

Augmenting mmWave Localization Accuracy Through Sub-6 GHz on Off-the-Shelf Devices

Alejandro Blanco
IMDEA Networks Institute
Universidad Carlos III de Madrid
Madrid, Spain
alejandro.blanco@imdea.org

Francesco Gringoli
University of Brescia/CNIT
Brescia, Italy
francesco.gringoli@unibs.it

Pablo Jiménez Mateo
IMDEA Networks
Madrid, Spain
pablo.jimenezmateo@imdea.org

Joerg Widmer
IMDEA Networks Institute
Madrid, Spain
joerg.widmer@imdea.org

ABSTRACT

Millimeter-wave (mmWave) technology is an important element to increase the throughput and reduce latency of future wireless networks. At the same time, its high bandwidth and highly directional antennas allow for unprecedented accuracy in wireless sensing and localization applications. In this paper, we thoroughly analyze mmWave localization and find that it is either extremely accurate or has a very high error, since there is significant mmWave coverage via reflections and even through walls. As a consequence, sub-6 GHz technology can not only provide (coarse) localization where mmWave is not available, but is also critical to decide among multiple candidate antennas and APs for accurate mmWave localization.

Based on these insights, we design a high-accuracy joint mmWave and sub-6 GHz location system. We enable CSI-based angle estimation and FTM-based ranging on off-the-shelf mmWave devices to implement our mechanism and carry out an extensive measurement campaign. Our system is the first to achieve 18 cm median location error with off-the-shelf devices under their normal mode of operation. We further release the location system (and in particular the CSI and FTM functionality) as well as the trace data from the measurement campaign to the research community.

CCS CONCEPTS

• **Networks** → **Location based services.**

KEYWORDS

Indoor localization, CSI, mmWave, AoA, ToF, Wireless Networks

ACM Reference Format:

Alejandro Blanco, Pablo Jiménez Mateo, Francesco Gringoli, and Joerg Widmer. 2022. Augmenting mmWave Localization Accuracy Through Sub-6 GHz on Off-the-Shelf Devices. In *The 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '22)*, June 25–July

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiSys '22, June 25–July 1, 2022, Portland, OR, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9185-6/22/06...\$15.00

<https://doi.org/10.1145/3498361.3538920>

1, 2022, Portland, OR, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3498361.3538920>

1 INTRODUCTION

The rapid increase in traffic demand both in mobile as well as wireless local area networks has boosted interest in networking at mmWave frequencies, where the huge available bandwidth enables significantly higher data rates than the ones that can be achieved at sub-6 GHz frequencies. 5G mmWave networks were rolled out a few years ago in several countries with user data rates of up to 2.5 Gb/s [26]. For Wireless Local Area Networks (WLANs), the IEEE 802.11ad amendment [27] operates at 60 GHz with a bandwidth of 2.16 GHz to deliver data rates of up to 4.6 Gbit/s, and its recently standardized successor 802.11ay [9] can reach up to 100 Gb/s through MIMO, higher order modulation schemes, and channel bonding. To achieve larger link distances, mmWave devices use directional antenna arrays to focus the emitted energy towards the intended receiver, since the range of a single omnidirectional mmWave antenna element is extremely limited. In addition, mmWave signals are blocked more easily by obstacles than sub-6 GHz signals, making it difficult to achieve ubiquitous coverage. For these reasons, wireless devices typically support multiple bands and mmWave connectivity is usually deployed together with sub-6 GHz networks to be able to fall back to the latter whenever the mmWave link is unavailable. Furthermore, to deal with the limited aperture (field-of-view) of mmWave antennas, some off-the-shelf devices integrate up to eight mmWave antenna arrays and switch between them for better angular coverage [15, 25].

The large number of elements in directional mmWave phased antenna arrays enables very accurate estimation of the angle of the signal and the high temporal resolution from the multi-GHz bandwidth provides excellent ranging accuracy. These characteristics make mmWave technology an ideal candidate to design highly accurate wireless location systems [8, 30, 47]. A location system estimates the Angle of Arrival (AoA) and Time of Flight (ToF) of the path between an Access Point (AP) and a client to determine the position of the client. In case this path is in direct Line of Sight (LOS), the angle and range information from AoA and ToF will lead to an accurate estimate of the position of the client. While obstructed LOS paths that pass through obstacles are often too weak to be detected, they may still provide accurate localization in some cases. However,

if the path to the client is Non Line of Sight (NLOS) via one or more reflections, both angle and range information correspond to the reflection rather than the actual client position. This can either occur when 1) only a reflected path exists, 2) the reflected path is much stronger than the LOS path (since the phased array antenna is steered towards the reflection), or 3) the device orientation is such that the LOS path is outside the aperture of the antenna and cannot be captured. Thus, a mmWave location system is either very accurate, or has a high error.

Unfortunately, the mmWave multi-path channel is typically sparse with only few available paths, and due to the limited range there is coverage from only one or very few APs. This makes it exceedingly difficult to distinguish between LOS and NLOS cases using mmWave information alone. However, operating a sub-6 GHz location system [18, 48, 49] along with the mmWave location system allows to distinguish reflections from (obstructed) LOS paths with the help of the lower frequency information. Sub-6 GHz antenna elements are omnidirectional and are thus less affected by device orientation, and the lower frequency leads to better coverage and more visible APs. In addition to improving mmWave localization, lower frequency location information can be used whenever there is no mmWave coverage, although it provides lower accuracy due to the lower bandwidth and lower number of antennas.

Only few prior works study the joint use of mmWave and sub-6 GHz bands. The most common use case is to reduce the mmWave beam-training overhead with the help of sub-6 GHz information [11, 28, 36, 41] to infer the optimal steering direction of the mmWave antenna array. The only system that combines mmWave and sub-6 GHz information for localization purposes “fuses” mmWave ToF-based ranging with sub-6 GHz AoA [23] but does not in fact provide an independent location system at either band.

In this paper, we propose a Multi-band Location System (MultiLoc) that uses both mmWave and sub-6 GHz channel information to achieve high-accuracy device localization. Despite significant research interest in localization, to the best of our knowledge no such multi-band location system has been proposed. MultiLoc uses the sub-6 GHz location estimate to decide whether to use mmWave information at all, and if yes which mmWave AP to use for the localization, whenever the client is covered by one or more mmWave APs. If there is no mmWave coverage or all mmWave paths are estimated to be NLOS, the sub-6 GHz location estimate is used directly. Otherwise, the location estimate from a *single* mmWave AP is used. Our system does not merge location estimates from multiple APs as the mmWave estimates are either very accurate or quite inaccurate. Combining accurate estimates only marginally improves localization performance, but merging even a single inaccurate estimate significantly degrades it. In case the device integrates multiple mmWave antenna arrays with different orientations to improve coverage, MultiLoc uses the array that provides the shortest mmWave ToF estimate to a given AP, since LOS paths are characterized by a shorter length than NLOS paths.

In order to test our system at scale, we implement MultiLoc on Commercial Off-The-Shelf (COTS) mmWave devices [25]. To this end, we port the OpenWRT embedded Linux to these devices and modify the mmWave device drivers in order to extract Channel State Information (CSI) information and enable Fine Time Measurement (FTM) ranging. This is the first open platform that supports




Figure 1: Estimated angles and path lengths when the client has a 45° misalignment with the AP

both mmWave CSI extraction and FTM, and we have made the code and tools freely available [1]. We deploy MultiLoc in indoor and outdoor environments and compare its performance against several baselines, including a multiband (MB), a mmWave, and a sub-6 GHz (Spotfi) system. We also evaluate MultiLoc as well as baseline system performance for different numbers of mmWave antenna arrays on the client. In a benign scenario with a dense deployment of APs and a client with several mmWave arrays facing in different directions, both MultiLoc and the mmWave baseline achieve a median error of 13 cm, whereas the MB baseline and the sub-6 GHz Spotfi have errors of 45 cm and 2.8 m, respectively. We find that only two independent arrays on the client are sufficient to guarantee an error below 20 cm with MultiLoc. Even in sparse and complex office deployments, and for outdoor and mixed indoor/outdoor scenarios, MultiLoc significantly outperforms the baseline schemes. The full MultiLoc implementation as well as the datasets of our measurement campaigns are available at [1].

2 MOTIVATION

We examine the behavior of mmWave localization and its dependence on device orientation in more detail in an indoor scenario. We place a client at different measurement points directly in front of the AP within LOS, as shown in Fig. 1. The five measurement points are equally spaced (with step size 1.2 m) on a line originating at the AP, with the closest being at 4.2 m. Parallel to that line at a distance of 3.4 m there is a wall. On the AP, we measure both ToF and AoA to the client to estimate its position relative to the AP (we provide further details of the ToF and AoA estimation algorithms in Section 3). According to the manufacturer, the aperture of the AP and client antennas is $\pm 30^\circ$, but we observe partial connectivity also beyond that angular range.

If the client has a rotation of 0° with respect to the AP, i.e., it points directly at the AP, the location estimates are very accurate with errors between 8 cm and 13 cm for the measurement points (not shown in the figure). For comparison, the baseline 802.11ac location system operating at 5 GHz obtains a median error of 0.3 m. Fig. 1 shows the same scenario when the client is rotated by 45° counterclockwise. The length and angle of each line represent the distance and the angle estimated for mmWave. For reflected paths, the client location estimated by the AP is given by following the solid line through the wall for the corresponding distance (i.e., mirroring the dotted lines at the wall). The location errors are 0.12, 0.13, 5.7, 4.4 and 6.4 m. The first two points still provide accurate localization despite the LOS path being outside of the antenna aperture, since the AP does not emit a signal at an angle steep enough to reach the client via a reflection, due to its own limited aperture. For the last three points, however, the chosen path is a

reflection off the wall towards which the client antenna is pointing, which significantly decreases the accuracy of the system. In a more extreme case with a rotation of 90 degrees and the client directly facing the wall, none of the antenna patterns of the array have any significant gain directly towards the AP, and for all five positions we observe reflected paths and location errors above 4 m.

These examples highlight how unreliable mmWave localization can be when AP and client are not aligned. This problem is exacerbated for larger AP-client distances and more complex indoor environments. In contrast, the sub-6 GHz system is significantly less accurate but at the same time less dependent on client orientation. This motivates the design of our multi-band location system which uses sub-6 GHz information to distinguish LOS and NLOS mmWave paths. Finally, since the mmWave location system is either very accurate or has a high error, *it is not beneficial to fuse the data from both bands*.

3 MULTI-BAND LOCALIZATION

We now explain the details of the localization techniques. We first give some background on ToF and path parameter estimation. We then explain the details of the MultiLoc AoA estimator for both, mmWave and sub-6 GHz. Finally, we discuss how to combine the measurements from several APs and for both frequency bands.

3.1 Background on FTM ranging

IEEE P802.11-REVmc [7] describes the FTM procedure that can be used by two FTM capable devices to estimate their distance without clock synchronization. One device, the Initiator, starts the FTM session by sending an FTM request frame. The other device, the Responder, acknowledges the request and confirms that the FTM session has started. After the FTM request has been acknowledged, both sides start exchanging packets. The Responder sends a first packet and records the timestamp t_1 when it is sent. Once this packet reaches the other side, the Initiator records the time and then sends back an ACK with two timestamps, the timestamp t_2 when the packet was received and t_3 when the ACK is sent. Once this ACK reaches the Responder, it again records the current time t_4 and extracts t_2 and t_3 from the ACK. The Round Trip Time (RTT) is then estimated as follows, where n is the number of measurements:

$$RTT = \frac{1}{n} \sum_{m=1}^n ((t_4(m) - t_1(m)) - (t_3(m) - t_2(m)))$$

For both mmWave and sub-6 GHz, the distance is then given by $\hat{d} = c \cdot RTT/2$, where c is the speed of light.

3.2 Background on path parameter estimation

Consider a wireless multipath channel where each path is defined by the parameters of the azimuth and the elevation angles at the transmitter and the receiver, denoted by α_{tx} , β_{tx} , α_{rx} and β_{rx} , respectively, the attenuation γ , the path delay τ and the Doppler shift λ . We call $v = [\alpha_{tx}, \beta_{tx}, \alpha_{rx}, \beta_{rx}, \gamma, \tau, \lambda]$ the vector of path parameters. Then, the multipath channel is expressed as:

$$H = \sum_{l=1}^L H(v_l) + W, \quad (1)$$

where L is the number of paths and W is L -dimensional white Gaussian noise.

For active localization purposes, we need to extract the path parameters of the direct path from the multipath channel. However, estimating them is not trivial because all the multipath components interact with each other once they arrive at the receiver. Hence, an accurate estimation of the direct path requires separating all the multipath components. To this end, we follow the approach of mD-track [48] which is a powerful path parameter decomposition algorithm originally designed for passive localization with a Uniform Linear Array (ULA) at sub-6 GHz. The elevation angle cannot be extracted due to the geometry of the array. The vector of path parameters of mD-track is defined as $v = [\alpha_{tx}, \alpha_{rx}, \gamma, \tau, \lambda]$. In this subsection, we provide an overview of this algorithm and in the next subsection we discuss our modifications of the mD-track algorithm for mmWave channels.

The mD-track algorithm estimates the path parameters iteratively by reconstructing the strongest path and its parameters and then subtracting it from the received signal, so that successively weaker paths can be estimated. Once it obtains an initial estimate for all the paths, it applies several rounds of refinement.

Initial estimation. mD-track iteratively determines the contribution of the strongest path, estimates its parameters, reconstructs it, and then subtracts it from the received channel. The result of the subtraction is the channel residual, which is then used to estimate the second strongest path's contribution, and so on. mD-track iterates multiple times until the parameters of all the significant paths are estimated. To this end, mD-track formulates a multi-dimensional estimator denoted as $z(\alpha_{tx}, \alpha_{rx}, \gamma, \tau, \lambda)$ which aims at maximizing the power of the received channel according to a set of possible path parameters. Therefore, the path parameters that maximize the z -function are the ones from the strongest signal as follows:

$$\hat{\alpha}_{tx}, \hat{\alpha}_{rx}, \hat{\tau}, \hat{\gamma} = \arg \max_v |z(\alpha_{tx}, \alpha_{rx}, \gamma, \tau, \lambda)|, \quad (2)$$

and the attenuation is estimated as follows:

$$\hat{\lambda} = z(\hat{\alpha}_{tx}, \hat{\alpha}_{rx}, \hat{\tau}, \hat{\gamma}) / (P \cdot T), \quad (3)$$

where P is the total number of antennas and T is the signal duration. Hence, the estimated vector of path parameters is \hat{v} and the reconstructed channel is $H(\hat{v})$.

For the l -th path, the channel residual is expressed as follows:

$$H_l^r = H - \sum_{l'=1}^{l-1} H(\hat{v}_{l'}), \quad (4)$$

and the residual for $l = 1$ is the original channel $H_0^r = H$.

After estimating the path parameters of the strongest path, mD-track reconstructs it and removes it from the received channel. The initial estimation phase ends once all the paths are reconstructed.

Iterative path parameter refinement. The initial estimates may contain errors since the subtraction step may leak information from stronger paths to weaker paths and vice versa. To improve the estimate, mD-track applies several rounds of refinement by summing the residual to the reconstructed path so that we can iterate over this path and re-estimate its parameters another time. By adding the reconstructed path to the residual we get the raw




Figure 2: URA layout

path as follows:

$$\hat{\mathbf{H}}_l^n = \hat{\mathbf{H}}_l + \mathbf{H}_l^r \quad (5)$$

The raw signal now contains less interference since other paths were already estimated in a prior iteration. However, the residual might still contain useful information that was not extracted correctly before, since the paths were interfering with each other in the initial estimate. mD-track then re-estimates the path parameters for the raw path as in the initial estimation rounds to obtain more accurate parameters. mD-track also re-estimates the residual for every path so that the residual contains less interference as well. mD-track continues iterating until the difference between the path parameters obtained in the last two rounds of iterations is lower than a configurable threshold.

3.2.1 Azimuth estimation at sub-6 GHz. We directly use the mD-track algorithm to estimate the azimuth at the transmitter and receiver, the attenuation and the path delay for every path with the sub-6 GHz system. For simplicity, we do not extract velocity via the Doppler effect since our scenarios do not involve mobility. For localization, we then extract the direct path as the one with the lowest delay, i.e., that arrives first in time. For details we refer the interested reader to [48].

3.3 mmWave angle estimator

We consider a wireless system where the transmitter and receiver have a single RF-chain connected to a Uniform Rectangular Array (URA) with $M \times N$ antenna elements. The signal sent by the transmitter propagates through the multipath channel along L different paths and arrives at the receiver. In the following we refer to Fig. 2 where the transmitter (the client) is located at p and the center of URA of the a -th AP is at \mathbf{x}_a , with the URA lying on the $X = 0$ plane. We characterize the paths from the source as they appear at the centre of the URA using the following parameters:

- **Complex attenuation** γ of the signal over the path.
- **Azimuth angle** α_{rx} between the path and the plane $Y = 0$.
- **Elevation angle** β_{rx} between the path and the plane $Z = 0$.

In order to extract the three parameters, we first need to characterize how the phase rotation in the elements of the URA changes over z and y . We denote by θ_z and θ_y the angles that characterize the antennas in the Z and Y axis, see Fig. 2. We formulate the phase rotation introduced by the signal on every antenna element with respect to the reference antenna element as:

$$a_{m,n} = md_z \sin(\theta_z) + nd_y \sin(\theta_y), \quad (6)$$

where d_z and d_y are the antenna spacing and m goes from 0 to $M - 1$, n goes from 0 to $N - 1$, and $(m, n) = (0, 0)$ is selected as the

reference antenna. Thus, the vector that contains the phases for the antennas across the Z direction is defined as:

$$\boldsymbol{\phi}(\theta_z) = [1, e^{-j2\pi d_z \sin(\theta_z)/\lambda}, \dots, e^{-j2\pi d_z (M-1) \sin(\theta_z)/\lambda}], \quad (7)$$

where λ is the wavelength of the RF signal. Similarly, for the antennas across the Y direction we have:

$$\boldsymbol{\phi}(\theta_y) = [1, e^{-j2\pi d_y \sin(\theta_y)/\lambda}, \dots, e^{-j2\pi d_y (N-1) \sin(\theta_y)/\lambda}] \quad (8)$$

Both angles θ_z and θ_y depend on the the direction of the path, which is characterized by the elevation and azimuth angles. From Fig. 2 we can see that θ_z corresponds to elevation α , so that $\sin(\theta_z)$ is equal to $\sin(\alpha_{rx})$. Angle θ_y , instead, depends on both the elevation and azimuth according to the relation $\sin(\theta_y) = \sin(\alpha_{rx}) \cos(\beta_{rx})$.

With the parameters above, we express the channel by summing all the contributions given by the L paths as follows:

$$\mathbf{H} = \sum_{l=1}^L \boldsymbol{\phi}^T(\theta_{z,l}) \gamma_l \boldsymbol{\phi}(\theta_{y,l}) + W, \quad (9)$$

where W is a white Gaussian noise matrix of size $M \times N$ and $(\cdot)^T$ is the Transpose operator.

3.3.1 AoA estimation at mmWave. The channel described above contains the path parameters for all the paths, but for localization we are interested in the direct path between the client and the AP. To this end, we follow the mD-track approach and we extend it to adapt it to a URA and the mmWave channel characteristics.

In our case, the z -function is the matching projection to the channel residual based on the path parameters, which for the l -th path is defined as follows:

$$z_l(\theta_z, \theta_y) = \boldsymbol{\phi}^*(\theta_z) \mathbf{H}_l^r \boldsymbol{\phi}^H(\theta_y), \quad (10)$$

where $(\cdot)^*$ and $(\cdot)^H$ are the conjugate and Hermitian operator, respectively.

Once the iterative refinement finishes, we extract the path parameters of the direct path. Since the path parameters for a mmWave single-carrier PHY do not contain any temporal information as in an OFDM system, we select the direct path as the one with highest power: we express the power of every path as $power_p = \|\gamma\|^2$. We denote the direct path as p_{dp} . The estimated azimuth is extracted from the $\theta_{z,p_{dp}}$ as follows:

$$\hat{\alpha}_{rx} = \text{asin}(\sin(\theta_{z,p_{dp}}) / \cos(\theta_{y,p_{dp}})), \quad (11)$$

and the estimated elevation is $\hat{\beta}_{rx} = \theta_{z,p_{dp}}$.

3.4 Location estimation

We now describe how MultiLoc determines the final location. First, each AP a computes the location of the client using its own known position \mathbf{x}_a as

$$\hat{\mathbf{p}}_a = \mathbf{x}_a + \hat{d}_a \begin{bmatrix} \cos(\hat{\alpha}_{rx,a}) \cos(\hat{\beta}_{rx,a}) \\ \sin(\hat{\alpha}_{rx,a}) \cos(\hat{\beta}_{rx,a}) \\ \sin(\hat{\beta}_{rx,a}) \end{bmatrix}. \quad (12)$$

We denote by $\hat{\mathbf{p}}_{a,lf}$ and $\hat{\mathbf{p}}_{a,hf}$ the estimated location using the sub-6 GHz and mmWave information, respectively. Note that the elevation angle cannot be estimated by the sub-6 GHz system, and it thus uses $\hat{\beta}_{rx} = 0$ for the estimation of $\hat{\mathbf{p}}_{a,lf}$. All the location information is gathered at a central controller that computes the client position. Before showing how MultiLoc merges the mmWave

and sub-6 GHz location estimates, we first explain how we localize at mmWave and sub-6 GHz bands individually.

3.4.1 Standalone mmWave location system. At typical AP densities, it is unlikely that several mmWave APs provide good coverage at a given location, due to the large path attenuation when traversing obstacles. Many client devices also include more than one antenna array on each device, which helps to overcome the limited angular aperture of a single array and makes the system much more robust to rotation. The first challenge is thus to pick the best array out of the available ones. For a client with K arrays, we obtain K position estimates. We then compute $\hat{\mathbf{p}}_{a,hf,k}$ which is the position estimated by AP a for client array k .

For each of these estimates, we compare the distance and pick the combination of array and AP which minimizes the estimated distance. Since mmWave behaves quasi-optically, the shortest distance is most likely to point to the LOS path and to the best AP if there is more than one within range. The position estimate at mmWave is $\hat{\mathbf{p}}_{hf} = \hat{\mathbf{p}}_{a',hf,k'}$, where a' and k' are the AP and array that result in the lowest distance estimate.

3.4.2 Standalone sub-6 GHz localization. The sub-6 GHz band has better coverage than mmWave and is more likely to penetrate through obstacles. Therefore, the sub-6 GHz localization system has to deal with more position estimates since more APs are able to localize the client than with mmWave. Obstructed paths may provide useful position information but the location estimates from the farthest APs to the client are more likely to be unreliable. We thus compensate the unreliable estimates by weighting the location estimate from every AP based on the estimated distance. We give the weight values of 1 and 0 to the APs which have the minimum and maximum estimated distance, respectively. For the remaining APs, we compute their weights through a linear interpolation between the estimated distances and the minimum and maximum distances (e.g., for three APs and ToF estimates of 10, 12 and 20 m, the corresponding weights are 1, 0.8 and 0). The position at sub-6 GHz is then computed using the weighted centroid approach:

$$\hat{\mathbf{p}}_{lf} = \frac{\sum_{a=0}^{A-1} \hat{\mathbf{p}}_{a,lf} \cdot w_a}{\sum_{a=0}^{A-1} w_a}, \quad (13)$$

where w_a is the weight given by the a -th AP and A is the number of available APs.

3.4.3 MultiLoc localization. For best performance under different conditions, MultiLoc combines the excellent localization precision of mmWave with the availability of sub-6 GHz. MultiLoc uses mmWave localization only if it is deemed reliable, and if not MultiLoc replaces it by sub-6 GHz positioning information. To do so, MultiLoc applies three steps. (1) MultiLoc takes the best array from the client for localization. However, the best array at the client may be still unreliable (in case it is pointing to an obstacle instead of the AP), and thus (2) MultiLoc applies a sub-6 GHz and mmWave discrimination. Finally, (3) MultiLoc selects the best AP. We explain in details the MultiLoc steps below.

(1) Best array selection. As explained before, the client may have several mmWave antenna arrays pointing in different directions. We are interested in determining the array that points in the direction of the AP, to avoid the large bias in the localization that arises

when MultiLoc takes the angle and ToF of a reflection. To do so, we assume that the array with the lowest estimated distance is the correct array since the direct LOS path is the shortest path from the client to the AP.

(2) Sub-6 GHz/mmWave discrimination. This is the most critical step of MultiLoc. Selecting the mmWave information from a reflection leads to high localization errors, but at the same time, whenever choosing the sub-6 GHz location even though the mmWave one is correct, MultiLoc loses accuracy. To this end, MultiLoc measures two metrics, (i) the difference between sub-6 GHz and mmWave ToF estimates, expressed as $\Delta_{tof} = \hat{d}_{hf} - \hat{d}_{lf}$ and (ii) the Euclidean distance between the sub-6GHz and mmWave location estimates, denoted as Δ_l . If the ToF difference Δ_{tof} is larger than 1.5 m (to provide some margin in case of inaccurate sub-6 GHz ToF estimates), the mmWave ToF is longer than the sub-6GHz one which implies that the mmWave path is likely to be NLOS. Thus, MultiLoc should use sub-6GHz positioning. We also observe cases where the sub-6 GHz azimuth is coarsely pointing to the client and the mmWave azimuth to a reflector while there is no significant difference in the ToF estimates. Here, the strongest mmWave path is NLOS but the direct path is strong enough so that the mmWave ToF is a combination of the ToF of these two paths. To overcome this issue, we use the Euclidean distance of the location estimates, Δ_l . In particular, a large Δ_l means that the two bands highly disagree concerning the position of the client, and the most likely reason is that the mmWave AoA is pointing to a reflector. We therefore set a threshold th_d for Δ_l above which, MultiLoc also uses sub-6 GHz localization. We experimentally determine $th_d = 4.5m$ as the value that minimizes the Root Mean Square Error (RMSE) of MultiLoc localization in several example scenarios, since RMSE is sensitive to outliers and thus helps to avoid high errors caused by faulty mmWave positioning. In summary, if $\Delta_{tof} > 1.5$ m or $\Delta_l > th_d$, MultiLoc uses sub-6 GHz and otherwise mmWave. We denote as $\hat{\mathbf{p}}_a$ the position estimate after applying discrimination for the a -th AP.

(3) Best AP selection. Having chosen between sub-6 GHz and mmWave for each individual AP estimate, MultiLoc has to determine the best AP for localization from all the AP candidates. Note that there may be APs with NLOS in the set of candidates, and taking one of those APs to localize will create large errors. Therefore, MultiLoc has to take the AP with the best channel conditions. We observe that the lowest values of the estimated distance are highly correlated with the most accurate localization estimates. However, we also observe an accuracy degradation for azimuth estimates closer to extreme angles ($\pm 90^\circ$), in particular a degradation by a factor 1.5 for cases where the estimated azimuth is either larger than $+45^\circ$ or lower than -45° . Thus, we formulate an objective function in terms of the sub-6GHz estimated distance and azimuth to measure the reliability of the location estimates giving more weight to the distance. Then, we select the best AP (denoted as a^*) as the one that minimizes the following objective function:

$$a^* = \arg \min_a \begin{cases} \mu \hat{d}_a + (1 - \mu) |\hat{\alpha}_{rx}| & \text{if } |\hat{\alpha}_{rx}| \geq 45^\circ \\ \hat{d}_a & \text{otherwise} \end{cases}, \quad (14)$$

where $|\cdot|$ refers to the absolute value.

The μ parameter controls the trade-off between distance and azimuth and is set to 0.85 if $|\hat{\alpha}_{rx}| \geq 45^\circ$ and to 1 otherwise. We base this selection on sub-6 GHz because it is agnostic to the client rotation since the mmWave ToF and azimuth of the AP may contain outliers due to an unsuitable rotation of the client. Thus, the final position estimate of MultiLoc is given by \hat{p}_{a^*} .

Sensitivity analysis. We did an in-depth sensitivity analysis for the parameters used in the sub-6 GHz/mmWave discrimination and best AP selection procedures. For all parameters, minor changes (on the order of 10%) do not significantly impact localization performance. Larger changes for these parameters may, in several cases, cause MultiLoc to select sub-6 GHz when mmWave is more accurate or vice versa, and thus degrade performance. Fortunately, we found these parameter settings to be very robust across different environments and deployment types, and we verify that they work well in the very diverse set of scenarios we investigate.

4 IMPLEMENTATION

Before discussing the details of wireless localization and the design of our multi-band algorithm, we present the COTS hardware on which we implement our system and explain how we modify the device driver to access CSI and FTM.

4.1 mmWave system

For the experimental evaluation we use COTS devices, specifically the MikroTik wAP 60G and MikroTik wAP 60Gx3 shown in Fig. 3. Since its proprietary operating system provides very limited access, we ported OpenWRT to this platform. This gives us the ability to customize the open source wil6210 Linux driver and enable CSI extraction functionality as well as FTM measurements. Our custom fork of OpenWRT and the mmWave driver are available at [1].

4.1.1 Hardware. The MikroTik wAP 60G devices come with a URA of 6 by 6 antenna elements and 60° aperture. The MikroTik wAP 60Gx3 uses three of such arrays arranged at different angles to provide a combined aperture of 180° . Unfortunately, as of now it is not possible to select which antenna to use and the CSI and FTM data are always gathered from the antenna in the middle, and we thus use these devices interchangeably. The same antenna has been partially reverse engineered [54] in order to have further control over it by using a custom FPGA implementation for the control path. Instead, our approach does not require additional hardware modifications which makes it inexpensive and easily accessible. The first challenge we faced was to map each logical antenna element to its physical position on the antenna array, which is important to estimate the AoA. To this end, we establish a link between a pair of MikroTik devices while capturing CSI measurements and we successively cover each antenna element with tin foil. When covered, an antenna element becomes noisy, and we can find the logical element by observing the changes in CSI magnitude of the antenna elements. The mapping is shown in Fig. 3. Since Qualcomm uses 32 bits to define the Beam Patterns (BP) of their devices, only 32 antennas can be addressed, and for this reason the 4 antennas in the corners are unused. Old versions of the MikroTik wAP 60G use the Sparrow B0 chipset, whereas recent versions come with a Sparrow D0 chipset.

The antenna spacing is 0.58 times the wavelength at 60.48 GHz [54]. This antenna spacing creates spatial ambiguity since it is more than half the wavelength. Here, several AoAs may produce the same phase delays at the antennas. To remove the ambiguity and select the correct AoA, we first identify the possible AoAs and we then choose the correct one based on the AoA estimate at sub-6 GHz. We choose as mmWave AoA the one that minimizes the angular difference with respect to the sub-6 GHz AoA. Note that at sub-6 GHz there are no ambiguities because the antenna spacing is exactly half of the wavelength.

4.1.2 Software. MikroTik devices use a proprietary operating system called RouterOS based on a heavily modified 3.3.5 Linux Kernel with a unique binary called *nova* that bundles all functionality. This OS is very limited, and only very basic operations are enabled. In order to perform research with these devices, we created our own custom OpenWRT fork with 60 GHz support. Our original development was done with first generation devices which have a Sparrow B0 chipset for which an official Qualcomm firmware exists. However, the current MikroTik wAP 60G and 60Gx3 APs come with a Sparrow D0 (QCA9500) chipset, for which no official Qualcomm firmware is available. After some trial and error we found that it is possible to extract the required firmware from the package wireless-6.40.9-arm.npk available on the MikroTik website, which contains the wil6210-d0.fw firmware. Newer packages are encrypted and are not compatible with the Linux driver. Furthermore, not all the firmwares have CSI and FTM functionality enabled.

All the communication from the operating system to the chipset and vice versa is done through Wireless Module Interface (WMI) calls. Most of those calls are defined in a file called wmi.h, distributed with the official sources of the wil6210 Linux driver developed by QCA. This file hints at CSI and FTM commands, but provides no information on how to invoke them. By analyzing the different WMI structures, we were able to add new calls for executing commands in the wil6210 driver and return results to user space, as shown in Fig. 4. More specifically, we implemented the necessary vendor subcommands, the missing events, the NLA (Network Level Authentication) policies for FTM, and implemented the required FTM functions and the parameters to those functions. All these modifications are included in our modified wil6210 driver [1]. Additionally, we ported the CSI extraction functionality available for the old chipset on the Talon AD7200 devices [30] to our OpenWRT image. FTM functionality was previously not available, and to the best of our knowledge this is the first time that FTM has been enabled in mmWave COTS devices. A CSI measurement only requires a simple WMI command with few parameters, whereas for FTM we need to start a specific FTM session with a specific set of parameters, which we determined through automated parameter testing.

Since the FTM functionality is experimental and was never publicly released by the vendor, stable long-term measurements are unfortunately not possible with the current firmware. We thus have to disconnect and then reconnect the client to the AP for each new FTM measurement. As a consequence, FTM measurements take 1.5 s on average. However, the FTM transaction itself (including the handshake and packet transmissions) only takes around 10 ms. We expect that with new firmware releases or some firmware hacking, reliable FTM will be possible in real time with these devices.




Figure 3: Hardware.

4.2 Sub-6 GHz system

For the sub-6 GHz localization implementation based on IEEE 802.11ac, we use two COTS devices as was done in other prior work [34], one to extract CSI and another for FTM. Fig. 3 shows the devices we use. The *CSI device* is an ASUS AC2900 RT-AC86U with a Broadcom BCM4366E chipset. We use the Nexmon CSI extractor tool from [12] to obtain CSI for 4x4 MIMO and 80 MHz IEEE 802.11ac frames. In this configuration, the extracted CSI matrix has a size of 4x4x256 complex values, where 256 corresponds to the number of subcarriers of the 80 MHz channel. As *FTM device* we use an ASUS AC1300 RT-ACRH13 router which supports FTM requests out of the box and a laptop with an Intel Dual Band Wireless-AC 8260 wireless interface as initiator, as in [14]. We apply some modifications to the Linux driver of the card to start the FTM session and extract the FTM measurements.

4.3 System calibration

Unfortunately, a range of preliminary measurements reveal that some calibration is needed to handle hardware imperfection that affect the CSI and ToF values, which would considerably reduce the accuracy of the AoA and ToF estimates.

4.3.1 60 GHz CSI. A CSI measurement contains a phase and amplitude value for each of the 32 antenna elements. Note that for all devices, antenna elements 31 and 32 only report noise, which leaves a total of 30 antenna elements to work with. The value reported by the driver is an integer between 0 and 1024 which corresponds to a phase between 0 and 2π . For every CSI measurement, we convert the 30 CSI values into a matrix form of 6x6 with a layout as shown in Fig. 3, to convert the logical to the corresponding physical antenna elements in the rectangular array. Further analysis of the CSI data reveals that the AP has to account for the phase added by the oscillator to every antenna in the array and some systematic phase shifts that are WiFi frame dependent.

• **Phase offsets.** Since the devices have one radio-frequency (RF) chain, there is one common oscillator to tune to the central frequency. However, the internal circuitry and layout of the device introduce phase offsets at each antenna element, which need to be removed for accurate AoA estimation. To this end, we mount an AP/client pair in a shielded environment at 0° azimuth and elevation angles, as seen in Fig. 5. This ensures that the direct path itself does not introduce any phase delays and the shielding removes any multipath effects that may distort the phase. We then measure the reference CSI that corresponds to the phases offsets. To remove

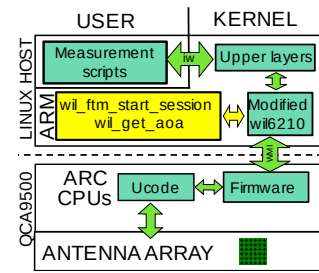


Figure 4: System modifications




Figure 5: Calibration setup.




Figure 6: Phase histogram

them from actual CSI measurements, we perform the element-wise Hadamard division between measured CSI data and the reference CSI. These offsets are constant over power-cycles and for different devices with the same antenna type, and thus this calibration has to be done only once.

We verify the calibration using a rotation table that rotates the device in steps of 5° both in the azimuth within $\pm 60^\circ$ and elevation angles within $\pm 30^\circ$. The estimated angles closely follow the ground truth, with very low angle errors below 2° . This measurement also shows that under ideal conditions, accurate angle estimation is possible far beyond the official antenna aperture of $\pm 30^\circ$.

• **Systematic phase shifts.** We also observe systematic phase shifts that affect all antenna elements and that appear randomly over the WiFi packets. These phase jumps remain the same as can be seen from the two Gaussian distributions in Fig. 6, one of which corresponds to the correct phase centered at 0 and another one centered between $-3\pi/8$ and $-\pi/2$. We can thus compensate the phase shifts as they follow a clear pattern, rather than discard them. To do so, we apply a Gaussian mixture model to the phase values over a set of WiFi packets to model the two Gaussian distributions, one for the correct CSI phases and the other for the phase-rotated CSI. We then shift the mean of the less populated Gaussian distribution to center it on the correct phase value.

4.3.2 60 GHz ToF. We further observe an offset for FTM measurements as was also reported in [14] for sub-6 GHz. To calibrate the FTM, we use the same setup as for the CSI and set up the routers at 2.5 m of distance. We then take 100 measurements and use the difference to the true distance to correct the offset. The FTM offset is the same for all the mmWave devices and is constant over time.

4.3.3 Sub-6 GHz CSI. The sub-6 GHz oscillator introduces a phase uncertainty at the four RF chains, which we have to remove to obtain accurate AoA at the AP side. We connect the four AP transmit ports to the corresponding four receive ports of the client using identical cables [48] to ensure that the transmitted signals arrive at the same time. The CSI values thus contain the constant phase offsets that we use for the phase correction. Unfortunately, this calibration has to be done per device and for each power cycle.

4.3.4 Sub-6 GHz ToF. Finally, the sub-6 GHz FTM measurements have an offset which we determine and correct by taking outdoor measurements at known distances. The sub-6 GHz FTM calibration is also device-dependent, but is constant over time.

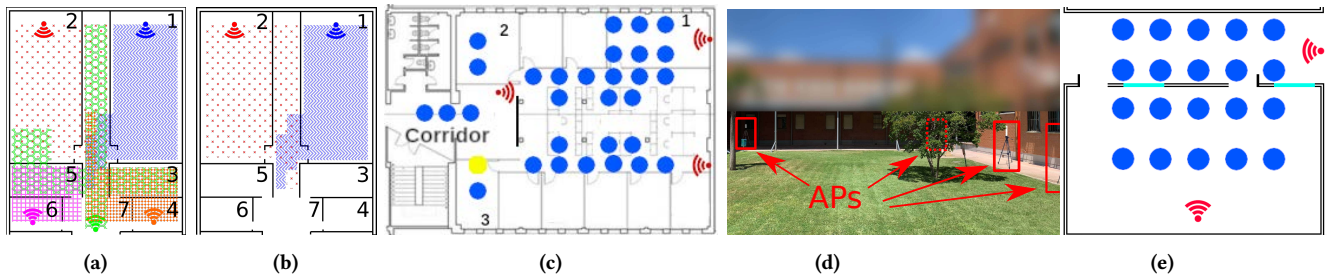


Figure 7: Scenarios and coverage maps: (a) indoor 5 APs; (b) indoor 2 APs; (c) office; (d) outdoor; (e) indoor/outdoor

5 EXPERIMENTAL EVALUATION

We run an extensive measurement campaign to test the localization accuracy of MultiLoc in several scenarios and compare it with different baseline location systems. We also analyze the performance of the key components of MultiLoc ToF, azimuth and elevation angles as a function of the rotation of the client.

At each measurement point we collect CSI and FTM for both sub-6 GHz and 60 GHz using the devices shown in Fig. 3. The mmWave antennas are highly directional and the device orientation significantly impacts localization performance. To emulate antennas with different orientations, as client we use two MikroTik wAP 60G devices mounted back-to-back on a rotation table model PTU-D47. This table allows for rotations of up to 200° , and thus covering 360° requires two devices. We then measure 4 different rotations at 45° steps and let each client connect to each AP in turn, for a total of 8 client orientations. This allows us to emulate clients with different numbers of mmWave sub-arrays. Having 8 arrays on the client ensures that it can reliably cover the whole azimuth. Note that systems with 8 arrays are already commercially available [44].¹ Currently, only the highly directional MikroTik devices are FTM-capable, but for a client with a more omnidirectional antenna array such as the TP-Link Talon AD7200 router, only one array would be needed. We take 300 CSI measurements and 10 FTM measurements per rotation and AP. For sub-6 GHz we use 2 different routers, one to capture the CSI and a second one to capture the ToF.

5.1 Scenarios

We deploy MultiLoc in four different scenarios in order to comprehensively evaluate its performance indoors as well as outdoors, and for LOS and NLOS settings.

Indoor scenario: This scenario has a size of 15×20 m and contains six rooms, one corridor, and up to five APs as shown in Fig. 7(a). We collect measurements in 110 different positions. For a deployment of five mmWave APs that we label with the number of the rooms, at most one AP is in LOS for each measurement position, even though some rooms are covered by more than one AP: for instance, the bottom-left corner of room 2 is covered by AP 2 and 7 (shown with red and green patterns, respectively). In this scenario, we have ubiquitous mmWave coverage even in rooms without AP. Our goal is to validate the robustness of MultiLoc against unreliable localization estimates that come from NLOS APs. The second scenario shown in Fig. 7(b) uses the same indoor environment but

with a more sparse deployment of only two APs that ensure at least sub-6 GHz coverage everywhere (sub-6 GHz coverage is not shown in the figures). In general, we observe a surprising degree of NLOS mmWave coverage even through walls. We also evaluate realistic cases where mmWave coverage is not available and MultiLoc needs to meet localization requirements using sub-6 GHz information.

Office scenario: We then test MultiLoc in a more crowded indoor office environment shown in Fig. 7(c), where people are working at their desks or move around. This environment contains a central open area with smaller offices and conference rooms on either side. It has a size of 25×12 m and is covered by 3 APs. We collect measurements in 31 different locations, including the open area, a conference room (room 1) and two offices (2 and 3). The central area is covered by 2 APs, but walls, pillars and other obstacles create obstructed LOS at some of the points. Moreover, one AP provides coverage in the conference room and also provides obstructed LOS to some points in the central area through a glass wall. This allows to evaluate how MultiLoc deals with partially and fully obstructed LOS settings and environment dynamics.

Outdoor scenario: This scenario comprises 4 APs and it is depicted in Fig. 7(d). It has a size of 18×21 m. In general, the APs have a clear LOS but trees and pillars may cause NLOS in certain areas. This scenario represents the benign conditions under which baseline location systems are usually tested, with the majority of the APs having LOS and good channel quality, and unreliable location estimates coming from spurious NLOS can easily be filtered out.

Indoor/Outdoor scenario: In more complex deployments such as a campus or multi-building office premises, devices may move between indoor and outdoor settings. This makes accurate localization particularly challenging since indoor APs provide partial coverage at close by outdoor locations, and vice versa. Therefore, in Fig. 7(e) we explore a hybrid scenario of size 25×18 m with an indoor and an outdoor AP. We collect measurements at 20 points (blue dots), half of them indoors and half outdoors.

5.2 Comparison with other systems

We compare MultiLoc against the following benchmarks:

- **Spotfi** [18] is a well-known sub-6 GHz location system based on AoA, which is estimated together with the path delay by a 2D MUSIC version using spatial smoothing. It merges information from multiple APs to obtain the user location.
- **High frequency baseline (HF baseline):** To the best of our knowledge, the only other mmWave location system [17] that is based on angle and distance estimates of ToF from Received Signal Strength Indicator (RSSI), which is highly inaccurate. For a fair

¹We were unfortunately unable to acquire an Airfide AP that directly integrates 8 sub-arrays. However, there should be no difference between the Airfide and our emulated 8 antenna array setup, other than causing a longer measurement duration.

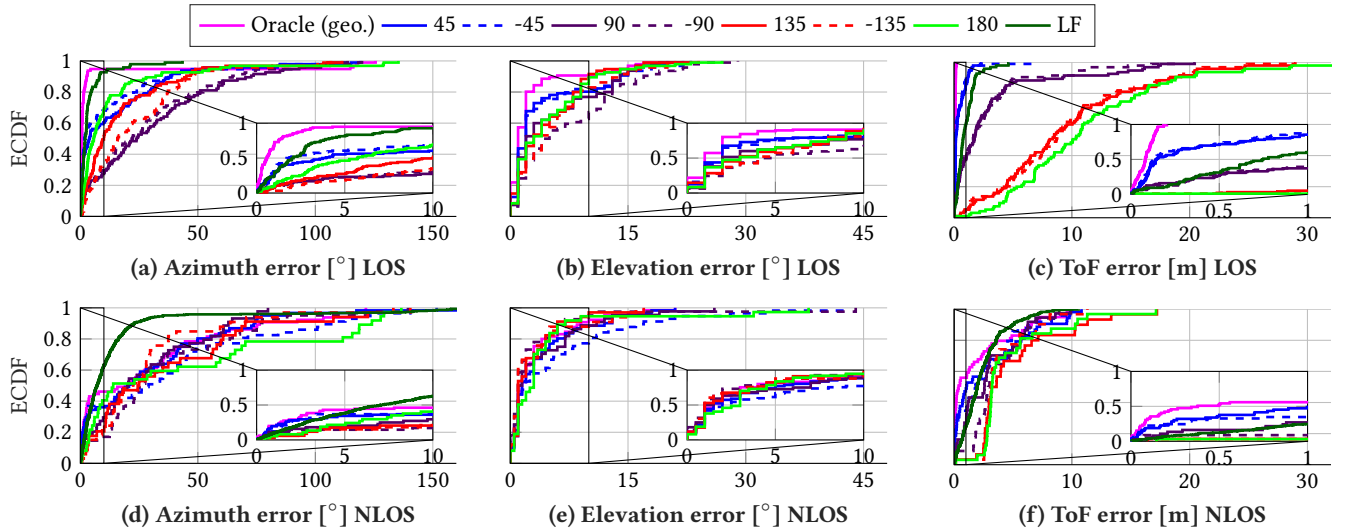


Figure 8: Azimuth, elevation and ranging errors for the indoor scenario with different rotations at the client side.

comparison, we use the ToF from our mmWave FTM and the AoA obtained by Multiple Signal Classification (MUSIC) [38] as improved 60 GHz baseline.

• **Multiband baseline (MB baseline):** The only multi-band system in the literature [23] fuses 60 GHz ToF information and AoA from sub-6 GHz extracted with MUSIC to localize targets. To replicate this system, we combine the mmWave FTM-ranging taken from the MikroTik and the AoA obtained with MUSIC from the ASUS AC2900 sub-6 GHz CSI.

Both the HF and MB baselines were designed for a single APs. Averaging information from all (including unreliable) APs degrades performance. For a fair comparison, we thus apply to these baselines the same AP and array selection strategies as for MultiLoc, and we verify that this indeed improves their performance.

5.3 Rotation analysis

In order to expand on the intuitions of Section 2, we perform an extensive rotation analysis of the client. Fig. 7(a) shows the coverage map for that scenario. We have divided the analysis of the indoor scenario in two separate sets to better explain the challenges to overcome, i.e., one set with positions and rotations where one of the APs is in LOS, and one set with the positions and rotations where no AP is in LOS for azimuth and elevation angles and ToF ranging. As explained in Section 2, mmWave is very susceptible to antenna misalignment when the client array is not directly facing the AP. Therefore, we quantify the estimation error according to array orientation, where the geometrical oracle indicates the array that is most closely aligned with the direct line between AP and client, and the misalignment in degrees indicates by how much the other orientations deviate from that direct line. Fig. 8 shows the empirical CDF for all the cases. In particular, subfigures (a) and (d) show the results for azimuth in LOS and NLOS, respectively, (b) and (e) show elevation results for LOS and NLOS and (c) and (f) show ToF results in LOS and NLOS. The figures include all estimates from all APs, not just the ones that would be chosen by MultiLoc.

From Fig. 8(a) with azimuth error for indoor LOS, we observe that if we choose the optimal rotation, there is less than 2.75° of

angle estimation error for 90% of the cases. The next best rotation is 180° with an error of less than 33° for 90% of the cases since it receives the direct path via the backplane and estimates a good angle with the sign changed. (However, 45° has a better median error of 1.8° .) The worst estimates are obtained when the antenna arrays are perpendicular to each other, i.e., $\pm 90^\circ$. The default beam patterns of the MikroTik devices favor pointing to the right side of the device, which can be seen also for the two rotations of $\pm 45^\circ$. As a comparison, results for sub-6 GHz have less than 10° of error for 90% of the cases. It is also noticeable that the azimuth errors for the oracle rotation jump from less than 5° to more than 120° , which is due to the spatial ambiguity as it may select the incorrect azimuth angle. MultiLoc resolves this ambiguity using sub-6 GHz information and obtain a maximum error of 6.04° . The LOS elevation errors are depicted in Fig. 8(b). As for the azimuth case, the oracle orientation provides the most accurate elevation estimates as the errors are below 5° for the 90% of the cases. We also observe a jump from 5° to 10° due to the influence of a first order reflection path that bounces off the floor of the room. This is an extremely challenging case since this path has a similar azimuth angle but a slightly different elevation angle than the direct path. Therefore, the MultiLoc AoA estimator might not discriminate these two paths and estimates only one path with a correct azimuth but with a slight elevation error. The other orientations behave similarly without an excessive degradation in the elevation errors, which implies that MultiLoc estimates either the elevation from the direct path or from the path that reflects off the floor. Unfortunately, sub-6GHz estimates the azimuth only due to the geometry of the antenna array, therefore elevation cannot be extracted. Fig. 8(c) shows the estimated ToF ranging error for LOS. The oracle rotation has an error of less than 0.21 m, and even for the deviations of $\pm 45^\circ$ we obtain less than 1.2 m of error for 90% of the cases. The worst case here is for 180° , as it always captures a very long reflected path. In comparison, sub-6 GHz has less than 2 m error for 90% of the cases. We next summarize the NLOS case. In Fig. 8(d) are depicted the NLOS azimuth errors, we see that the azimuth estimated for mmWave is affected by huge errors for all rotations, while sub-6 GHz always performs better

and reports less than 23° of error for 90% of the cases. The NLOS elevation errors are found in Fig. 8(e). It is noticeable that all the rotations perform at the same level achieving an elevation error of 7° for the 90% of the cases. Finally, we show in Fig. 8(f) the estimated distance error for the NLOS case. mmWave still yields better results than sub-6 GHz for 75% of the cases, with a median error of 0.4 m for the oracle orientation. Although the corridor has only one 60 GHz AP in LOS, all of the other NLOS APs still provide coverage to some parts of it. Next to the doors to the rooms, for example the door in room number 1, the coverage mainly comes from reflections, but the top part of the corridor also has coverage from AP 2 via obstructed LOS directly through the wall.

For brevity we omit the figures for the rest of the scenarios and we summarize the results for the outdoor one. Different from the indoor scenario, there is no significant NLOS. We observe less than 3.1° of azimuth error for 90% of the cases. Here, the two offsets of $\pm 45^\circ$ perform better than 180° as there are no close walls that can reflect that signal. Once more, the two offsets $\pm 90^\circ$ perform worst. With sub-6 GHz we get less than a 7° error for 90% of the cases. Regarding the elevation angle, all the errors are below 4° for the oracle orientation since the path that reflects off the ground is considerably weaker than in the indoor scenario as the ground of the outdoor scenario is grass. Hence the interference that this path adds to the estimation of the direct path is almost negligible. As for the distance, the oracle rotation has less than 0.55 m of error for all the points, followed by offset -45° with an error of less than 0.42 m for 90% of the cases. For sub-6 GHz we obtain a maximum error of 4.8 m for all the points.

In summary, mmWave performance is excellent under ideal conditions but degrades even for small changes in device orientation. The sub-6 GHz system has less attenuation when traversing obstacles and walls and the omnidirectional antennas are insensitive to rotation, but it is less accurate.

5.4 Location results

MultiLoc overcomes the limitations of standalone mmWave systems by identifying when mmWave estimates are accurate and when they are not. We now analyze the joint location accuracy for clients with 1, 2, 4 and 8 mmWave antenna arrays. For the case of fewer than 8 arrays, we report the location errors per measurement point for each of the (8, 4 and 2) possible rotations. We compare MultiLoc and its stand-alone subsystems, i.e., the mmWave module (HF MultiLoc) and the sub-6 GHz module (LF MultiLoc), to the baselines described in Section 5.2. Due to limited space, we only show in the empirical CDF of the performance for the 8 arrays and 1 array cases. While in the figures we report the location accuracy for a 2D system, the elevation angle allows to also compute the 3D coordinates of the client, and in the text we also report the accuracy of the 3D version of MultiLoc.

5.4.1 Indoor scenario with 5 APs. Fig. 7(a) shows the scenario with 5 multi-band APs, located in rooms 1, 2, 4 and 6 and in corridor 7. In this dense deployment, NLOS propagation happens only in rooms 3 and 5, where we collect 14 measurement points. Results can be found in Fig. 9(a) for the cases of 8 and 1 arrays and in Fig. 11 for all the cases. For the 8 arrays case, MultiLoc, its HF component and the HF baseline exhibit a median error of 13 cm that grows

up to 45 cm for the MB baseline because of the lower accuracy of the sub-6 GHz AoA. Moreover, while the maximum error achieved by MultiLoc is as low as 0.5 m, it grows to 5 m for the HF baseline that can not cope with spatial ambiguity. Regarding the sub-6 GHz results, LF MultiLoc achieves a median error of 1 m whereas Spotfi obtains 2.8 m: this indicates that weighting the APs by their distance considerably improves the performance. We emphasize that for the 1 array case, MultiLoc correctly selects either 60 GHz or sub-6 GHz, while the MB baseline, HF MultiLoc and HF baseline perform poorly with median errors of 1.9 m, 3.15 m and 3.7 m. With few arrays, MultiLoc and the MB baseline perform better than mmWave-only systems since the sub-6 GHz information is less dependent on the rotation of the user. Interestingly, the accuracy for the sub-6 GHz system is lower in the corridor where multipath is more significant. In general, mmWave sees much less multipath due to the narrow antenna beams and the higher signal attenuation, which leads to higher accuracy. We also observe that the median error remains stable between 8 and 4 arrays (13 cm) and it slightly degrades for 2 arrays (23 cm), but for 1 array, there is a significant increase in the median error to 65 cm. This indicates that a 2 arrays setting is an appropriate trade-off between accuracy and cost. Finally, even though there are pure NLOS areas (rooms 3 and 5), MultiLoc is able to achieve an error of less than 0.5 m in all cases. For 3D MultiLoc, the median error slightly degrades from 13 cm (2D) to 19 cm (3D) for the 8 array case. However, MultiLoc accuracy for the bottom 10% of the cases degrades more significantly from 50 cm (2D) to 2.5 m (3D). This is caused by the influence of the ground reflection on the elevation angle estimation as explained in Section 5.3. Instead of 8, when the array size is changed to 4, 2 and 1 arrays, the median errors of 3D MultiLoc increase to 18 cm, 37 cm and 71 cm, respectively.

5.4.2 Indoor scenario with 2 APs. This scenario mimics a sparse deployment where only the large rooms contain 60 GHz APs, whereas the other rooms are covered by sub-6 GHz APs, as shown in Fig. 7(b). We measure the same 110 points as in the 5 APs scenario, with the difference that there is only mmWave LOS in rooms 1 and 2, and obstructed LOS in the corridor, whereas the rest of the rooms have NLOS coverage from sub-6 GHz.

We summarize results in Fig. 9(b). For the 8 arrays case, the median errors are 16 cm for MultiLoc, HF MultiLoc and the HF baseline while they are 21 cm, 53 cm, 1.25 m and 3 m for 3D MultiLoc, MB baseline, LF MultiLoc and Spotfi. We note that mmWave-only systems do not work in many positions: almost 20% of the points are without 60GHz coverage. Here, MultiLoc provides reasonable accuracy thanks to sub-6 GHz with a maximum error of 4.5 m. It is worth highlighting that Spotfi gets poor accuracy for the bottom 20% of the cases with errors above 6 m, which indicates that it also suffers from the lack of APs. Regarding results with 1 array, MultiLoc achieves one meter median error while the others perform poorly achieving results that are worse than Spotfi. Impressively, with 4 and 8 antenna arrays, MultiLoc is able to find a direct path in all the LOS rooms and even in some parts of the corridor where there is through-wall coverage from APs 1 and 2. Overall, MultiLoc takes advantage of mmWave in (obstructed) and in LOS areas and sub-6 GHz in NLOS, reducing the median error by a factor of 5.8 between MultiLoc and a mmWave-only system.

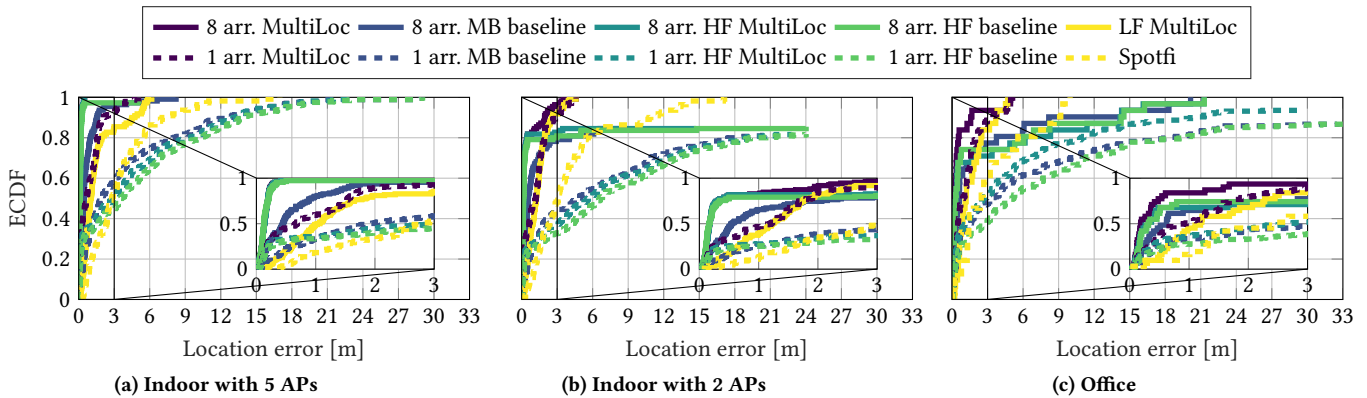


Figure 9: Location error for all indoor scenarios with 1 and 8 antenna arrays configurations.

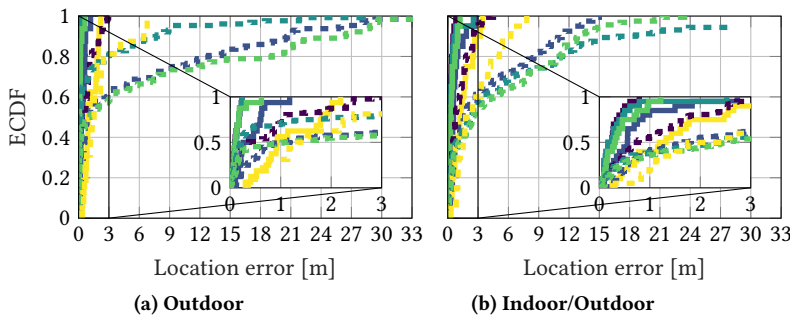


Figure 10: Location error for the outdoor and indoor/outdoor scenarios with 1 and 8 antenna array configurations.




Figure 11: Location accuracy for different systems and antenna array configurations in the indoor scenario.

5.4.3 *Office scenario.* The previous scenarios are devoid of obstacles, and an entire room either had good LOS coverage by one AP, or had no LOS coverage at all. We now study a more realistic scenario Fig. 7(c), with more complex obstructed LOS cases from several APs. For instance, the central area contains obstacles that block the direct path at some points and it also has obstructed coverage by the AP in the conference room (room 1). The results of this scenario are shown in Fig. 9(c). For the 8 array case, MultiLoc is the most accurate one achieving a median error of 18 cm while 3D MultiLoc, HF MultiLoc and HF and MB baselines achieve 35 cm, 31 cm, 31 cm and 60 cm, respectively. Although the scenario is much richer in multipath components than the previous one, the median error does not degrade which implies that the mmWave directional links make localization more robust to reflections from surrounding objects. However, the maximum errors are higher, up to 5 m, due to the difficult conditions for the points in the corridor and in rooms 2 and 3. The corridor has mmWave coverage via first order reflection from a wall (black wall in Fig. 7(c)), which makes mmWave localization unreliable. The measurement points in rooms 2 and 3 also have mmWave coverage via reflections since the walls of the rooms obstruct the LOS path. For all these points, MultiLoc filters the unreliable mmWave locations and uses sub-6GHz information. As a result, the MultiLoc errors are below 5 m while HF MultiLoc, HF and MB baselines get errors above 15 m. Interestingly, the yellow point exhibits accurate mmWave localization since one AP in the central area provides partially obstructed LOS mmWave coverage. Regarding the 1 array case, MultiLoc provides sub-meter median accuracy (90 cm) while 3D MultiLoc, HF MultiLoc and

HF and MB baselines have median errors of 1 m, 2.5 m, 5 m and 3.3 m, respectively. Concerning the sub-6GHz-only systems, the LF MultiLoc outperforms Spotfi, as expected, with a median error of 1.2 m compared to 2.5 m.

5.4.4 *Outdoor scenario.* The outdoor scenario has fewer reflections than the indoor one since walls are much further away. For most of the cases we either have a direct path or a very large reflection, except when there is a tree in the middle. This leads to an improvement compared to the indoor scenario, with the median error decreasing to 10 cm as depicted in Fig. 10(a). For sub-6 GHz, also LF MultiLoc and Spotfi achieve better performance with errors of 90 cm and 1.5 m. Every time the mmWave system sees a reflection, MultiLoc is able to correctly identify it and fall back to sub-6 GHz. For 1 array, MultiLoc’s errors are below 4 m while the other mmWave approaches have errors above 15 m. Here, 3D MultiLoc achieves a median error of 20 cm and a maximum error of 40 cm, since MultiLoc estimates more accurately the elevation angle as the reflections from the grass-covered ground are weaker.

5.5 Indoor/Outdoor scenario

The last scenario contains two APs, one indoors and one outdoors. In this case, selecting the correct AP is critical to avoid large location errors. We show results in Fig. 10(b). Regarding the 8 array case, both MultiLoc and HF MultiLoc exhibit a median error of 17 cm while 3D MultiLoc and the HF and multiband baselines obtain 25 cm, 30 cm and 55 cm, respectively. Again, LF MultiLoc achieves better performance than Spotfi. For the 1 array setting, MultiLoc is the

only system achieving a submeter median error (75 cm). Interestingly, there are a few points where MultiLoc selects the outdoor AP when the client is inside and vice versa: this happens when the direct path crosses a window (we depict windows in cyan in Fig. 7(e)). In these cases we observe similar positioning accuracies as with LOS conditions. This indicates that in buildings with glass walls, the positioning accuracy of outdoor clients can be as good as indoor ones. We also note that the accuracy of LF MultiLoc is slightly worse in indoor conditions due to the rich multipath environment, while mmWave for both, indoor and outdoor, produces similar accuracy due to multipath sparsity.

6 RELATED WORK

Location systems at different frequencies have attracted a lot of interest, as we summarize next.

- *Sub-6GHz*: Localization using sub-6GHz channels has been extensively studied over the past two decades. In the early 2000s, the first localization systems were using RSSI to estimate distance [4, 6, 51, 52]. However, due to the unpredictable effects of the environment on the RSSI, overall accuracy was limited. Huge improvements were demonstrated with the arrival of IEEE 802.11n enabling new approaches for AoA and ToF estimation [10, 18, 35, 39, 43, 49, 50]. In 2016, IEEE introduced FTM ranging protocol for WiFi, which was thoroughly evaluated both in indoor and outdoor environments [14, 16]. More recently, [33] presented a sub-6 GHz location system with sub-meter accuracy in NLOS settings.

- *mmWave*: Localization with such high frequencies is appealing due to the high spatial and temporal resolution. The first mmWave location system [42] estimated the position of the target directly using RSSI. This approach was later extended to RSSI fingerprinting [13, 29] and was also combined with AoA measurements [17]. AoA and Angle difference of Arrival (AdoA) techniques were also evaluated in [5, 20, 31] and were shown to achieve good performance. In addition, mmWave localization can be used to drive handover decisions [30]. Different from these angle based systems, which multiple APs per room, MultiLoc works with just one. Similar to MultiLoc, the ToF-based localization system in [21, 22] achieve outstanding accuracy but they were implemented on an experimental platform. MultiLoc, instead, is implemented on COTS devices, which makes its deployment more straightforward. We would like to highlight that MultiLoc is the first system implemented on COTS devices that enables estimation of AoA and ToF. Finally, also passive mmWave localization was investigated in [8, 45, 46].

- *Multiband*: To the best of our knowledge, the only multiband location system other than MultiLoc is reported in [23]. That system, however, merges sub-6 GHz AoA and 60 GHz ranging to estimate the position of the target device. MultiLoc, instead, combines complete position estimates from the two bands to improve localization performance. For this reason, MultiLoc is much more robust since it continues to work when 60 GHz is not available and can filter unreliable mmWave estimates.

Multi-connectivity systems combine sub-6 GHz and mmWave bands for other purposes, e.g., to reduce mmWave beam training overhead using sub-6 GHz information. It was demonstrated that there is a strong correlation between high and low frequency bands [24], which enables the design of sub-6 GHz assisted mmWave beam

steering mechanism. The authors of [11] show how to improve mmWave beam search using sub-6 GHz information, simulating a sub-6 GHz antenna with 4 elements and a 32-element mmWave antenna. A similar approach is taken in [36] which focuses on sub-6 GHz assisted mmWave beam steering and link maintenance. The authors experimentally demonstrate how to use sub-6 GHz ToF and an external gyroscope to estimate the position of the AP. In [2], the spatial correlation matrix is transformed from sub-6 GHz to 60 GHz in order to reduce mmWave training overhead. Such transformation can also be determined using deep-learning techniques [3].

Testbeds are essential for mmWave experimentation. OpenMili [53] and X60 platform [37] are both FPGA-based 60 GHz testbeds with custom antenna arrays. M-Cube [54] reverse engineers the control path of COTS 60 GHz RF frontends to build mmWave multi-array and MIMO systems with up to 8 antenna arrays. The MIMORPH platform [19] is an open source FPGA-based multiband system that supports real time operation. Several efforts exist to instrument COTS devices for mmWave experimentation. The first work in this direction is [40] which extracts signal strength of received frames from a 60 GHz TP-Link Talon AD7200 to directly determine angle information from only a few beam training packets and thus reduce beam-training overhead. For the same COTS devices, it is even possible to extract the CSI to then design custom beam patterns that maximize the signal-to-noise ratio [32]. From a practical perspective, such COTS testbeds are highly important since they can be deployed easily and at the same time show the constraints of realistic consumer hardware. Besides localization, they can be an important platform for research on wireless sensing, imaging, mapping and backscatter. We note that while extracting the CSI from the aforementioned Talon routers is possible, the Talon's physical shape of the antenna array makes AoA estimation more difficult and less accurate, since the antenna elements are not uniformly distributed as in the MikroTik we use here.

7 CONCLUSIONS

In this paper, we demonstrated that the accuracy of stand-alone mmWave location systems highly depends on the antenna array orientation. To address this problem, we introduced MultiLoc, a dual band location system which uses sub-6 GHz technology to i) decide which of the mmWave APs and which of the available antenna arrays to choose when localizing a target and to ii) directly determine the position of the client when the mmWave location estimate is unreliable. We carried out an extensive measurement campaign for both bands to evaluate MultiLoc. The results show that our system is the first to achieve 18 cm median location error with off-the-shelf devices in complex indoor environments. To inspire future research on mmWave and multi-band localization, we publicly released the location system and the CSI and FTM extraction tools, as well as the datasets of our measurement campaigns.

ACKNOWLEDGMENT

This research work was sponsored in part by the European Union's Horizon 2020 research and innovation programme under Grant No.871249 (LOCUS), the Spanish Ministry of Science and Innovation (MICIU) grant RTI2018-094313-B-I00 (PinPoint5G+) and the Region of Madrid through TAPIR-CM (S2018/TCS-4496).

REFERENCES

- [1] alejandroBlancoPizarro. Multiloc. <https://github.com/IMDEANetworksWNG/MultiLoc>, 2022.
- [2] Anum Ali, Nuria Gonzalez-Prelcic, and Robert W. Heath. Estimating millimeter wave channels using out-of-band measurements. *2016 Information Theory and Applications Workshop (ITA)*, 2016.
- [3] M. Alrabeiah and A. Alkhatieb. Deep learning for mmwave beam and blockage prediction using sub-6 ghz channels. *IEEE Transactions on Communications*, 68(9):5504–5518, 2020.
- [4] P. Bahl and V.N. Padmanabhan. Radar: an in-building rf-based user location and tracking system. In *Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064)*, volume 2, pages 775–784 vol.2, 2000.
- [5] Guillermo Bielsa, Joan Palacios, Adrian Loch, Daniel Steinmetzer, Paolo Casari, and Joerg Widmer. Indoor localization using commercial off-the-shelf 60 GHz access points. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 2384–2392. IEEE, 2018.
- [6] Paul Castro, Patrick Chiu, Ted Kremenek, and Richard Muntz. A probabilistic room location service for wireless networked environments. In *International conference on ubiquitous computing*, pages 18–34. Springer, 2001.
- [7] IEEE Approved Draft Standard for Information technology–Telecommunications and information exchange between systems - Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE P802.11-REVmc/D8.0, August 2016*, pages 1–3774, 2016.
- [8] Dolores Garcia, Jesus Omar Lacruz, Pablo Jiménez Mateo, and Joerg Widmer. POLAR: Passive object localization with IEEE 802.11 ad using phased antenna arrays. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pages 1838–1847. IEEE, 2020.
- [9] Yasaman Ghasempour, Claudio R. da Silva, Carlos Cordeiro, and Edward W. Knightly. IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine*, 55(12):186–192, 2017.
- [10] Domenico Giustiniano and Stefan Mangold. Caesar: Carrier sense-based ranging in off-the-shelf 802.11 wireless lan. In *Proceedings of the Seventh Conference on emerging Networking Experiments and Technologies*, pages 1–12, 2011.
- [11] Nuria Gonzalez-Prelcic, Anum Ali, Vutha Va, and Robert W. Heath. Millimeter-Wave Communication with Out-of-Band Information. *IEEE Communications Magazine*, 55(12):140–146, 2017.
- [12] Francesco Gringoli, Matthias Schulz, Jakob Link, and Matthias Hollick. Free your CSI: A channel state information extraction platform for modern Wi-Fi chipsets. In *Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, pages 21–28, 2019.
- [13] Pei-Yuan Hong, Chi-Yu Li, Hong-Rong Chang, YuanHao Hsueh, and Kuochen Wang. Wbf-ps: Wigig beam fingerprinting for uav positioning system in gps-denied environments. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pages 1778–1787. IEEE, 2020.
- [14] Mohamed Ibrahim, Hansi Liu, Minitha Jawahar, Viet Nguyen, Marco Gruteser, Richard Howard, Bo Yu, and Fan Bai. Verification: Accuracy evaluation of WiFi fine time measurements on an open platform. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 417–427, 2018.
- [15] Airfide Inc. <http://airfidenet.com>, accessed 2021.
- [16] Kevin Jiokeng, Gentian Jakllari, Alain Tchana, and André-Luc Beylot. When FTM Discovered MUSIC: Accurate WiFi-based Ranging in the Presence of Multipath. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pages 1857–1866. IEEE, 2020.
- [17] Ojas Kanhere and Theodore S Rappaport. Position locationing for millimeter wave systems. In *2018 IEEE Global Communications Conference (GLOBECOM)*, pages 206–212. IEEE, 2018.
- [18] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. Spotfi: Decimeter level localization using wifi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 269–282, 2015.
- [19] Jesus O. Lacruz, Rafael Ruiz Ortiz, and Joerg Widmer. A Real-Time Experimentation Platform for sub-6 GHz and Millimeter-Wave MIMO Systems. *The 19th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys 2021)*, 2021.
- [20] Jesus Omar Lacruz, Dolores Garcia, Pablo Jiménez Mateo, Joan Palacios, and Joerg Widmer. mm-FLEX: an open platform for millimeter-wave mobile full-bandwidth experimentation. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, pages 1–13, 2020.
- [21] Nebojsa Maletic, Vladica Sark, Marcus Ehrig, Jesús Gutiérrez, and Eckhard Grass. Experimental evaluation of round-trip tof-based localization in the 60 ghz band. In *2019 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–6. IEEE, 2019.
- [22] Nebojsa Maletic, Vladica Sark, Marcus Ehrig, Jesus Gutierrez, and Eckhard Grass. Performance evaluation of los round-trip tof localization: A 60ghz band case study. In *WSA 2020; 24th International ITG Workshop on Smart Antennas*, pages 1–6. VDE, 2020.
- [23] Nebojsa Maletic, Vladica Sark, Jesús Gutiérrez, and Eckhard Grass. Device localization using mmWave ranging with sub-6-assisted angle of arrival estimation. In *2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, pages 1–6. IEEE, 2018.
- [24] Pablo Jimenez Mateo, Alejandro Blanco Pizarro, Norbert Ludant, Matteo Marugan Borelli, Amanda Garcia-Garcia, Adrian Loch, Zhenyu Shi, Yi Wang, and Joerg Widmer. A Comprehensive Study of Low Frequency and High Frequency Channel Correlation. *2019 International Conference on Computing, Networking and Communications (ICNC)*, 2019.
- [25] MikroTik. <https://mikrotik.com/products/group/60-ghz-products>, accessed 2021.
- [26] Arvind Narayanan, Eman Ramadan, Jason Carpenter, Qingxu Liu, Yu Liu, Feng Qian, and Zhi-Li Zhang. A First Look at Commercial 5G Performance on Smartphones. *Proceedings of The Web Conference 2020*, 2020.
- [27] Thomas Nitsche, Carlos Cordeiro, Adriana Flores, Edward Knightly, Eldad Perahia, and Joerg Widmer. IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi. *IEEE Communications Magazine*, 52(12):132–141, 2014.
- [28] Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer. Steering with eyes closed: Mm-Wave beam steering without in-band measurement. *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015.
- [29] Alain Olivier, Guillermo Bielsa, Irene Tejado, Michele Zorzi, Joerg Widmer, and Paolo Casari. Lightweight indoor localization for 60-ghz millimeter wave systems. In *2016 13th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pages 1–9. IEEE, 2016.
- [30] Joan Palacios, Paolo Casari, Hany Assasa, and Joerg Widmer. LEAP: Location Estimation and Predictive Handover with Consumer-Grade mmWave Devices. *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, 2019.
- [31] Joan Palacios, Paolo Casari, and Joerg Widmer. JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems. *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, 2017.
- [32] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. Adaptive codebook optimization for beam training on off-the-shelf ieee 802.11 ad devices. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 241–255, 2018.
- [33] Alejandro Blanco Pizarro, Joan Palacios Beltrán, Marco Cominelli, Francesco Gringoli, and Joerg Widmer. Accurate ubiquitous localization with off-the-shelf IEEE 802.11ac devices. *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021.
- [34] Maurizio Rea, Traian Emanuel Abrudan, Domenico Giustiniano, Holger Clausen, and Veli-Matti Kolmonen. Smartphone positioning with radio measurements from a single wifi access point. In *Proceedings of the 15th International Conference on Emerging Networking Experiments and Technologies*, pages 200–206, 2019.
- [35] Maurizio Rea, Aymen Fakhreddine, Domenico Giustiniano, and Vincent Lenders. Filtering noisy 802.11 time-of-flight ranging measurements from commoditized wifi radios. *IEEE/ACM Transactions on Networking*, 25(4):2514–2527, 2017.
- [36] Maurizio Rea, Domenico Giustiniano, Guillermo Bielsa, Danilo De Donno, and Joerg Widmer. Beam Search Strategy for MillimeterWave Networks with Out-of-Band Input Data. *2020 Mediterranean Communication and Computer Networking Conference (MedComNet)*, 2020.
- [37] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Roshan Shyamsunder, Owen Torres, et al. X60: A programmable testbed for wideband 60 ghz channels with phased arrays. *Computer Communications*, 133:77–88, 2019.
- [38] Ralph Schmidt. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation*, 34(3):276–280, 1986.
- [39] Souvik Sen, Jeongkeun Lee, Kyu-Han Kim, and Paul Congdon. Avoiding multipath to revive inbuilding wifi localization. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*, pages 249–262, 2013.
- [40] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. Compressive millimeter-wave sector selection in off-the-shelf ieee 802.11 ad devices. In *Proceedings of the 13th International Conference on emerging Networking Experiments and Technologies*, pages 414–425, 2017.
- [41] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. WiFi-Assisted 60 GHz Wireless Networks. *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, 2017.
- [42] Marco Vari and Dajana Cassioli. mmwaves rssi indoor network localization. In *2014 IEEE International Conference on Communications Workshops (ICC)*, pages 127–132. IEEE, 2014.
- [43] Deepak Vasishth, Swarun Kumar, and Dina Katabi. Decimeter-level localization with a single WiFi access point. *NSDI'16: Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, page 165–178, Mar 2016.
- [44] Song Wang, Jingqi Huang, Xinyu Zhang, Hyoil Kim, and Sujit Dey. X-Array: Approximating Omnidirectional Millimeter-Wave Coverage Using an Array of Phased Arrays. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020.
- [45] Teng Wei and Xinyu Zhang. mtrack: High-precision passive tracking using millimeter wave radios. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 117–129, 2015.

- [46] Chenshu Wu, Feng Zhang, Beibei Wang, and KJ Ray Liu. mmtrack: Passive multi-person localization using commodity millimeter wave radio. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pages 2400–2409. IEEE, 2020.
- [47] Chenshu Wu, Feng Zhang, Beibei Wang, and K. J. Ray Liu. mmTrack: Passive Multi-Person Localization Using Commodity Millimeter Wave Radio. *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 2020.
- [48] Yaxiong Xie, Jie Xiong, Mo Li, and Kyle Jamieson. mD-Track: Leveraging Multi-Dimensionality in Passive Indoor Wi-Fi Tracking. *The 25th Annual International Conference on Mobile Computing and Networking*, 2019.
- [49] Jie Xiong and Kyle Jamieson. Arraytrack: A fine-grained indoor location system. In *10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13)*, pages 71–84, 2013.
- [50] Jie Xiong, Karthikeyan Sundaresan, and Kyle Jamieson. Tonetrack: Leveraging frequency-agile radios for time-based indoor wireless localization. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 537–549, 2015.
- [51] M.A. Youssef, A. Agrawala, and A. Udaya Shankar. Wlan location determination via clustering and probability distributions. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications, 2003. (PerCom 2003)*, pages 143–150, 2003.
- [52] Moustafa Youssef and Ashok Agrawala. The horus wlan location determination system. In *Proceedings of the 3rd international conference on Mobile systems, applications, and services*, pages 205–218, 2005.
- [53] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. OpenMili: A 60 GHz software radio platform with a reconfigurable phased-array antenna. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 162–175, 2016.
- [54] Renjie Zhao, Timothy Woodford, Teng Wei, Kun Qian, and Xinyu Zhang. M-cube: A millimeter-wave massive MIMO software radio. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020.