

Angle of Arrival Estimation based on Channel Impulse Response Measurements

Anton Ledergerber*, Michael Hamer and Raffaello D’Andrea

Abstract—In recent years, ultra-wideband radio technology has become increasingly popular as a space- and cost-effective solution to the problem of indoor localization. This paper demonstrates how measurements of the channel impulse response can be used to estimate a signal’s angle of arrival at a receiving antenna. This novel method requires no additional hardware, uses only a single antenna, and works with unsynchronized clocks and one-way communication. We evaluate our method on a real-dataset, and experimentally demonstrate how a mobile robot can localize itself by measuring angles to multiple ultra-wideband anchors.

I. INTRODUCTION

The design of small, omnidirectional ultra-wideband (UWB) antennas is challenging, since the frequency response of the antenna should ideally be constant over different angles and for a wide range of frequencies [1]. Furthermore, these antennas are often mounted on devices (e.g. robots), whose physical construction causes signal reflection and distortion. Non-idealities in the antenna and its surroundings affect the channel impulse response (CIR) and reduce the accuracy of timestamps obtained via leading edge detection. As noted in [2], [3], this leads to systematic inaccuracies in time-of-flight-based UWB localization systems. Existing research has attempted to compensate for these inaccuracies via tailored antenna designs [4], [5], or by using models to predict and correct for systematic timestamp inaccuracies [6], [7]. Instead of trying to reduce the effects of antenna non-idealities, this paper proposes the opposite: use, and where possible amplify, these non-idealities to estimate a signal’s angle of arrival (AOA).

Historically, AOA estimation has required either a rotating, directional antenna, or the measurement of a signal at multiple locations (e.g. using an array of antennas or multiple, synchronized receivers). The AOA estimation method presented in this work requires no additional hardware, enabling the estimation of a signal’s AOA using a single receiving antenna. This method can either be used in a standalone fashion, requiring no clock synchronization between the transmitter and the receiver, or can be used to augment existing time-of-flight or received-signal-strength UWB localization systems without requiring hardware changes.

A. Related Work

An overview of how the environment affects the channel impulse response via fundamental propagation processes

This work was supported by the Swiss National Science Foundation
The authors are with the Institute for Dynamic Systems and Control,
ETH Zurich, Switzerland

* Corresponding author: antonl@ethz.ch

such as reflection and diffraction is given in [8]. The environment can be used in several ways in UWB localization systems. As discussed in [9], if it is possible to obtain ground truth measurements for a given environment, fingerprinting techniques that make use of the environment’s channel impulse response can be used for localization. Alternatively, and as investigated by [10], [11], if a floor plan of reflective surfaces is available, it is possible to localize by resolving multipath components in the measured channel impulse response and mapping them to known reflective surfaces.

In addition to environmental effects, the signal’s propagation channel is affected by the frequency response of both the transmitting and receiving antenna. The effects of the antennas’ frequency responses on the accuracy of UWB systems is investigated in [12], [13] and [14].

AOA estimation is at the core of most radar applications [15]. Overviews of AOA estimation techniques are given in [16], [17]. Generally, the AOA is determined via two or more antennas in an array (e.g. [18]), by rotating antennas with a directive antenna pattern (e.g. [19]) or by collecting measurements at multiple points in space (e.g. [20]). The work presented herein demonstrates AOA estimation using a single antenna, attached to a low-cost and small size UWB radio. This work can be combined with time of flight localization approaches using identical hardware; however, we do not discuss this further in this paper. To the best of the authors’ knowledge, this is the first time that AOA was estimated using only one transmitting and one receiving antenna and without relative motion between the two antennas.

B. Outline

The paper is structured as follows. In Section II, we discuss how CIR measurements are obtained and investigate distortions caused by the antenna and by objects in the antenna’s vicinity. In Section III, we explain how knowledge of these systematic distortions can be used to estimate the AOA of a UWB signal and how a CIR to AOA mapping can be learned with a neural network. The quality of this mapping is evaluated experimentally in Section IV, and the method’s applicability to robot localization is demonstrated in Section V, where a robot is localized solely based on AOA estimates obtained via the CIR.

II. CHANNEL IMPULSE RESPONSE MEASUREMENTS

A UWB communication channel is characterized by its channel impulse response function, which is extremely dif-




Fig. 1. This figure shows a transmitting antenna sending a signal to a receiving antenna, which is mounted on a rigid body with a reflective surface in its vicinity. As shown, a signal propagates along multiple pathways of different lengths on its way to the receiver. The signal’s multipath components are recombined upon reception with various phase-shifts. This recombination can be noticeable in the CIR and can reduce the accuracy of leading-edge detection methods. Each pathway in the diagram is labeled according to its contribution to the CIR, as in (1). The ellipse as given in (2) is also shown.

difficult to model accurately since many influences need to be considered [8]. Of these influences, we explicitly discuss two: the antenna transfer function, and objects in the environment.

The transfer functions of the transmitting and receiving antennas are generally a function of the angle of departure (AOD) α_{tx} and AOA α_{rx} , respectively. These functions can have a large influence on the CIR, particularly if directional antennas are used [14].

Objects in the environment which interact with the electromagnetic signals also strongly affect the CIR. We group these objects into those objects belonging to the same rigid body as the transmitting or receiving antennas (hereinafter the antenna’s “local environment”), and objects belonging to the larger environment, which are typically further away from the antennas. This is depicted in Fig. 1.

We combine the impulse response of the transmitting and receiving antennas with the impulse response caused by objects in each antenna’s local environment. We denote these combined impulse responses as $h_{tx}(t, \alpha_{tx})$ and $h_{rx}(t, \alpha_{rx})$, respectively, where α_{tx} and α_{rx} denote AOD and AOA. Assuming a cascaded, linear, time invariant model, and a UWB channel with N_{MP} different multipath components, the measured channel impulse response can be written as

$$h_{CIR}(t) = \sum_{n=1}^{N_{MP}} h_{tx}(t, \alpha_{tx, n}) * h_{env, n}(t) * h_{rx}(t, \alpha_{rx, n}), \quad (1)$$

where $*$ is the convolution operator, where the AOD and the AOA of the n -th multipath component are denoted by $\alpha_{tx, n}$ and $\alpha_{rx, n}$ respectively, and where $h_{env, n}$ denotes the transfer function of the environment’s n -th multipath component. This is illustrated in Fig. 1.

To investigate the effects of the antennas’ transfer functions and of objects in their local environments, a UWB module is placed at a fixed position and instructed to transmit, while a receiving module is placed 5 m from the transmitter on a rotating platform, allowing measurement of the CIR for different AOAs around the antenna’s azimuth axis. The elevation angle is kept at zero throughout the paper as we only consider a two-dimensional setup; however,

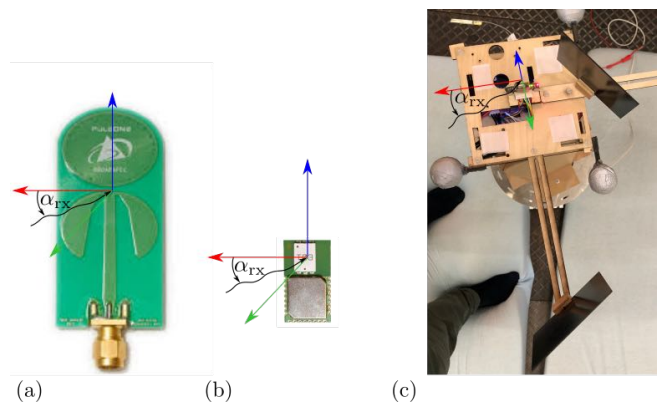


Fig. 2. The different receiver configurations for which the CIR is measured. (a) Time Domain Broadspec Antenna, (b) Partron dielectric chip antenna, (c) Partron dielectric chip antenna with carbon plates attached in close proximity. Receiving antennas are placed on a rotating platform in order to measure their CIR over a large range of AOAs.

the techniques presented herein are applicable to three-dimensional setups. DW1000 UWB radio modules [21] are used for both the transmitter and receiver.

The transmitting module is outfitted with a Time Domain Broadspec antenna, and three different receiver configurations are tested: a Time Domain Broadspec antenna, a Partron dielectric surface-mount antenna, and a Partron dielectric surface-mount antenna with reflective carbon plates mounted in its vicinity. These three configurations are shown in Fig. 2. The position and attitude of each antenna is recorded using a motion capture system.

Measurement of the CIR is enabled by the DW1000, which gives access to its measurement of the complex-valued impulse response envelope with a resolution of $T_s = 1/(2f_c) \approx 1$ ns, where $f_c = 499.2$ MHz is the chipping frequency. The magnitude of this complex-valued impulse response envelope is the envelope of the physical (real-valued) CIR [15, p.281]. For simplicity, we hereinafter use CIR to mean the complex-valued impulse response envelope. To achieve a sub-nanosecond resolution, multiple CIRs sampled at different times can be aligned and combined into a single,

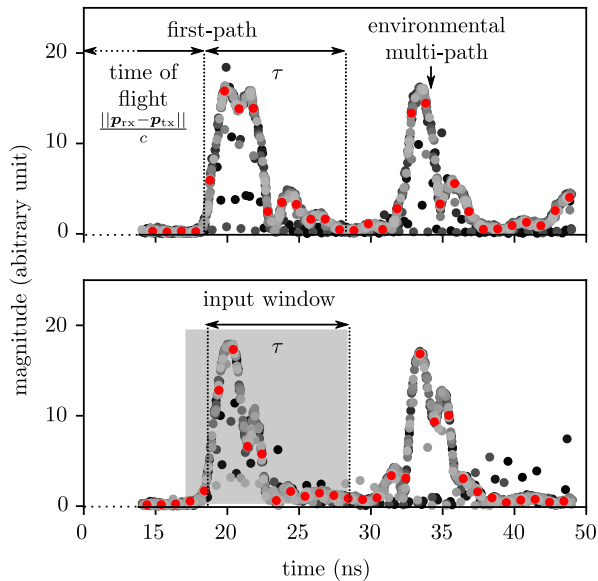


Fig. 3. The magnitude of two accumulated CIRs, each taken at a different AOA and based on 50 CIR measurements. In both plots, the different CIR measurements are aligned by the first-path estimate provided by Decawave’s proprietary leading edge detection algorithm. Different measurements are shown in different tones of gray, with one measurement in each plot highlighted in red. Each accumulated CIR is shown starting four samples before the estimated first path. On each plot we note a secondary peak caused by an environmental multi-path. When comparing the two plots, we note that the first peak is of a slightly different shape: this is caused by an AOA-dependent antenna transfer function, and by the overlapping of different reflections caused by the receiving antenna’s local environment. We tune the variable τ to trim multi-path components caused by the environment, and estimate the AOA based only on samples of a single measurement (e.g. shown in red) within a window of the first path.

high resolution CIR. The alignment can be done based on Decawave’s proprietary leading edge detection algorithm’s estimate of the first path, or by solving an optimization problem, or using ground truth measurements [22]. The magnitude of such an accumulated CIR is shown in Fig. 3, with tones of gray representing different measurements of the CIR, and with one such measurement highlighted in red.

The magnitude of the accumulated CIR for the three different configurations and across different AOAs is shown in Fig. 4. We note that the CIR obtained with the larger Time Domain Broadspec antenna (Fig. 4(a)) is more uniform than the CIR obtained with the smaller Partron dielectric antenna (Fig. 4(b)). If reflective surfaces are placed close to the antenna, the measured CIR is further distorted (Fig. 4(c)). Although it may at first appear contrived, in the majority of applications antennas are mounted on or are integrated in devices made of reflective materials such as metals or carbon fibers, and this third situation is therefore representative of a realistic use-case. As noted in, for example [6], [7], leading-edge detection algorithms can be sensitive to such AOA-dependent changes in the CIR, resulting in systematic biases or increased noise in reception timestamps. For timestamp-based localization, it is therefore important to minimize these effects. To the contrary, our method demonstrates that these AOA-dependent distortions in the CIR provide useful




Fig. 4. The magnitude of the accumulated CIR is plotted as a function of AOA (vertical axis) for each of the different receiver configurations shown in Fig. 2. A large magnitude is shown in white, and a small magnitude is shown in black. The first path is at $t = 0$ ns.

information for estimating a signal’s AOA.

III. LEARNING THE CIR TO AOA MAPPING

A. CIR Measurement and Windowing

Considering Fig. 3, we note that measurements of the CIR taken at different AOAs show differences in the shape of the first peak. For line-of-sight conditions, these differences in shape can be attributed to an AOA-dependent antenna transfer function and to changes in the delay of reflections and disturbances caused by objects in the antenna’s local environment.

As shown in Fig. 1 and Fig. 3, environmental multi-path components have an additional influence on the measured CIR. We tune the variable τ to include reflections from the antenna’s local environment, and to exclude paths more than a certain distance $c\tau$ longer than the direct path. This defines an ellipse given by

$$\left\| \mathbf{p}_{\text{tx}} - \begin{pmatrix} x \\ y \end{pmatrix} \right\| + \left\| \mathbf{p}_{\text{rx}} - \begin{pmatrix} x \\ y \end{pmatrix} \right\| \leq \|\mathbf{p}_{\text{rx}} - \mathbf{p}_{\text{tx}}\| + c\tau, \quad (2)$$

where \mathbf{p}_{tx} and \mathbf{p}_{rx} are the positions of the transmitter and receiver respectively, $\|\cdot\|$ is used to denote the Euclidean norm, c is the speed of light, and (x, y) are the x, y coordinates of an arbitrary point within the ellipse. We assume that the ellipse formed between each transmitting anchor and each position within the robot’s operating area




Fig. 5. A standard neural network with three hidden layers is used to predict the probability distribution of the AOA given a CIR. As shown in Fig. 3, a single measurement of the complex CIR envelope is windowed starting two samples before the estimated first path and finishing ten samples ($\tau = 10$ ns) after the first path. The magnitude and phase of these thirteen samples are fed to the input layer of the neural network. The outputs of the neural network are the unnormalized log probabilities of the received signal's AOA corresponding to each of the N_{bin} bins.

is free from environmental objects, and thus assume that a CIR measurement trimmed after τ seconds of the first path is entirely a function of the impulse responses of the transmitting and receiving antennas' local environment: $h_{\text{tx}}(t, \alpha_{\text{tx}})$ and $h_{\text{rx}}(t, \alpha_{\text{rx}})$, respectively.

Unfortunately, as shown in Fig. 4 the mapping of CIR to AOA is not always one-to-one: a similar CIR might be obtained for different AOAs due to these AOAs having similar antenna transfer functions. Despite this, it is possible to learn the probability that the measured CIR is a result of convolving $h_{\text{tx}} * h_{\text{env}}$ with $h_{\text{rx}}(t, \alpha_{\text{rx}})$. We approach this modeling using a neural network trained on a large dataset of CIR measurements paired to their ground-truth AOA which were obtained by a motion capture system.

B. The Network Structure

Given these CIR AOA pairs, the conditional probability distribution $p(\alpha_{\text{rx}} | h_{\text{CIR}}(t))$ is learned in a supervised learning framework by minimizing the cross-entropy between the training data and the model distribution [23, p. 173]. The AOA is discretized into $N_{\text{bin}} = 256$ bins and a neural network is trained to predict the unnormalized log probabilities that a signal is received with an AOA corresponding to a certain bin. The neural network consists of three, fully connected, hidden layers of size 800, 400 and 400, and layer outputs are passed through the rectified linear unit activation function (ReLU, [23, p. 168]). The network's structure is shown in Fig. 5. The input to the neural network is a 13-sample window of the complex-valued CIR envelope, starting two samples before the location of the first-path sample and ending 10 samples ($\tau = 10$ ns) after the first path sample. This corresponds to an ellipse with $c\tau = c \cdot 10$ ns = 3 m as defined in (2). The magnitude and phase of the 13 complex-valued samples (contained in the vector $\mathbf{h}_{\text{CIR}} \in \mathbb{C}^{13}$) are fed to the input layer of the neural network. Denoting with $\mathbf{z} \in \mathbb{R}^{N_{\text{bin}}}$ the output of the neural network, the normalized probability that the signal's AOA belongs to bin

$i \in \{1, 2, \dots, N_{\text{bin}}\}$ is

$$p(\text{bin} = i | \mathbf{h}_{\text{CIR}}) = \frac{\exp(\mathbf{z}[i])}{\sum_{j=1}^{N_{\text{bin}}} \exp(\mathbf{z}[j])}. \quad (3)$$

We train the network to minimize the cross-entropy loss

$$J(\alpha_{\text{rx}}, \mathbf{h}_{\text{CIR}}) = -\log p(\text{bin} = \text{bin}(\alpha_{\text{rx}}) | \mathbf{h}_{\text{CIR}}), \quad (4)$$

where $\text{bin}(\alpha_{\text{rx}})$ denotes the ground-truth bin of the training sample.

C. Training

The network is implemented using Tensorflow [24] and trained on $N_{\text{train}} = 453000$ $(\alpha_{\text{rx}}, \mathbf{h}_{\text{CIR}})$ pairs using the ADAM optimizer [23, p. 301] and the cost function given in (4). This data is collected with a setup as shown in Fig. 6, where transmitters are placed around a rectangle measuring approximately 4 m \times 3 m. A receiver mounted on a moving platform (a Roomba robot) drives around in a random fashion within the space. The transmitters use the Time Domain Broadspec antenna (Fig. 2(a)), and the receiver uses a Partron dielectric chip antenna with two carbon plates placed in its local environment (Fig. 2(c)). Each transmitter sends messages at random times to the receiver, and the receiver records the associated CIR measurement. The random transmission interval allows receivers to measure the CIR at an average frequency of 360 Hz. The slow data logging routine limits this frequency, which could be substantially increased if position estimation was performed directly on the robot.

The $(\alpha_{\text{rx}}, \mathbf{h}_{\text{CIR}})$ pairs are randomly shuffled and split into a training and a validation dataset. Training is stopped once the accuracy in the validation datasets stops improving. Additionally, an independent test dataset is collected in which the moving platform covers positions and orientations not visited in the training dataset.



Fig. 6. In order to measure the CIR for different AOAs, different AODs and different distances, a UWB receiver is mounted on a moving, rotating platform (Roomba). Transmitters send signals in a random fashion to the receiver, which measures the associated CIR. Both receiver and transmitter are equipped with markers which can be seen by the overhead motion capture system for purposes of gathering ground-truth data.

IV. RESULTS

This section describes the performance of the neural network described in the previous section on the test data set. If the estimated AOA is defined to be the maximum a-posteriori estimate predicted by the network, i.e.

$$\widehat{\text{bin}} := \arg \max_{i \in \{1, \dots, N_{\text{bin}}\}} p(\text{bin} = i | \mathbf{h}_{\text{CIR}}) \quad (5)$$

$$\widehat{\alpha}_{\text{rx}} := 2\pi \frac{\widehat{\text{bin}} - 1}{N_{\text{bin}}}, \quad (6)$$

a cumulative distribution plot of the error in the AOA can be made as shown in blue in Fig. 7. In this plot, we see that roughly 80% of the maximum a-posteriori estimates show an absolute error of less than 25° . By collecting 10 consecutive measurements sampled at different times $\mathbf{h}_{\text{CIR}}(1), \dots, \mathbf{h}_{\text{CIR}}(10)$ from the same transmitter, and predicting the AOA as

$$\widehat{\text{bin}} := \arg \max_{i \in \{1, \dots, N_{\text{bin}}\}} \sum_{j=1}^{10} p(\text{bin} = i | \mathbf{h}_{\text{CIR}}(j)) \quad (7)$$

we see an increase in prediction accuracy. The cumulative distribution of this error is shown in orange in Fig. 7, and shows that more than 90% of the predictions have an absolute error of less than 25° .

The output of the neural network is, however, not a single probability or bin estimate, but rather a probability distribution across all bins. This distribution is visualized in Fig. 8 for measurements taken by a robot moving in the environment shown in Fig. 9. We note that the distribution is bimodal for the CIR measurements obtained around $t = 71$ s, and that although the maximum a-posteriori estimate of the AOA may be incorrect, the probability assigned by the neural network to the correct bin is still quite high.

V. APPLICATION TO LOCALIZATION SYSTEMS

Any UWB localization system based on time of flight measurements, or based on received signal strength can be augmented by AOA estimation. Such hybrid UWB localization approaches are for example discussed in [25]. In the following, and in order to show the potential of the proposed AOA estimation technique for localization, we localize a robot based solely on the a-posteriori AOA probability distribution $p(\alpha_{\text{rx}} | \mathbf{h}_{\text{CIR}})$, provided by the trained neural network.




Fig. 7. The cumulative probability of the error in the estimated, maximum a-posteriori AOA. In blue, the maximum a-posteriori estimate given a single CIR measurement; in orange, the maximum a-posteriori estimate given 10 consecutive CIR measurements.

A. A Simple Localization Problem

We consider the problem of localizing the mobile robot (Roomba) used for collecting the training and test data as described in Section III-C. The robot moves in a 2-D environment with randomly transmitting anchors placed at known positions. The robot's state in the inertial reference frame is given by $\mathbf{x} = (x_R, y_R, \theta_R)$, where x_R and y_R are the robot's Cartesian coordinates, and θ_R is the angle describing its orientation (see Fig. 9). Using wheel odometry, the robot measures the distance travelled Δp and the change of its orientation $\Delta \theta$ every 15 ms. Using these measurements as system inputs $\mathbf{u} = (\Delta p, \Delta \theta)$, the robot's discrete time process model $\mathbf{x}(k+1) := q(\mathbf{x}(k), \mathbf{u}(k), \boldsymbol{\eta}(k))$ with process noise $\boldsymbol{\eta} = (\eta_x, \eta_y, \eta_\theta)$ for a sampling period of 15 ms is described by the following equations

$$x_R(k+1) = x_R(k) - \sin(\theta_R(k))\Delta p(k) + \eta_x(k) \quad (8)$$

$$y_R(k+1) = y_R(k) + \cos(\theta_R(k))\Delta p(k) + \eta_y(k) \quad (9)$$

$$\theta_R(k+1) = \theta(k) + \Delta \theta(k) + \eta_\theta(k), \quad (10)$$

where $k = 1, 2, \dots$, and where the process noise is assumed to have a zero mean normal distribution

$$\boldsymbol{\eta}(k) \sim \mathcal{N}(0, \boldsymbol{\Sigma}), \boldsymbol{\Sigma} = \begin{pmatrix} (2 \text{ mm})^2 & 0 & 0 \\ 0 & (2 \text{ mm})^2 & 0 \\ 0 & 0 & (1.5^\circ)^2 \end{pmatrix}. \quad (11)$$

The robot is equipped with a UWB receiver, which measures the CIR whenever a signal is received from one




Fig. 8. This figure shows the probability distribution over the AOA given by the neural network for a single transmitter during one minute of the test dataset. The ground truth AOA is indicated with a red dashed line. It is visible that at $t = 71$ s the distribution is bi-modal. The position and orientation of the receiver for this part of the test dataset is labeled in Fig. 9, and a circle used to mark the transmitter from which the CIR is measured.

of the randomly transmitting anchors at time k . This CIR measurement is passed through a trained neural network to obtain the estimated binned probability density function of the AOA $p(\alpha_{rx}(k)|\mathbf{h}_{CIR}(k))$.

B. Particle Filter

Due to the potential multi-modality of the a-posteriori AOA probability distribution, a particle filter is used to track the robot's state estimate. The particle filter framework is briefly outlined in the following, and the reader is referred to [26] for a more in depth introduction.

1) *Initialization*: The particle filter is initialized with $N_{PF} = 1000$ particles $\mathbf{x}^p, p \in \{1, 2, \dots, N_{PF}\}$ whose initial orientations are drawn from the uniform distribution $U_\theta(-\pi, \pi)$ and whose initial coordinates are drawn from $U_{x_R}(-0.5 \text{ m}, 4.5 \text{ m})$ and $U_{y_R}(-0.5 \text{ m}, 3.5 \text{ m})$.

2) *Prediction step*: At each iteration, the process model of (8)-(10) is used to update to each particle \mathbf{x}^p as

$$\mathbf{x}^p(k+1) = q(\mathbf{x}^p(k), \mathbf{u}(k), \boldsymbol{\eta}^p(k)) \quad (12)$$

3) *Measurement update*: When a UWB signal is received, each particle is weighted according to the a-posteriori probability distribution, given by the neural network, i.e.

$$w^p(k) = p(\alpha_{rx}^p(k)|\mathbf{h}_{CIR}(k)), \quad (13)$$




Fig. 9. This figure shows the transmitter positions (▼) and the path travelled by the robot in the test dataset (gray line). Only 60 seconds of this test is shown in black to improve plot clarity and facilitate comparison with the estimate of the particle filter over this same period (blue line). The tracking results across the entire 600 second experiment are plotted in Fig. 10. The AOA α_{rx} with respect to one transmitter is labeled at 71 s, and is additionally marked in Fig. 8 for purposes of comparison over the same 60 seconds.

where the expected AOA α_{rx}^p of each particle p is

$$\alpha_{rx}^p(k) = \text{atan2}(y_{tx} - y_R^p(k), x_{tx} - x_R^p(k)) - \theta_R^p(k), \quad (14)$$

wherein x_{tx} and y_{tx} are the x and y coordinates of the transmitter from which a signal was received. After the particle weights have been calculated, the particles are resampled to get N_{PF} posterior particles, all with equal weights.

C. Results

The performance of the particle filter was evaluated using the test dataset described in Section IV. Fig. 9 shows a 60 second segment of the path travelled by the robot in black (only 60 seconds are plotted to increase clarity of the plot, the full 600 second experiment is shown in Fig. 10). The path estimated by the particle filter is shown in blue (calculated as the mean of the x, y coordinates of the particles). The root-mean-squared error in the estimated position and the estimated orientation are 0.26 m and 2.9°, respectively.

To investigate whether it is also possible to estimate the state of an object with unknown dynamics using the AOA estimate, a second particle filter was run using a random walk process model, i.e.

$$x_R(k+1) = x_R(k) + \eta_x(k) \quad (15)$$

$$y_R(k+1) = y_R(k) + \eta_y(k) \quad (16)$$

$$\theta_R(k+1) = \theta(k) + \eta_\theta(k). \quad (17)$$




Fig. 10. The robot's path is shown in black, the mean of the state estimate generated by the particle filter with the process model given in (8)-(10) is shown in blue, and the mean of the state estimate of the particle filter employing the process model given in (15)-(17) is shown in orange.

Since this process model is driven solely by the zero mean process noise, the noise's covariance is increased to

$$\Sigma = \text{diag}((12 \text{ mm})^2, (12 \text{ mm})^2, (5.4^\circ)^2), \quad (18)$$

in order to track the robot's motion. The mean of the state distribution generated by this particle filter is shown in orange in Fig. 10. It is clear the evolution of the estimated state is not as smooth as for the case with the correct process model, particularly in positions where the AOA distribution is multi-modal (e.g. at approx. $t = 71$ s). Nevertheless, the filter is able to track the robot's state with a root-mean-squared error in the estimated position of 0.37 m and in the estimated orientation of 3.6° .

VI. CONCLUSION

This paper discusses how the AOA of a UWB signal can be estimated based on a measurement of the CIR, which is a function of the angle-dependent antenna transfer function. To the best knowledge of the authors this approach is a novel one, with this paper serving as a proof of concept. However, the results presented in this paper only scratch the surface of this novel sensor modality and more research is needed to evaluate its applicability to real-world problems.

ACKNOWLEDGMENT

The authors would like to thank Michael Egli, Marc-Andr  Corzillius and Karl Marcus Aalton for their technical support. Many people have contributed to the Flying Machine Arena in which this research was conducted. A list of these people can be found at: <http://flyingmachinearena.org/people>.

REFERENCES

- [1] S. B. T. Wang, A. M. Niknejad, and R. W. Brodersen, "Circuit modeling methodology for UWB omnidirectional small antennas," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 4, pp. 871–877, Apr. 2006.
- [2] Y. Shen and M. Z. Win, "Effect of path-overlap on localization accuracy in dense multipath environments," in *2008 IEEE International Conference on Communications*, May 2008, pp. 4197–4202.
- [3] M. Hamer and R. D'Andrea, "Self-calibrating ultra-wideband network supporting multi-robot localization," *IEEE Access*, vol. 6, pp. 22 292–22 304, 2018.
- [4] R. Cicchetti, E. Miozzi, and O. Testa, "Wideband and UWB antennas for wireless applications: A comprehensive review," *International Journal of Antennas and Propagation*, vol. 2017, 2017, Article ID 2390808.
- [5] A. Dumoulin, "A study of integrated UWB antennas optimised for time domain performance," Doctoral Thesis, Dublin Institute of Technology, 2012.
- [6] K. Bregar and M. Mohor i , "Improving indoor localization using convolutional neural networks on computationally restricted devices," *IEEE Access*, vol. 6, pp. 17 429–17 441, 2018.
- [7] A. Ledergerber and R. D'Andrea, "Calibrating away inaccuracies in ultra wideband range measurements: A maximum likelihood approach," *IEEE Access*, vol. 6, pp. 78 719–78 730, 2018.
- [8] A. F. Molisch, "Ultra-wide-band propagation channels," *Proceedings of the IEEE*, vol. 97, no. 2, pp. 353–371, Feb. 2009.
- [9] C. Steiner, "Location fingerprinting for ultra-wideband systems: The key to efficient and robust localization," Doctoral Thesis, ETH Zurich, 2010.
- [10] P. Meissner and K. Witrisal, "Multipath-assisted single-anchor indoor localization in an office environment," in *2012 19th International Conference on Systems, Signals and Image Processing (IWSSIP)*, Apr. 2012, pp. 22–25.
- [11] J. Kulmer *et al.*, "Using DecaWave UWB transceivers for high-accuracy multipath-assisted indoor positioning," in *2017 IEEE International Conference on Communications Workshops (ICC Workshops)*, May 2017, pp. 1239–1245.
- [12] F. Munier *et al.*, "On the effect of antennas on uwb systems," in *28th URSI General Assembly, New Delhi, India*, Oct. 2005, pp. 128–137.
- [13] V. Sipal, B. Allen, and D. Edwards, "Effects of antenna impulse response on wideband wireless channel," Nov. 2010, pp. 129–132.
- [14] A. Muqaiel, A. Safaai-Jazi, B. Woerner, and S. Riad, "UWB channel impulse response characterization using deconvolution techniques," in *The 2002 45th Midwest Symposium on Circuits and Systems, 2002. MWSCAS-2002.*, vol. 3, Aug 2002, pp. III–605.
- [15] J. Sachs, *Handbook of ultra-wideband short-range sensing: Theory, sensors, applications*. Weinheim: Wiley-VCH, 2012.
- [16] S. Barua, a. P. Ghosa, and a. K. Sandrasegaran, "A survey of direction of arrival estimation techniques and implementation of channel estimation based on SCME," in *2015 12th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology (ECTI-CON)*, June 2015, pp. 1–5.
- [17] X. Zhang and R. Cao, *Direction of Arrival Estimation: Introduction*. Wiley Encyclopedia of Electrical and Electronics Engineering, 2017, pp. 1–22.
- [18] I. Dotlic *et al.*, "Angle of arrival estimation using decawave DW1000 integrated circuits," in *14th Workshop on Positioning, Navigation and Communications (WPNC)*, Oct. 2017, pp. 1–6.
- [19] B.-C. Min *et al.*, "A directional antenna based leader–follower relay system for end-to-end robot communications," *Robotics and Autonomous Systems*, vol. 101, pp. 57–73, Mar. 2018.
- [20] C. J. Lowrance and A. P. Lauf, "Direction of arrival estimation for robots using radio signal strength and mobility," in *2016 13th Workshop on Positioning, Navigation and Communications (WPNC)*, Oct. 2016, pp. 1–6.
- [21] "Product Documentation | DecaWave." [Online]. Available: <https://www.decawave.com/product-documentation>
- [22] A. Moschevikin *et al.*, "Investigations on passive channel impulse response of ultra wide band signals for monitoring and safety applications," in *2016 3rd International Symposium on Wireless Systems within the Conferences on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS-SWS)*, Sept. 2016, pp. 97–104.
- [23] Ian Goodfellow and Yoshua Bengio and Aaron Courville, *Deep Learning*. MIT Press, 2016.
- [24] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [25] A. Alarifi *et al.*, "Ultra wideband indoor positioning technologies: Analysis and recent advances," *Sensors (Basel, Switzerland)*, vol. 16, no. 5, 2016.
- [26] D. Simon, *Optimal state estimation : Kalman, H ∞ , and nonlinear approaches*. Hoboken, N.J: Wiley-Interscience, 2006.