task of following the human.

One of the areas of possible improvement is the feedback provided to the

user. The proposed solution uses a speaker mounted on the robot to report the state of the system. This is acceptable if the human can hear the audio output of the robot. In case the environment is particularly noisy or the human wears protective equipment, no feedback would be available. It is possible to use Bluetooth to transmit the status messages to the tag carried by the human. The tag could contain a speaker that plays the status messages, or the status messages could be played into headphones.

Another option for understandable feedback is to add lights to the robot. These could be simple status LEDs that show the state of the system, but light can also be used to convey more complex messages. The robot could use lights to signal its intentions – the desired direction of movement, the need to stop because of an obstacle, the estimated position of the human. Some of these capabilities are already present on commercially available robotic platforms[1].

Staying behind the human is a great solution for situations where visual information is not available. In such a case, the human acts as a "path planner" — he selects a traversable collision-free path that the robot should be able to follow. If the robot is able to sense its environment through LiDAR or cameras, collision avoidance and path planning can be done onboard, meaning that the robot can be anywhere around the robot. The best feedback for the human is if he can see that the robot is moving together with him. If the robot is staying behind the human, it is outside of the field of view of the human and the human has to turn around to see the robot; this leads to decreased comfort of the user. If the robot is supposed to assist the human (for example, by carrying some equipment), it is beneficial if the robot stays in a position that is easy to access, for example if the robot stays to the side of the human, it is possible to interact with it even during movement.

The comfort of the user could be further increased by designing a custom hardware for the tag. The current solution uses two separate boards that cannot be stacked, but a single board containing both transmitters could have a significantly smaller footprint.

[1] `https://bostondynamics.com/blog/the-next-step-in-safe-autonomous-robotic-inspection/`

Bibliography

[1] Frauenhofer Institute for Communication, Information Processing and Ergonomics, "Scientifically supported Bundeswehr test series puts transport robots through their paces for the infantry." https://www.fkie.fraunhofer.de/en/press-releases/bw-test-series-ugv.html. Accessed: 2024-3-25.

[2] I. Ocskay, "Az integrált, moduláris, vezető nélküli szárazföldi járműrendszer bemutatója Németországban [*Exhibitions of integrated, modular unmanned ground vehicle systems in Germany*]," *Haditechnika*, vol. 57, no. 3, pp. 27–32, 2023.

[3] Rheinmetall, "The Rheinmetall Mission Master family." https://www.rheinmetall.com/en/products/unmanned-vehicles/unmanned-vehicles/mission-master-a-ugv. Accessed: 2023-11-22.

[4] Rheinmetall, "Path autonomy kit." https://www.rheinmetall.com/en/products/unmanned-vehicles/unmanned-vehicles/path-autonomy-kit. Accessed: 2023-11-22.

[5] Diehl Defence, "Keeping pace with the infantry." https://www.diehl.com/defence/en/press-and-media/news/keeping-pace-with-the-infantry/. Accessed: 2023-11-22.

[6] Elbit Systems, "Probot, a flexible platform to support the infantry." https://elbitsystems-de.com/en/probot-flexible-platform-to-support-the-infantry/. Accessed: 2023-11-22.

[7] Milrem Robotics, "The THeMIS UGV." https://milremrobotics.com/defence/. Accessed: 2023-11-22.

[8] Milrem Robotics, "Autonomy." https://milremrobotics.com/autonomy/. Accessed: 2023-11-22.

[9] J. Wang and E. Olson, "AprilTag 2: Efficient and robust fiducial detection," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2016.

[10] M. Manz, T. Luettel, F. von Hundelshausen, and H.-J. Wuensche, "Monocular model-based 3D vehicle tracking for autonomous vehicles in unstructured environment," in *2011 IEEE International Conference on Robotics and Automation*, pp. 2465–2471, IEEE, 2011.

[11] C. Fries and H.-J. Wuensche, "Monocular template-based vehicle tracking for autonomous convoy driving," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2727–2732, IEEE, 2014.

[12] J. L. Giesbrecht, H. K. Goi, T. D. Barfoot, and B. A. Francis, "A vision-based robotic follower vehicle," in *Unmanned Systems Technology XI*, vol. 7332, pp. 451–462, SPIE, 2009.

[13] M. Himmelsbach, A. Mueller, T. Lüttel, and H.-J. Wünsche, "LIDAR-based 3D object perception," in *Proceedings of 1st international workshop on cognition for technical systems*, vol. 1, 2008.

[14] C. Fries and H.-J. Wuensche, "Autonomous convoy driving by night: The vehicle tracking system," in *2015 IEEE International Conference on Technologies for Practical Robot Applications (TePRA)*, pp. 1–6, IEEE, 2015.

[15] A. Barrera, C. Guindel, J. Beltrán, and F. García, "Birdnet+: End-to-end 3d object detection in lidar bird's eye view," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1–6, IEEE, 2020.

[16] S. Pang, D. Morris, and H. Radha, "CLOCs: Camera-LiDAR object candidates fusion for 3D object detection," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 10386–10393, IEEE, 2020.

[17] H. Schneiderman, M. Nashman, A. J. Wavering, and R. Lumia, "Vision-based robotic convoy driving," *Machine Vision and Applications*, vol. 8, no. 6, pp. 359–364, 1995.

[18] V. Walter, N. Staub, A. Franchi, and M. Saska, "UVDAR System for Visual Relative Localization With Application to Leader–Follower Formations of Multirotor UAVs," *IEEE Robotics and Automation Letters*, vol. 4, pp. 2637–2644, July 2019.

[19] Y. Cheng and T. Zhou, "UWB Indoor Positioning Algorithm Based on TDOA Technology," in *2019 10th International Conference on Information Technology in Medicine and Education (ITME)*, pp. 777–782, 2019.

[20] R. Mazraani, M. Saez, L. Govoni, and D. Knobloch, "Experimental results of a combined TDOA/TOF technique for UWB based localization systems," in *2017 IEEE International Conference on Communications Workshops (ICC Workshops)*, pp. 1043–1048, 2017.

[21] J. Tiemann and C. Wietfeld, "Scalable and precise multi-UAV indoor navigation using TDOA-based UWB localization," in *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–7, 2017.

[22] S. Güler, M. Abdelkader, and J. S. Shamma, "Peer-to-Peer Relative Localization of Aerial Robots With Ultrawideband Sensors," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 5, pp. 1981–1996, 2021.

[23] T.-M. Nguyen, A. Hanif Zaini, C. Wang, K. Guo, and L. Xie, "Robust Target-Relative Localization with Ultra-Wideband Ranging and Communication," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2312–2319, 2018.

[24] K. Guo, X. Li, and L. Xie, "Ultra-Wideband and Odometry-Based Cooperative Relative Localization With Application to Multi-UAV Formation Control," *IEEE Transactions on Cybernetics*, vol. 50, no. 6, pp. 2590–2603, 2020.

[25] J. P. Queralta, Q. Li, F. Schiano, and T. Westerlund, "VIO-UWB-Based Collaborative Localization and Dense Scene Reconstruction within Heterogeneous Multi-Robot Systems," in *2022 International Conference on Advanced Robotics and Mechatronics (ICARM)*, pp. 87–94, 2022.

[26] C. C. Cossette, M. Shalaby, D. Saussié, J. R. Forbes, and J. Le Ny, "Relative Position Estimation Between Two UWB Devices With IMUs," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4313–4320, 2021.

[27] Z. Jianyong, L. Haiyong, C. Zili, and L. Zhaohui, "RSSI based Bluetooth low energy indoor positioning," in *2014 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 526–533, 2014.

[28] L. Bai, F. Ciravegna, R. Bond, and M. Mulvenna, "A Low Cost Indoor Positioning System Using Bluetooth Low Energy," *IEEE Access*, vol. 8, pp. 136858–136871, 2020.

[29] S. Chai, R. An, and Z. Du, "An Indoor Positioning Algorithm using Bluetooth Low Energy RSSI," in *Proceedings of the 2016 International Conference on Advanced Materials Science and Environmental Engineering*, pp. 274–276, Atlantis Press, 2016/04.

[30] G. Pau, F. Arena, Y. E. Gebremariam, and I. You, "Bluetooth 5.1: An Analysis of Direction Finding Capability for High-Precision Location Services," *Sensors*, vol. 21, no. 11, 2021.

[31] S. Monfared, T.-H. Nguyen, L. Petrillo, P. De Doncker, and F. Horlin, "Experimental Demonstration of BLE Transmitter Positioning Based

on AOA Estimation," in *2018 IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pp. 856–859, 2018.

[32] Z. Hajiakhondi-Meybodi, M. Salimibeni, K. N. Plataniotis, and A. Mohammadi, "Bluetooth Low Energy-based Angle of Arrival Estimation via Switch Antenna Array for Indoor Localization," in *2020 IEEE 23rd International Conference on Information Fusion (FUSION)*, pp. 1–6, 2020.

[33] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726–1743, 2013.

[34] P. Teja Singamaneni, A. Favier, and R. Alami, "Human-Aware Navigation Planner for Diverse Human-Robot Interaction Contexts," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5817–5824, 2021.

[35] A. Mateus, D. Ribeiro, P. Miraldo, and J. C. Nascimento, "Efficient and robust Pedestrian Detection using Deep Learning for Human-Aware Navigation," *Robotics and Autonomous Systems*, vol. 113, pp. 23–37, 2019.

[36] "IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs): Amendment 1: Add Alternate PHYs," *IEEE Std 802.15.4a-2007 (Amendment to IEEE Std 802.15.4-2006)*, pp. 1–210, 2007.

[37] "IEEE Standard for Low-Rate Wireless Networks–Amendment 1: Enhanced Ultra Wideband (UWB) Physical Layers (PHYs) and Associated Ranging Techniques," *IEEE Std 802.15.4z-2020 (Amendment to IEEE Std 802.15.4-2020)*, pp. 1–174, 2020.

[38] J. Krška, *Hyperbolic positioning in UWB networks with non-transmitting tag.* Master Thesis, CTU Faculty of Electrical Engineering, 2021.

[39] V. Navrátil, J. Krška, and F. Vejražka, "Concurrent bidirectional tdoa positioning in uwb network with free-running clocks," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 4434–4450, 2022.

[40] P. Šimek, *Precision of Position Determination Radio Systems.* Bachelor's Thesis, CTU Faculty of Electrical Engineering, 2022.

[41] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, *et al.*, "ROS: an open-source Robot Operating System," in *ICRA workshop on open source software*, vol. 3, p. 5, Kobe, Japan, 2009.

[42] M. A. Branch, T. F. Coleman, and Y. Li, "A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems," *SIAM Journal on Scientific Computing*, vol. 21, no. 1, pp. 1–23, 1999.

[43] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, p. 226–231, AAAI Press, 1996.

[44] Qorvo, "DWM1001 Firmware API Guide." `https://www.qorvo.com/products/d/da007975`. Accessed: 2023-11-22.

[45] u-blox, "u-locateEmbed AT Commands Manual." `https://content.u-blox.com/sites/default/files/documents/u-connectLocate-ATCommands-Manual_UBX-22025586.pdf`. Accessed: 2024-04-1.

Appendix A

Data Format of the Serial Communication

A.1 Qorvo DWM1001c

When the module is set to tag mode, it can send the measured distances over UART to a connected computer. Each anchor has an ID, and it can transmit its position (has to be set manually for each anchor). The messages received from the tag have the following form:

```
1151[5.00,8.00,2.25]=6.48 0CA8[0.00,8.00,2.25]=6.51
111C[5.00,0.00,2.25]=3.18 1150[0.00,0.00,2.25]=3.16
le_us=2576 est[2.57,1.98,1.68,100]
```

The individual fields of the message are separated by a space. The first four fields of the example message are the measurements of four anchors. The first four characters (for example, `1151`) are the ID of the anchor. The three floating point numbers in the brackets (for example `[5.00,8.00,2.25]`) are the coordinates of the anchor (provided during the setup of the anchor). The number after the equal sign is the measured distance between the anchor and the tag (for example `6.48`).

If at least three anchors are available, the position of the tag is estimated. The tag reports the time in microseconds required to compute the estimate (in the above example, it is the field `le_us=2576`) and the estimated position with the quality of the estimate (`est[2.57,1.98,1.68,100]`, the quality is 100). These two fields are optional; if less than three anchors are available, the message will not contain them. More details about the messages and possible settings of the module are described in the API documentation [44].

A.2 u-blox XPLR-AOA1 and ANT-B10

The antenna uses UART to send the measured angles to a connected computer. The messages have the following form:

```
+UUDF:CCF9578E0D8A,-42,20,0,-43,37,"CCF9578E0D89","",15869,23
```

`+UUDF:CCF9578E0D8A` means that the message is the angle calculation event for tag with ID `CCF9578E0D8A`. The measured values are in the second (RSSI of the first polarisation — `-42`), third (azimuth — `20`) and fourth (elevation — `0`) field. All of these values are integers. The fifth field is currently unused (`-43`), the sixth field contains the channel (`37`), the seventh field is the ID of the anchor (`"CCF9578E0D89"`), the eighth field is a user-defined message, the ninth field the timestamp (`15869`) and the tenth is the value of the counter contained in the packet sent by the tag (`23`).

The antenna receives AT commands over the serial interface, the available commands can be found in the AT Commands Manual [45]. In the case of the antenna board, the commands can mostly be used to read information about the board; it is possible to set the board to track a specific tag (each tag has a unique ID, and the board measures angles to all available tags by default). The active tag uses the same set of AT commands; it is, for example, possible to change the transmit power.

Appendix B

Custom ROS Messages

The driver for the u-blox XPLRAOA kits takes the information sent over serial communication and transforms it into a ROS message (Listing B.1). The measured azimuth and elevation are published in radians and have a value between $-\pi$ and π, RSSI is a negative floating point number with a value usually between -50 and -80.

Listing B.1: xplraoa1_ros/Angles

```
float64 azimuth
float64 elevation
float64 RSSI
```

Each range measurement corresponds to one anchor, the tag sends a list of these measurements. The ROS message (Listing B.2) that the driver publishes therefore has the same structure — it is a list of individual measurements. Each measurement is another ROS message (shown in Listing B.3) contains all the available information — the ID of the anchor, the position of the anchor (set in the firmware of the anchor during configuration) and the measured range.

Listing B.2: dwm1001c_ros/UWBMeas

```
dwm1001c_ros/Anchor[] measurements
```

Listing B.3: dwm1001c_ros/Anchor

```
string id
geometry_msgs/Point location
float64 dist
```

The DWM1001 module set to the tag mode is capable of computing its own position if at least three anchors are available. This work uses a single anchor, but the driver maintains this ability of the modules. The message (Listing B.4) contains the computed position, its quality (integer in range 0 to 100), and the computation time in microseconds.

Listing B.4: dwm1001c_ros/TagLocation

```
geometry_msgs/Point location
int32 quality
int32 computation_time
```

Appendix C

An Example of Deployment of the System

The process of deploying the system on a new robot will be shown in the example of a Spot robot with 3 TWR modules and 4 AoA antennas.

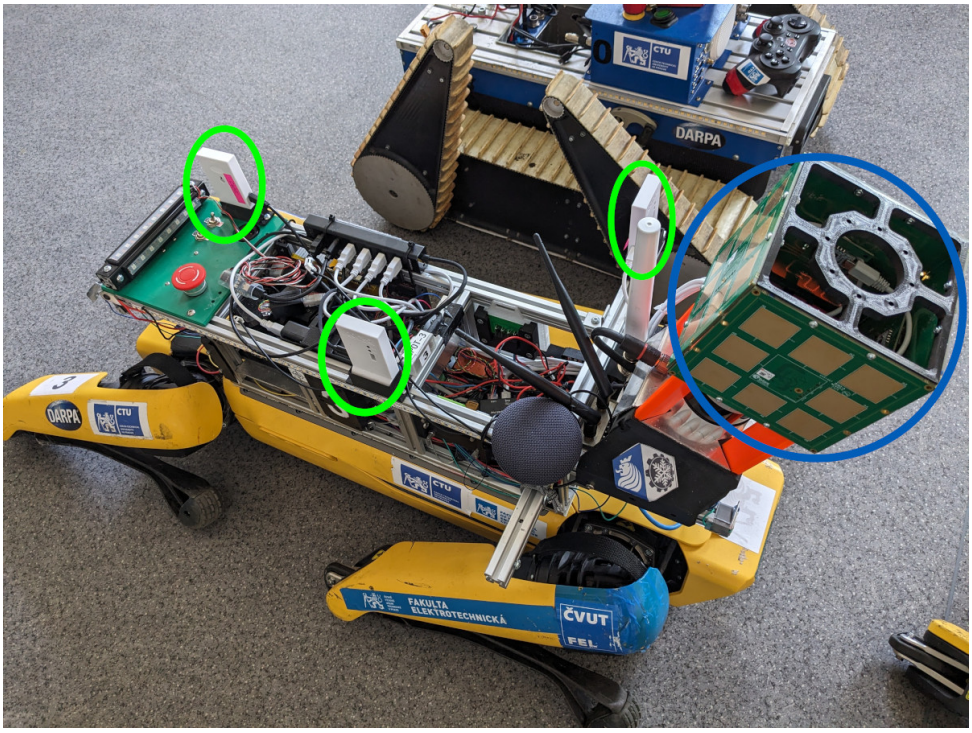

Figure C.1: The Spot robot equipped with three UWB TWR modules (marked with green colour and four Bluetooth AoA antenna boards (in blue circle). The reference point on the payload frame is marked with purple circle.

The first step in deploying the system is to select an appropriate number of sensors and their placement on the body of the robot. Three TWR modules are enough to localise the human, but the addition of AoA measurements can increase the usable range.

Once the sensors are mounted, it is necessary to measure their position. A single reference point has to be used for each system. This point can be seen in Figure C.1, there is a middle bar on top of the payload frame on the Spot robot, the position of this bar is known, and it is easy to measure the

position of each TWR module with respect to this bar. The AoA antennas are mounted as a single cube, and the position of each antenna with respect to the centre of the cube is known. The position of the bracket to which the AoA cube is mounted is known too, which means that no additional measurements are necessary. The position of the sensors should be measured to their reference frame, the AoA antenna measures the angles from the centre of the board, and TWR modules have the antenna placed in the upper left corner.

The measured positions must be written into the corresponding configuration files. The ROS package `follow_me` contains folder `config` with a separate configuration file for each system, an example of the config file for UWB TWR is shown in Listing C.1.

Listing C.1: The config file for UWB TWR

```
fixed_frame: "locator"
human_frame: "twr_position"
use_3d: false

ids: [
  "5722",
  "5A84",
  "C4B5"
]

positions: [
  #x,y,z
  0.175, 0.075, 0.115,
  -0.07, -0.04, 0.093,
  -0.30, 0.08, 0.098,
]

calibration: [
  0, 0, 1, 0,
  0, 0, 1, 0,
  0, 0, 1, 0,
]
```

The last step is to configure the launch files. It is necessary to set the ports to which the sensors are connected (e.g. `/dev/ttyACM0`). This is done in a separate file for each sensor system. It is advisable to create a udev rule for each sensor (example is shown in Listing C.2 for Qorvo UWB modules and Listing C.3 for u-blox ANT-B10 antenna boards connected through the FTDI UMFT234XF board), as Micro USB can temporarily lose connection due to heavy vibrations. The device then reconnects to the same port, and the driver can continue publishing the measurements.

Listing C.2: udev rule for the Qorvo DWM-1001-dev module

```
ACTION=="add", ENV{ID_VENDOR_ID}=="1366", ENV{ID_MODEL_ID}==
"0105", ENV{ID_SERIAL_SHORT}=="000760166508", SYMLINK+=
"ttyACM_uwb5722"
```

Listing C.3: udev rule for u-blox ANT-B10 antenna boards connected through FTDI UMFT234XF board

```
ACTION=="add", ENV{ID_VENDOR_ID}=="0403", ENV{ID_MODEL_ID}==
"6015", ENV{ID_SERIAL_SHORT}=="FT5TXNW4", SYMLINK+=
"ttyUSB_aoa1"
```

Appendix D

Source Code of the ROS Nodes

The source code for all the nodes developed as part of this work is available on GitHub. The nodes are split into multiple ROS packages:

- UWB TWR driver: `https://github.com/brozb/dwm1001_ros`
- Bluetooth AoA kit driver: `https://github.com/brozb/xplraoa_ros`
- UWB TDoA driver: `https://github.com/brozb/uwb_tdoa`
- Localisation nodes: `https://github.com/brozb/follow_me`
- Node for tracking the human in point clouds: `https://github.com/brozb/lidar_detector`