SALMA: UWB-based Single-Anchor Localization System using Multipath Assistance

Bernhard Großwindhager[†], Michael Rath[†], Josef Kulmer, Mustafa S. Bakr,
Carlo Alberto Boano, Klaus Witrisal, and Kay Römer
Faculty of Electrical and Information Engineering, Graz University of Technology, Austria
{grosswindhager, mrath, kulmer, mustafa.bakr, cboano, witrisal, roemer}@tugraz.at
[†] Authors contributed equally to this work

ABSTRACT

Setting up indoor localization systems is often excessively time-consuming and labor-intensive, because of the high amount of anchors to be carefully deployed or the burdensome collection of fingerprints. In this paper, we present SALMA, a novel low-cost UWB-based indoor localization system that makes use of only one anchor and that does neither require prior calibration nor training. By using only a crude floor plan and by exploiting multipath reflections, SALMA can accurately determine the position of a mobile tag using a single anchor, hence minimizing the infrastructure costs, as well as the setup time. We implement SALMA on off-the-shelf UWB devices based on the Decawave DW1000 transceiver and show that, by making use of multiple directional antennas, SALMA can also resolve ambiguities due to overlapping multipath components. An experimental evaluation in an office environment with clear line-of-sight has shown that 90% of the position estimates obtained using SALMA exhibit less than 20 cm error, with a median below 8 cm. We further study the performance of SALMA in the presence of obstructed line-of-sight conditions, moving objects and furniture, as well as in highly dynamic environments with several people moving around, showing that the system can sustain decimeter-level accuracy with a worst-case average error below 34 cm.

CCS CONCEPTS

• **Computer systems organization** → *Embedded and cyber-physical systems*; • **Networks** → Location based services;

KEYWORDS

Indoor localization, ultra-wideband, multipath, single-anchor.

1 INTRODUCTION

Localizing people and objects in a precise and accurate way is a key requirement for future location-aware Internet of Things (IoT) applications such as assisted living [54], health care [20], and robot navigation [14, 24]. As of today, achieving an accurate position estimation is still a grand challenge especially *indoors*, where global navigation satellite systems such as GPS, Galileo, GLONASS, and Beidou are not applicable due to the limited signal reception [4].

Challenges of indoor positioning. When estimating the position of a device indoors, one needs to deal with severe *attenuation, multipath, and scattering* of signals due to walls, furniture, or other surrounding objects. This is, for example, a major challenge for localization systems making use of narrowband RF technologies: solutions based on IEEE 802.15.4 [18, 29], Bluetooth [1, 3], and Wi-Fi [9, 16] are indeed highly susceptible to multipath fading, and can hardly achieve a sub-meter accuracy in these settings [38].

Indoor environments are also *highly dynamic*: moving people and objects may obstruct the line-of-sight (LOS) path between a source and a receiver. This is, for example, particularly challenging for localization systems based on optical technologies. Furthermore, the unpredictable presence of interference sources (e.g., co-located wireless devices using the same frequency band) can cause loss of information and fluctuations in the received signal strength that drastically affect the accuracy of positioning algorithms.

A practical localization system needs to achieve a high positioning accuracy despite these inherent properties of indoor environments. This task is further complicated by the fact that an ideal indoor localization system should maximize the accuracy, efficiency, and responsiveness of position estimation, while *minimizing deployment efforts and costs*. Although a plethora of distinct approaches has been proposed in the literature, none of them can yet achieve a high accuracy at minimal costs and is thus widely accepted [35].

Deployment overhead still too high. After comparing the performance of more than 100 state-of-the-art indoor localization systems under the same settings, Lymberopoulos and Liu [35] have concluded that the set-up procedure of existing solutions is excessively *time-consuming* and *labor-intensive*. As a consequence, the use of most systems is still impractical in real-world deployments.

This state of affairs represents a serious problem, because recent solutions based on Ultra-wideband (UWB) could easily achieve accuracies in the order of decimeters [24, 31]. However, one cannot fully exploit this outstanding positioning accuracy, because of the high overhead in deploying the required infrastructure. These systems employ indeed multiple anchors (e.g., at least 8 [52], 9 [23],

or 15 [24]), each of which needs to be carefully placed [19] in order to maximize the system's performance – a burden that is not sustainable on a large scale. Similarly, localization systems based on RSS profiling [16, 43] require a laborious offline data collection process before deployment to acquire the radio maps (fingerprints). Even worse, fingerprinting needs to be frequently repeated to cope with environmental changes, such as furniture setup and human motion: this makes the installation costs prohibitively high.

Minimizing the deployment effort of localization systems while still allowing to sustain a high positioning accuracy is hence a long-due fundamental step towards the creation of solutions that are viable for real-world IoT applications.

Contributions. In this paper we present SALMA, a novel UWB-based indoor localization system that can sustain a decimeter-level accuracy despite the use of only a *single* anchor. SALMA removes the need of multiple anchors by exploiting multipath propagation, i.e., specular reflections originating from static objects. The system works out of the box without any time-consuming setup phase, as it does not require any prior calibration, training, or position estimates (i.e., SALMA is based on neither fingerprinting nor other learning algorithms). All that is needed is a crude floor plan showing the geometry of the building in which the system is installed. The map includes static objects such as walls and windows only, to avoid adaptations if furniture or other objects are moved.

Starting from this floor plan and the known location and orientation of the anchor, SALMA models the theoretical multipath propagation and compares it with the estimated channel impulse response (CIR) derived by the anchor node, as shown in Sect. 3. Exploiting the position-related information embedded in the CIR allows to unambiguously determine the position of a tag using a single anchor with an accuracy comparable to the one achieved by common multi-anchor UWB systems. This way, SALMA reduces the infrastructure costs and setup time, hence addressing the omni-present trade-off between accuracy and deployment costs.

We implement SALMA on off-the-shelf UWB devices based on the popular Decawave DW1000 transceiver, building – to the best of our knowledge – the *first* low-cost single-anchor UWB-based indoor localization system. In particular, as shown in Sect. 4, we support multiple tags simultaneously and shift the burden of position estimation to the anchor node. This allows to keep the design of the mobile tag simple, so to preserve its limited battery capacity.

We specifically implement SALMA for two-dimensional settings in order to support map-based navigation and tracking applications such as locating patients in hospitals [21], assistance for visually impaired, disabled, and elderly people [7, 40, 47, 54], as well as monitoring sport events [33, 45]. The applicability to three-dimensional settings and the resulting challenges are discussed in Sect. 8.

We further show how overlapping multipath components (MPCs) may limit the performance of SALMA when using a single omni-directional antenna. To alleviate this problem, we illustrate in Sect. 5 how to improve the robustness of SALMA using *multiple directional antennas*. The latter enable the exploration of the angular information of MPCs to enhance the system's performance remarkably.

In Sect. 6, we carry out a thorough experimental evaluation[1] of the performance of SALMA in an office environment with clear LOS

conditions. Among others, our results reveal that 90% of position estimates obtained with SALMA exhibit less than 20 cm error, with a median below 8 cm. This performance was obtained with a single measurement snapshot from four directional antennas. We further simulate how the accuracy of SALMA can be improved with a higher number of antennas with narrower bandwidth.

In Sect. 7, we study the performance of SALMA in the presence of *obstructed LOS*, showing that 90% of position estimates exhibit less than 30 cm error, with a median below 15 cm. Furthermore, we evaluate SALMA in a challenging setting (stockroom) reaching a 90% error of 44.5 cm, and show that moving objects and furniture has a limited effect on the performance. We also deploy SALMA in an office crowded with tens of people moving in/out across 24-hours, and show that – even in such a highly dynamic environment – SALMA sustains a worst-case average error below 34 cm. Therefore, by exploiting the redundancy offered by multipath reflections, SALMA achieves a high accuracy even with obstructed LOS, addressing an inherent vulnerability of traditional systems.

In summary, this paper makes the following contributions:

- We present SALMA, a UWB-based indoor localization system making use of only a single anchor and requiring neither prior profiling nor calibration (Sect. 3);
- We implement SALMA on off-the-shelf UWB devices and support multiple tags simultaneously (Sect. 4);
- We increase the robustness of the system to overlapping MPCs by using multiple directional antennas (Sect. 5);
- We evaluate the performance of SALMA experimentally in different scenarios with clear LOS and show that 90% of position estimates exhibit less than 20 cm error (Sect. 6);
- We show that SALMA is resilient to obstructed LOS situations and that it sustains a high accuracy even in dynamic environments with objects and people moving around (Sect. 7).

2 SALMA: OVERVIEW

Fig. 1a shows a sketch of SALMA's design. The system consists of a single anchor (fixed infrastructure) and multiple battery-powered mobile tags (devices to be localized). The anchor is connected to and powered by a central notebook running a *localization engine* that computes the position of each tag.

Every tag initiates a double-sided two-way ranging (DS-TWR) with the anchor node, following a time division multiple access scheme. The two-way ranging process allows the anchor to estimate the distance $d_0 = \|\mathbf{p} - \mathbf{a}\|$, with $\mathbf{p}$ and $\mathbf{a}$ being the tag and the anchor location, respectively (Fig. 1b). Upon completion of the DS-TWR process, the anchor records the estimated distance $\hat{d}_0$, as well as an estimate of the channel impulse response (CIR) provided by the UWB transceiver, and forwards this info to the localization engine.

Exploiting multipath propagation. The CIR embeds information about the multipath propagation consisting of reflections from walls. Traditional UWB localization systems employ the CIR to estimate the distance d_0, which is related to the path delay τ_0 as follows: $d_0 = \tau_0 \cdot c_0$, with c_0 being the speed of light. Therefore, these systems only use the path delay τ_0, and forgo remaining multipath components (MPCs). SALMA, instead, additionally uses delays of reflected multipath components, which contain additional geometric information (cf. τ_k and d_k in Fig. 1b for $k = 1, \ldots, 4$).

---

[1]All datasets are publicly accessible under http://www.iti.tugraz.at/SALMA [12].
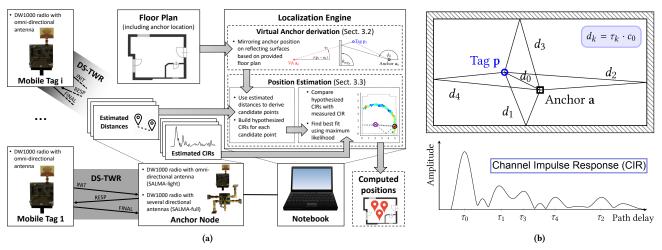
Figure 1: Overview of SALMA's design (a): the system makes use of the multipath propagation between a single anchor and a tag i. The multipath propagation is characterized by the estimated CIR containing position-related information (b).

Localization engine. Starting from a floor plan showing the geometry of the building in which the system is installed[2], and the known location of the anchor, SALMA models the theoretical multipath propagation by employing the concept of virtual anchors [39] and by building a hypothesized CIR for several candidate positions. The latter are selected on a circle of radius $\hat{d}_0$ centered in **a**, with $\hat{d}_0$ being the estimated distance derived from the DS-TWR. The localization engine then compares the hypothesized CIR of each candidate point with the one measured through the DS-TWR process, and returns the best fit using maximum likelihood estimation.

As we will show in the next sections, exploiting the position-related information encoded in the MPCs allows to unambiguously and accurately determine the position of a tag using a single anchor.

3 SALMA: DESIGN PRINCIPLES

We describe next the mathematical principles behind the functional stages of SALMA, showing how the system can leverage the information that is contained in the observed CIR to accurately narrow down the tag position. First, we present a model of the observed CIR including the multipath component (MPC) parameters in Sect. 3.1. We then explain in Sect. 3.2 how to use the known anchor position and floor plan to determine *virtual anchors* that can relate the tag position to parameters embedded in the CIR. Third, we describe in Sect. 3.3 how these parameters are used in combination with the observed CIR to obtain a position estimate. Sect. 4 then outlines how these methods are implemented on off-the-shelf hardware.

3.1 Signal model

Taking advantage of multipath propagation requires its proper modeling. In the following, we introduce the signal model relating the effective system response (i.e., the observed CIR) and the parameters of multipath components. We assume that a tag is equipped with a single omni-directional antenna, while the anchor can carry

---

[2]While furniture and other objects do affect the performance of the system, the impact stays in reasonable bounds as demonstrated in Sect. 6 and 7. Thus, there is no need to keep track of whether tables, shelves, or other furniture have moved.

M antennas. Each antenna with index $m = 1, \ldots, M$ is characterized with its beampattern $b_m(\phi)$. The observed CIR $r_m(t)$ between a single tag and the anchor's m^{th} antenna can be modeled as:

$$r_m(t) = \sum_{k=0}^{K} \alpha_k b_m(\phi_k) s_{\text{DW}}(t - \tau_k) + w_m(t). \tag{1}$$

The first term on the right-hand-side describes K specular MPCs, i.e., dominant reflections, of the transmitted signal $s_{\text{DW}}(t)$. The latter includes de-spreading and filtering at the receiver. Each MPC is characterized by its complex-valued amplitude α_k, angle of departure ϕ_k and delay τ_k. These MPCs are resulting from reflections at flat surfaces such as walls, windows or doors and will be further discussed in Sect. 3.2. The last term $w_m(t)$ denotes zero-mean white Gaussian measurement noise with variance σ_w^2. Note that the proposed signal model in (1) can model single omni-directional as well as multiple directional antenna measurements.

The signal $r_m(t)$ is sampled with frequency $f_s = 1/T_s$ and N_s samples are acquired. Hence, we use vector notation [26, 27] to compactly describe the signal model in (1) as:

$$\mathbf{r} = \mathbf{X}(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha} + \mathbf{w} \tag{2}$$

with

$$\begin{bmatrix} \mathbf{r}_1 \\ \vdots \\ \mathbf{r}_M \end{bmatrix} = \begin{bmatrix} \mathbf{X}_1(\boldsymbol{\tau}, \boldsymbol{\phi}) \\ \vdots \\ \mathbf{X}_M(\boldsymbol{\tau}, \boldsymbol{\phi}) \end{bmatrix} \boldsymbol{\alpha} + \begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_M \end{bmatrix} \tag{3}$$

and

$$\mathbf{r}_m = [r_m(0 \cdot T_s), \ldots, r_m([N_s - 1] \cdot T_s)]^{\text{T}}$$
$$\mathbf{X}_m(\boldsymbol{\tau}, \boldsymbol{\phi}) = [b_m(\phi_0)\mathbf{s}(\tau_0) \ldots b_m(\phi_K)\mathbf{s}(\tau_K)]$$
$$\mathbf{s}(\tau_k) = [s_{\text{DW}}(0 \cdot T_s - \tau_k), \ldots, s_{\text{DW}}([N_s - 1] \cdot T_s - \tau_k)]^{\text{T}}$$
$$\mathbf{w}_m = [w_m(0 \cdot T_s), \ldots, w_m([N_s - 1] \cdot T_s)]^{\text{T}}$$
$$\boldsymbol{\phi} = [\phi_0, \ldots, \phi_K]^{\text{T}}; \ \boldsymbol{\tau} = [\tau_0, \ldots, \tau_K]^{\text{T}}; \ \boldsymbol{\alpha} = [\alpha_0, \ldots, \alpha_K]^{\text{T}}.$$

Thus, the proposed signal model connects the MPC parameters $(\alpha_k, \phi_k, \tau_k)$ with the expected CIR. In Sect. 3.2, we relate these parameters to the tag position.
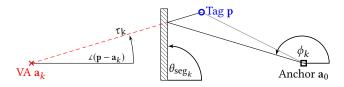
Figure 2: The concept of virtual anchors (VA) and its use in calculating the angle of departure ϕ_k and delay τ_k.

3.2 Geometric model and virtual anchors

The MPC parameters of the CIR contain position-related information regarding the tag location as well as the environment [28]. Following Euclidean geometry, simple relations can be obtained for ϕ_k and τ_k. In particular, we employ the concept of virtual anchors (VAs) [39] in order to relate MPC parameters to the tag positions (see Fig. 2). To obtain the positions of the virtual anchors $\mathbf{a}_k$ ($k > 0$), the position of the physical anchor $\mathbf{a}_0 \triangleq \mathbf{a}$ is mirrored at each reflective flat surface. Fig. 2 illustrates the top-view of a single reflection. A specular MPC (black solid) originates at the wall segment. Assignment of the specular MPC to a virtual anchor (red cross) enables an efficient calculation of the MPC parameters, delay, and angle. The delay τ_k follows as geometric distance between tag and VA, divided by the speed of light c_0, according to

$$\tau_k = \frac{1}{c_0}\|\mathbf{p} - \mathbf{a}_k\|. \quad (4)$$

We describe the angle of departure ϕ_k via the azimuth angle between tag and VA $\angle(\mathbf{p} - \mathbf{a}_k)$ according to

$$\phi_k = 2\,\theta_{\mathrm{seg}_k} - \angle(\mathbf{p} - \mathbf{a}_k). \quad (5)$$

Here, θ_{seg_k} denotes the angle of the involved reflective surface that was used to generate the VA $\mathbf{a}_k$ (see Fig. 2).

Note that, in this work, we limit the multipath propagation to *single-bounce* reflections, i.e., only a single reflective object is bounced during the path's propagation. Hence, the number of considered surfaces also determines the number of used MPCs K and VAs, e.g., for the floorplan shown in Fig. 1b we set $K = 4$ resulting in four VAs. In principle, the virtual anchor model can be extended to cover higher-order reflections as well. However, higher-order reflections are attenuated strongly, due to their increased path length and additional reflection losses. It should be also noted that, for each tag position $\mathbf{p}$, the visibility of the VAs has to be taken into account. This means that we have to check the direct path from $\mathbf{p}$ to the VA position $\mathbf{a}_k$ for intersections with any obstacles or wall segments. Only if there is a single involved intersection with the correct wall segment, we can use the k-th MPC in the signal model.

While the parameters τ_k and ϕ_k can be directly derived from the geometric model using the known VAs, a proper model for the MPC amplitudes α_k is difficult to obtain [27]. Hence, we propose to treat α_k as nuisance parameter, estimated directly from the observation $\mathbf{r}$.

3.3 Position estimation

In the following, we present a position estimator based on the CIR measurements. We aim for a maximum likelihood (ML) estimator, derived from the signal model in (2). To allow efficient computations, we assume complex-valued white Gaussian measurement noise $\mathbf{w}$.

The likelihood $p(\mathbf{r}|\mathbf{p})$ of observation $\mathbf{r}$ conditioned on tag position $\mathbf{p}$ follows as:

$$p(\mathbf{r}|\mathbf{p}) = \left(\frac{1}{\pi\sigma_w^2}\right)^{MN_s} \exp\left\{-\frac{1}{\sigma_w^2}\|\mathbf{r} - X(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha}\|^2\right\} \quad (6)$$

where $\boldsymbol{\tau}$ and $\boldsymbol{\phi}$ are related to the tag position via (4) and (5). Taking the log of (6) results in the log-likelihood function

$$\log p(\mathbf{r}|\mathbf{p}) = -MN_s \log(\pi\sigma_w^2) - \frac{1}{\sigma_w^2}\|\mathbf{r} - X(\boldsymbol{\tau}, \boldsymbol{\phi})\boldsymbol{\alpha}\|^2. \quad (7)$$

This function depends on MPC amplitudes $\boldsymbol{\alpha}$. We propose to estimate $\boldsymbol{\alpha}$ as least squares solution [27] according to

$$\hat{\boldsymbol{\alpha}} = (X^{\mathrm{H}}(\boldsymbol{\tau}, \boldsymbol{\phi})X(\boldsymbol{\tau}, \boldsymbol{\phi}))^{-1}X^{\mathrm{H}}(\boldsymbol{\tau}, \boldsymbol{\phi})\mathbf{r} \quad (8)$$

with $(\cdot)^{\mathrm{H}}$ denoting the conjugate and transposed. The position estimate $\hat{\mathbf{p}}$ maximizing the log-likelihood function can be formulated as a non-linear optimization problem:

$$\hat{\mathbf{p}} = \arg\max_{\mathbf{p}\in\mathcal{P}} \log p(\mathbf{r}|\mathbf{p}) = \arg\min_{\mathbf{p}\in\mathcal{P}} \|\mathbf{r} - X(\boldsymbol{\tau}, \boldsymbol{\phi})\hat{\boldsymbol{\alpha}}\|^2. \quad (9)$$

The parameters $\boldsymbol{\tau}$ and $\boldsymbol{\phi}$ are determined via the geometry, and these parameters in turn build the hypothesized CIR $X(\boldsymbol{\tau}, \boldsymbol{\phi})\hat{\boldsymbol{\alpha}}$, which is compared to the observed CIR $\mathbf{r}$. The position for which the hypothesis comes closest to the observation (and thus maximizes the likelihood) is chosen as the position estimate $\hat{\mathbf{p}}$. Searching for a global maximum requires to evaluate (9) at each feasible tag position $\mathcal{P}$, i.e., all positions within the communication range to the anchor. As shown in [26], this exhaustive search can be limited to potential candidate points that are located along a circle around $\mathbf{a}$ with radius $\hat{d}_0$. We consider N_C candidate points $\mathcal{P} = \{\mathbf{p}^{(j)}\}_{j=1}^{N_C}$ where each point is drawn independently with Gaussian distributed radius $d^{(j)} \sim \mathcal{N}(\hat{d}_0, \sigma_{\mathrm{DW}}^2)$ and uniformly distributed angle $\phi^{(j)} \sim \mathcal{U}(0, 2\pi)$ [26]. Candidate points lying outside of the room are discarded. These can be determined with simple line equation tests using the given floor plan. The number of candidate points N_C has a direct impact on the accuracy of the found estimate (9) and will be studied in Sect. 6.3.

4 IMPLEMENTATION ON OFF-THE-SHELF DEVICES

We implement SALMA on off-the-shelf UWB devices. After introducing the hardware in Sect. 4.1, we sketch the scheme used to derive the distance between the tag and anchor as well as the CIR $\mathbf{r}_m$ in Sect. 4.2. We then illustrate how the system can support multiple tags in Sect. 4.3 and describe the implementation of the position estimation in Sect. 4.4.

4.1 Hardware

The system consists of Decawave EVB1000 platforms used for both anchor and tags (Fig. 3a). These platforms employ the low-cost IEEE 802.15.4-compliant UWB transceiver DW1000 [5]. The tags are battery-powered and can be moved around freely. The anchor, instead, is located at a fixed position $\mathbf{a}$ and is connected to a notebook running MATLAB. The antenna at the tag is a self-made linearly polarized omni-directional dipole antenna (Fig. 3b), but any off-the-shelf omni-directional UWB antenna is suitable. At the anchor, instead, we employ either a single omni-directional antenna (Sect. 4.4) or multiple directional antennas (Sect. 5.2).
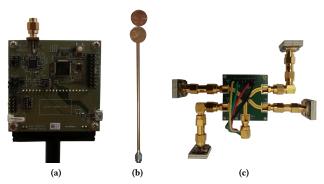
(a)   **(b)**   **(c)**

Figure 3: Decawave EVB1000 node (a) with self-made omnidirectional dipole antenna (b) and switchable directional antenna system employed in Sect. 5.2 (c).

Transmitted pulse shape. The proposed signal model in (1) requires a known transmitted pulse shape $s_{DW}(t)$. The IEEE 802.15.4-2015 standard allows the generation of an arbitrary pulse shape, as long as it fulfills certain requirements on its cross-correlation with a standard reference signal, a root raised cosine pulse with a roll-off factor of $\beta = 0.5$ [50]. Decawave follows the IEEE 802.15.4-2015 standard, but does not provide information regarding the transmitted signal of the DW1000. Therefore, we identify $s_{DW}(t)$ in a measurement campaign. We place a transmitter and receiver 1 m apart from each other in clear LOS conditions. The receiver logs 1000 CIRs. In a post-processing step, we separate the LOS from the CIR and calculate an average over these signals, which defines the transmitted pulse shape $s_{DW}(t)$ of the DW1000.

4.2 Acquiring CIR and ranging

As shown in Fig. 1a, the localization engine of SALMA requires to estimate the distance $\hat{d}_0$ between the tag and the anchor, and to derive information about the multipath propagation by acquiring the CIR provided by the DW1000.

Two-way ranging. Due to the missing synchronization between anchor and tags, we employ a double-sided two-way ranging scheme (DS-TWR) to estimate the distance $\hat{d}_0 = \|\mathbf{p} - \mathbf{a}\|$ between each tag and the anchor (see Fig. 4). The DS-TWR scheme consists of three messages, each of which contains an 11-byte MAC header embedding source and destination address, as well as a 16-bit checksum. The payload of the first message (*INIT*) and the second message (*RESP*) is 1 byte long (MSG_ID). The last message (*FINAL*) is 16 bytes long and contains the message ID as well as three 5-byte timestamps [13]. The uncalibrated distance $\hat{d}_{TWR}$ is calculated in the DS-TWR scheme with [6, pp. 213]:

$$\hat{d}_{TWR} = \frac{T_{round1} \cdot T_{round2} - T_{reply1} \cdot T_{reply2}}{T_{round1} + T_{round2} + T_{reply1} + T_{reply2}} \quad (10)$$

To calibrate the distance estimate, we perform 5000 DS-TWR trials between the anchor and a tag placed 2 m apart from each other. The derived variance and mean of the difference between the reported distance $\hat{d}_{TWR}$ and the true distance $d_0 = 2$ m is $\sigma_{DW}^2 = (0.054\,\text{m})^2$ and $\mu = 0.26$ m, respectively. Hence, the calibrated distance estimate follows as $\hat{d}_0 = \hat{d}_{TWR} - \mu$. The distance $\hat{d}_0$ and the variance σ_{DW}^2 define the distribution of the candidate points around the anchor, as shown in Sect. 3.3.
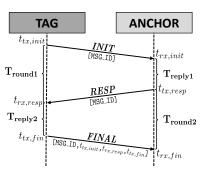

Figure 4: Double-sided two-way ranging scheme.

Acquisition of CIR. Besides deriving the distance $\hat{d}_0$ between anchor and tag, the anchor acquires the CIR r_m from the *FINAL* message received from the tag. Fig. 1b illustrates an exemplary r_m. The sampling period is set to $T_s = 1/f_s = 1/(2 \cdot 499.2\,\text{MHz}) = 1.0016$ ns. Each sample consists of a 16-bit real integer and a 16-bit imaginary integer resulting in a total size of 4048 Bytes. To reduce the amount of data read via SPI from the DW1000, we limit the length of the CIR to $N_s = 100$ samples.

Having the tags initiate the DS-TWR lets the anchor receive the required information to run the localization algorithm (i.e., the *INIT* and *FINAL* message). At the same time, it also allows to shift the burden on the anchor, which is typically static and much more powerful than tags, as it is line-powered and connected to a backbone localization engine that performs the CPU-intensive calculations. This is advantageous in real-world deployments, as tags are able to control the position update rate based on their energy budget. For example, by equipping a tag with an accelerometer, one can initiate a position update only in case of a movement, and remain in low-power mode otherwise.

4.3 Supporting multiple tags

We have so far considered only a single tag placed at an unknown position $\mathbf{p}$. SALMA can support up to N_t tags placed at positions $\mathbf{p}_i$ (with $i = 1...N_t$) by employing a slotted ALOHA scheme. The duration of a time-slot is related to the computation time necessary to obtain a position estimate (evaluated experimentally in Sect. 6.4) plus a guard interval of 1 ms at the start/end of each time-slot to overcome mis-alignments due to clock drifts.

In our current implementation, the anchor periodically broadcasts beacon messages embedding information about the time-slots' occupancy every 30 seconds[3]. Tags are not assigned to specific time-slots, but have instead the freedom to use any of the unoccupied ones: this enables a tag to use several time-slots in case it requires a higher update rate. In principle, this scheme may lead to collisions if two tags pick the same time-slot. This is, however, a well-known issue that has been largely studied in RFID systems where the reader sends a request and tags pick a random slot to answer [8, 42]. Existing anti-collision schemes can be readily applied also in SALMA. For example, the anchor can monitor the number of occupied slots and adjust their number accordingly, or adapt the slot duration by changing the number of candidate points.

---

[3]Due to the stable clock of the EVB1000 board (10 ppm), an even higher interval between beacon messages can be safely selected.

4.4 SALMA-light: Position estimation using omni-directional antennas

After the anchor has acquired the estimated distance $\hat{d}_0$ and the CIR r_m, SALMA needs to carry out the position estimation as described in Sect. 3.3. We provide a first implementation of such a position estimation by equipping the anchor node with a single omni-directional antenna: we call this implementation *SALMA-light*. When using a single antenna, only a single CIR observation is available, which greatly simplifies the signal model from (3) with $M = 1$.

Obtaining a position estimate. We use $\hat{d}_0$ to obtain candidate points, as described in Sect. 3.3. At each candidate point, only the MPC delays τ_k are calculated using (4), since the beampattern has no effect on the estimate. The amplitude estimate from (8) requires a computationally demanding matrix inversion, and, in the case of overlapping MPCs, the matrix might not even be invertible. Hence, we approximate the log-likelihood value from (9) iteratively [26]:

$$\text{init} : r^{(0)} = r$$

$$\text{for } k = 1 \ldots K : \alpha_k = s^{\mathrm{H}}(\tau_k) r^{(k-1)} \tag{11}$$

$$r^{(k)} = r^{(k-1)} - \alpha_k s(\tau_k) \tag{12}$$

Essentially, we take the observed CIR r and sequentially subtract sub-hypotheses ($\alpha_k s(\tau_k)$) by using pulses shifted to the respective τ_k and weighted by single amplitude estimates α_k. The resulting $r^{(K)}$ is then the left-over 'residual' signal. The latter represents how similar the hypothesized and measured CIRs are, and is thus used as an approximation of the log-likelihood. This procedure is repeated for each candidate point and the one with highest log-likelihood value is chosen to be the tag position estimate $\hat{p}$.

Limitation: multipath ambiguities. While this method is simple, the non-accessible beampattern restricts the algorithm to delay information only. This restriction makes the algorithm sensitive to overlapping MPCs, as well as to ambiguities in the delay times of MPCs, which may degrade the positioning performance significantly, as shown in Sect. 6.2.

5 TACKLING MULTIPATH AMBIGUITIES

As discussed in Sect. 4.4, SALMA-light uses measurements from a single antenna only, which makes the algorithm sensitive to overlapping MPCs and ambiguities. In this section, we introduce *SALMA-full*: an enhanced version of the system in which the anchor makes use of multiple switchable directional antennas[4]. Hence, we may now take advantage of the full signal model from (3), where each antenna m is characterized by its beampattern $b_m(\phi)$ covering one sector of the azimuth plane.

The combined observations of the antennas enable the system to separate closely-arriving MPCs in the spatial domain. However, the combination of several antenna measurements requires phase-coherency between the measurements, which is not given by low-cost transceivers. In the following, we tackle the phase-coherency issue (Sect. 5.1), describe how to carry out position estimation using directional antennas (Sect. 5.2), and highlight the key differences in the employed hardware compared to SALMA-light (Sect. 5.3).

___
[4]This system was showcased at SenSys'17 [11].

5.1 Non-phase-coherent amplitude estimates

Phase-coherency demands accurate radio clocks, which are not provided by off-the-shelf UWB transceivers like the DW1000. In our case, this affects the implementation of the presented amplitude estimates in (8). Inaccurate clocks between consecutive measurements are perceived as a phase change in the baseband-equivalent CIR. Thus, amplitude estimates from consecutive measurements differ in their complex-valued phase $\alpha_{k,m} \approx e^{j\varphi} \alpha_{k,m'}$, where φ denotes the unknown phase offset. However, the unkown phase offset φ is required for the position estimate in (9).

To overcome the necessity of phase coherency, we follow the approach presented in [27]. Assuming non-overlapping MPCs $(s(\tau_k)^{\mathrm{H}} s(\tau_k) \approx 0)$, an MPC amplitude α_k can be estimated independently as projection of the normalized signal $\frac{s^{\mathrm{H}}(\tau_k)}{s^{\mathrm{H}}(\tau_k) s(\tau_k)}$ onto the m-th measurement r_m according to

$$\alpha_{k,m} = \frac{s^{\mathrm{H}}(\tau_k) r_m}{s^{\mathrm{H}}(\tau_k) s(\tau_k)}. \tag{13}$$

Furthermore, the amplitude estimate in (8) can be written as complex-valued average. Relaxing the complex-valued weighted average by an absolute-valued average [41] results in an estimate of the k-th MPC amplitude α_k^{avg} according to

$$\alpha_k^{\mathrm{avg}} = \frac{\sum\limits_{m=1}^{M} |\alpha_{k,m}| \cdot |b_m(\phi_k)|^2}{\sum\limits_{m'=1}^{M} |b_{m'}(\phi_k)|^2}. \tag{14}$$

The remaining phase $\angle \alpha_{k,m}$ is extracted from the individual antenna measurements and the amplitude estimate $\hat{\alpha}_{k,m}^{\mathrm{avg}}$ of the m-th antenna and k-th MPC results in

$$\hat{\alpha}_{k,m}^{\mathrm{avg}} = \alpha_k^{\mathrm{avg}} \exp(j \angle \alpha_{k,m}). \tag{15}$$

This approximation combines MPC amplitudes from non-phase-coherent measurements, taking into account the directivity of the M antennas.

5.2 SALMA-full: position estimation using directional antennas

In contrast to SALMA-light, SALMA-full collects observations from M directional antennas. The antennas are physically separated (see Fig. 3c), which results in different range estimates from tag to each antenna. Since this difference is smaller than the standard deviation of the DW1000 ranging, this error can be neglected. However, to create the candidate points as described in Sect. 3.3, we use the mean value of all ranges.

For each candidate point, the MPC parameters τ and ϕ are calculated using the VA positions in (4) and (5), respectively. For the amplitude estimates, the same iterative approach is followed as in Sect. 4.4, but it is adapted to use the stacked observed CIRs $r = [r_1^{\mathrm{T}}, \ldots, r_M^{\mathrm{T}}]^{\mathrm{T}}$ and to take the non-coherent amplitude estimates from the previous section into account. For this, in the iteration step (11), we use α_k^{avg} from (14), and for step (12), we use

$$r_m^{(k)} = r_m^{(k-1)} - b_m(\phi_k) \hat{\alpha}_{k,m}^{\mathrm{avg}} s(\tau_k).$$

This gives us a (stacked) residual $r^{(K)} = \left[(r_1^{(K)})^{\mathrm{T}}, \ldots, (r_M^{(K)})^{\mathrm{T}}\right]^{\mathrm{T}}$ representing the similarity between the hypothesized and measured CIR that is used as an approximation for the log-likelihood. The final estimate $\hat{p}$ is obtained by evaluating the log-likelihood for each candidate point and by picking the one achieving the maximum.

Two aspects are worth of note regarding the beampatterns $b_m(\phi)$ of the antennas: first, we use 36 sampled values $b_m(i \cdot \Phi_s)$ with a spacing of $\Phi_s = 10°$ obtained from a measurement campaign. Second, when the anchor is employed, it can be oriented with $\Phi_o = j \cdot 10°$, where j might be chosen as desired. For the implementation, this value has to be known. SALMA then uses $b_m\left(\lfloor \frac{\phi_k + \Phi_o + 5}{10} \rfloor \bmod 36\right)$ to approximate $b_m(\phi_k)$, also taking the orientation into account.

5.3 Hardware differences

In contrast to SALMA-light, in SALMA-full we exploit four self-made and low-cost directional antennas with a half-power beamwidth of about $150°$ (see Fig. 3c). The antennas are mounted such that each one points in a different cardinal direction. The evaluation in Sect. 6.2 shows that, even with this wide beamwidth, SALMA achieves an error below 20 cm for 90% of the estimated positions. The higher number of antennas increases the acquisition time of CIRs and distance estimates. Hence, the acquisition duration is higher than that of SALMA-light, as discussed in Sect. 6.4.

6 EVALUATION

We evaluate the positioning capabilities of SALMA in challenging indoor environments: an office (Room A, see Fig. 5a and 6a), and a stockroom (Room B, see Fig. 5b and 6b). After describing the experimental setup in Sect. 6.1, we answer the following questions:

- What is the benefit of using SALMA-full over SALMA-light? (Sect. 6.2);
- What is the *accuracy* achieved by SALMA, and at which computational costs? (Sect. 6.3);
- How long does it take to estimate a position, and what are the implications on *scalability*? (Sect. 6.4);
- Do more (and better) antennas improve the performance of SALMA? (Sect. 6.5).

We answer all these questions in Room A under clear LOS conditions. In Sect. 7, we will then specifically evaluate how SALMA performs in more challenging environments with obstructed LOS (both rooms), and a dynamic environment due to moving objects and people (Room A).

6.1 Experimental setup

We carry out the evaluation in an office containing obstacles and scattering objects such as desks, chairs, shelves, and PC monitors, as shown in Fig. 5a and 6a (Room A). We place the tag in $N_{\mathrm{EP}} = 35$ evenly distributed evaluation points, while fixing the anchor next to the table. We mount both anchor and tag on a tripod at a height of 1.50 m, i.e., well above the obstacles, so to have clear LOS conditions. The anchor is connected to a Lenovo ThinkPad T450s notebook running MATLAB. The tag, instead, is battery-powered and can move freely. The only pre-processing required by SALMA is to enter the anchor location and orientation as well as the coordinates of the surrounding four wall segments. The following settings are used
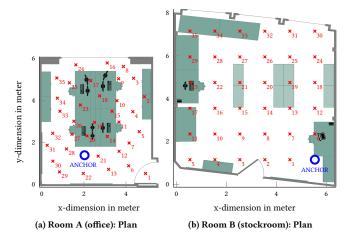

(a) Room A (office): Plan   (b) Room B (stockroom): Plan

Figure 5: Evaluation setup (2D-plan): we consider 35 evaluation points (red crosses) in two different environments.

(a) Room A (office): Picture

(b) Room B (stockroom): Picture

Figure 6: Evaluation setup (Picture): the white dashed line marks the measurement height under obstructed LOS.

by the DW1000: maximum data rate (6.8 Mbps), pulse repetition frequency of 64 MHz, and a preamble symbol repetition of 1024. Channel 7 is used due to its high bandwidth (900 MHz) and since our self-made directional antennas are optimized for this band. At each evaluation point (denoted by p_{EP}), we perform 100 position estimates, hence carrying out 3500 evaluations in total for both SALMA-light and SALMA-full. We denote the i-th position estimate by $\hat{p}_i$, and obtain the absolute position error with

$$\mathrm{Err}_i = \|\hat{p}_i - p_{\mathrm{EP}}\|. \qquad (16)$$

Statistically, we look at the cumulative distribution function (CDF) over the errors Err_i using all evaluation points (i.e., 3500 estimates) for different configurations, as illustrated in the upcoming sections.

6.2 SALMA-light vs. SALMA-full

In this section we examine the performance of both SALMA implementations comparatively.

Handling ambiguities. SALMA-light relies solely on the position information contained in the arrival times of MPCs. Hence, the
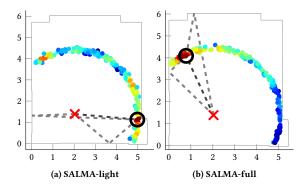
(a) SALMA-light   (b) SALMA-full

Figure 7: Evidence of multipath ambiguities.

resulting likelihood for the positions is highly multimodal, or in other words, there are multiple regions that seem to best fit the observed signal. This is demonstrated in Fig. 7, which shows the positioning result using SALMA-light (Fig. 7a) and SALMA-full (Fig. 7b) for one estimation run on position 34 (cf. Fig. 5a). The colored dots indicate the candidate point positions, where the color represents the likelihood values (red=high, blue=low). SALMA-light has three regions showing similarly high likelihood values (red and orange dots), caused by similarly long MPC paths, which results in a completely wrong estimate. In contrast, SALMA-full can narrow the estimate down to find the true position of the tag. This is possible due to the combined directional antenna observations, where wrong candidate points have low likelihood values because amplitude values do not fit to the antenna patterns.

Quantitative comparison. We show the improvement quantitatively by accounting for all 3500 estimates via the CDF of the absolute position error. Fig. 8 shows the CDF for SALMA-light (dashed blue line) and SALMA-full (solid orange line). With SALMA-light, 67.3% of all evaluations have a position error below 25 cm. On the other hand, 21.7% of the evaluations have an error above 1 m: these outliers are caused by the multipath ambiguities, as just explained. By using directional antennas, SALMA-full can mitigate these outliers: 90% of all evaluations are below 20.17 cm, whilst 99% of the evaluations are below 29.72 cm.

We can hence conclude that SALMA-full clearly outperforms SALMA-light thanks to the additional angular information. Hence, we focus the next evaluations on SALMA-full only.

6.3 Localization accuracy

With the promising results shown in Sect. 6.2, we investigate the accuracy of SALMA in more detail, and focus also on the computational costs.

Role of candidate points. We examine the impact of the number of candidate points used by SALMA-full. To this end, we perform 3500 estimates for different number of candidate points $N_C \in \{50, 100, 200, 500, 1000\}$. Fig. 9 shows the resulting CDF: even when using only 50 candidate points (blue, star), 90% of the estimates have an error below 30 cm. However, there are outliers for about 5% of the estimates. Increasing the number of candidate points removes the outliers and improves the performance to a "saturation point" at about 200 candidate points (i.e., a higher number of points gives negligible improvements). Hence, we make N_C=200
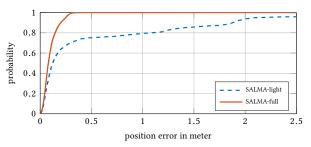

Figure 8: SALMA-light sustains an accuracy below 30 cm only in 70% of the cases due to multipath ambiguities. By exploiting the angular domain, SALMA-full exhibits an error below 30 cm in 99% of the cases (Room A).
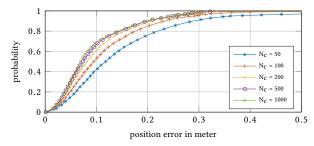

Figure 9: Impact of the number of candidate points on SALMA's accuracy: N_C=200 acts a good trade-off.

our preferred setup and use it for all further evaluations, unless stated otherwise. The number of candidate points increases the computation time linearly, so N_C can act as a trade-off between computational costs and accuracy, as illustrated in Sect. 6.4.

Individual evaluation points. A more detailed display of the accuracy is shown in Fig. 10a. For the 100 estimates at each position of Room A, the mean (blue circle) and the 3-fold standard deviation (black error ellipse) are shown. The former indicates that there is little estimation bias (distance to ground truth). As for the standard deviation, with the good ranging precision of UWB, the ranging deviation is small (facing the LOS), while info gained by the MPCs determines the angle deviation (perpendicular to the LOS). Overall, the good performance is reinforced, while there are certain positions (e.g., 9, 16, and 18) with a slightly higher bias. Of special note is that the accuracy of SALMA does not degrade at higher ranges: pos. 8, 16, and 24 are placed more than 4 meters apart from the anchor, but their estimates are as accurate as the ones obtained at positions much closer to the anchor. This is highlighted in Fig. 11 showing the average position error with respect to the real distance between anchor and tag. This stands in contrast to many other indoor positioning techniques, where the inaccuracy increases quickly with the range, e.g., visual systems [34].

6.4 Scalability

The number of supported tags by SALMA is limited by (i) the computation time of the position estimation, (ii) the duration of the DS-TWR, as well as (iii) the time needed to stream the CIR and additional info to the notebook via USB. The duration of a DS-TWR is mainly defined by the packet length of its three packets, which is 3.49 ms. Streaming one CIR to MATLAB takes 4.62 ms.
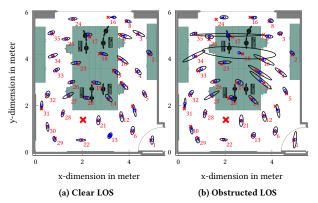
(a) Clear LOS **(b) Obstructed LOS**

Figure 10: Error ellipses showing position bias and three-fold standard deviation (Room A).

The computation time of the position estimation depends on the notebook's performance and on the number of candidate points. In the evaluations, the algorithms are performed in MATLAB on a Lenovo ThinkPad T450s with 2.59 GHz clock and 8 GB RAM. An increase in the number of candidate points scales the computation time linearly. Thus, we evaluate the time needed per candidate point. SALMA-light takes 174.77±12.2 μs and SALMA-full requires 955.13±23.5 μs per candidate point. Thus, with $N_C = 200$ candidate points, the algorithms take 34.95 ms and 191.03 ms, respectively. The overall duration of a position estimation with SALMA-light when using $N_C = 200$ is hence 43.06 ms, resulting in an update rate of 23 Hz. Using $N_C = 50$, instead, gives an update rate of 60 Hz.

When using SALMA-full, for each antenna a DS-TWR trial is performed and four CIRs are acquired: this reduces the achievable update rate. In total, SALMA-full requires 223.5 ms for $N_C = 200$ and 79.8 ms for $N_C = 50$, resulting in an update rate of about 4.5 Hz and 12.5 Hz, respectively. Thus, SALMA can easily compete with comparable solutions and outdoor positioning systems like GPS.

6.5 The role of the antenna

We examine next how SALMA would perform when using more antennas with more directive beampatterns. As such antennas are not yet commercially available, we simulate artificial CIRs (r) and ranging ($\hat{d}_0$) for the same tag positions shown in Fig. 5a. For any tag position **p**, we can create an artificial CIR in two steps:

(1) Specular part: we shift and add pulses $s_{\text{DW}}(t - \tau_k)$ weighted by beampattern $b(\phi_k)$ and amplitude α_k using the known delays τ_k and angles ϕ_k. The amplitude exhibits free-space pathloss and each reflection halves the magnitude. We consider MPCs up to order two.

(2) Scattering part: we simulate diffuse multipath by drawing realizations of a Gaussian random process whose variance is defined by a double exponential power delay profile according to equation (9) from [22]. Additionally we simulate AWGN measurement noise with an SNR of 29.5 dB at 1 m.

This simulation setup allows us to adjust the half-power beamwidth (HPBW) of the antennas and to recreate the effect of clutter by setting a signal-to-interference[5] ratio (SIR). The latter is defined by the ratio between LOS and scattering energy. We determined empirically that an SIR of 3 dB properly describes the environment.
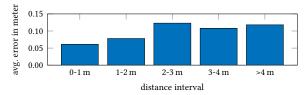
[5]Interference, in this case, refers to self-interference due the scattering part.

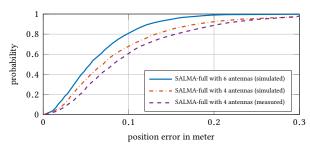**Figure 11: Average position error w.r.t. the real distance.**

Figure 12: Simulated performance of SALMA when using multiple antennas with more directive beampatterns.
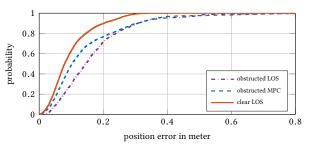

Figure 13: Performance of SALMA in clear LOS, obstructed MPC, and obstructed LOS situations in Room A.

To recreate the performance behavior from the SALMA-full measurement runs described in Sect. 6.3, we set the HPBW to 150°, matching the properties of the used antennas. Additionally, we carry out simulations using six antennas with a HPBW of 90°, reflecting a higher quality implementation. Fig. 12 shows the results. On the one hand, we can see that the SALMA-full simulation (red curve) fits the measured results (purple curve) closely. A slightly better performance is achieved in the simulation, because the impact of bias due to floor plan inaccuracies is not present. On the other hand, we can notice that, when simulating six antennas (blue curve), the performance of SALMA improves significantly: the 90% error decreases by about 10 cm to almost reach the sub-decimeter mark, whilst 99% of the estimates achieve an error below 20 cm.

7 ROBUSTNESS TO NON-LINE-OF-SIGHT AND DYNAMIC ENVIRONMENTS

Indoor environments are inherently highly dynamic due to moving humans and objects. Thus, the value of a localization system strongly depends on (i) its performance under obstructed LOS, (ii) its behavior in different environments and (iii) its robustness in crowded settings. In this section, we discuss the performance of SALMA under non-line-of-sight (NLOS) conditions (Sect. 7.1), when furniture is moved without updating the map (Sect. 7.2) and in the case of a highly-dynamic and crowded environment (Sect. 7.3).

7.1 Performance under NLOS conditions

In situations of a blocked LOS, range-based systems suffer from a positive bias [36]. This is either caused by the lower propagation speed in case the signal propagates through the obstacle, or, in case of a fully blocked LOS, due to the misinterpretation of a reflection as the direct path.

Distance bias. SALMA requires the distance estimate $\hat{d}_0$ between anchor and tag to distribute the candidate points on a circle around the anchor as described in Sect. 3.3. Thus, an obstructed LOS causes an increase in the radius of the circle. First, we analyze the impact of different objects on $\hat{d}_0$. We place tag and anchor 2 m apart from each other and perform 1000 DS-TWR trials with different objects blocking the LOS. The objects included: a metal plate (800×450×3 mm), PC monitors, and humans. The threshold-based mechanism of the DW1000 was able to detect a leading edge corresponding to the LOS in each of the trials. But, indeed, the obstructed LOS leads to a positive bias in the range estimate. Metal plate and PC monitor caused a range bias of 11 cm and 13.1 cm, respectively. Even in the case of humans blocking the LOS, the leading edge was successfully detected, but two humans blocking the LOS already led to a range bias of 41.8 cm. Further evaluations will show that SALMA is robust even in the case of a range bias.

Accuracy evaluation. SALMA is not just making use of the LOS component, but also of specular MPCs. Thus, we evaluate SALMA also in situations of blocked MPCs. We repeat the evaluation in Room A described in Sect. 6.1: this time, however, we mount the tag and the anchor at a height of 1.20 m, corresponding to the height of monitors, shelves, and people in the room. Depending on the position of the evaluation points, this results in obstructed LOS for twelve of these points $EP_{OLOS} = \{3, 8, 9, 10, 15, 16, 18, 19, 25, 26, 33, 34\}$, which results in 1200 evaluations. Note that, for all these evaluation points, there were also specular MPCs blocked by objects. For twenty points, instead, the LOS was still clear but specular MPCs were blocked $EP_{OMPC} = \{1, 2, 4 - 7, 11 - 14, 17, 20 - 24, 27, 28, 32, 35\}$. In total, 2000 evaluations were acquired in these situations. The remaining 300 evaluations are still in clear LOS with no blocked MPCs, thus, they are ignored. Fig. 13 (magenta dash dotted line) shows the CDF of all evaluations under blocked LOS (EP_{OLOS}). The median is at 14.5 cm and the error for 90% of the estimates is still below 30.7 cm. The blue dashed line in Fig. 13 shows the CDF just considering evaluations where significant multipath components are blocked by obstacles or humans (EP_{OMPC}). The median is at 10.25 cm and the error for 90% of the estimates is below 30.52 cm. This shows that SALMA remains robust even in the case of blocked MPCs.

Qualitative evaluation. In Fig. 10b, we see again the accuracy for individual evaluation points, now for the obstructed LOS case. While the position bias (distance blue circles to red crosses) did not increase significantly, we can see that the variance in the angular direction increases for most of the evaluation points. The error ellipses shown in Fig. 10b indicate two evaluation points with significantly higher variances (no. 17 and 18) as the other points. The reason is the unfortunate position of the anchors in this case. The PC monitors and obstacles at the left and right wall block the respective MPCs, thus, the position information obtained at these positions comes only from the LOS and the reflection from the
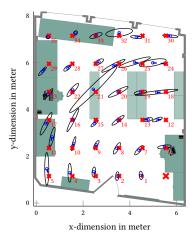

Figure 14: Error ellipses showing position bias and threefold standard deviation with empty racks in Room B.
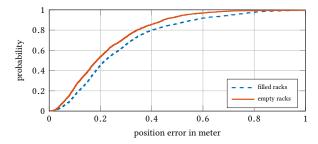

Figure 15: Performance of SALMA in Room B in the case of moving obstacles (storage racks are empty or filled).

window. Since these two reflections are arriving at the same angle, SALMA suffers from a poor geometric configuration. This results in ambiguities similar to the ones shown in Sect. 6.2. Due to significant MPCs from the left and right wall, this situation was not evident in the clear LOS case (see Fig. 10a).

7.2 Performance in stockroom with moving obstacles

In Sect. 6 and 7.1, we have performed all the measurements in Room A. To prove the capabilities of SALMA also in more challenging environments and in the presence of moving obstacles, we have evaluated its performance also in Room B (see Fig. 5b and 6b).

Performance in more challenging environments. To challenge SALMA, we chose a stockroom that is larger than Room A (46.7 m^2 vs. 31.6 m^2) and cluttered with desks, storage racks (bright rectangles in Fig. 5b) and several other metal objects (see Fig. 6b). We have mounted anchor and tag at a height of 1.20 m. Fig. 15 (solid orange line) shows the CDF of all evaluations in Room B. The median is at 18.6 cm and 90% of all estimates obtain an error below 44.5 cm. Thus, compared to the evaluation in Room A, the performance of SALMA is slightly worse due to the larger room with more clutter and wall materials with unfavorable reflective properties (see Sect. 8). Fig. 14 shows the accuracy for individual evaluation points. Similar to Room A (see Sect. 7.1), some positions
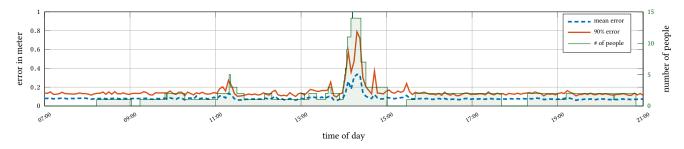
Figure 16: Snippet of a 24-hours experiment in dynamic environments. The dashed blue line depicts the mean error of 50 position estimates over time, whilst the solid orange line shows the 90% error. Despite the people moving in/out of the room (green line), SALMA can sustain a decimeter-level position accuracy.

(e.g., 4, 5, 12, 18, 24) suffer from an unfavorable anchor placement as the LOS is arriving from the same angle as strong reflections.

Moving obstacles. The performance of localization systems based on RSS profiling and fingerprinting is highly affected by moving obstacles. Thus, changing the furniture in a room often requires to update or repeat measurements. To evaluate the performance of SALMA in the case of moving obstacles or furniture, we have stocked up the storage racks in Room B with full beer crates and other objects (see Fig. 6b). As the goal of SALMA is to minimize the setup effort, we do not model reflections from obstacles such as the full storage racks. Fig. 15 shows that the position error (dashed blue line), while higher due to the range bias introduced by obstructed LOS, still stays in reasonable bounds, relatively unaffected by the additional reflections. Thus, SALMA only slightly loses accuracy in favor of practicability and setup time.

7.3 Performance in a crowded environment

For a final stress test of SALMA and to evaluate its behavior in a dynamic environment including NLOS situations, we employed our system again in the office scenario (Room A) for a non-stop 24 hours run. The system was exposed to the usual ongoing work flow that involves multiple people passing by the system, thereby blocking the LOS or MPCs, hence creating a dynamic environment. During the 24 h experiment, SALMA localized three tags at representative positions (positions 3, 15, and 33) simultaneously. The positions were deliberately chosen to be under obstructed LOS. Every five seconds we estimated the tag positions resulting in 51840 position estimates. We evaluate the performance of the system on multiple levels. Fig. 16 shows the mean error (dashed blue line) and the 90% error (solid orange line) over 50 position updates from 07:00 - 21:00 o'clock. Additionally, we track the number of present people in the room during the experiment (green stairstep graph). The figure focuses on daytime, since over night no one was in the room and the performance remained constant. It can be seen that the usual working environment (with the two designated working people present) does not impair the performance of the system providing an average error below 11.2 cm. In terms of present people, there are two events prominent in Fig. 16: at 11:00 o'clock there was a meeting with five people and at 14:00 o'clock we have presented SALMA to thirteen people making it in total fifteen people in the room simultaneously. We asked people to move around the room freely during the presentation, thus, the LOS and the MPCs were

obstructed in a dynamic fashion. Even though the error increases during these periods, still, when the room was completely filled with people, the average error was below 34 cm and the 90% error below 79 cm. The latter indicates that SALMA is robust also in a highly dynamic environment and under NLOS conditions.

Comparison to other multi-anchor systems. Comparing the accuracy of SALMA with other UWB-based systems is difficult, as they are either evaluated in mobile 2D [15, 31, 52] or static 3D [23, 24] scenarios. Silva et al. [49] report a 2D static LOS mean error of 16.6 cm. SALMA instead achieves an average error of just 9.85 cm in clear LOS. Kempke et al. [23, 24] report a 90% error of 77 cm and 50 cm in static 3D, respectively. In contrast, SALMA achieves a 90% error of 50 cm between 14:00-14:30 o'clock, thus, under obstructed LOS and when up to fifteen people were walking around. Therefore, it is fair to say that SALMA can compete and even outperform existing systems, despite using just a single anchor.

8 DISCUSSION

Our evaluation demonstrates the capabilities of SALMA to perform accurate positioning in typical indoor environments. However, it has also highlighted a number of challenges and open questions that we will elaborate in detail in this section.

Sensitivity to chosen anchor position. Due to SALMA's principle, just one anchor per room is required. In our evaluations, we examined two typical choices for anchor positions, namely, in the vicinity of the room center (Room A) and in the corner of the room (Room B). Both variants have pros and cons: in Room A we have a full candidate point circle for many ranges, which increases the risk of ambiguities, especially for SALMA-light. However, SALMA-full can take full advantage of the beampatterns in all directions to stay relatively unaffected (as we have demonstrated in Sec. 6.2). In Room B, we have, at most, a quarter circle of candidate points: on the one hand, this reduces possible ambiguities. However, on the other hand, this results in higher ranges with reduced signal strength and reduced benefit from the angular information. We also pointed out some difficult positions in both rooms where LOS and the strongest reflection come from the same direction, resulting in a plateau in the likelihood which leads to a dilution of precision. These cases exist no matter what anchor position is chosen.

The anchor orientation can be set arbitrarily, but it has to be fixed and known to correctly weight the amplitudes.

Including the third dimension. SALMA is designed specifically to perform 2D positioning. This choice is rather pragmatic: physical and algorithmic setup of the system are simplified dramatically, enabling a practical implementation with short setup time and efforts, while only using a single anchor. Also, many applications (e.g., navigation tasks) do not require any height information. In principle, the methods can be extended to the third dimension: (i) determining the VAs can be done by mirroring at plane surfaces, (ii) for the MPC angle one needs to take the elevation beampatterns into account and (iii) the candidate points are placed on a range sphere rather than circle. However, this drastically increases the computational complexity and makes the position likelihood even more multimodal. A 3D model of the environment could help to avoid ambiguities due to floor or ceiling reflections, however, our antennas exhibit a fairly narrow elevation pattern, hence, the impact of ceiling, floor, and other reflections is limited significantly.

Effect of wall materials. The main setup effort for SALMA is the determination of reflecting surfaces in the considered environment. However, additional care has to be taken with regard to the material of the surfaces. Preferably, materials such as glass and metal enable good reflectors and including them in the models enhances the position estimate. On the other hand, plaster boards or wooden surfaces, even if they are flat and smooth, give little to no contribution in terms of specular reflections and can in fact decrease the performance. For example, in Room B, the eastern wall, even though close to the anchor, is made out of plasterboard and does not contribute with a specular multipath component. Thus, it should not be included in the signal model as a source of a virtual anchor.

9 RELATED WORK

Indoor localization technologies. Many RF technologies have been investigated for indoor localization, such as Wi-Fi [9], Bluetooth [1, 3], and IEEE 802.15.4 [29, 44]. However, these systems hardly achieve an accuracy below 1 m, require a high amount of reference nodes, and typically come with a high deployment effort. Optical systems are among the most accurate indoor localization systems, but cannot inherently operate in NLOS conditions [34]. SALMA, instead, reaches a median error of 15 cm and a 90% error of 30 cm even in obstructed LOS conditions. Acoustic systems can also achieve decimeter-level accuracy, but their biggest enemy – multipath propagation – is SALMA's best friend [30, 37].

UWB indoor localization systems. UWB-based systems can also achieve decimeter-level accuracy [35, 57]. Recently, several systems have been implemented using low-cost UWB transceivers [15, 23, 24, 31, 49, 52]. However, these systems require a high amount of anchors, typically between eight [31, 52] and fifteen [24]. SALMA, instead, uses a single anchor and – to the best of our knowledge – no comparable solution exists. In terms of accuracy, as discussed in Sect. 7.3, it is fair to say that SALMA can compete and even outperform existing systems, despite using just a single anchor.

Multipath-assisted localization systems. Theoretical works have discussed the performance bounds of multipath-assisted indoor localization via simulation [10, 17, 55] and using very expensive, bulky and wired-synchronized equipment [25, 32, 39]. Instead,

with SALMA, we are the first to enable the exploitation of multipath reflections for low-cost, low-power wireless localization systems.

Directional antennas to enable single-anchor systems. Several works have exploited electronically steerable or switchable antenna systems to enable single-anchor localization using narrow-band technologies [2, 48]. However, SALMA outperforms all these systems due to the exploitation of the position-related information provided by the MPCs. Sun et al. [51] presented a UWB-based system claimed to achieve decimeter-level accuracy. However, their measurement setup is vague and a thorough analysis of the system performance is missing. Quing et al. [46] and Zhang et al. [56] presented similar systems, but solely based on simulation.

Although not exploiting directional antennas, also Chronos [53] requires just a single access point to estimate the position of another device. In particular, Chronos uses an omni-directional antenna array and emulates a wideband radio on commodity Wi-Fi systems. Still, due to the position-related information provided by the MPCs, SALMA outperforms Chronos in terms of accuracy. Furthermore, by using the license-free ISM bands, Chronos interferes and is prone to the interference of other devices using the 2.4 GHz band.

10 CONCLUSIONS AND FUTURE WORK

In this paper, we present SALMA, a low-cost UWB-based indoor localization system that exploits multipath reflections to tear down the position estimation to a unique solution while only using a single anchor. Besides a crude floor plan and the position of the anchor, the system does not need any prior calibration or training phase. By using directional antennas, we increased the robustness of SALMA against overlapping MPCs. We extensively evaluated the performance of SALMA under LOS and NLOS conditions, as well as during a 24 h stress-test to challenge SALMA in dynamic settings. Under LOS, SALMA achieved a median error below 8 cm and an error below 20 cm for 90% of the position estimates. Even under obstructed LOS and in a highly dynamic environment SALMA sustains a high accuracy.

Our aim in this paper was to show the outstanding capabilities of SALMA without using a tracking filter and solely utilizing single-shot single-anchor measurements. In future work, we will combine SALMA with a particle filter and an inertial measurement unit to benefit from past position estimates. Moreover, we will perform an exhaustive evaluation of SALMA in mobile environments.

ACKNOWLEDGMENTS

REFERENCES

[1] M.S. Bargh and R. de Groote. 2008. Indoor Localization based on Response Rate of Bluetooth Inquiries. In *Proc. of the 1^{st} ACM Int. Workshop on Mobile Entity Localization and Tracking in GPS-less Environments (MELT)*.

[2] A. Cidronali, S. Maddio, G. Giorgetti, and G. Manes. 2010. Analysis and Performance of a Smart Antenna for 2.45-GHz Single-Anchor Indoor Positioning. *IEEE Transactions on Microwave Theory and Techniques* 58, 1 (Jan. 2010).

[3] G. Conte et al. 2014. BlueSentinel: a First Approach using iBeacon for an Energy Efficient Occupancy Detection System. In *Proc. of the 1^{st} ACM Conf. on Embedded Systems for Energy-Efficient Buildings (BuildSys)*.

[4] D. Dardari, P. Closas, and P. M. Djurić. 2015. Indoor Tracking: Theory, Methods, and Technologies. *IEEE Transactions on Vehicular Technology* 64, 4 (April 2015).

[5] Decawave Ltd. 2016. *DW1000 Datasheet. Version 2.12*.

[6] Decawave Ltd. 2016. *DW1000 User Manual. Version 2.10*.

[7] D. Espes, A. Daher, Y. Autret, E. Radoi, and P. Le Parc. 2013. Ultra-Wideband Positioning for Assistance Robots for Elderly. In *Proc. of the 10^{th} International Conference on Signal Processing, Pattern Recognition and Applications (SPPRA)*.

[8] K. Finkenzeller. 2008. *RFID Handbuch – Grundlagen und praktische Anwendungen von Transpondern, kontaktlosen Chipkarten und NFC*. Hanser.

[9] S. Gansemer et al. 2010. RSSI-based Euclidean Distance Algorithm for Indoor Positioning Adapted for the use in Dynamically Changing WLAN Environments and Multi-level Buildings. In *Proc. of the 2010 IPIN Conference*.

[10] C. Gentner, T. Jost, W. Wang, S. Zhang, A. Dammann, and U. Fiebig. 2016. Multipath Assisted Positioning with Simultaneous Localization and Mapping. *IEEE Transactions on Wireless Communications* 15, 9 (Sept 2016), 6104–6117.

[11] B. Großwindhager et al. 2017. Demo Abstract: UWB-based Single-anchor Low-cost Indoor Localization System. In *Proc. of the 15^{th} ACM Int. Conf. on Embedded Networked Sensor Systems (SenSys), demo session*.

[12] B. Großwindhager et al. 2018. Dataset: Single-Anchor Indoor Localization with Decawave DW1000 and Directional Antennas. In *Proc. of the 1st Workshop on Data Acquisition To Analysis (DATA'18)*.

[13] B. Großwindhager et al. 2018. Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications. In *Proc. of the 19th IEEE Intern. Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*.

[14] K. Guo, Z. Qiu, C. Miao, A. H. Zaini, C.-L. Chen, W. Meng, and L. Xie. 2016. Ultra-Wideband-Based Localization for Quadcopter Navigation. *Unmanned Systems Journal* 4, 1 (Jan. 2016).

[15] F. Hartmann et al. 2015. Design of an Embedded UWB Hardware Platform for Navigation in GPS Denied Environments. In *Proc. of the IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT)*.

[16] S. He and S-H G. Chan. 2016. Wi-Fi Fingerprint-based Indoor Positioning: Recent Advances and Comparisons. *IEEE Comm. Surveys & Tutorials* 18, 1 (Aug. 2016).

[17] B. Hu, Z. Shi, and Y. Wang. 2018. Single-Sensor Based Indoor Localisation by Exploiting Multipath Propagation. *Electronics Letters* 54, 3 (Feb. 2018), 179–181.

[18] X. Huang, W. Zhu, and D. Lu. 2010. Underground Miners Localization System based on ZigBee and WebGIS. In *Proc. of the 18^{th} Int. Conf. on Geoinformatics*.

[19] J. T. Isaacs, D. J. Klein, and J. P. Hespanha. 2009. Optimal Sensor Placement for Time Difference of Arrival Localization. In *Proceedings of the 48^{th} IEEE Conference on Decision and Control (CDC)*.

[20] SM R. Islam et al. 2015. The Internet of Things for Health Care: a Comprehensive Survey. *IEEE Access* 3 (June 2015).

[21] L. Jiang, L. N. Hoe, and L. L. Loon. 2010. Integrated UWB and GPS Location Sensing System in Hospital Environment. In *Proc. of the 5^{th} Int. Conf. on Industrial Electronics and Applications (ICIEA)*.

[22] J. Karedal, S. Wyne, P. Almers, F. Tufvesson, and A. F. Molisch. 2007. A Measurement-Based Statistical Model for Industrial Ultra-Wideband Channels. *IEEE Transactions on Wireless Communications* 6, 8 (Aug. 2007), 3028–3037.

[23] B. Kempke et al. 2016. SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization. In *Proc. of the 14^{th} Int. Conf. on Embedded Network Sensor Systems (SenSys)*.

[24] B. Kempke, P. Pannuto, and P. Dutta. 2015. PolyPoint: Guiding Indoor Quadrotors with Ultra-Wideband Localization. In *Proc. of the 2^{nd} Int. Workshop on Hot Topics in Wireless (HotWireless)*.

[25] Y. Kuang, K. Åström, and F. Tufvesson. 2013. Single Antenna Anchor-Free UWB Positioning Based on Multipath Propagation. In *Proceedings of the International Conference on Communications (ICC)*.

[26] J. Kulmer et al. 2017. Using Decawave UWB Transceivers for High-accuracy Multipath-assisted Indoor Positioning. In *Proc. of the 5^{th} Int. Workshop on Advances in Network Localization and Navigation (ANLN)*.

[27] J. Kulmer, S. Grebien, M. Rath, and K. Witrisal. 2018. On the Unimportance of Phase-Coherent Measurements for Beampattern-Assisted Positioning. In *Proc. of the Int. Conf. on Wireless Comm. and Networking (WCNC)*.

[28] J. Kulmer, E. Leitinger, S. Grebien, and K. Witrisal. 2018. Anchorless Cooperative Tracking using Multipath Channel Information. *IEEE Transactions on Wireless Communications* PP, 99 (Jan. 2018).

[29] J. Larranaga, L. Muguira, J. M. Lopez-Garde, and J. I. Vazquez. 2010. An Environment Adaptive ZigBee-based Indoor Positioning Algorithm. In *Proc. of the Int. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*.

[30] P. Lazik and A. Rowe. 2012. Indoor Pseudo-Ranging of Mobile Devices using Ultrasonic Chirps. In *Proc. of the 10^{th} Int. Conf. on Embedded Network Sensor Systems (SenSys)*.

[31] A. Ledergerber, M. Hamer, and R. D'Andrea. 2015. A Robot Self-Localization System using One-Way Ultra-Wideband Communication. In *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*.

[32] E. Leitinger, M. Fröhle, P. Meissner, and K. Witrisal. 2014. Multipath-Assisted Maximum-Likelihood Indoor Positioning using UWB Signals. In *Proc. of the Int. Conf. on Comm. Workshops (ICC)*.

[33] R. Leser, A. Schleindlhuber, K. Lyons, and A. Baca. 2014. Accuracy of an UWB-based Position Tracking System used for Time-Motion Analyses in Game Sports. *European Journal of Sport Science* 14, 7 (Feb. 2014).

[34] S. Liu and T. He. 2017. SmartLight: Light-weight 3D Indoor Localization Using a Single LED Lamp. (Nov. 2017).

[35] D. Lymberopoulos and J. Liu. 2017. The Microsoft Indoor Localization Competition: Experiences and Lessons Learned. *IEEE Signal Processing Magazine* 34, 5 (Sept. 2017).

[36] S. Maranò, W. M. Gifford, H. Wymeersch, and M. Z. Win. 2010. NLOS Identification and Mitigation for Localization Based on UWB Experimental Data. *IEEE Journal on Selected Areas in Communications* 28, 7 (Sept. 2010).

[37] R. Mautz. 2009. The Challenges of Indoor Environments and Specification on some Alternative Positioning Systems. In *Proc. of the 6^{th} Int. Workshop on Positioning, Navigation and Communication (WPNC)*.

[38] R. Mautz. 2012. *Indoor positioning technologies*. Ph.D. Dissertation. ETH Zürich.

[39] P. Meissner, C. Steiner, and K. Witrisal. 2010. UWB Positioning with Virtual Anchors and Floor Plan Information. In *Proc. of the 7^{th} Int. Workshop on Positioning, Navigation and Communication (WPNC)*.

[40] R. Mirza, A. Tehseen, and AV J. Kumar. 2012. An Indoor Navigation Approach to Aid the Physically Disabled People. In *Proc. of the 18^{th} Int. Conf. on Computing, Electronics and Electrical Technologies (ICCEET)*.

[41] P. Mowlaee et al. 2016. *Single Channel Phase-Aware Signal Processing in Speech Communication: Theory and Practice*. John Wiley & Sons.

[42] V. Namboodiri et al. 2012. An Extensive Study of Slotted Aloha-based RFID Anti-collision Protocols. *Computer Communications* 35, 16 (Sept. 2012).

[43] C. Nerguizian, C. Despins, and S. Affes. 2006. Geolocation in Mines with an Impulse Response Fingerprinting Technique and Neural Networks. *IEEE Transactions on Wireless Communications* 5, 3 (March 2006).

[44] C. W. Ou et al. 2017. A ZigBee Position Technique for Indoor Localization Based on Proximity Learning. In *Proc. of the International Conference on Mechatronics and Automation (ICMA)*.

[45] B. Perrat et al. 2015. Quality Assessment of an Ultra-Wide Band Positioning System for Indoor Wheelchair Court Sports. *Proc. of the Inst. of Mechanical Engineers, Part P: Journal of Sports Eng. and Techn.* 229, 2 (April 2015).

[46] X. Qing, Zhi Ning Chen, and T. S. P. See. 2009. Sectored Antenna Array for Indoor Mono-Station UWB Positioning Applications. In *Proc. of the 3^{rd} European Conf. on Antennas and Propagation*. 822–825.

[47] T. H. Riehle, P. Lichter, and N. A. Giudice. 2008. An Indoor Navigation System to Support the Visually Impaired. In *Proc. of the 30^{th} Int. Conf. of the IEEE Engineering in Medicine and Biology Society*.

[48] M. Rzymowski, P. Woznica, and L. Kulas. 2016. Single-Anchor Indoor Localization Using ESPAR Antenna. *IEEE Antennas and Wireless Prop. Letters* 15 (Nov. 2016).

[49] B. Silva, Z. Pang, J. Åkerberg, J. Neander, and G. Hancke. 2014. Experimental Study of UWB-based High Precision Localization for Industrial Applications. In *Proc. of the Int. Conf. on Ultra-WideBand (ICUWB)*.

[50] IEEE Computer Society. 2015. Standard for Low-Rate Wireless Networks. (2015).

[51] X. Sun, Y. Ma, J. Xu, J. Zhang, and J. Wang. 2008. A high Accuracy Mono-Station UWB Positioning System. In *Proc. of the Int. Conf. on Ultra-WideBand (ICUWB)*.

[52] J. Tiemann, F. Eckermann, and C. Wietfeld. 2016. ATLAS - An Open-Source TDOA-based Ultra-Wideband Localization System. In *Proc. of the Int. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*.

[53] D. Vasisht, S. Kumar, and D. Katabi. 2016. Decimeter-level Localization with a Single WiFi Access Point. In *Proc. of the 13^{th} Usenix Conf. on Networked Systems Design and Implementation (NSDI)*.

[54] K. Witrisal et al. 2016. High-Accuracy Localization for Assisted Living: 5G systems will Turn Multipath Channels from Foe to Friend. *IEEE Signal Processing Magazine* 33, 2 (March 2016).

[55] K. Witrisal and P. Meissner. 2012. Performance Bounds for Multipath-Assisted Indoor Navigation and Tracking (MINT). In *Proc. of the Int. Conf. on Comm. (ICC)*.

[56] H. Zhang, X. Cui, B. An, and T. A. Gulliver. 2013. A Distance and Angle Estimated Method based on Single UWB Station. In *Proc. of the Int. Conf. on Sig. Proc., Comm. and Computing (ICSPCC)*.

[57] J. Zhang, P. V. Orlik, Z. Sahinoglu, A. F. Molisch, and P. Kinney. 2009. UWB Systems for Wireless Sensor Networks. *Proc. IEEE* 97, 2 (Feb. 2009).