

Single- and Multiple-Access Point Indoor Localization for Millimeter Wave Networks

Joan Palacios, Guillermo Bielsa, Paolo Casari, Joerg Widmer

Abstract—Millimeter wave (mmWave) location systems not only provide accurate positioning for location-based services, but can also help optimize network operations, for example through location-driven beam steering and access point association. In this paper, we design and evaluate localization schemes that exploit the characteristics of mmWave communication systems. We propose two range-free algorithms belonging to the broad classes of triangulation and angle difference-of-arrival. The schemes work both with multiple anchors and with as few as a single anchor, under the only assumption that the floor plan and the positions of the mmWave access points are known. Moreover, they are designed to be lightweight, so that even computationally-constrained devices can run them. We evaluate our proposed algorithms against two benchmark approaches based on fingerprinting and angles of arrival, respectively. Our results, obtained both by means of simulations and through measurements involving commercial 60-GHz mmWave devices, show that sub-meter accuracy is achieved in most of the cases, even in the presence of only a single access point. The availability of multiple access points substantially improves the localization accuracy, especially for large indoor spaces.

Index Terms—Millimeter wave; localization; simulation; ray tracing; measurements; commercial 60 GHz hardware

I. INTRODUCTION

Millimeter-wave (mmWave) communications in the 30 to 300 GHz range are a disruptive technology for future high-rate wireless communications [2]–[4]. mmWaves are characterized by crisp reflections carrying significant power with little energy scattering around the direction predicted by the law of reflection [5], [6], and a predominance of the line-of-sight (LoS) component over reflected components in the absence of blockage [7]. These channel effects have been studied in great detail, e.g., in [6], [8], and the above characteristics imply a quasi-optical propagation behavior, which enables indoor localization methods based on geometrical propagation assumptions [9].

Indoor localization is considered an enabling technology for a range of applications [10], including people and asset track-

ing, large-scale virtual reality, next-generation RFID tags [11], assisted living [12], tracking of patients in protected areas, as well as other functions related to social proximity and environment-dependent customization of mobile devices [13]. This makes it a hot topic for current theoretical and practical research, as proven by the broad participation to yearly localization competitions such as [14]. Recently, some indoor localization algorithms have achieved remarkable accuracy. For example, Chronos [15] achieves a median error of about 70 cm and 1 m in LoS and non-LoS (NLoS) environments, respectively, using multi-channel WiFi signals. The competition report in [16] also shows that a very good average accuracy is achieved by some approaches such as [17] (infrastructure-based, mixing angle and distance measurements) and [18] (infrastructure-free, based on fingerprinting in dense WiFi deployments). Other systems achieve sub-meter accuracy by integrating WiFi signal processing with inertial sensors [19], or by deploying ultra-wideband devices [20], [21].

Importantly, accurate localization and tracking can serve as a proxy for physical mmWave communication functions, such as beam training [22], handover and context switching. While the degree of accuracy required for indoor localization may vary widely depending on the specific application, we remark that no extreme accuracy is required to assist, for example, handover and beam training/tracking in mmWave networks. In fact, the typical beamwidth of current mmWave hardware ranges between 7 and 10 degrees or more [23]. Such apertures translate into a horizontal span between 60 and 90 cm at 5 meters from the access point (AP), or between 90 and 180 cm at 10 meters. In this context, sub-meter accuracy would be more than sufficient to help choose the right sector for link establishment, or to help detect sector changes as a client moves.

Still, several characteristics of mmWave links require to rethink the typical localization methods employed in indoor environments. For example, according to the Friis equation, due to the short wavelength the path loss can easily exhibit 30 to 40 dB higher attenuation compared to microwave signals over typical link distances. Moreover, in some bands the spectrum has further absorption peaks (e.g., at 60 GHz, due to oxygen absorption) [9], [24]. While this implies that a LoS propagation path is easily distinguishable from NLoS paths, it also limits the range of indoor mmWave communications, which is often confined to a single room or a portion thereof [25]. For communications, these effects are typically compensated for via antenna arrays that permit to form narrow and very directional beams, implying that LoS and NLoS components can be discriminated in the angle domain.

It is desirable that mmWave localization algorithms can run on terminals with limited computational capabilities or

Manuscript received March 28, 2018; revised July 27, 2018 and October 16, 2018; accepted February 05, 2019. Date of publication XXX XX, XXX; date of current version XXX XX, XXX. This work has been supported in part by the ERC project SEARCHLIGHT (grant no. 617721), the Ramon y Cajal grant RYC-2012-10788, the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919) and the Spanish Ministry of Science, Innovation and Universities the DiscoEdge grant (TIN2017-88749-R). The associate editor coordinating the review of this paper and approving it for publication was P. Salvo Rossi.

A preliminary version of this work has been presented at the IEEE SECON 2016 conference [1].

All authors (e-mail: first.name.last.name@imdea.org) are with the IMDEA Networks Institute, 28918 Madrid, Spain. Joan Palacios and Guillermo Bielsa are also with University Carlos III of Madrid, 28918 Madrid, Spain.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier xx.xxxx/TWC.xxxx.xxxxxxx

energy constraints. This requires to achieve sufficient accuracy without complex signal processing through a low-overhead design. The localization process can be aided in part by the anchor APs, which can simplify the environment estimation by broadcasting constraints such as room boundaries (walls, ceiling height, etc.) as well as their own position. Besides this information, the localization algorithms should rely only on the information that is available locally on the device, i.e., provided by the device’s RF interface.

In this paper, we propose two localization schemes for mmWave systems that work solely based on the angle measurements implicitly provided by the beam training process of mmWave devices. One of the schemes uses these measurements for a form of triangulation, and the other one uses them to compute angle difference of arrival (ADoA) values. The two algorithms offer different tradeoffs between accuracy and environment awareness requirements. They are designed to work with as many APs as are available in the indoor mmWave network, and leverage indoor multipath propagation in order to improve the localization accuracy. In conditions of particularly low coverage, these characteristics make it possible to localize a node with only a single AP. Note that, even in case the PHY layer of the mmWave devices cannot be easily modified to directly make angle of arrival (AoA) information available, sector information of the highly directional antenna arrays is available at the MAC layer [26], making very it easy to pass such angle-related information to higher-layer protocols. Our algorithms assume to know the indoor area’s floor plan and the location of the mmWave APs. This is a common assumption for indoor localization [27]–[31] and can be realized by means of map-based services [32]. However, we do not require the knowledge of temporary obstructions, such as human bodies blocking mmWave propagation paths. As long as such obstructions do not lead to a very low number of visible paths for the client, the loss of a few propagation paths does not hamper the functionality of the algorithms. Additionally, spurious paths originating from temporary reflections are very likely to be filtered out by the validation or optimization steps of the algorithms. The feasibility and practicality of our schemes is demonstrated through both simulations and experiments with commercial mmWave systems operating in the 60 GHz band, in two different indoor scenarios.

This paper substantially extends several aspects of our previous work [1]. First, while [1] is concerned with single-anchor localization, this work focuses on realistic mmWave deployments with multiple APs and extends all localization algorithms to work with multiple anchors. In this sense, the algorithms in [1] are a special case of those presented in this paper. When designing these improved algorithms, we adhere to the initial design requirement that localization should be achieved using only angle information as provided by the mmWave hardware (e.g., through sector information) and that the algorithms should have limited computational complexity. Moreover, we present several additional simulation results both for multiple and for a single anchor AP. These simulations are complemented by two sets of experimental results obtained with commercial mmWave hardware, including a multi-AP deployment in a realistic office space.

The remainder of this paper is organized as follows. Section II surveys related work. Section III describes the proposed algorithms. Section IV discusses simulation results and draws initial conclusions on the performance of the algorithms, and Section V presents an experimental campaign involving actual mmWave hardware, which characterizes the performance of our localization algorithms in real-world environments. Finally, Section VI draws concluding remarks.

II. RELATED WORK

We subdivide related work on localization into a general review of classical schemes (Section II-A) and more recent advances (Section II-B).

A. Classical localization schemes

Trilateration and multilateration are among the most common range-based localization methods [33]. Ranging can be provided via the evaluation of the Time of Arrival (ToA) or of the received signal strength (RSS) [34], if a path loss model is available. However, range-based approaches may suffer from inaccurate ranging in typical scenarios. Triangulation [33] requires the knowledge of the distance between two anchor nodes and of the angles of departure from the anchors. Here, angle estimation errors negatively affect the accuracy, especially if the distance between the client and the anchors is large. Fingerprinting-based methods [33], [35] rely on the assumption that measurable radio fingerprints are distinct for a given location. While the definitions of fingerprint vary broadly depending on the scenario and radio technology [36], this method has the significant drawback that a very high measurement effort is required for fingerprint collection (the greater the desired accuracy, the more measurements must be collected). The mathematical analysis of the error performance of schemes based on TOA, TDoA, RSS, and AoA in an NLoS environment is provided in [37] via Cramer-Rao bound expressions, observing that the AoA distance error tends to increase faster than TDoA or RSS errors. The fusion of ToA and RSS measurements via an adaptive likelihood metric in WiFi scenarios is advocated in [38]. The proposed algorithm works with the empirical distribution of such measurements and is robust to non-Gaussian noise, achieving a median error of about 2 m.

The choice of range-based or range-free methods depends on the availability of a precise path loss model: in general, range-based methods achieve better accuracy, but require to accurately tune the parameters of the path loss model in a way that is environment-dependent, and can lead to large errors if such parameters are wrongly estimated or updated too infrequently. In these cases, angle-based methods are preferable.

B. Recent advances in indoor localization

Recent work significantly advanced WiFi-based indoor localization. For example, ArrayTrack [39] implements an FPGA-based AoA localization system that estimates locations via a maximum likelihood (ML) approach. The system is

shown to be accurate but sensitive to changes in orientation and antenna height. CUPID [40] is a range-based system that extracts the direct propagation path in multipath WiFi environments and locates a user using multiple RSSI measurements. Centaur [41] fuses acoustic and radio-frequency ranging using Bayesian inference to achieve better accuracy than with either system alone. SAIL [42] exploits a single AP for localization, and improves the accuracy of position estimates through dead reckoning enabled by a smartphone's sensors. The system needs to estimate the gait of the user carrying the smartphone, and is sensitive to pose changes. iLocScan [43] harnesses the multipath propagation of WiFi signals in order to localize a signal source and roughly estimate the surrounding boundaries. The approach is demonstrated using a USRP implementation of the algorithm and the related antenna array signal processing. The common aspect of these works is that multi-antenna signal processing is used to improve the performance of the localization algorithms, even in scenarios where only a single AP is available.

Especially mmWave technology is amenable to high-accuracy indoor localization [44]. The shorter mmWave wavelength, compared to microwave frequencies in the sub-6 GHz range, makes it possible integrate larger antenna arrays in a smaller space, enabling MIMO approaches for localization. By leveraging a sparse representation of the MIMO channel matrix, a localization technique based on ML was designed in [45]. However, such techniques are computationally intensive, and efficient implementations usually require programmable hardware such as FPGAs. The performance of RSS-, TDoA- and AoA-based localization schemes in mmWave systems is compared via simulations in [46], assuming the presence of several anchor nodes deployed over a circumference around the receiver. It is observed that the AoA approach achieves the smallest localization error because of the broad AoA spectrum diversity originating from the circular geometry. Under the assumption of no initial environment knowledge, JADE [47] localizes a client based solely on ADoA information, and maps the environment by leveraging the relationship between physical and virtual anchors, where virtual anchors correspond to the reflected paths of physical anchors. While the method obtains sub-meter localization accuracy in simulations, it requires several APs in order to achieve a realistic map of the environment. A single-anchor pseudo-lateration method is proposed in [48] that estimates the position of a client by measuring the distance from a physical anchor and from a corresponding virtual anchor. Like other range-based methods, pseudo-lateration requires environment-dependent calibration in order to achieve good results. mmWave has been combined with different technologies to localize a client. In [49], the authors consider uplink mmWave and downlink visible light communications, where the lowest error obtained was about 6 cm. For mmWave localization, the authors rely on a single multi-antenna anchor, and exploit reflected paths.

While mmWave systems make it possible to discriminate signals in the angular domain, ultra-wideband (UWB) systems discriminate different propagation paths in time, and enable accurate localization via time-of-arrival (ToA)-based ranging

techniques. In [50] the authors review UWB ranging techniques together with the primary sources of TOA error (e.g., clock drift, and interference). Fundamental ToA bounds are described in both ideal and multipath environments. UWB can potentially enable concurrent ranging [21], which has the potential to reduce the ranging time by allowing overlapping transmissions.

By observing that deterministic multipath components typically bear up to 90% of the UWB channel impulse response's energy, the authors in [51] employ a single anchor node, extract ToA information from the AoA spectrum, and estimate the distance between the client and physical or virtual anchors via an ML approach. The knowledge of the floor plan is assumed in order to identify the location of virtual anchors. In [52], the TDoA between different multipath components is used to compute range estimates via an ML approach. Filtering such estimates through a Kalman filter yields a median localization error below 5 m in simulations.

Cooperative localization is also a promising paradigm that improves the localization accuracy at the cost of exchanging information among the network nodes [53]. In the iterative algorithm DILOC [54], the nodes exchange location estimates and derive their position as the convex combination of the coordinated of a proper set of neighboring nodes. The process is initialized by $m+1$ anchors in an m -dimensional environment. DILAND [55] improves over DILOC by supporting noisy distance measurements and unstable communication links. In the network navigation framework [56], the nodes localize themselves using internal measurements (e.g., inertial sensors) and intra-node measurements (e.g., radio ranging). The paper presents theoretical results to establish how navigation information evolves in a cooperative network. In [57] and [58] the authors develop a framework to characterize the fundamental limits of localization accuracy for a single client and in a wideband cooperative localization scenario, respectively. In the latter, it is assumed that the nodes exchange whole waveforms, instead of summary information such as ToA or RSS.

Our work differs from the above literature in that our algorithms employ only angle information in mmWave scenarios unlike, e.g., [40], [46], [48]. While incorporating RSS-based ranging in the location system tends to achieve decimeter-level accuracy [44], angle information is readily available from the sector scanning carried out by the firmware of mmWave devices for beam training and link establishment, and does not require to build environment-specific path loss models. Therefore, a distinctive feature of our algorithms is that they do not require to change the communication standard (unlike [15], [40], [45]), to incorporate additional hardware or UWB systems (unlike [21], [39], [41], [49]–[51]), or to collect large fingerprinting databases (as in [18], [35], [36]). Additionally, we explicitly avoid the exchange of additional messages for the purpose of localization (besides the ones mandated by the mmWave standard), which is different from collaborative localization approaches [54]–[58]. We design our algorithms for low complexity, such that they can be implemented on simple end devices: this avoids more complex ML approaches [39], [45], [51], [52]. Finally, our algorithms are tailored to typical mmWave scenarios, which normally

Table I
NOTATION EMPLOYED IN THE PAPER

Name	Meaning	Name	Meaning
$\mathbf{p}, \mathbf{n}, \mathbf{p}_{\text{tx}}^n$	Point, normal vector, location of AP n	\mathcal{A}^n	Set of the virtual anchor locations
N_{AP}	Number of APs	$\tilde{\mathcal{A}}^n$	Partition of set $\mathcal{A}^n \cup \{\mathbf{p}_{\text{tx}}^n\}$
$P_{\mathbf{p}}^n(\alpha)$	AoA spectrum for AP n at point \mathbf{p}	\mathcal{A}_{μ}^n	Set of virtual anchors mirrored up to μ times for AP n
$N_{\mathbf{p}}^n$	Number of MPCs in $P_{\mathbf{p}}^n(\alpha)$	A_{μ}	Set defined as $\bigcup_{n=1}^{N_{\text{AP}}} \mathcal{A}_{\mu}^n$
$\mathbf{M}^{(n,\mathbf{P})}$	$2 \times N_{\mathbf{p}}^n$ matrix with MPCs in $P_{\mathbf{p}}^n(\alpha)$	v_{max}	Number of anchors used for validation in TV
$\check{\mathbf{M}}^{(\mathbf{P})}$	Concatenation of all $\mathbf{M}^{(n,\mathbf{P})}$ matrices	β	Vector such that AP β_i generates MPC i in $\check{\mathbf{M}}^{(\mathbf{P})}$
δ_m	Angle difference $\check{\mathbf{M}}_{2,m}^{(\mathbf{P})} - \check{\mathbf{M}}_{2,1}^{(\mathbf{P})}$	\mathcal{D}^n	Set of fingerprints for AP n
\mathcal{Z}	Set of obstacles	\mathcal{F}	Set of fingerprint measurement locations
α_0	Reference angle	$\mathbf{T}^{(n,\mathbf{P})}$	NLoS feature matrix (last $N_{\mathbf{p}}^n - 1$ rows of $\mathbf{M}^{(n,\mathbf{P})}$)
$\Delta\alpha$	Error on α_0 (compass bias)	σ	AoA estimation error

feature a high density of APs in order to guarantee sufficient coverage. Still, the algorithms also work under sparse deployment conditions, and even in the presence of only a single access point.

III. MILLIMETER WAVE LOCALIZATION SCHEMES

We now present the localization algorithms, which are designed such that the client can run them locally by measuring only angle of arrival information. Such information is inherently available in mmWave scenarios, which implement directional communications through phased arrays, and run beam training algorithms to identify the best steering direction for the array's main lobe [59], even in dynamic scenarios [60], [61]. The algorithms presented in the following are representative of two broad classes of approaches, those based on triangulation and those based on ADoA. Such classes have different advantages and disadvantages. The first algorithm, named Triangulate-Validate (TV), requires fewer anchors to localize a client successfully, but is sensitive to errors in the device orientation. The second one is based on ADoA information, and has the advantage that it is inherently immune to orientation biases, but requires a larger number of anchors than TV. These algorithms will be compared against two benchmark approaches from the literature: the first is based on fingerprinting (FP) [1], [36], which is lightweight for the client, but requires the setup of a fingerprint database, typically implying a very significant measurement collection effort; the second one is named CCAL and is based on the solution of ADoA localization as a convex quadratic problem [62]. All algorithms are designed to leverage the sparse AoA spectrum that results from the quasi-optical mmWave propagation, and can manage the multiple APs that are typically encountered in indoor mmWave deployments. However, even if the indoor space is illuminated by only a single AP (for example in case other APs are blocked by obstacles), the algorithms can leverage multipath propagation in order to localize a node using both LoS and NLoS paths.

We now present the required background and notation. The most important symbols and definitions are collected in Table I. With focus on an indoor scenario, we define a three-dimensional Cartesian coordinate system centered in one of the corners of the room. The system is described by the canonical vectors $\mathbf{e}_x = (1, 0, 0)$, $\mathbf{e}_y = (0, 1, 0)$ (oriented orthogonally along the floor sides), and $\mathbf{e}_z = (0, 0, 1)$ (oriented

along the height of the room), such that any point $\mathbf{q} = q_x \mathbf{e}_x + q_y \mathbf{e}_y + q_z \mathbf{e}_z$ can be mapped to a triple (q_x, q_y, q_z) . The room boundaries and any other permanent obstacles containing radio-reflective surfaces are grouped in the reflective objects set \mathcal{Z} . Obstacles are modeled as three-dimensional polyhedra with flat polygonal faces, straight edges and sharp vertices. We treat each face as an oriented surface S , represented by its normal vector

$$\mathbf{n} = \frac{(\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)}{\|(\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)\|}, \quad (1)$$

where \mathbf{p}_1 , \mathbf{p}_2 and \mathbf{p}_3 are points of S , $\|\cdot\|$ denotes the Euclidean norm, and \times the cross-product. We assume that one or more mmWave APs $1, \dots, n, \dots, N_{\text{AP}}$ are installed in the room at the locations $\mathbf{p}_{\text{tx}}^1, \dots, \mathbf{p}_{\text{tx}}^n, \dots, \mathbf{p}_{\text{tx}}^{N_{\text{AP}}}$. These APs act as reference nodes for localization purposes. We also assume that each node trying to localize itself is aware of the location of the APs and of the obstacles in \mathcal{Z} . This is in line with our assumption that localization must be attained even by devices with limited computational power or hard energy constraints. The algorithms are designed to leverage the reflections of the AP signals off indoor boundaries in order to compute location estimates. For every AP n , the input to our algorithms is the AoA spectrum $P_{\mathbf{p}}^n(\alpha)$, which records the distribution (over the azimuthal plane) of the amplitude of multipath components (MPCs) of the signal from AP n , as seen at a given location \mathbf{p} . The AoA spectrum is a function of the azimuth α , measured relative to a reference angle α_0 . For each AP n , $P_{\mathbf{p}}^n(\alpha)$ is processed to yield a compact representation of different MPCs at \mathbf{p} . In particular, a peak in the reception pattern is identified with an MPC [63]. An MPC can be either a LoS or an NLoS path. We remark that, with respect to the classical triangulation and angle-difference-of-arrival localization algorithms, our schemes are unaware of the correspondence between the APs and the MPCs, and do not know a priori whether an MPC corresponds to a LoS or an NLoS propagation path. Enacting procedures to estimate this correspondence and localize a node is a distinctive feature of our algorithms. Such procedures have been designed to have low complexity, thus avoiding exhaustive search or maximum likelihood approaches in the design of both TV and ADoA. We collect the MPCs related to every AP n in a $2 \times N_{\mathbf{p}}^n$ matrix $\mathbf{M}^{(n,\mathbf{P})}$, where $N_{\mathbf{p}}^n$ is the number of detected MPCs in $P_{\mathbf{p}}^n(\alpha)$. The first row of $\mathbf{M}^{(n,\mathbf{P})}$ contains the amplitude of




Figure 1. Example scenario, including the geometry of angle difference-of-arrival localization.

each MPC sorted in decreasing order; the second row contains the AoA of the MPC, relative to the reference α_0 . In this way, each column $\mathbf{M}_{:,k}^{(n,\mathbf{p})}$ of $\check{\mathbf{M}}^{(n,\mathbf{p})}$ (where the colon notation $:,k$ denotes all elements of the corresponding dimension, and in this case it means all rows of column k) can be seen as a vector in polar coordinates, departing from \mathbf{p} , where $M_{1,k}^{(n,\mathbf{p})}$ and $M_{2,k}^{(n,\mathbf{p})}$ are the amplitude and phase of the vector relative to α_0 , respectively. Finally, we define $\check{\mathbf{M}}^{(\mathbf{p})}$ as the $2 \times \sum_{n=1}^{N_{\text{AP}}} N_{\mathbf{p}}^n$ concatenation of all matrices $\check{\mathbf{M}}^{(n,\mathbf{p})}$, re-sorted in order of decreasing MPC amplitude, and we call $\boldsymbol{\beta}$ the $1 \times \sum_{n=1}^{N_{\text{AP}}} N_{\mathbf{p}}^n$ vector whose i th element β_i indicates the index of the AP that generates the i th MPC in $\check{\mathbf{M}}^{(\mathbf{p})}$. For simplicity, we describe the algorithms developed in this paper for the case of node localization on the azimuthal plane, but their extension to the 3D case is straightforward.

A. Virtual anchors

The signals emitted by a given AP located at \mathbf{p}_{tx}^n generate an AoA spectrum $P_{\mathbf{p}}^n(\alpha)$ at location \mathbf{p} . Each MPC in this spectrum can be modeled as emitted by a virtual anchor that would be the source of a LoS signal reaching \mathbf{p} along the same AoA of the MPC. The position of the virtual anchor can be determined by mirroring the location of AP n with respect to the surfaces where its signal incurs a reflection. Call $\mathcal{A}^n = \{\mathbf{a}_0^n, \mathbf{a}_1^n, \dots\}$ the set containing the positions of the possible virtual anchors for AP n , and call $\bar{\mathcal{A}}^n = \{A_0^n, A_1^n, A_2^n, \dots\}$ a partition of the set $\mathcal{A}^n \cup \{\mathbf{p}_{\text{tx}}^n\}$. We let $A_0^n = \mathbf{p}_{\text{tx}}^n$, whereas each set A_i^n , $i = 1, 2, \dots$ contains the virtual anchors that have been mirrored i times with respect to any surface of the objects in \mathcal{Z} . Note that there is no limit to the number of times the AP can be mirrored. However, in practice the mmWave signal of the AP will fade quickly as it propagates and reflects off multiple surfaces: this is a substantially different aspect with respect to, e.g., UWB systems at lower frequencies, and allows us to truncate \mathcal{A}^n by considering a maximum reflection order μ [63]. To this end, we define $\mathcal{A}_\mu^n = \bigcup_{i=0}^{\mu} A_i^n$.

Fig. 1 shows an indoor scenario with two walls (thick gray lines) and two physical APs at locations \mathbf{p}_{tx}^1 and \mathbf{p}_{tx}^2 , corresponding to anchors \mathbf{a}_1 and \mathbf{a}_2 , respectively. Two virtual

anchors \mathbf{a}_3 and \mathbf{a}_4 correspond to a first- and a second-order reflection, respectively, i.e., $\mathbf{a}_3 \in A_1^1$ and $\mathbf{a}_4 \in A_2^1$.

B. The Triangulate-Validate (TV) algorithm

With this algorithm, a node at an unknown position \mathbf{p} estimates its position via a number of triangulation steps followed by a validation of the estimated locations. The algorithm assumes that the mmWave client has measured the AoA spectrum $P_{\mathbf{p}}^n(\alpha)$ for each AP n , and has derived the matrix $\check{\mathbf{M}}^{(\mathbf{p})}$ and the vector $\boldsymbol{\beta}$ defined above. We recall that the knowledge of the AP locations $\{\mathbf{p}_{\text{tx}}^n\}_{n=1}^{N_{\text{AP}}}$, of the obstacle set \mathcal{Z} and of the reference angle α_0 is also assumed. If the association between the anchors in each set \mathcal{A}^n and the MPCs in $\check{\mathbf{M}}^{(\mathbf{p})}$ were known, it would be possible to directly triangulate the position of \mathbf{p} . However, we only assume to know the ID of the APs from which the MPCs in $\check{\mathbf{M}}^{(\mathbf{p})}$ originate, but not the physical or virtual anchor they are associated with. Hence, we have to estimate this association via the procedure explained in the following, which is designed to be significantly less complex than an ML approach [51].

With reference to the pseudo-code in Algorithm 1, we start by considering virtual anchors from all APs up to a reflection order μ (line 3). While a high value of μ would yield a richer virtual anchor set \mathcal{A}_μ , a low value is more meaningful for triangulation: in fact, reflections weaken the signal, and virtual anchors of higher order can be quite far from the receiver. In turn, this distance would translate into a large triangulation error, even for small angle errors in the AoA spectrum. In addition, a signal at mmWave frequencies is usually so weak after more than two reflections that it cannot be decoded. Hence, we set $\mu=2$ in the following.

We start by considering $\check{\mathbf{M}}_{:,1}^{(\mathbf{p})}$ and $\check{\mathbf{M}}_{:,2}^{(\mathbf{p})}$ which, due to the sorting of $\check{\mathbf{M}}^{(\mathbf{p})}$, correspond to the MPCs with highest amplitude. Before using them to triangulate a position, we need to transform these entries into vectors *departing* from the position of any anchor, expressed relative to the reference Cartesian coordinate system of the room. This yields two vectors $\mathbf{u}_1 = -\mathbf{Q}\check{\mathbf{M}}_{:,1}^{(\mathbf{p})}$ and $\mathbf{u}_2 = -\mathbf{Q}\check{\mathbf{M}}_{:,2}^{(\mathbf{p})}$, where \mathbf{Q} is the coordinate transformation (line 4). We now make an initial guess that the anchors from which \mathbf{u}_1 and \mathbf{u}_2 emanate are two points $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{A}_\mu$. Given that there exist N_{AP} APs in the environment, each of which can be associated to multiple virtual anchors depending on the cardinality of set \mathcal{Z} , we exploit the ID of the AP generating a virtual anchor in $\boldsymbol{\beta}$ by imposing that $\mathbf{a}_i \in \mathcal{A}_\mu^{\beta_i}$ and $\mathbf{a}_j \in \mathcal{A}_\mu^{\beta_j}$ (line 6). We then triangulate a location by solving the following linear system in two unknowns $t_1 \geq 0$ and $t_2 \geq 0$ (line 7):

$$\mathbf{a}_i + \mathbf{u}_1 t_1 = \mathbf{a}_j + \mathbf{u}_2 t_2. \quad (2)$$

Call \mathbf{p}_k , $k \geq 0$ the position found. If \mathbf{p}_k is valid with respect

to the floor map constraints (line 8),¹ we validate the position by measuring how compatible the remaining MPCs are with the positions of other virtual anchors in \mathcal{A}_μ . We assign a weight $w_k > w_\ell$ to all anchors of partition subsets A_k^n and A_ℓ^n , $0 \leq \ell < k \leq \mu$, $A_k^n, A_\ell^n \subset \mathcal{A}_\mu$, for all APs n . This implements the consideration that the validations involving virtual anchors closer to \mathbf{p} should be given greater importance. We now choose v_{\max} further MPCs $\check{\mathbf{M}}_{:,m}^{(\mathbf{p})}$, $3 \leq m \leq v_{\max} + 2$, to be involved in the validation process. For each MPC m , we consider all virtual anchors in $\mathcal{T} = \mathcal{A}_\mu \setminus \{\mathbf{a}_i, \mathbf{a}_j\}$ and associate a cost c_k to \mathbf{p}_k as follows:

$$c_k = \sum_{m=3}^{v_{\max}+2} \min_{\mathbf{a} \in \mathcal{T} \cap \mathcal{A}_\mu^{\beta_m}} \left[\cos^{-1} \left(-\mathbf{u}_m \cdot \frac{\mathbf{a} - \mathbf{p}_k}{\|\mathbf{a} - \mathbf{p}_k\|} \right) \right]^2 w_{\omega(\mathbf{a})}, \quad (3)$$

where $\mathbf{u}_m = -\check{\mathbf{Q}}\check{\mathbf{M}}_{:,m}^{(\mathbf{p})}$ and $\omega(\mathbf{a}) = \ell$ if $\mathbf{a} \in A_\ell^n$ for some AP index n . We note that for a given MPC m , the argument of the sum is identically zero if there exists an anchor $\mathbf{a} \in \mathcal{T}$ that lies exactly on the line leaving \mathbf{p}_k with direction $\check{\mathbf{M}}_{:,m}^{(\mathbf{p})}$, whereas it increases if the minimum angle, centered in \mathbf{p}_k , between any anchor $\mathbf{a} \in \mathcal{T}$ and $\check{\mathbf{M}}_{:,m}^{(\mathbf{p})}$ increases. The squaring operation penalizes large discrepancies (lines 9 to 15). After computing the argument of the min in (3), the anchor \mathbf{a} that minimizes the argument is removed from \mathcal{T} (line 14).

The TV steps are repeated for all possible associations of $\check{\mathbf{M}}_{:,1}^{(\mathbf{p})}$ and $\check{\mathbf{M}}_{:,2}^{(\mathbf{p})}$ to the anchors in the set \mathcal{A}_μ , returning a total of K estimates \mathbf{p}_k , $k = 1, \dots, K$, and their related costs c_k (line 6). Note that $K \leq |\mathcal{A}_\mu|(|\mathcal{A}_\mu| - 1)$ since the algorithm fails if a triangulated position is found to be outside the room, or if the ensemble of the AoA spectra received from all APs, $\check{\mathbf{M}}^{(\mathbf{p})}$, contains fewer than 3 MPCs. The final estimate of \mathbf{p} returned by the TV algorithm is the one with the minimum associated cost (line 16). In lines 13 and 14, the function $\text{COST}(\mathbf{p}_k, \mathbf{a}, \mathbf{u}_m)$, returns the argument of the min function in (3). We remark that if $N_{\text{AP}} = 1$, Algorithm 1 falls back to the triangulate-validate algorithm presented in [1].

C. The Angle Differences-of-Arrival (ADoA) algorithm

The TV algorithm requires the knowledge of the reference angle α_0 , or equivalently, of the coordinate transformation matrix \mathbf{Q} introduced in Section III-B. As this assumption cannot always be met, and the measurement of α_0 (e.g., as provided by a smartphone's digital compass) may be affected by a significant error, we develop a second algorithm based on the Angle Differences-of-Arrival (ADoA) among MPCs. This algorithm is slightly more complex than TV, but is immune both to errors in α_0 and to variations thereof across the room area.

We start by defining the angles $\delta_1 = \check{\mathbf{M}}_{2,2}^{(\mathbf{p})} - \check{\mathbf{M}}_{2,1}^{(\mathbf{p})}$ and $\delta_2 = \check{\mathbf{M}}_{2,3}^{(\mathbf{p})} - \check{\mathbf{M}}_{2,1}^{(\mathbf{p})}$. With reference to Fig. 1, that depicts

¹The function ISVALID can be computed based on the winding number $\mathcal{W}_z(\mathbf{p})$, which equals n if curve z revolves around \mathbf{p} n times in the counterclockwise direction, and 0 if it does not enclose \mathbf{p} . For indoor localization, we can assume without loss of generality that the curves $z \in \mathcal{Z}$ are traveled in the counterclockwise direction by their parametric equations $z = z(t)$, $0 \leq t \leq 1$, and that they have no loops. The logical function ISVALID can then be defined as $\text{ISVALID}(\mathbf{p}, \mathcal{Z}) = \mathbb{1}[\sum_{z \in \mathcal{Z}} \mathcal{W}_z(\mathbf{p}) = 1]$, where $\mathbb{1}[\mathbf{p}]$ is True if predicate \mathbf{p} is true, and the sum is identically 1 only if the client is inside the room, but outside any obstacle.

Algorithm 1: The Triangulate-Validate algorithm.

```

1 Function TRIANGVAL ( $P_{\mathbf{p}}^n(\alpha) \forall n, \{\mathbf{p}_{\text{tx}}^n\}_{n=1}^{N_{\text{AP}}}, \mathcal{Z}, \alpha_0, \mu,$ 
 $v_{\max}, \beta)$ 
2    $\mathcal{A}_\mu^n \leftarrow \bigcup_{\ell=0}^{\mu} \text{DETERMINEANCHORS}(\ell, \mathbf{p}_{\text{tx}}^n, \mathcal{Z})$ 
3    $\mathcal{A}_\mu \leftarrow \bigcup_{n=1}^{N_{\text{AP}}} \mathcal{A}_\mu^n$ 
4   Map  $\check{\mathbf{M}}^{(\mathbf{p})}$  to canonical base, compute  $\mathbf{u}_1, \mathbf{u}_2$ 
5    $k \leftarrow 0$ 
6   foreach pair  $(\mathbf{a}_i, \mathbf{a}_j), \mathbf{a}_i \in \mathcal{A}_\mu^{\beta_i}, \mathbf{a}_j \in \mathcal{A}_\mu^{\beta_j}, \mathbf{a}_i \neq \mathbf{a}_j$  do
7      $\mathbf{p}_k \leftarrow$  point such that  $\mathbf{a}_i + \mathbf{u}_1 t_1 = \mathbf{a}_j + \mathbf{u}_2 t_2$ 
8     if ISVALID( $\mathbf{p}_k, \mathcal{Z}$ ) then
9        $\mathcal{T} \leftarrow \mathcal{A}_\mu \setminus \{\mathbf{a}_i, \mathbf{a}_j\}$ 
10       $c_k \leftarrow 0$ 
11      for  $m = 3$  to  $v_{\max} + 2$  do
12         $\mathbf{u}_m \leftarrow -\check{\mathbf{Q}}\check{\mathbf{M}}_{:,m}^{(\mathbf{p})}$ 
13         $c_k \leftarrow c_k + \min_{\mathbf{a} \in \mathcal{T} \cap \mathcal{A}_\mu^{\beta_m}} \text{COST}(\mathbf{p}_k, \mathbf{a}, \mathbf{u}_m)$ 
14         $\mathcal{T} \leftarrow$ 
15           $\mathcal{T} \setminus \{\arg \min_{\mathbf{a} \in \mathcal{T} \cap \mathcal{A}_\mu^{\beta_m}} \text{COST}(\mathbf{p}_k, \mathbf{a}, \mathbf{u}_m)\}$ 
16       $k \leftarrow k + 1$ 
17   return  $\hat{\mathbf{p}} = \mathbf{p}_{\arg \min_k (c_k)}$ 

```

a typical ADoA localization scenario, the ADoA algorithm is described as follows: given two points \mathbf{a}_i and \mathbf{a}_j , find the locus of the points \mathbf{p} such that the angle $\widehat{\mathbf{a}_i \mathbf{p} \mathbf{a}_j}$ (where \mathbf{p} is the corner), is constant and equal to the angle δ_1 defined above. This locus is a circumference, of which the segment $\widehat{\mathbf{a}_i \mathbf{a}_j}$ is a chord. The client is then located on this circumference. We assume the angle $\widehat{\mathbf{a}_i \mathbf{p} \mathbf{a}_j}$ to be positive if \mathbf{a}_i follows \mathbf{a}_j in the counterclockwise direction within the space of a semi-circumference and $|\widehat{\mathbf{a}_i \mathbf{p} \mathbf{a}_j}| < \pi$, where the $|\cdot|$ operator represents the absolute value of the angle measure. In Fig. 1, consider the circumference C_1 . Given the angle $\widehat{\mathbf{a}_1 \mathbf{p} \mathbf{a}_2}$, we have $\widehat{\mathbf{a}_1 \mathbf{o}_1 \mathbf{a}_2} = 2 \widehat{\mathbf{a}_1 \mathbf{p} \mathbf{a}_2}$, where $\widehat{\mathbf{a}_1 \mathbf{o}_1 \mathbf{a}_2}$ is the central angle that stands on the same chord $\widehat{\mathbf{a}_1 \mathbf{a}_2}$. If $\widehat{\mathbf{a}_1 \mathbf{p} \mathbf{a}_2} > \pi/2$, $\widehat{\mathbf{a}_1 \mathbf{p} \mathbf{a}_2}$ is concave, hence $\widehat{\mathbf{a}_1 \mathbf{o}_1 \mathbf{a}_2} > \pi$. In this case, we wrap the angle back into the interval $[-\pi, \pi)$ via the operator $W(\cdot)$. The radius r_1 of C_1 can be found as

$$r_1 = \frac{\|\mathbf{a}_1 - \mathbf{a}_2\|}{2 \sin \left(|W(\widehat{\mathbf{a}_1 \mathbf{o}_1 \mathbf{a}_2})|/2 \right)}. \quad (4)$$

Finally, given $\zeta = r_1 \cos(W(\widehat{\mathbf{a}_1 \mathbf{o}_1 \mathbf{a}_2})/2)$ and $\mathbf{u} = (\mathbf{a} - \mathbf{b}) \times (0, 0, 1)$, the center of the circumference C_1 is found as

$$\mathbf{o}_1 = \mathbf{m}_1 + \text{sgn} \left(W(\widehat{\mathbf{a}_1 \mathbf{o}_1 \mathbf{a}_2}) \right) \frac{\zeta \mathbf{u}}{\|\mathbf{u}\|}, \quad (5)$$

where sgn returns the angle sign, and $\mathbf{m}_1 = (\mathbf{a}_1 + \mathbf{a}_2)/2$ is the middle point of $\widehat{\mathbf{a}_1 \mathbf{a}_2}$. The vector

$$\mathbf{w}_1 = (\mathbf{o}_1 - \mathbf{m}_1) \text{sgn} \left(W(\widehat{\mathbf{a}_1 \mathbf{o}_1 \mathbf{a}_2}) \right) \quad (6)$$

points to the section of C_1 defined by the chord $\widehat{\mathbf{a}_1 \mathbf{a}_2}$ where the constant angle requirement is satisfied. We now consider a second chord $\widehat{\mathbf{a}_1 \mathbf{a}_3}$ and the circumference C_2 as the locus of the points \mathbf{q} where the angle $\widehat{\mathbf{a}_1 \mathbf{q} \mathbf{a}_3} = \delta_2$. The radius r_2 and the center \mathbf{o}_2 of C_2 can be computed in the same way as

Algorithm 2: The Angle Difference-of-Arrival algorithm.

```

1 Function ADOA( $P(\alpha)$ ,  $\{\mathbf{p}_{\text{tx}}^n\}_{n=1}^{N_{\text{AP}}}$ ,  $\mathcal{Z}$ ,  $\mu$ ,  $v_{\text{max}}$ ,  $\beta$ )
2    $\mathcal{A}_\mu^a \leftarrow \bigcup_{\ell=0}^\mu \text{DETERMINEANCHORS}(\ell, \mathbf{p}_{\text{tx}}^a, \mathcal{Z})$ 
3    $\mathcal{A}_\mu \leftarrow \bigcup_{n=1}^{N_{\text{AP}}} \mathcal{A}_\mu^n$ 
4   Map  $\check{\mathbf{M}}^{(\mathbf{p})}$  to canonical base, compute  $\delta_1, \delta_2$ 
5    $\ell \leftarrow 0$ 
6   foreach triple
   ( $\mathbf{a}_i, \mathbf{a}_j, \mathbf{a}_k$ ),  $\mathbf{a}_i \in \mathcal{A}_\mu^{\beta_i}, \mathbf{a}_j \in \mathcal{A}_\mu^{\beta_j}, \mathbf{a}_k \in \mathcal{A}_\mu^{\beta_k}, i \neq j \neq k$ 
7   do
    $C_1 \leftarrow \text{DETERMINECIRC}(i, j, \delta_1)$ 
    $C_2 \leftarrow \text{DETERMINECIRC}(i, k, \delta_2)$ 
    $\mathbf{p}_\ell \leftarrow \text{DETERMINEINTERSEC}(C_1, C_2) \setminus \{\mathbf{p}_{\text{tx}}^n\}_{n=1}^{N_{\text{AP}}}$ 
   if ISVALID( $\mathbf{p}_\ell, \mathcal{Z}$ ) then
    $\mathcal{T} \leftarrow \mathcal{A}_\mu \setminus \{\mathbf{a}_i, \mathbf{a}_j, \mathbf{a}_k\}$ 
    $c_\ell \leftarrow 0$ 
   for  $m = 4$  to  $v_{\text{max}} + 3$  do
    $c_\ell \leftarrow$ 
    $c_\ell + \min_{\mathbf{a} \in \mathcal{T} \cap \mathcal{A}_\mu^{\beta_m}} \text{COST}(\mathbf{p}_\ell, \mathbf{a}, \mathbf{a}_i, \delta_m)$ 
    $\mathcal{T} \leftarrow \mathcal{T} \setminus$ 
    $\{\arg \min_{\mathbf{a} \in \mathcal{T} \cap \mathcal{A}_\mu^{\beta_m}} \text{COST}(\mathbf{p}_\ell, \mathbf{a}, \mathbf{a}_i, \delta_m)\}$ 
    $\ell \leftarrow \ell + 1$ 
8   return  $\hat{\mathbf{p}} = \mathbf{p}_{\arg \min_\ell (c_\ell)}$ 

```

above. The intersection \mathbf{p} between C_1 and C_2 is considered a feasible location estimate whenever it is located on the sections of C_1 and C_2 pointed to by the orientation vectors \mathbf{w}_1 and \mathbf{w}_2 . This means that, given the two chord centers \mathbf{m}_1 and \mathbf{m}_2 , it must hold that $\mathbf{w}_1 \cdot (\mathbf{p} - \mathbf{m}_1) > 0$ and $\mathbf{w}_2 \cdot (\mathbf{p} - \mathbf{m}_2) > 0$.

The pseudo-code of the procedure that provides an estimate $\hat{\mathbf{p}}$ of the location of a node based on ADoA is given in Algorithm 2. The algorithm starts by determining $\mathcal{A}_\mu, \check{\mathbf{M}}^{(\mathbf{p})}, \delta_1$ and δ_2 (lines 3 and 4). As with the Triangulate-Validate algorithm, we are initially unable to map each detected MPC to its related anchor node. Therefore, the ADoA algorithm collects a set of eligible positions which are the result of the intersection between the circumferences determined by the angle differences δ_1 and δ_2 and the chords $\overline{\mathbf{a}_i \mathbf{a}_j}$ and $\overline{\mathbf{a}_j \mathbf{a}_k}$, $\mathbf{a}_i, \mathbf{a}_j, \mathbf{a}_k \in \mathcal{A}_\mu$. Lines 7 to 9 determine the center and radius of C_1 and C_2 via (4) and (5), and compute their intersection \mathbf{p}_ℓ by checking that it is feasible based on the orientation vectors \mathbf{w}_1 and \mathbf{w}_2 (see Eq. (6)). At this point, as for the TV algorithm, we pick v_{max} additional anchors $m = 4, \dots, v_{\text{max}} + 3$ from $\check{\mathbf{M}}^{(\mathbf{p})}$ (line 13), and for each of them we compute $\delta_m = \check{\mathbf{M}}_{2,m}^{(\mathbf{p})} - \check{\mathbf{M}}_{2,1}^{(\mathbf{p})}$. The contribution of MPC m to the cost of location \mathbf{p}_ℓ (lines 14 and 15) then depends on how well an anchor among those in set $\mathcal{T} = \mathcal{A}_\mu \setminus \{\mathbf{a}_i, \mathbf{a}_j, \mathbf{a}_k\}$ approaches the angle difference δ_m . This cost is defined as follows:

$$\text{COST}(\mathbf{p}_\ell, \mathbf{a}, \mathbf{a}_i, \delta_m) = (|\widehat{\mathbf{ap}_k \mathbf{a}_i}| - \delta_m)^2 \quad (7)$$

The ADoA procedure is repeated for different triples $(\mathbf{a}_i, \mathbf{a}_j, \mathbf{a}_k) \in \mathcal{A}_\mu$ (line 6), and the position with the minimum associated cost is returned as the node location estimate (line 17).

Algorithm 3: ADoA fingerprinting-based localization.

```

1 Function FINGERPRINT( $P(\alpha)$ ,  $\mathcal{D}^n$ ,  $\beta$ )
2   Compute  $\mathbf{M}^{(n,\mathbf{p})}$ , extract  $\mathbf{T}^{(n,\mathbf{p})}$ 
3   foreach  $f \in \mathcal{F}$  do
4      $c_f \leftarrow 0$ 
5     foreach  $1 \leq n \leq N_{\text{AP}}$  do
6        $\mathcal{P}^{(n)} \leftarrow$ 
        $\text{CLOSESTPAIRS}(\mathbf{T}_{2,:}^{(n,\mathbf{p})}, \mathbf{T}_{2,:}^{(n,\mathbf{f})} - \mathbf{T}_{2,1}^{(\beta_1,\mathbf{f})})$ 
7       foreach pair  $(i, j) \in \mathcal{P}^{(n)}$  do
8          $c_f \leftarrow$ 
          $c_f + \text{MPCCOST}(\mathbf{T}_{:,i}^{(n,\mathbf{p})}, \mathbf{T}_{:,j}^{(n,\mathbf{f})} - \mathbf{T}_{2,1}^{(\beta_1,\mathbf{f})})$ 
9   return  $\hat{\mathbf{p}} \leftarrow \arg \min_f c_f$ 

```

D. Benchmark: Localization based on ADoA fingerprinting (FP)

We now present a fingerprint-based localization algorithm. Unlike TV and ADoA, FP does not strictly require that \mathbf{p}_{tx}^n , $n = 1, \dots, N_{\text{AP}}$ and \mathcal{Z} are known to the node to be localized. We assume that the area has been previously characterized by creating a database \mathcal{D}^n of AoA spectra for all APs, measured at a set of different locations $\mathcal{F} = \{\mathbf{f}_1, \mathbf{f}_2, \dots\}$. Given a spectrum $P_{\mathbf{p}}^n(\alpha)$ related to AP n and measured at location \mathbf{p} , the algorithm looks up the most similar spectrum in \mathcal{D}^n (according to some proximity measure) and returns its corresponding location. In order to be fair to the TV and ADoA schemes, we assume we can identify the LoS path, and define the fingerprint of $P_{\mathbf{p}}^n(\alpha)$ in terms of the ADoA of any NLoS path, computed relative to AoA of the LoS path. This choice is based on the fact that the strongest MPC is typically the LoS path [7], [63], and that we wish to make the feature matching process resilient to position-dependent errors in the reference angle α_0 . Given the pattern $P_{\mathbf{p}}^n(\alpha)$ and its MPC matrix $\mathbf{M}^{(n,\mathbf{p})}$, call $\mathbf{T}^{(n,\mathbf{p})}$ the NLoS feature matrix that collects the last $N_{\mathbf{p}}^n - 1$ rows of $\mathbf{M}^{(n,\mathbf{p})}$, where $N_{\mathbf{p}}^n$ is the number of rows of $\mathbf{M}^{(n,\mathbf{p})}$. From all elements of $\mathbf{T}_{2,:}^{(n,\mathbf{p})}$, we finally subtract the angle of the LoS path $\mathbf{M}_{2,1}^{(n,\mathbf{p})}$, in order to create a vector of ADoA NLoS features, which is finally used as a fingerprint. Note that with this definition, fingerprints are also immune to a non-zero bias on the reference angle α_0 . In the pseudo-code in Algorithm 3, the above step corresponds to line 2.

The AoA spectra received at any two different points typically have one LoS feature, and may have a different number of NLoS features. For each entry of \mathcal{D}^n measured at location \mathbf{f} (line 3), we execute an adapted instance of the closest point algorithm [64] which returns the indices i and j of the NLoS features in $\mathbf{T}^{(n,\mathbf{p})}$ and $\mathbf{T}^{(n,\mathbf{f})}$, respectively, whose angle differences are most similar. The similarity of the patterns at \mathbf{p} and \mathbf{f} is finally conveyed by the cost

$$c_f = - \sum_{n=1}^{N_{\text{AP}}} \sum_{\forall (i,j) \in \mathcal{P}^{(n)}} |\mathbf{T}_{2,i}^{(n,\mathbf{p})} - (\mathbf{T}_{2,j}^{(n,\mathbf{f})} - \mathbf{T}_{2,1}^{(\beta_1,\mathbf{f})})|, \quad (8)$$

The function MPCCOST repeatedly called within the cycle at




Figure 2. Simulation scenario.

lines 7 and 8 computes the inner term of the sum in (8). These steps are repeated separately for each AoA spectrum corresponding to each of the APs. Finally, the algorithm returns the estimated position $\hat{\mathbf{p}}$ as the the location of the entry $\mathbf{f} \in \mathcal{F}$ that has minimum cost (line 9).

E. Benchmark: a Convex-Combination based AoA Localization (CCAL)

We now describe the Convex-Combination based AoA Localization (CCAL) scheme [62]. This algorithm works based only on AoA information and exploits the presence of virtual anchors, hence it constitutes a good benchmark for our localization algorithms. CCAL develops from the observation that, in the presence of noisy angle measurements, client localization can be casted as an optimization problem with a non-convex objective function. The authors then propose a linear approximation to the objective function that makes the localization problem amenable to be solved. The use of virtual anchors makes it possible to bound the area where the source can be located in the presence of angle errors.

F. Complexity of the algorithms

The TV algorithm loops twice over the set of anchors and once over the total combination of AoAs and anchors in the MPC set. Therefore its complexity is $\mathcal{O}(\max_n |\mathcal{A}^n|^3 \cdot v_{\max})$, where $\max_n |\mathcal{A}^n|$ is the maximum number of physical and virtual anchors associated to any AP. The ADoA algorithm loops three times over the number of anchors and once over the number of AoAs and anchors in the MPC set. Its complexity is $\mathcal{O}(\max_n |\mathcal{A}^n|^4 \cdot v_{\max})$. FP cycles once over the measurements data set, once over the APs, and once to find the closest pairs of MPCs, which leads to a complexity of $\mathcal{O}(|\mathcal{D}| \cdot \max_n |\mathcal{A}^n|^2)$. Finally, CCAL needs to solve a convexified quadratic problem with $\mathcal{O}(N_{\text{AP}}^2)$ variables and $\mathcal{O}(N_{\text{AP}}^2)$ constraints, for which there exists a polynomial-time interior point algorithm. CCAL's complexity is thus $\mathcal{O}(N_{\text{AP}}^{12})$. We remark that the algorithmic complexity of TV, ADoA and FP is perfectly acceptable in typical indoor mmWave deployments.




Figure 3. Statistics of the localization error for all algorithms as a function of the number of APs, N_{AP} . The solid boxes indicate the quartiles of the error distribution, the dashed error bars convey the 10th and 90th percentiles. $\sigma = 2^\circ$, $\Delta\alpha = 0$.

IV. NUMERICAL RESULTS

A. Simulation setup and scenario

In this section, we evaluate the performance of the proposed localization algorithms by means of simulation. We proceed by first describing an indoor space representing the network area as the input to a ray tracer; this includes the location of the APs and of the client, as well as the location of all walls. Second, for each position of the client, we run the ray tracer to simulate indoor mmWave propagation, and thereby retrieve the AoA spectra. Third, we run the localization algorithms of Section III using these AoA spectra and evaluate the localization error with respect to the true position of the client. We remark that no change is required to the algorithms with respect to the description in Section III.

The ray tracer is the same as described in [65, Section VI-B], and deterministically reproduces all relevant mmWave propagation phenomena, including path loss, material- and frequency-dependent reflection losses, shadowing and blockage. The output of the ray tracer collects the rays that lead to an actual received power contribution at the receiver (or eigenrays). Such output is translated into the AoA spectrum $P_{\mathbf{p}}^n(\alpha)$ measured by the receiver at its location \mathbf{p} . The AoA spectrum is derived for every transmitting AP n , and is the input to all localization algorithms.

The simulation scenario consists of an L-shaped room, as depicted in Fig. 2. The vertical leg on the left is 5×15 m², whereas the top-right section is 5×5 m². We set $\alpha_0 = \pi/2$, i.e., angles are measured in a counterclockwise direction starting from the south-to-north direction. Up to four APs are located in the room at coordinates $\mathbf{p}_{\text{tx}}^A = (2, 13)$ m, $\mathbf{p}_{\text{tx}}^B = (8, 13)$ m, $\mathbf{p}_{\text{tx}}^C = (2, 2)$ m, and $\mathbf{p}_{\text{tx}}^D = (3.5, 8.5)$ m. In the following, when we set $N_{\text{AP}} = 1$, we activate only the AP at \mathbf{p}_{tx}^A . For $N_{\text{AP}} = 2$ we activate the APs at \mathbf{p}_{tx}^B and \mathbf{p}_{tx}^C ; for $N_{\text{AP}} = 3$ those at \mathbf{p}_{tx}^A , \mathbf{p}_{tx}^B , and \mathbf{p}_{tx}^C . Finally, for $N_{\text{AP}} = 4$, all of the four APs in the room will be active.

For all algorithms, we truncate \mathcal{A} to $\mathcal{A}_\mu = \bigcup_{n=1}^{N_{\text{AP}}} \mathcal{A}_\mu^n$ by considering only the virtual anchor nodes that correspond to up to $\mu = 1$ reflections. We recall that different APs and virtual anchors become visible to the user as it moves. The fingerprinting algorithm has been trained through a reasonably dense set of measurements \mathcal{D}^n for $n = 1, \dots, N_{\text{AP}}$, which

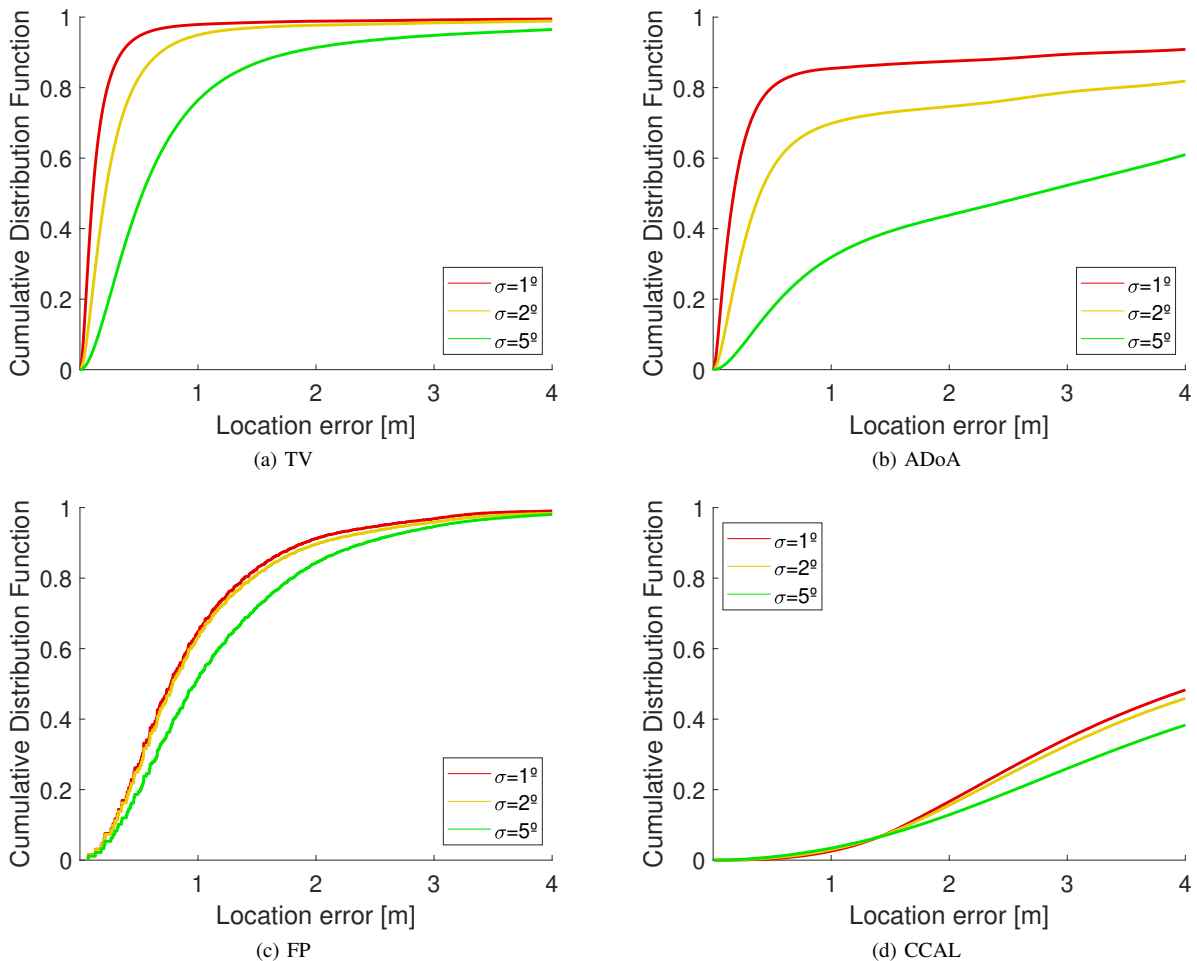


Figure 4. CDF of the localization error for TV, ADoA and FP as a function of the standard deviation of the MPC angle estimation error, σ . $N_{\text{AP}} = 2$, $\Delta\alpha = 0$.

includes a uniform matrix of 8×4 points in the vertical room section on the left, and 4×4 points over the top-right section.

We simulate the fact that the measurement of the AoA spectrum by the node to be localized is not ideal, but rather may be affected by errors $\Delta\alpha$ on the reference angle α_0 as well as on the AoA estimates of the MPCs, e.g., as a consequence of small changes in the propagation environment, or due to imperfect beamforming of the phased antenna array at the receiver. This non-ideal behavior is modeled for a given position \mathbf{p} by first estimating the true AoA spectrum, by extracting the MPC matrix $\mathbf{M}_{2,:}^{(n,\mathbf{p})}$ for any AP n , and by adding a different random Gaussian-distributed displacement of standard deviation σ to the elements of $\mathbf{M}_{2,:}^{(n,\mathbf{p})}$. To derive realistic MPC angle estimation errors, we synthesize beam patterns for a uniform linear antenna array with 8, 16 and 32 elements. For typical signal-to-noise ratios, these correspond to zero-mean Gaussian-distributed errors of standard deviation $\sigma = 5^\circ$, $\sigma = 2^\circ$, and $\sigma = 1^\circ$, respectively.

B. Simulation Results

We characterize the performance of the algorithms through the statistics of the localization error, as well as via the median location error experienced at different locations throughout the room. In our analysis, we explore the effect of different values

of N_{AP} , σ and of the bias on the reference angle α_0 , denoted as $\Delta\alpha$ and called compass bias for short in the following. In Fig. 3 we show a box plot for the localization error incurred by ADoA, FP and TV for different values of N_{AP} , for the case of $\sigma = 2^\circ$ and $\Delta\alpha = 0$.

We observe that even with one AP, TV, ADoA and FP achieve a sub-meter localization accuracy in at least 50% of the cases. However, since for $N_{\text{AP}} = 1$ the illumination of the room is sub-optimal, the dispersion of the error is significant for both ADoA and FP, whereas 95% of TV's errors are within less than 1.5 m. Conversely, CCAL does not perform well. While some positions experience small localization errors, these constitute a minor fraction of the cases, and the median error is about 8 m. The main reason behind this is that CCAL is not very robust to angle errors larger than those presented in [62], i.e., up to $\pm 0.1^\circ$. Increasing N_{AP} improves the mmWave illumination across the room and some of the algorithms benefit from this. TV's localization error decreases for increasing N_{AP} , and levels off for $N_{\text{AP}} \geq 3$. Since ADoA requires a larger number of MPCs in order to fix the client's location, it only reaches the good performance of TV for 4 APs. FP's performance is limited by the number of locations at which a training fingerprint is measured or, more precisely, by the fingerprint density. For this reason,

the statistics of FP’s localization error improve marginally for increasing N_{AP} . CCAL’s performance improves for $N_{\text{AP}} = 2$, but then the error increases if additional APs are included. The reason is that, with an increasing number of APs, the set of virtual anchors grows, breaking the assumptions behind the linear approximation of the angles in [62]. We remark that the above results cannot be achieved by the TV, ADoA, and FP algorithms originally presented in [1], as they cannot leverage the presence of multiple APs.

We now consider the effect of MPC estimation errors on the performance of the algorithms. Fig. 4 presents four graphs, one for each of the TV, ADoA, FP and CCAL algorithms. In each panel, we plot the cumulative distribution function (CDF) of the localization error, for $\sigma = 1^\circ$, $\sigma = 2^\circ$ and $\sigma = 5^\circ$. We remark that the latter represents a very unfavorable case, with 97% of the angle estimation errors comprised in an interval of $\pm 15^\circ$ around the true AoA of each MPC. For this analysis, we set $N_{\text{AP}} = 2$ (which is also the best case for CCAL) and $\Delta\alpha = 0$. From panel 4a, we observe that while the performance of TV decreases for increasing σ , the localization error statistics remain very good even for $\sigma = 5^\circ$: in this case, the probability of achieving a sub-meter localization error is about 0.8, and decreases to 0.93 and 0.99 for $\sigma = 2^\circ$ and $\sigma = 1^\circ$, respectively. The situation is slightly different with ADoA (panel 4b), as an error of variance σ^2 on the MPC angle estimation translates into an error of variance $2\sigma^2$ on angle differences. This adds to ADoA’s need to observe a higher number of MPCs than TV, and results in a higher maximum error. The probability of sub-meter accuracy for ADoA is 0.88 for $\sigma = 1^\circ$, but decreases to about 0.35 for $\sigma = 5^\circ$. The FP algorithm (panel 4c) behaves differently: depending on the value of σ , the probability of achieving a sub-meter localization accuracy ranges between 0.5 and 0.62. Specifically, for low values of σ , the limiting factor is the density of points in FP’s training data set; for $\sigma = 5^\circ$, instead, the effect of wrong MPC angle measurements dominates, and makes the AoA spectra measured by the client more similar to fingerprint database entries associated to farther positions. The results of CCAL (Fig. 4d) are aligned with those in Fig. 3, as the algorithm achieves sub-meter accuracy in less than 10% of the cases, and is characterized by a very high maximum error, above 4 m in at least 50% of the simulated estimates. Considering that FP requires to actually collect and maintain a training database, and given that FP’s performance is worse than TV’s and ADoA’s even for $\sigma = 1^\circ$, we conclude that FP is not a competitive algorithm for indoor mmWave localization. The same conclusion can be drawn for CCAL due to its unsatisfactory performance.

Finally, in Fig. 5 we evaluate the effect of a non-zero value of the compass bias $\Delta\alpha$ on the performance of TV for the case $N_{\text{AP}} = 3$, $\sigma = 2^\circ$. We remark that ADoA and FP are immune to non-zero values of $\Delta\alpha$ by design. Compared to Fig. 4a, the average and maximum errors achieved by TV increase. While 98% of the estimated locations are within 1 m of the true client location for $\Delta\alpha = 1^\circ$, for higher values of $\Delta\alpha$ we observe that the minimum error increases to about 30 cm, and the fraction of location estimates with sub-meter accuracy decreases to about 82%.

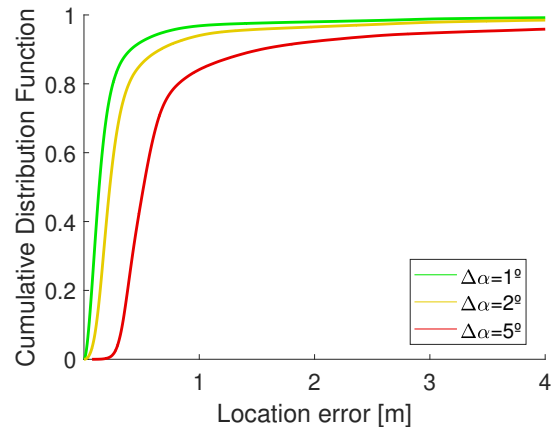


Figure 5. CDF of the localization error for TV as a function of the client’s compass bias, $\Delta\alpha$. $N_{\text{AP}} = 3$, $\sigma = 2^\circ$.

In order to investigate the source of the errors observed for the TV, ADoA, FP, and CCAL localization methods in our simulation scenarios, we depict in Fig. 6 a heat map of the median localization error throughout the simulation area for $N_{\text{AP}} = 2$ and $\Delta\alpha = 0^\circ$. Dark-blue hues represent near-zero error, whereas yellow hues represent high error. Fig. 6 reveals that TV achieves the lowest error throughout the map. A few areas of uncertainty exist for TV, mainly close to the walls. In these regions, the location of the client is estimated to be outside the boundaries of the area, therefore the validation step of the TV algorithm discards the estimate. For ADoA, we observe a generally very good performance, but high-error areas around both APs. These are the areas where the three strongest anchors and the client’s location tend to lie on the same circumference. As a result, the ADoA localization problem becomes ill-conditioned, and the localization error increases. The FP algorithm (Fig. 6c) also achieves a remarkably good error in this case. In some areas, this error increases significantly due to the similarity among arrival patterns throughout the room. Further results obtained with a four times-denser training set (not included due to lack of space) show that FP’s performance improves, but not substantially, leaving areas affected by significant errors close to the bottom and leftmost walls. CCAL’s error (Fig. 6d) shows that acceptable errors concentrate near the bottom-right and top-right corners of the L-shaped room: in these areas, the assumptions behind CCAL are verified and the linear approximation to its cost function is accurate. Otherwise, large errors occur as also noted above. These results confirm our preliminary conclusion that FP and CCAL are not preferred solutions for indoor mmWave localization.

While TV shows generally good performance under most conditions considered here and tends to outperform both ADoA and FP in a number of cases, it remains very sensitive to a biased reference angle α_0 , which instead does not affect ADoA and FP. We show this through the TV error maps in Fig. 7. As expected, we observe that the performance of TV decreases noticeably for increasing σ and $\Delta\alpha$: in fact, erroneous AoA MPCs translate into higher localization error for any triangulation-based method. For $\sigma = 2^\circ$, even a

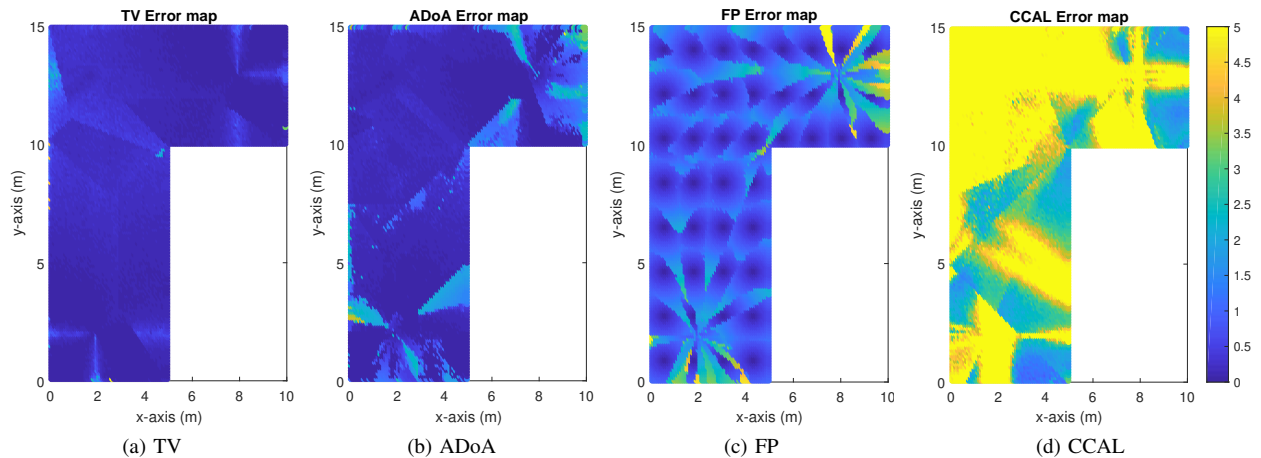


Figure 6. Median localization error maps for all methods, $\sigma = 2^\circ$, $N_{AP} = 2$, $\Delta\alpha = 0$.

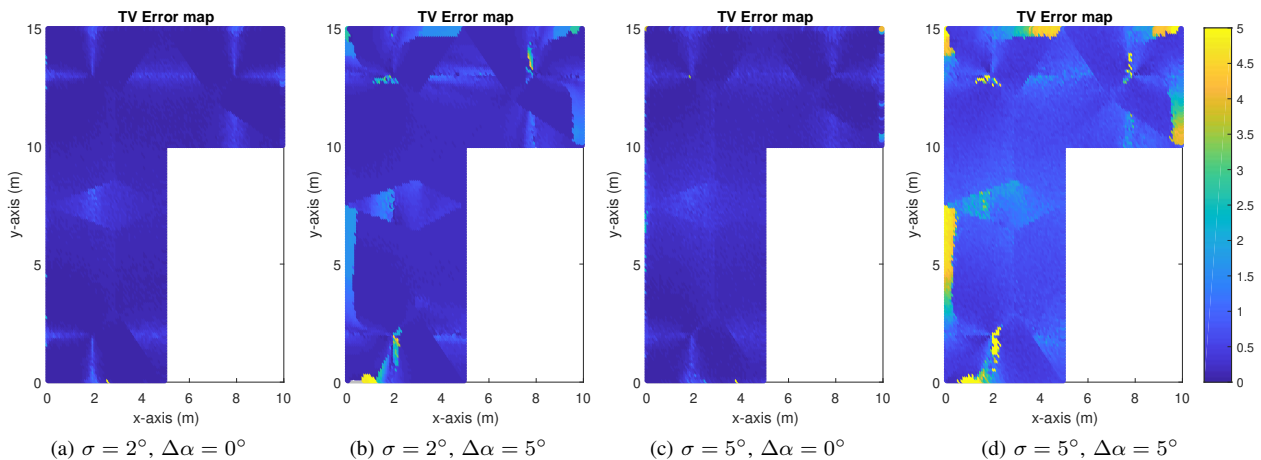


Figure 7. Median localization error maps for TV, for different values of σ and $\Delta\alpha$. $N_{AP} = 3$.

significant reference angle bias $\Delta\alpha = 5^\circ$ leads to acceptable errors, mostly below 1 m (Fig. 7b). For $\sigma = 5^\circ$, instead, the superposition of inaccurate AoA estimates and a high $\Delta\alpha$ implies that the error remains generally high, with small regions of sub-meter errors, mainly close to the three APs.

V. EXPERIMENTAL RESULTS

We now perform experiments with actual mmWave hardware in realistic propagation scenarios, in order to validate the simulation results described in Section IV. We focus only on TV and ADoA for our experimental results, due to the worse performance shown by CCAL and FP. Furthermore, collecting and updating a training database in a fully operational and time-varying environment involves a significant effort and is very cumbersome for practical deployments.

We consider two scenarios: an empty room of size 6.3×8.9 m with a single AP installed (also considered in [1] and re-evaluated here to compare our new algorithms); and a more complex, fully functional L-shaped office area with typical open-space furniture, including workstations and screens, as well as building elements such as columns and glass walls with metal frames. In the L-shaped room, we install up to 5 APs in order to achieve good coverage. The latter is

representative of dense mmWave AP deployments as are envisioned in the literature to provide adequate indoor mmWave performance [24]. The empty room was chosen so that the LoS mmWave signal can be received clearly at all locations in the room, whereas the L-shaped room is representative of an actual working environment with signal blockage, such that the localization algorithms have to rely more heavily on both the real and virtual anchors they can observe at a given time.

A. Empty room with a single AP

We start our evaluation with the empty room. The room is unfurnished, but does contain fixtures and appliances which cause reflections. This includes metal plumbing, heating radiators, wall-mounted rack cabinets, and metal floor-mounted cable leads. In one of the corners of the room, we install a Dell D5000 wireless docking station. This consumer-grade mmWave equipment operates according to the WiGig protocol, and creates a 60-GHz mmWave data link between a nearby computer and the peripherals attached to the station. To establish the link, the D5000 periodically carries out a device discovery process, involving the transmission of 32 discovery frames, each corresponding to a different configuration of the station's 2×8 phased antenna array. An example of the signal received during this process is shown in Fig. 8. These device




Figure 8. Dell D5000’s device discovery frames. Each frame corresponds to a different configuration of the station’s 2×8 phased antenna array [1], [66].

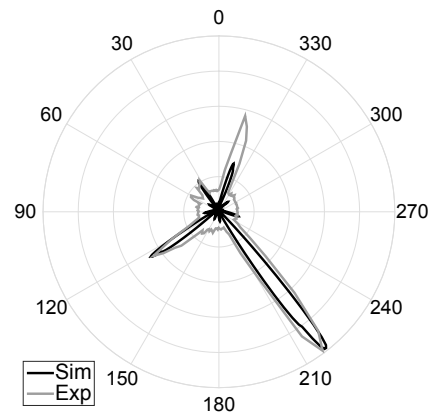


Figure 9. AoA spectrum in the empty room at coordinates (1.85, 5.75) m. (Adapted from [1].)

discovery frames are sent every 102.4 ms. For additional details, we refer the interested reader to [66].

For the receiver we use a 60-GHz VubIQ development system, connected to an Agilent MSO-X 3034 oscilloscope and equipped with a standard 25 dBi-gain horn antenna mounted on an Arduino-controlled rotating stage. The VubIQ reveals LoS paths as well as any MPC reflected off the boundaries of the room at any given position. We input the VubIQ’s down-converted, 1.8 GHz-bandwidth, analog modulated I/Q output signals directly into the oscilloscope, which digitizes and stores the traces for analysis. In this way, we collect the amplitude of each of the 32 discovery frames over the azimuthal plane. We control the rotating stage to take measurements at orientations spaced 3° apart. For every orientation, the receiver captures all 32 frames as in Fig. 8, where each frame is received with a different amplitude in each position, depending on the antenna patterns used by the D5000 for each frame and the mmWave signal propagation throughout the room. From these traces, an AoA spectrum is extracted by measuring the amplitude of the strongest frame. This replicates the behavior of a typical mmWave beam training handshake, where two devices would try several combinations of the beam pattern configurations available to them, until the best transmit and receive beam patterns are found. Fig. 9 shows one of the experimental patterns (in gray), measured at position (1.85, 5.75) m (see also Fig. 10 for a reference about the location). We observe that the LoS path between the D5000 and the VubIQ has an AoA of 215° . Strong reflections are detected with AoAs of 125° and 340° , and a weaker reflection has an AoA of 30° . For comparison, we also plot an AoA spectrum simulated via ray tracing (in black), which proves the very good agreement between ray tracing simulations and measurements.

We carried out AoA spectrum measurements at a total of 44 points throughout the room. Figs. 10a and 10b show the localization performance of TV and ADoA, respectively, for all of the measurement points. The color of each circle conveys the accuracy of the position estimate: blue hues correspond to low error, whereas green and yellow hues correspond to increasingly higher errors. A red cross in the bottom-right

corner of the room represents the location of the D5000 docking station, i.e., the mmWave AP. The validity of each position estimate is checked against the room boundaries (this corresponds to the `ISVALID(.)` function in Algorithms 1 and 2). The $30 \text{ cm} \times 30 \text{ cm}$ area surrounding the location of the AP is also marked as an invalid location.

In general, TV achieves better localization performance than ADoA, with most estimated positions within 1 m of the true client location. The largest localization errors are primarily due to the large widths of some arrivals in the measured AoA spectra, which make it difficult to accurately estimate the corresponding AoA. Specifically, for the locations near the transmitter, this is the case for the LoS AoA. For TV, wrong AoA estimates result in wrong triangulated positions, which are found to be outside valid room boundaries. For this reason, TV is unable to localize the node at seven locations out of 44: two close to the AP, two concentrating around the measurement point at (3.8, 2.5), and three at the opposite end of the room. As ADoA requires at least three MPCs in the AoA spectrum to work properly, the weakness of the received signal and the presence of mainly low-SNR second-order reflections in some measurements throughout the room explains why ADoA incurs a large error at such locations. Fig. 10c summarizes the results by showing the cumulative distribution function (CDF) of the localization error, for the points for which a valid position can be estimated. We observe that both TV and ADoA achieve a sub-meter median error. While the probability to compute a very accurate location estimate is similar for both algorithms, overall TV has significantly higher location accuracy: its probability to obtain a sub-meter location error is about 0.8, against about 0.5 for ADoA.

B. L-shaped working area with multiple APs

This second scenario is an L-shaped work place with an open area, six small offices and two labs, all of which are in active use. The walls are made of brick and have glass windows. Six columns are present in the open area. Glass panels divide the offices from the open area. Office and lab furniture (including tables, chairs, screens, as well as metal door frames and cupboards) make the propagation setting realistically complex.

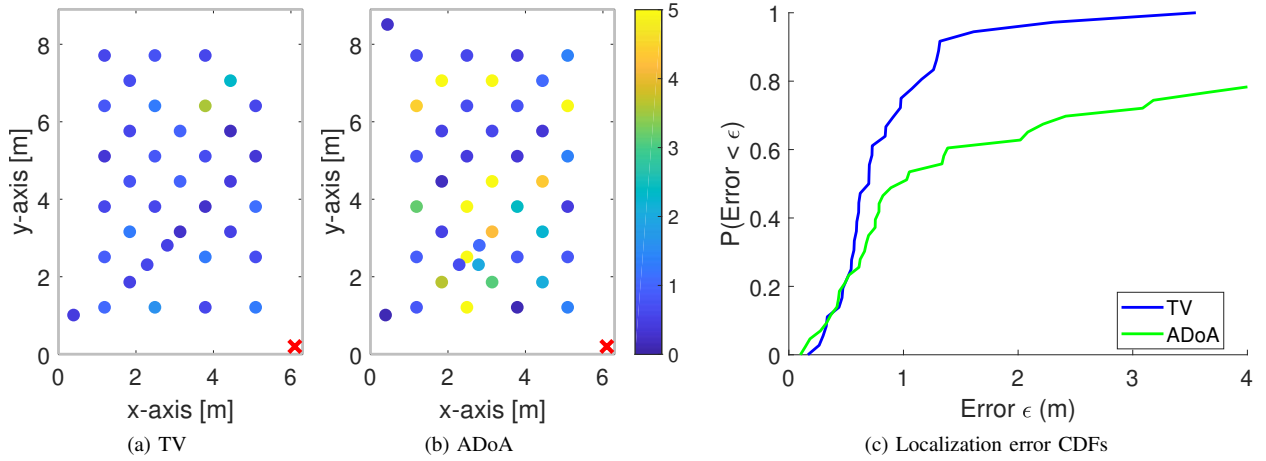


Figure 10. Map and CDF of the localization error of TV and ADoA in the empty room with a single AP.




Figure 11. Photo of the open space measurement setup of Fig. 12, showing the client, AP2 and AP5.




Figure 12. Multi-AP measurement setup in the L-shaped room area. Five different APs are deployed, each using one of three different antenna types (omnidirectional, 120°-aperture and 80°-aperture).

As the room is larger and contains different types of obstacles, a single AP cannot cover the whole area. We therefore investigate scenarios including up to five APs deployed throughout the room. We remark that using D5000 docking stations would make it difficult to differentiate among the signals transmitted by each AP. We solve this issue by using the following five transmitters: one Pasternak VubIQ, two SiversIMA DC1005V/00 and two SiversIMA CO2201A. In order to be able to separate each platform's signal at the receiver, we tune the local oscillator of each device to a different frequency within the 60 GHz band. Each transmitter is equipped with antennas of different beam widths, depending on the location of the transmitter, with the general objective of covering as much space as possible, so that the maximum number of transmitter signals can be seen and simultaneously sampled from every location in the room.

To emulate the user, we use the same Pasternak VubIQ receiver as before, equipped with a 7°, 25 dBi-gain horn antenna. However, instead of connecting the VubIQ output to the oscilloscope, we use an Agilent EXA N9010A signal analyzer, capable of recording the mmWave signal and of differentiating between the different frequencies of the trans-

mitters. We take measurements at angular steps of 0.45°, which covers the azimuthal plane with 800 measurements in total. We use pseudo omni-directional transmitters in order to speed up the collection of the measurements and capture all available paths of a given AP at the same time. As a general rule-of-thumb, we equipped APs located in open areas with an omni-directional antenna; where omni-directionality is not needed we employed wide-beam antennas (e.g., 80°-aperture horn antennas for APs near corners, and open waveguide terminations translating into a 120° antenna aperture for APs located near walls). After collecting the measurements for the different locations, we retrieve the AoA patterns for each AP by isolating the corresponding frequencies throughout all rotation steps. Different peaks in each pattern correspond to the LoS arrival (when available) and to one or more NLoS arrivals. Fig. 11 shows a view of the deployment.

Fig. 12 shows the deployment setup superimposed to the floor plan. Each AP has a unique identifier, and is depicted as a circle arc representing the aperture and direction of its antenna. The omni-directional transmitters were placed centrally in the open area, whereas the 80° horn antennas were positioned in the two west corners of the open area, and the 120°




Figure 13. CDF of the localization error for TV and ADoA for the multi-AP deployment in the L-shaped room, for different numbers of active APs.

transmitter in the lower east corner. We measured the AoA pattern from all visible APs at 66 different positions, with an average separation of 1.3 m between nearest measurement points. Note that the furniture constrained the measurements, which were taken around the work stations, in empty areas, and along the corridors. The opacity of glass to mmWaves and the resulting lack of coverage prevented us to do measurements in the offices and labs surrounding the open area.

We evaluate the performance of TV and ADoA in the above scenario for different numbers of active APs between 1 and 5. For the cases of 1 to 3 APs, we test the outcome of all possible choices of active and inactive APs, and select the combination that yields the best results. For example, for a single active AP, the best choice is AP 2 in Fig. 12, which provides the maximum area coverage compared to all other APs. For the case of two active APs, the best APs are AP 2 and AP 4, where the latter provides additional coverage especially around the top-left corner of the open area.

Fig. 13 shows the CDF of the localization error for the best combination of 1, 2, 3, and 5 active APs. As expected, when $N_{AP} = 1$ (only AP 2 is active), the localization error is good in some cases, but still significant: sub-meter accuracy is achieved in roughly 30% of the estimates for TV and 5% of the cases for ADoA, and the maximum error of both algorithms is large as well. Increasing the number of APs improves the mmWave coverage throughout the room, so that the accuracy of both TV and ADoA increases. For example, sub-meter accuracy is achieved by TV in about 60% of the cases with 2 active APs, and up to 85% of the cases with 5 APs. ADoA's performance increases similarly with additional active APs, up to a 70% probability of sub-meter accuracy with 5 active APs. This confirms that, in a realistic deployment, ADoA needs additional anchors (hence better mmWave coverage) in order to be able to compute the client location correctly.

In those locations where the localization outcome is inaccurate, the cause of the errors is typically traced back to insufficient mmWave illumination, as is the case in the

top left and rightmost sections of the room. In some cases, the validation step fails because localization errors make the predicted location exceed the room boundaries. Other possible sources of error are the alignment of anchors with the location of the users (which makes the triangulation problem ill-conditioned), or the case where the three strongest anchors and the client's location lie on the same circumference (which makes the ADoA problem ill-conditioned).

VI. CONCLUSIONS

In this paper, we developed two localization algorithms tailored to the characteristics of mmWave propagation: one based on a Triangulate-Validate (TV) procedure, and one based on Angle Difference-of-Arrival (ADoA) information. We characterized the performance through measurements involving commercial 60-GHz mmWave hardware, both in an empty room and in a fully operational work environment that was in use during the measurements. Our results show that unless a large error affects AoA estimates, our algorithms achieve a very good localization performance. Specifically, they achieve sub-meter accuracy with high probability in the presence of sufficient mmWave illumination. An extensive set of simulations that compare TV and ADoA against an algorithm based on fingerprinting and an ADoA and virtual anchor-based algorithm from the literature show that TV is in general a better choice, especially when the orientation of the client can be accurately estimated. Conversely, the ADoA algorithm is immune to any compass bias, and is thus a better choice if the device orientation is affected by large errors.

This work has been supported in part by the ERC project SEARCHLIGHT (grant no. 617721), the Ramon y Cajal grant RYC-2012-10788, the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919) and the Spanish Ministry of Science, Innovation and Universities the DiscoEdge grant (TIN2017-88749-R).

REFERENCES

- [1] A. Olivier, G. Bielsa, I. Tejado, M. Zorzi, J. Widmer, and P. Casari, "Lightweight indoor localization for 60 GHz millimeter wave systems," in *Proc. IEEE SECON*, 2016, pp. 1–9.
- [2] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-wave cellular wireless networks: Potentials and challenges," *Proc. IEEE*, vol. 102, no. 3, pp. 366–385, 2014.
- [3] A. Osseiran, F. Boccardi, V. Braun, K. Kusume, P. Marsch, M. Maternia, O. Queseth, M. Schellmann, H. Schotten, H. Taoka *et al.*, "Scenarios for 5G mobile and wireless communications: the vision of the METIS project," *IEEE Commun. Mag.*, vol. 52, no. 5, pp. 26–35, 2014.
- [4] F. Boccardi, R. W. Heath, A. Lozano, T. L. Marzetta, and P. Popovski, "Five disruptive technology directions for 5G," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 74–80, 2014.
- [5] L. Subrt, P. Pechac, and S. Zvanovec, "New approach to modeling of diffuse reflection and scattering for millimeter-wave systems in indoor scenarios," *PIERS Online*, vol. 6, no. 8, pp. 719–722, 2010.
- [6] K. Venugopal and R. W. Heath, "Millimeter wave networked wearables in dense indoor environments," *IEEE Access*, vol. 4, pp. 1205–1221, 2016.
- [7] A. Maltsev, R. Maslennikov, A. Sevastyanov, A. Khoryaev, and A. Lomayev, "Experimental investigations of 60 GHz WLAN systems in office environment," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1488–1499, Oct. 2009.
- [8] S. Geng, J. Kivinen, X. Zhao, and P. Vainikainen, "Millimeter-wave propagation channel characterization for short-range wireless communications," *IEEE Trans. Veh. Technol.*, vol. 58, no. 1, pp. 3–13, Jan. 2009.
- [9] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.
- [10] A. Yassin, Y. Nasser, M. Awad, A. Al-Dubai, R. Liu, C. Yuen, R. Raulefs, and E. Aboutanos, "Recent advances in indoor localization: A survey on theoretical approaches and applications," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 2, pp. 1327–1346, Apr. 2017.
- [11] K. Wu, P. Burasa, T. Djerafi, and N. Constantin, "Millimeter-wave identification for future sensing, tracking, positioning and communicating systems," in *Proc. G5MM*, Aalto, Finland, Jun. 2016, pp. 1–4.
- [12] K. Witrals, P. Meissner, E. Leitinger, Y. Shen, C. Gustafson, F. Tufveson, K. Haneda, D. Dardari, A. F. Molisch, A. Conti, and M. Z. Win, "High-accuracy localization for assisted living: 5G systems will turn multipath channels from foe to friend," *IEEE Signal Process. Mag.*, vol. 33, no. 2, pp. 59–70, Mar. 2016.
- [13] L. Barkhuus and A. K. Dey, "Location-based services for mobile telephony: a study of users' privacy concerns," in *Proc. Interact*, vol. 3, 2003, pp. 702–712.
- [14] D. Lymberopoulos and J. Liu, "The Microsoft indoor localization competition: Experiences and lessons learned," *IEEE Signal Process. Mag.*, vol. 34, no. 5, pp. 125–140, Sep. 2017.
- [15] D. Vasisht, S. Kumar, and D. Katabi, "Decimeter-level localization with a single WiFi access point," in *Proc. USENIX NSDI*, vol. 16, Santa Clara, CA, Mar. 2016, pp. 165–178.
- [16] D. Lymberopoulos, J. Liu, X. Yang, R. R. Choudhury, V. Handziski, and S. Sen, "A realistic evaluation and comparison of indoor location technologies: Experiences and lessons learned," in *Proc. IPSN*, 2015, pp. 178–189.
- [17] R. Reimann, A. Bestmann, and M. Ernst, "Locating technology for AAL applications with direction finding and distance measurement by narrow bandwidth phase analysis," in *Evaluating AAL Systems Through Competitive Benchmarking*, S. Chessa and S. Knauth, Eds. Springer, 2013, pp. 52–62.
- [18] C. Beder and M. Klepal, "Fingerprinting based localisation revisited: A rigorous approach for comparing rssi measurements coping with missed access points and differing antenna attenuations," in *Proc. IPIN*, Sydney, Australia, Nov. 2012, pp. 1–7.
- [19] H. Zou, Z. Chen, H. Jiang, L. Xie, and C. Spanos, "Accurate indoor localization and tracking using mobile phone inertial sensors, WiFi and iBeacon," in *Proc. IEEE INERTIAL*, Kauai, HI, 2017, pp. 1–4.
- [20] P. Meissner, E. Leitinger, M. Fröhle, and K. Witrals, "Accurate and robust indoor localization systems using ultra-wideband signals," in *Proc. ENC*, Vienna, Austria, Apr. 2012.
- [21] P. Corbalán and G. P. Picco, "Concurrent ranging in ultra-wideband radios: Experimental evidence, challenges, and opportunities," in *Proc. EWSN*, 2018, pp. 55–66.
- [22] Z. Xiao, P. Xia, and X.-G. Xia, "Enabling UAV cellular with millimeter-wave communication: potentials and approaches," *IEEE Commun. Mag.*, vol. 54, no. 5, pp. 66–73, May 2016.
- [23] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in o-the-shelf IEEE 802.11 ad devices," in *Proc. ACM CoNEXT*, 2017, pp. 414–425.
- [24] T. S. Rappaport, F. Gutierrez, E. Ben-Dor, J. N. Murdock, Y. Qiao, and J. I. Tamir, "Broadband millimeter-wave propagation measurements and models using adaptive-beam antennas for outdoor urban cellular communications," *IEEE Trans. Antennas Propag.*, vol. 61, no. 4, pp. 1850–1859, Apr. 2013.
- [25] F. Giannetti, M. Luise, and R. Reggiani, "Mobile and personal communications in the 60 GHz band: A survey," *Springer Wireless Pers. Commun.*, vol. 10, no. 2, pp. 207–243, 1999.
- [26] T. Nitsche and A. B. Flores and E. W. Knightly and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *Proc. IEEE INFOCOM*, Hong Kong, China, Apr. 2015, pp. 2416–2424.
- [27] N. Gutierrez, C. Belmonte, J. Hanvey, R. Espejo, and Z. Dong, "Indoor localization for mobile devices," in *Proc. IEEE ICNSC*, Miami, FL, Apr. 2014, pp. 173–178.
- [28] S. Ito, F. Endres, M. Kuderer, G. D. Tipaldi, C. Stachniss, and W. Burgard, "W-RGB-D: Floor-plan-based indoor global localization using a depth camera and WiFi," in *Proc. IEEE ICRA*, Hong Kong, China, May 2014, pp. 417–422.
- [29] Y. Wang, A. K. S. Wong, and R. S. K. Cheng, "Adaptive room-level localization system with crowd-sourced WiFi data," in *Proc. SAI IntelliSys*, London, UK, Nov. 2015, pp. 463–469.
- [30] Y. Guo, Y. Sun, Y. Li, T. Y. Wu, and M. S. Obaidat, "Accurate indoor localization with crowd sensing," in *Proc. IEEE ICC*, Kuala Lumpur, Malaysia, May 2016, pp. 1–7.
- [31] X. Lin, X.-W. Chang, and X. Liu, "LocMe: Human locomotion and map exploitation based indoor localization," in *Proc. IEEE PerCom*, Kona, HI, Mar. 2017, pp. 131–140.
- [32] D. Zeinalipour-Yazti, C. Laoudias, K. Georgiou, and G. Chatzimilioudis, "Internet-based indoor navigation services," *IEEE Internet Comput.*, vol. 21, no. 4, pp. 54–63, Jul. 2017.
- [33] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of wireless indoor positioning techniques and systems," *IEEE Trans. Syst., Man, Cybern. C*, vol. 37, no. 6, pp. 1067–1080, Nov. 2007.
- [34] B. Cook, G. Buckberry, I. Scowcroft, J. Mitchell, and T. Allen, "Indoor location using trilateration characteristics," in *Proc. London Commun. Symp.*, London, UK, Sep. 2005, pp. 147–150.
- [35] K. Kaemarungsi and P. Krishnamurthy, "Modeling of indoor positioning systems based on location fingerprinting," in *Proc. IEEE INFOCOM*, vol. 2, Hong Kong, China, Mar. 2004, pp. 1012–1022.
- [36] Z. Wei, Y. Zhao, X. Liu, and Z. Feng, "DoA-LF: A location fingerprint positioning algorithm with millimeter-wave," *IEEE Access*, vol. 5, pp. 22 678–22 688, Sep. 2017.
- [37] Y. Qi, H. Kobayashi, and H. Suda, "Analysis of wireless geolocation in a non-line-of-sight environment," *IEEE Trans. Wireless Commun.*, vol. 5, no. 3, pp. 672–681, 2006.
- [38] J. Prieto, S. Mazuelas, A. Bahillo, P. Fernandez, R. M. Lorenzo, and E. J. Abril, "Adaptive data fusion for wireless localization in harsh environments," *IEEE Trans. Signal Process.*, vol. 60, no. 4, pp. 1585–1596, 2012.
- [39] J. Xiong and K. Jamieson, "Arraytrack: A fine-grained indoor location system," in *Proc. USENIX*, Lombard, IL, Apr. 2013, pp. 71–84.
- [40] S. Sen, J. Lee, K.-H. Kim, and P. Congdon, "Avoiding multipath to revive inbuilding WiFi localization," in *Proc. ACM MobiSys*, Taipei, Taiwan, Jun. 2013, pp. 249–262.
- [41] R. Nandakumar, K. K. Chintalapudi, and V. Padmanabhan, "Centaur: Locating devices in an office environment," in *Proc. ACM Mobicom*, Istanbul, Turkey, Aug. 2012, pp. 281–292.
- [42] A. T. Mariakakis, S. Sen, J. Lee, and K.-H. Kim, "SAIL: Single access point-based indoor localization," in *Proc. ACM MobiSys*, Bretton Woods, NH, Jun. 2014, pp. 229–232.
- [43] C. Zhang, F. Li, J. Luo, and Y. He, "iLocScan: Harnessing multipath for simultaneous indoor source localization and space scanning," in *Proc. ACM SenSys*, Memphis, TN, Nov. 2014, pp. 91–104.
- [44] F. Lemic, J. Martin, C. Yarp, D. Chan, V. Handziski, R. Brodersen, G. Fettweis, A. Wolisz, and J. Wawrzyn, "Localization as a feature of mmWave communication," in *Proc. IWCWC*, Paphos, Cyprus, Sep. 2016, pp. 1033–1038.
- [45] H. Deng and A. Sayeed, "Mm-wave MIMO channel modeling and user localization using sparse beamspace signatures," in *Proc. IEEE SPAWC*, Toronto, Canada, Jun. 2014, pp. 130–134.

- [46] H. El-Sayed, G. Athanasiou, and C. Fischione, "Evaluation of localization methods in millimeter-wave wireless systems," in *Proc. IEEE CAMAD*, Athens, Greece, Dec. 2014, pp. 345–349.
- [47] J. Palacios, P. Casari, and J. Widmer, "JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems," in *Proc. IEEE INFOCOM*, Atlanta, GA, May 2017, pp. 1–9.
- [48] J. Chen, D. Steinmetzer, J. Classen, E. W. Knightly, and M. Hollick, "Pseudo lateration: Millimeter-wave localization using a single RF chain," in *Proc. IEEE WCNC*, San Francisco, CA, Mar. 2017, pp. 1–6.
- [49] P. Botsinis, D. Alanis, S. Feng, Z. Babar, H. V. Nguyen, D. Chandra, S. X. Ng, R. Zhang, and L. Hanzo, "Quantum-assisted indoor localization for uplink mm-wave and downlink visible light communication systems," *IEEE Access*, vol. 5, pp. 23 327–23 351, 2017.
- [50] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Z. Win, "Ranging with ultrawide bandwidth signals in multipath environments," *Proc. IEEE*, vol. 97, no. 2, pp. 404–426, 2009.
- [51] P. Meissner, C. Steiner, and K. Witrals, "UWB positioning with virtual anchors and floor plan information," in *Proc. WPNC*, Dresden, Germany, Mar. 2010, pp. 150–156.
- [52] C. Gentner and T. Jost, "Indoor positioning using time difference of arrival between multipath components," in *Proc. IPIN*, Montbeliard-Belfort, France, Oct. 2013, pp. 1–10.
- [53] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero, R. L. Moses, and N. S. Correal, "Locating the nodes: cooperative localization in wireless sensor networks," *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 54–69, 2005.
- [54] U. A. Khan, S. Kar, and J. M. Moura, "Distributed sensor localization in random environments using minimal number of anchor nodes," *IEEE Trans. Signal Process.*, vol. 57, no. 5, pp. 2000–2016, 2009.
- [55] —, "DILAND: An algorithm for distributed sensor localization with noisy distance measurements," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1940–1947, 2010.
- [56] Y. Shen, S. Mazuelas, and M. Z. Win, "Network navigation: Theory and interpretation," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 9, pp. 1823–1834, 2012.
- [57] Y. Shen and M. Z. Win, "Fundamental limits of wideband localization—Part I: A general framework," *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 4956–4980, Oct. 2010.
- [58] Y. Shen, H. Wymeersch, and M. Z. Win, "Fundamental limits of wideband localization—Part II: Cooperative networks," *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 4981–5000, 2010.
- [59] Y. Kim, H. Y. Lee, P. Hwang, R. K. Patro, J. Lee, W. Roh, and K. Cheun, "Feasibility of mobile cellular communications at millimeter wave frequency," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 3, pp. 589–599, Apr. 2016.
- [60] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng, "Noncoherent mmWave path tracking," in *Proc. ACM HotMobile*, Sonoma, CA, USA, Feb. 2017, pp. 13–18.
- [61] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "BeamSpy: Enabling robust 60 GHz links under blockage," in *Proc. USENIX NSDI*, Santa Clara, CA, Mar. 2016, pp. 193–206.
- [62] C. Wang, F. Qi, G. Shi, and X. Wang, "Convex combination based target localization with noisy angle of arrival measurements," *IEEE Trans. Wireless Commun.*, vol. 3, no. 1, pp. 14–17, Feb. 2014.
- [63] H. Xu, V. Kukshya, and T. Rappaport, "Spatial and temporal characteristics of 60 GHz indoor channels," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 3, pp. 620–630, Apr. 2002.
- [64] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*. MIT Press and McGraw-Hill, 2001.
- [65] J. Palacios, D. De Donno, and J. Widmer, "Tracking mm-wave channel dynamics: Fast beam training strategies under mobility," in *Proc. IEEE INFOCOM*, Atlanta, GA, May 2017, pp. 1–9.
- [66] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, "Boon and bane of 60 GHz networks: Practical insights into beamforming, interference, and frame level operation," in *Proc. ACM CoNEXT*, Heidelberg, Germany, Dec. 2015, pp. 17:1–17:13.



localization for quasi-optical communication systems.



developing algorithms to improve its performance.



Manager of the NAUTILUS and WISEWAI projects. In 2015, he joined the IMDEA Networks Institute, Madrid, Spain, where he leads the Ubiquitous Wireless Networks group. He served in the organizing committee of several international conferences. He serves on the editorial boards of the IEEE TRANSACTIONS ON MOBILE COMPUTING and of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, and has been guest editor of a special issue of IEEE ACCESS on "Underwater Acoustic Communications and Networking," as well as of an earlier special issue of the *Hindawi Journal of Electrical and Computer Engineering* on the same topic. He received two best paper awards. His research interests include many aspects of networked communications, such as channel modeling, network protocol design, localization, simulation, and experimental evaluations.



Switzerland on ultra-wide band communication and network coding. He was a visiting researcher at the International Computer Science Institute in Berkeley, USA, University College London, UK, and TU Darmstadt, Germany. Joerg Widmer authored more than 150 conference and journal papers and three IETF RFCs, and holds 13 patents. He serves on the editorial board of the IEEE TRANSACTIONS ON MOBILE COMPUTING and *Elsevier Computer Networks* and is a member of the program committees of several major conferences. He was awarded an ERC consolidator grant, the Friedrich Wilhelm Bessel Research Award of the Alexander von Humboldt Foundation, a Mercator Fellowship of the German Research Foundation, a Spanish Ramon y Cajal grant, as well as nine best paper awards.

Joan Palacios received his B.Sc. degree in mathematics in 2015 from the Universitat de Valencia, Spain, and his M.Sc. degree in multimedia and communications in 2016 from the Universidad Carlos III de Madrid, Spain. In September 2015, he joined the IMDEA Networks Institute, Madrid, Spain as a PhD student. His areas of interest lie in mm-Wave communications with main research focus on the development of efficient schemes for beam training and tracking, hybrid analog-digital beamforming, and Angle Difference of Arrival (ADoA) based

Guillermo Bielsa is currently a Research Assistant at IMDEA Networks in Madrid, Spain. He is also with Universidad Carlos III de Madrid (UC3M), where he is enrolled as a PhD student. He obtained both his bachelor's degree in Communication System Engineering and his master's degree in Multimedia and Communications from UC3M in 2015 and 2016, respectively. He developed a great interest for wireless communications during his studies and now he is working in the field of mmWave communications, studying its behavior and

Paolo Casari (S'05–M'08–SM'13) received the PhD in Information Engineering in 2008 from the University of Padova, Italy. He was on leave at the Massachusetts Institute of Technology in 2007, working on underwater communications and networks. He collaborated to several funded projects including CLAM (FP7), RACUN (EDA), and several US ARO, ONR and NSF initiatives. He is the PI of the NATO SPS project ThreatDetect, and the scientific coordinator of the EU H2020 RECAP and SYMBIOSIS projects. Previously, he was Technical