

Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points

Guillermo Bielsa^{†‡} Joan Palacios^{†‡} Adrian Loch[†] Daniel Steinmetzer^b Paolo Casari[†] Joerg Widmer[†]

[†]IMDEA Networks Institute, Madrid, Spain

[‡]Universidad Carlos III de Madrid, Spain

^bSecure Mobile Networking Lab, Technische Universität Darmstadt, Germany

Abstract—The very large bandwidth available in the 60 GHz band allows, in principle, to design highly accurate positioning systems. Integrating such systems with standard protocols (e.g., IEEE 802.11ad) is crucial for the deployment of location-based services, but it is also challenging and limits the design choices. Another key problem is that consumer-grade 60 GHz hardware only provides coarse channel state information, and has highly irregular beam shapes due to its cost-efficient design. In this paper, we explore the location accuracy that can be achieved using such hardware, without modifying the 802.11ad standard. We consider a typical 802.11ad indoor network with multiple access points (APs). Each AP collects the coarse signal-to-noise ratio of the directional beacons that clients transmit periodically. Given the irregular beam shapes, the challenge is to relate each beacon to a set of transmission angles that allows to triangulate a user. We design a location system based on particle filters along with linear programming and Fourier analysis. We implement and evaluate our algorithm on commercial off-the-shelf 802.11ad hardware in an office scenario with mobile human blockage. Despite the strong limitations of the hardware, our system operates in real-time and achieves sub-meter accuracy in 70% of the cases.

I. INTRODUCTION

Communications in the millimeter-wave (mmWave) band place strong requirements on the hardware. The very high frequency, bandwidth, and data rates require high processing power at the transceivers. To ensure that consumer-grade mmWave hardware is viable, the design of commercial off-the-shelf devices is often tailored to specific applications. For instance, IEEE 802.11ad access points (APs) use cost-efficient phased antenna arrays that result in extremely irregular beam shapes [1]. Further, their firmware operates on coarse Signal-to-Noise Ratio (SNR) values to select the best transmit sector. While sufficient for communications in scenarios with a small number of nodes, this strongly limits other mmWave use cases such as high accuracy positioning.

Related work shows that mmWave signals can provide sub-centimeter accuracy [2], but this requires specialized hardware such as motorized horn antennas. This allows location systems to precisely estimate the incidence angle of a signal, since the beam shapes of horn antennas are much more focused than the ones of consumer-grade phased antenna arrays. As a result, existing mmWave positioning systems are not integrated with mmWave communication systems [3], [4]. The difference in terms of hardware requirements is significant. Implementing

both communications and localization on a common cost-efficient consumer-grade platform is highly challenging, but still crucial to enable next-generation location-based services.

In this paper, we explore the accuracy of mmWave positioning that can be achieved using commercial off-the-shelf 60 GHz IEEE 802.11ad hardware. We consider a typical indoor 60 GHz network deployment with multiple APs per room. This ensures coverage despite the high attenuation in the 60 GHz band and the high probability of blockage. Further, we do not modify the operation of 802.11ad, but only collect information that is available at the APs as a natural by-product. Specifically, we collect the SNR with which the APs receive the sector sweep messages of the clients in a room.

According to the standard, clients transmit such messages periodically on each of their transmit sectors to cover all directions. The standard considers idealized sectors which are triangle-shaped, and thus cover a well-defined angular range. Locating the client based on such a sector model given the location of the APs is straightforward. However, the shape of the sectors used in consumer-grade commercial hardware is not only highly irregular, but also does not clearly point towards a specific direction since the underlying beam patterns often have two or more equally-strong lobes [1], [5]. As a result, relating a sector identifier to an angular direction is highly challenging. Furthermore, current 60 GHz hardware uses quasi-omnidirectional beam patterns for reception and thus does not perform receive beam training. This additionally limits the angular information that such consumer-grade hardware can provide for positioning. Finally, triangulation requires that at least three APs are in range in order to determine the location of the client. Since blockage in mmWave networks is common, this cannot always be ensured, and the positioning system should be designed to cope with missing information.

We address the above issues as follows. Instead of relating each sector identifier to a specific angle, we design a method to compute a sparse channel decomposition conveying the power and the angle of departure of each propagation path. We then merge the information of all APs and all beam patterns to estimate the location of the user. To this end we use linear programming along with Fourier analysis. In addition, we prevent location errors due to blockages by using a particle filter that implements a mechanism similar to “dead reckoning.” This limits the location error even if a client is in range

of less than three APs. Our approach operates in real-time since it updates the location estimation dynamically each time a client performs a sector sweep. We implement our system on commercial off-the-shelf 60 GHz APs with electronically steerable phased antenna arrays and a limited SNR accuracy of 0.25 dB. We modify the APs' firmware to obtain the SNR of each sweep message that the AP receives. Other than that, the modified firmware behaves exactly the same as the original one. This major effort results in a first-of-its-kind real-time testbed, that allows us to obtain unprecedented insight into practical mmWave location systems.

We deploy our testbed both in a large empty indoor space, as well as in an open-plan office space with mobile human blockage, furniture, and static obstacles such as pillars. On average, we achieve sub-meter location accuracy in 70% of the cases. For our best results, this value increases up to 87%. This is a remarkable result given the tremendous constraints that consumer-grade hardware imposes on mmWave channel decomposition methods and localization algorithms. We are the first to explore the limitations of a location system implemented on such hardware. Our contributions are as follows:

- We analyze the limitations of consumer-grade mmWave hardware for user positioning, and show how to tackle them using techniques such as Fourier analysis and particle filtering.
- We design and implement a mmWave location system for IEEE 802.11ad devices with zero overhead and no modification to the protocol.
- Our system operates in real-time on off-the-shelf 60 GHz APs in real-world scenarios. To this end, we modify the firmware of the APs while maintaining compliance to the IEEE 802.11ad standard.
- We achieve sub-meter location accuracy in 70% of the cases despite highly irregular beam shapes, 0.25 dB SNR quantization, and no receiver beam training.

The remainder of this paper is structured as follows. Section II surveys related work. In Section III, we explain the operation of our localization system. We then present our practical implementation in Section IV and our testbed results in Section V. Finally, Section VI concludes the paper.

II. RELATED WORK

Wireless localization is a very well investigated research area [6]. A large body of work studies localization using signals below 10 GHz such as WiFi at 2.4 and 5 GHz. The recent advances in mmWave communication systems have drawn attention to higher frequencies, where the large available bandwidth inherently allows for much higher location accuracy [7]. In this section, we first survey techniques for lower frequency bands and then study the techniques that are designed for mmWave scenarios [8].

A. Low Frequency Localization

Existing localization approaches are often based on triangulation or multilateration [9]. The key idea is to estimate the distance, the angle of arrival (AoA), or both from two or more

APs to the client that shall be localized. Systems based on distance ranging are strongly dependent on the accuracy of the ranging process [10]. A common approach is to estimate the distance based on the received signal strength (RSS), which typically results in very coarse results. Estimating the time of flight (ToF) can yield more accurate location [11] but requires precise timing and synchronization among devices. Systems based on AoA estimation [12], [13] require an array of antennas at the receiver but achieve high accuracy in networks with many APs. The key problem for both distance and AoA estimation is distinguishing the line-of-sight (LoS) path from reflected non-line-of-sight (NLoS) paths. This is particularly relevant at low frequencies due to the rich multipath environment. Still, sophisticated localization algorithms can exploit multipath to perform simultaneous localization and mapping (SLAM) [14]. In SLAM, the system tracks how reflections change as the user moves in order to map obstacles [15]. To improve the accuracy of localization systems at low frequencies, related work suggests using additional information such as gyroscope and sensor readings [16].

B. High Frequency Localization

Localization is an inherent feature of mmWave systems due to the sparseness of the channel and the use of directional communication [7]. Since the wavelength is in the order of millimeters, even mobile nodes can fit large antenna arrays. Related work shows how fully digital antenna array architectures allow for highly accurate localization mechanisms based on massive multiple-input-multiple-output (MIMO) [17]. However, commercial mmWave hardware typically features analog antenna arrays with a single transceiver chain, which only allows for coarse AoA estimation based on analog beamforming. Most existing work does not consider such limitations but instead studies theoretical performance assuming idealized hardware. For this ideal case, results show that mmWave localization can not only reveal the position of a device but also its orientation [18]. Further studies indicate that SLAM is feasible in mmWave [4], and that the use of Orthogonal Frequency-Division Multiplexing (OFDM) in mmWave can allow a multiple-input-single-output (MISO) system to estimate location [19].

Practical mmWave localization work is limited. Related work shows that triangulation [3] and ranging with multilateration [20] are feasible in practice. However, such approaches assume specialized hardware such as lab-grade horn antennas and accurate measurement equipment. Moreover, they often do not enable simultaneous communication and localization. In contrast, our solution allows for both, and is designed for consumer-grade devices with strong hardware impairments. Thus, our system clearly stands apart from all of the above body of work.

III. LOCALIZATION ALGORITHM

One of the motivations for our work is that making a mmWave network aware of the client locations opens up interesting opportunities to, for example, optimize client-AP

associations dynamically according to mobility, AP coverage, and traffic requirements. In turn, this makes the network much more robust to interference, and improves fairness by load balancing clients across available APs. It also increases the resilience to mutual signal blockage among different users, as this could be foreseen and remediated in a preemptive manner.

Client localization can be achieved at zero communication overhead thanks to the information collected by 802.11ad-compliant APs during the beam training phase. Our algorithm comprises two main components to localize a client despite the limitations of consumer-grade devices: *i*) a simple linear programming formulation that allows the network to estimate the angle of departures (AoDs) of the mmWave signal from the client, in a way that such AoDs are compatible with SNR values measured by the APs for different beam pattern choices of the client; and *ii*) a modified particle filter (PF) that obtains feasible estimates of the client's location while the client moves. Most importantly, step *i*) is not based on any simplifying assumption such as custom beam pattern design, "triangular" beam pattern shapes, or the availability of phase information; additionally, step *ii*) is based on low-complexity particle updating formulas, and on an informed way of creating new particles that have a higher probability of being generated at the actual location of the client, thus speeding up the convergence of the PF and substantially improving its accuracy. In the following, we provide the details of our localization algorithm.

A. Angle-of-Departure estimation

Call \mathbf{a}_i and \mathbf{x} the coordinates of AP i and client node to be localized (see Figure 1) and φ is the client orientation with respect to an absolute coordinate system. For transmission, the client can choose among a total of $B = 34$ beam patterns $p_b(\theta)$, for $b = 1, \dots, B$, where θ is the emission angle and $p_b(\theta)$ is the amplitude gain of pattern b along θ . Whenever a client performs beam training, each AP records the received signal strength indicator (RSSI) and the SNR corresponding to every beam pattern tested by the client that was detected by the AP. The APs then forward this information to a central server, where the location process runs. Call $\gamma_i^{(b)}$ the SNR (in dB) measured by AP i when the client transmits with beam pattern b , and let $P_{R_i}^{(b)} = 10^{\gamma_i^{(b)}/20}$ be the corresponding signal amplitude.

The most critical issues for the design of the localization algorithm involve the lack of phase information (which makes the problem non-linear and prohibits typical angle decomposition algorithms), the rough quantization of log-scale SNR values with a resolution of 0.25 dB, and a device firmware which is often too slow to log RSSI and SNR information for all client beam patterns. Ultimately, the latter issue results in incomplete measurements.

A typical assumption for angle estimation is that the energy carried by the LoS path from the client to any AP i exceeds the energy of NLoS paths. In order to find the LoS AoDs $\hat{\theta}_i$ that best match the SNR measurements of the APs, it would




Fig. 1. Reference scenario for the localization algorithm. The irregular beam patterns of a consumer-grade mmWave device may generate NLoS paths (dashed lines) with similar power as LoS paths (solid lines).

then be sufficient to solve the following minimum mean-square error (MMSE) problem:

$$\hat{\theta}_i = \arg \min_{\theta} \min_{\alpha} \sum_{b \in \mathcal{B}_i} \left(P_{R_i}^{(b)} - \alpha p_b(\theta) \right)^2. \quad (1)$$

However, the assumption that the LoS path predominates the RSSI and SNR measurements is not true in practice, as the transmit beam patterns of mmWave devices can emit significant power through secondary lobes, resulting in significant NLoS path energy reaching the receiving APs. An example using two beam patterns from the Talon AD7200 devices is shown in Figure 1. Assume that the client at location \mathbf{x} communicates to the AP at \mathbf{a}_3 through the beam pattern plotted in blue (and specifically through its large lobe pointing towards the bottom-left corner). A secondary lobe just above the main one would similarly amplify a NLoS path reaching \mathbf{a}_3 , making the NLoS path power non-negligible with respect to that of the LoS path.

To avoid making the above assumption, we take a different approach, and assume that the power received from mmWave signals from the client dominates noise in the APs' SNR measurements. Since we have neither phase information from the APs nor access to the phased antenna array weights, we cannot estimate how different multipath components of the same transmitted signals interfere in the complex domain. In the same vein, we cannot directly apply the non-coherent path estimation approach in [21], due to the AP's coarse dB-scale quantization of measured SNR values. Hence, we conservatively constrain the measured amplitude to be less than the sum of the amplitude of all paths, if such paths were interfering constructively. This makes it possible to formulate the AoD estimation problem as the following linear program with variables $\alpha_i(\theta)$:

$$\min_{\theta} \sum_{\theta} \alpha_i(\theta) \left(\sum_{b \in \mathcal{B}_i} p_b(\theta) \right)^{1/2} \quad (2a)$$

$$\text{s.t.} \sum_{\theta} \alpha_i(\theta) p_b(\theta) \geq P_{R_i}^{(b)}, \quad \forall b \in \mathcal{B}_i \quad (2b)$$

$$\alpha_i(\theta) \geq 0, \quad \forall \theta \in \Theta \quad (2c)$$

where \mathcal{B}_i is the set of beam patterns for which AP i was




Fig. 2. Different channel decompositions. Left: MMSE result from Eq. (1). Center: Linear programming result with variables $\alpha_i(\theta)$ in problem (2). Right: smoothed linear programming (using the variables $v_i(\theta)$ defined in Eq. (8)).

able to retrieve RSSI and SNR measurements,¹ and Θ is the set of the decision variables, whose cardinality depends on the resolution of the angular domain quantization. The latter is a design parameter and is set to 400 points in our implementation. The term $(\sum_{b \in \mathcal{B}_i} p_b(\theta)^2)^{1/2}$ in the objective function prevents giving excessive weight to AoDs that have a much higher relative gain along direction θ . Besides being very fast to solve, problem (2) has a solution bounded by 0. It is easy to prove that, as a consequence, the number of non-zero components of α_i is bounded by $|\mathcal{B}_i|$. In practice, the total number of non-zero terms $\alpha_i(\theta)$ tends to be very low.

As an example, in the left and center panels of Figure 2, we show the AoD estimation results obtained by four different APs using the MMSE approach in (1) and the linear programming approach in problem (2). The dots on the outer circle of the plots correspond to actual AoDs. We observe that the linear programming solution estimates a number of likely directions, among which at least one is in fact very close to the LoS path between the AP and the client. Conversely, the MMSE solution shows larger discrepancies between the estimated and the actual LoS angles.

B. Goodness function for the user location

We now leverage the terms $\alpha_i(\theta)$ estimated through (2) to find a position $\hat{\mathbf{x}}$ and orientation $\hat{\varphi}$ for the client that agree with the angular estimation by different APs. Define $\hat{\theta}_i^{\mathbf{x},\varphi} = \hat{\theta}_i^{\mathbf{x}} + \varphi$ as the AoD of the LoS mmWave path emanating from the client as estimated by anchor i , when the client has orientation φ , and the angles $\hat{\theta}_i^{\mathbf{x}}$ refer to the case where the user is aligned to the reference coordinate system, i.e., $\varphi = 0$. Call

$$G^{\mathbf{x}}(\varphi) = \sum_i H_i^{\mathbf{x}}(\varphi) = \sum_i \alpha_i(\hat{\theta}_i^{\mathbf{x},\varphi}), \quad (3)$$

where \mathbf{x} is the location of the user. We define an angular goodness function for the user location as

$$\mathcal{L}(\mathbf{x}) = \max_{\varphi} G^{\mathbf{x}}(\varphi). \quad (4)$$

The expression in (3) can be reformulated as follows:

$$G^{\mathbf{x}}(\varphi) = \sum_i \langle \alpha_i, \delta_{\hat{\theta}_i^{\mathbf{x},\varphi}} \rangle \quad (5)$$

¹Due to firmware inefficiencies, this is typically a subset of 65% to 75% of the beam patterns available to the client, and this subset changes across different measurements.

where $\delta_{\hat{\theta}_i^{\mathbf{x},\varphi}}(\theta)$ is the Dirac delta distribution centered on $\hat{\theta}_i^{\mathbf{x},\varphi}$ and $\langle \cdot, \cdot \rangle$ denotes the inner product. Since $\hat{\theta}_i^{\mathbf{x},\varphi} = \hat{\theta}_i^{\mathbf{x}} + \varphi$, we have $\delta_{\hat{\theta}_i^{\mathbf{x},\varphi}} = \delta_{\hat{\theta}_i^{\mathbf{x}}} \otimes \delta_{\varphi}$ where, \otimes denotes circular convolution over the angular domain. Therefore, using Fourier transform properties,

$$\begin{aligned} H_i^{\mathbf{x}}(\varphi) &= \langle \alpha_i, \delta_{\hat{\theta}_i^{\mathbf{x},\varphi}} \rangle = \langle \alpha_i, \delta_{\hat{\theta}_i^{\mathbf{x}}} \otimes \delta_{\varphi} \rangle \\ &= \langle \mathcal{F}(\alpha_i), \mathcal{F}(\delta_{\hat{\theta}_i^{\mathbf{x}}})\mathcal{F}(\delta_{\varphi}) \rangle = \langle \mathcal{F}(\alpha_i)\mathcal{F}(\delta_{\hat{\theta}_i^{\mathbf{x}}})^\dagger, \mathcal{F}(\delta_{\varphi}) \rangle \\ &= \langle \mathcal{F}^{-1}(\mathcal{F}(\alpha_i)\mathcal{F}(\delta_{\hat{\theta}_i^{\mathbf{x}}})^\dagger), \delta_{\varphi} \rangle, \end{aligned} \quad (6)$$

where $\mathcal{F}(\cdot)$ and $\mathcal{F}^{-1}(\cdot)$ are the discrete and inverse discrete Fourier transform (DFT) operators, respectively, and the dependence on φ of the terms on the right hand side of (6) has been dropped for clarity. Eq. (6) finally simplifies to

$$H_i^{\mathbf{x}} = \mathcal{F}^{-1}(\mathcal{F}(\alpha_i)\mathcal{F}(\delta_{\hat{\theta}_i^{\mathbf{x}}})^\dagger). \quad (7)$$

The above formulation is exact only if AoD estimates are not affected by any error. In order to compensate for measurement errors, we substitute $\alpha_i(\theta)$ with a cyclically convoluted version

$$v_i(\theta) = \alpha_i(\theta) \otimes g(\theta), \quad (8)$$

where $g(\theta) = \exp(-\theta^2/(2\sigma_e^2))/\sqrt{2\pi\sigma_e^2}$ is a Gaussian kernel of standard deviation $\sigma_e = 10^\circ$. The right panel in Figure 2 shows the result of this convolution. We then define

$$\bar{G}^{\mathbf{x}}(\varphi) = \sum_i \bar{H}_i^{\mathbf{x}}(\varphi), \quad (9)$$

where $\bar{H}_i^{\mathbf{x}} = \mathcal{F}^{-1}(\mathcal{F}(v_i)\mathcal{F}(\delta_{\hat{\theta}_i^{\mathbf{x}}})^\dagger)$, and

$$\bar{\mathcal{L}}(\mathbf{x}) = \max_{\varphi} \bar{G}^{\mathbf{x}}(\varphi). \quad (10)$$

The objective function in (10) can be used to evaluate the goodness of a given client location for a set of power measurements per client beam pattern, and can be computed even if the involved angles do not exactly coincide with the AP angle measurements. Notably, the formula is very fast to compute, as it involves only the computation of FFT/IFFT, products and the maximum operator. We use (10) to develop a PF that processes AP measurements and localizes a mmWave device in real time.

C. Distance-based SNR likelihood

We will now provide a function to determine how likely a SNR measurement is given that the user is located in a given position. To do this, we will consider that the maximum SNR (in dB) measured by AP i for the best client beam pattern follows a log-normal distribution, conditioned to a path loss model, i.e.,

$$\gamma_i^{(\max)} \sim \mathcal{N}(\bar{\gamma}_i^{(\max)}(\|\mathbf{x} - \mathbf{a}_i\|), \sigma_d^2), \quad (11)$$

where $\|\cdot\|$ is the Euclidean norm and $\bar{\gamma}_i^{(\max)}(\|\mathbf{x} - \mathbf{a}_i\|)$ is the expected value of the maximum SNR according to the path loss propagation model

$$\bar{\gamma}_i^{(\max)}(\|\mathbf{x} - \mathbf{a}_i\|) = \kappa - 10\eta \log_{10}(\|\mathbf{x} - \mathbf{a}_i\|). \quad (12)$$




Fig. 3. Path loss regression model for the computation of $\gamma_i^{(\max)}$.

We estimate the parameters κ and η of the model by measuring pairs of $\bar{\gamma}_i^{(\max)}(\|\mathbf{x} - \mathbf{a}_i\|)$ against $\log_{10}(\|\mathbf{x} - \mathbf{a}_i\|)$ values in an empty room, for random locations and orientations of the client and of the APs. We then perform a linear regression to estimate the values of the two parameters. The result is shown in Figure 3. The distance-based likelihood of the maximum SNR measurements can be finally written as

$$\mathcal{D}(x) = \prod_i \frac{1}{\sqrt{2\pi}\sigma_d} \exp\left(-\frac{(\gamma_i^{(\max)} - \bar{\gamma}_i^{(\max)}(\|\mathbf{x} - \mathbf{a}_i\|))^2}{2\sigma_d^2}\right) \quad (13)$$

In (13), $\sigma_d = 15$ dB, which is a sufficiently large value to take into account additional attenuation caused by partial or full signal blockage. We remark that this method is only a rough consistency check, and not a full ranging method based on the precise calibration of the SNR measurement provided by the hardware. In fact, the value output by the function $\mathcal{D}(x)$ is not used to determine the user location, but rather to help filter spurious solutions when more than one location is feasible from the AoD estimation.

D. Modified particle filter for client localization

In order to localize the user in real time, we employ a modified PF algorithm, shown in Figure 4. The top part of the flow chart behaves like a typical PF. Particles are generated at random for the initialization of the algorithm, grouped in the initial particle set, and evolved according to a mobility model. The fitness of the evolved position is updated based also on the objective functions in (10) and (13), which incorporates measurements collected by the APs. Evolved particles are grouped in the posterior set and become the past particle set for the next PF cycles.

We modify the particle generation procedure in order to incorporate a single ‘‘informed’’ particle in the initial particle set when a new AoD estimate is computed by the APs. The informed particle is positioned at a location computed using an angle difference of arrival (ADoA) algorithm applied to the three angle measurements with greatest associated $\alpha_i(\theta)$ among those computed by the different APs. This solution can be computed in closed form (it is basically an intersection of circumferences), which is much faster than computing, e.g., $\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \tilde{\mathcal{L}}(\mathbf{x})$. The informed particle has 0 speed and an initial fitness value equal to 1 in order to satisfy the fitness normalization criteria, and is evolved like all other particles. When creating the set of posterior particles, we apply a proper fitness renormalization, in order to avoid giving excessive weight to the offspring of the informed particle.

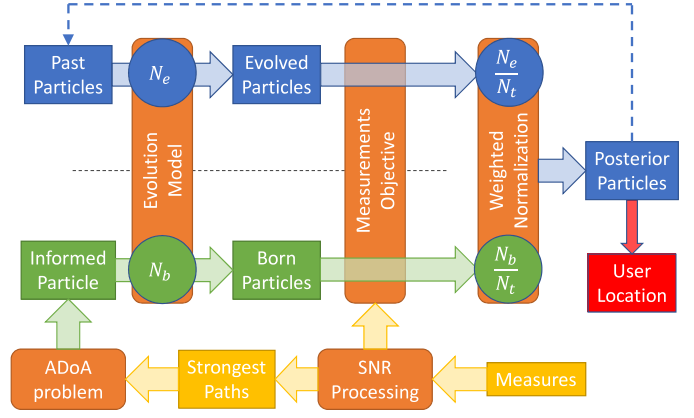


Fig. 4. Modified particle filter flowchart.

1) *Particle initialization and evolution*: Call \mathbf{x} , \mathbf{v} and f the location, the speed and the fitness of a particle. Together, these three quantities constitute the state of the particle, where the fitness values are normalized so that $\sum_{k=1}^{|\mathcal{P}|} f_k = 1$. As an initialization step, before any measurement, N_t particles are initialized at random according to a Gaussian distribution of mean and variance equal to the sample mean and numerical covariance of the AP location set, with speed equal to 0 and fitness N_t^{-1} . These particles are passed to the algorithm as the past particles set \mathcal{P}_p for the first iteration. At the beginning of each iteration, an informed particle is generated using an ADoA algorithm applied to the AoD measurements assuming they are LoS. Note that this may not be possible in some cases, e.g., due to lack of three AoDs or measurement errors preventing the computation of a feasible solution. If the informed particle can be computed, we evolve N_b particles from it, and $N_e = N_t - N_b$ from the initial particle set. Otherwise, we evolve $N_e = N_t$ only from the initial particles. Call \mathcal{P}_e and \mathcal{P}_b the sets of particles thus evolved, where $N_e = |\mathcal{P}_e|$ and $N_b = |\mathcal{P}_b|$.

To evolve a particle from the past particle set, we first choose a particle k at random with discrete probability distribution $\{f_\ell\}$, $\ell \in \mathcal{P}_p$. The procedure to evolve an informed particle is analogous, except that in this case all particles have the same parent. Call \mathbf{x}_k , \mathbf{v}_k and f_k the location, speed and fitness of particle k , respectively. We evolve a new particle through the following update equations:

$$\begin{aligned} \mathbf{x}' &= \mathbf{x}_k + \mathbf{v}_k \Delta t + \Delta \mathbf{x} + t \Delta \mathbf{v} \\ \mathbf{v}' &= \mathbf{v}_k + \Delta \mathbf{v}, \end{aligned} \quad (14)$$

where the apex $'$ denotes the state variables of the evolved particle. The model in (14) is a uniform linear motion model, with uncertainties $\Delta \mathbf{x} \sim \mathcal{N}(0, \sigma_x^2 \mathbf{I})$ and $\Delta \mathbf{v} \sim \mathcal{N}(0, \sigma_v^2 \mathbf{I})$ on the position and on the speed, respectively, and Δt is the time elapsed between the current and the previous measurement epochs. We set $\sigma_x = 2$ m and $\sigma_v = 1$ m/s. Define

$$\Delta \mathbf{f} = \mathbf{x}' - (\mathbf{x}_k + \mathbf{v}_k \Delta t) = \Delta \mathbf{x} + \Delta t \Delta \mathbf{v}, \quad (15)$$

so that $\Delta \mathbf{f} \sim \mathcal{N}(0, \sigma_x^2 + \Delta t^2 \sigma_v^2)$. To evolve the particle fitness values, we rely on $\tilde{\mathcal{L}}(\mathbf{x})$ in (10) and $\mathcal{D}(x)$ in (13). Since $\tilde{\mathcal{L}}(\mathbf{x})$




Fig. 5. Talon Router in our testbed (left) and in the anechoic chamber (right).

and $D(x)$ are independent of the particle speed, instead of generating $\Delta \mathbf{x}$ and $\Delta \mathbf{v}$ directly, we generate a sample value of $\Delta \mathbf{f}$ and compute $\Delta \mathbf{x}$ and $\Delta \mathbf{v}$ with maximum likelihood such that $\Delta \mathbf{f} = \Delta \mathbf{x} + t\Delta \mathbf{v}$. It can be easily shown that

$$\begin{aligned} \Delta \mathbf{x} &= \Delta \mathbf{f} \sigma_x^2 / (\sigma_x^2 + \Delta t^2 \sigma_v^2) \\ \Delta \mathbf{v} &= \Delta \mathbf{f} \Delta t^2 \sigma_v^2 / (\sigma_x^2 + \Delta t^2 \sigma_v^2) \end{aligned} \quad (16)$$

Therefore, the final update expression for \mathbf{x}_k becomes

$$\begin{aligned} \mathbf{x}' &= \mathbf{x}_k + \mathbf{v}_k \Delta t + \Delta \mathbf{f} \cdot \\ \mathbf{v}' &= \mathbf{v}_k + \Delta \mathbf{v}. \end{aligned} \quad (17)$$

In order to update the fitness of particle k based on the AP measurements, we compute

$$f'_k = f_k \frac{1}{2\pi(\sigma_x^2 + \Delta t^2 \sigma_v^2)} \exp\left(-\frac{\|\Delta \mathbf{f}\|^2}{2(\sigma_x^2 + \Delta t^2 \sigma_v^2)}\right) \bar{\mathcal{L}}(\mathbf{x}) D(x), \quad (18)$$

Multiplying by $\bar{\mathcal{L}}(\mathbf{x})$ and $D(x)$ includes the fitness update due to the angle estimation and the SNR of the measurements.

2) *Normalization*: Usually, the fitness of evolved particles is normalized such that $\sum_{k \in \mathcal{N}_t} f_k = 1$. However, in our modified PF this would result in excessive weight being for particles in set \mathcal{P}_b , since all of these particles evolve from the same informed particle which has fitness equal to 1. We correct this by imposing that the probability of evolving from sets \mathcal{P}_e and \mathcal{P}_b on the next measurement iteration is proportional to the number of elements of each set as

$$\sum_{k \in \mathcal{P}_e} f_k = \frac{N_e}{N_t} \quad \text{and} \quad \sum_{k \in \mathcal{P}_b} f_k = \frac{N_b}{N_t}. \quad (19)$$

In the following, we set $N_t = 1024$ and $N_b = 256$.

IV. IMPLEMENTATION

We base the practical implementation of our location system on the TP-Link Talon AD7200 router (see Figure 5, left), which is the first 60 GHz device that fully implements the IEEE 802.11ad standard. The Talon router uses a Qualcomm QCA9500 60 GHz chipset that comes with a phased antenna array of 32 antenna elements, whose phase and magnitude can be individually controlled. The router has 34 predefined antenna configurations hard-coded in the firmware, that can be selected to steer the antenna array. We use the updated firmware with version number “3.3.3.7759”.




Fig. 6. Examples of beam patterns used by the TP-Link Talon AD7200 router.

A. Firmware Modification

As the default firmware running on the QCA9500 chip does not allow to access the signal strength of received frames, we modified the firmware to integrate this feature [5], [22]. To this end, we used the Nexmon firmware patching framework [23], which enables the development of binary firmware extensions in C. By matching the patterns of IEEE 802.11ad sector sweep frames with the memory inside the chip, we identified the parts of the firmware that were responsible for handling the sector sweep frames. At these memory addresses, we patched the firmware to extract the measured signal strength [5]. For each sector that is probed during the sector sweep, we reveal the SNR and RSSI value. As the sector sweep is performed periodically and probes all available sectors with different beam-patterns, this effectively allows us to sense the environment.

B. Beam pattern Measurement

To obtain the beam patterns of all the different sectors, we conduct measurements using our modified firmware in an anechoic chamber (Figure 5, right) [5]. This allows us to avoid reflections that would reduce the accuracy of the measurements. We place one device on a custom rotation head that steers mechanically in different directions. A second device is placed at a distance of 3m facing the rotating one. To measure a beam pattern, we force the firmware to only use that specific pattern and establish a connection between both devices. We then record the SNR and RSSI while transmitting ping messages to maintain the connection active. We adjust the direction of the rotation head in steps of 1° , and thus measure the beam patterns with high accuracy. Figure 6 shows three of the 34 beam patterns of the Talon router as an example. As it is evident from the figure, their highly irregular shape and lack of a clear main lobe [5] are a major challenges for the localization algorithm.

V. EVALUATION

In this section we present our localization results. First we describe our scenarios, then explain the metrics we use, and finally show our results.

A. Testbed

We use eight Talon routers as APs and one as a station. To this end, we install a variant of LEDE [24] ported to this architecture along with the latest `wil6210` device driver from the Linux kernel and configure them accordingly. Due to the inner workings of the firmware, we can only extract the SNR




Fig. 7. Measurement setups

of the station’s sector sweep messages at the AP to which it is associated. To circumvent this hardware limitation (without very substantial firmware changes), we invert the role of the station and the APs. As a result, all fixed nodes can connect simultaneously to the node that the system is locating, enabling us to obtain all of the SNR values. For clarity, we stick to the same nomenclature regarding the APs and the station as in previous sections.

We deploy our testbed in two different setups. The first is an empty auditorium which is 11×21 meters large and has no furniture, as shown in Figure 7a. This is a controlled environment with no blockage or movement, and thus allows us to study the accuracy of our system in an ideal setting. We deploy the routers such that we obtain the best coverage in the center of the room. We measure the location accuracy at 32 different positions on a grid of eight rows and three columns. At each position, we consider four different orientations of the station, pointing towards each of the walls of the room.

Figure 7b depicts our second setup, which is a real-world office environment. The size of the room is 7.4×13.5 meters. This scenario includes office furniture such as tables, screens, and chairs. We test our system during office hours, that is, humans are present and move according to their regular activities. We deploy the APs such that we maximize coverage in the area, taking into account the impact of columns. We measure the location accuracy at 40 different positions, and consider the same four different orientations at each position.

B. Metrics and Device Configurations

We evaluate the performance of our system in terms of the localization error at each of the measurement positions in each of our scenarios. We obtain 100 measurements per position and device orientation. Thus, we have a total of 400 measurements per position. This allows us to compute the location error for a number of different device configurations. Specifically, we consider three cases:

- **Single antenna:** Our first case is the standard device configuration with a single antenna array. We include measurements for all four possible orientations and consider each orientation to be equally probable.
- **Four antennas:** Our second case considers a station with four antenna arrays, since future mmWave systems are likely to include more antenna arrays to combat blockage. As a result, the location system obtains four different location estimates. We implement an objective function that computes a score for each of the four estimates based on the likelihood of each estimate being correct. Our system chooses the estimate with the highest score.
- **Upper bound:** Our last case is an upper bound. At each location, we use the device antenna that results in the smallest location error. For instance, for the antenna facing a wall, the location error is much larger than for one facing multiple APs. Note that the system cannot know a priori which orientation is the most accurate one, and thus this estimate provides an upper bound on the performance that our system can theoretically achieve.

We show the error distributions as cumulative distribution functions (CDFs). For the error maps, we depict the median error of the measurements at each location, taking into account the corresponding device configuration.

C. Results

In this section, we present the evaluation results of our localization system. First, we discuss the behavior of the system in a dense network deployment with a total of eight APs. After that, we study the performance that we can achieve with fewer APs. To this end, we discard the measurements of some of the APs to obtain a sparse deployment scenario.

1) *Dense Deployment:* Figures 8 and 10 depict our measurement results in the auditorium. In Figure 8, we show a map of the median error for all of the measured positions and orientations for the single antenna configuration including some examples of the estimated locations. We observe that our system achieves high accuracy throughout the room, with slight degradations along the walls between APs. The underlying reason is that these areas have slightly worse coverage from the rest of the APs. In Figure 10 we show the cumulative distribution function for all of the measured positions. With a single antenna system, we can obtain sub-meter accuracy in roughly 70% of the cases. Further, the error is below $3m$ in more than 87% of the occasions. For the upper bound, we achieve sub-meter accuracy in 85% of the occasions while the remaining 15% have an error below $3m$. As expected, the performance of a system with four antennas lies in between the two above cases, also having a maximum error of $3m$.

Figures 9 and 11 show the corresponding results for the office environment. We observe a similar performance as for the auditorium. However, since the room contains obstacles, the system does not have full AP coverage at all of the positions that we consider. The error map in Figure 9 shows that, similarly to our previous results in Figure 8, we obtain the highest accuracy in the central part of the scenario due to the

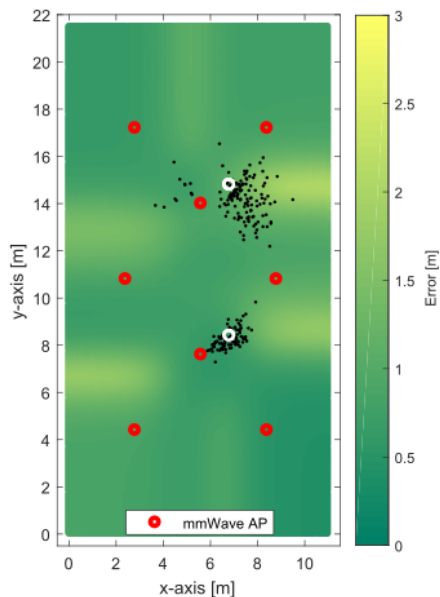


Fig. 8. Auditorium error map. The white circles represent two measured positions while the black dots show their corresponding location estimations.




Fig. 9. Office error map. The white circles represent two measured positions while the black dots show their corresponding location estimations.




Fig. 10. CDFs: Auditorium results




Fig. 11. CDFs: Office results

better coverage of the APs. The above performance is remarkably good given that we perform the measurements during office hours with the resulting intermittent human blockage. Figure 11 shows that the system provides sub-meter accuracy for 60% of the cases for both single and multiple antenna configurations. The upper bound achieves sub-meter accuracy 80% of the times. The three different metrics converge at an error of $3m$ for roughly 90% of the cases. As expected, the accuracy is lower than in the auditorium case, but still close to the upper bound. Future 60 GHz devices featuring more directional beam patterns and more accurate SNR sampling will provide even better results.

2) *Sparse deployment*: In Figures 12 and 13, we show the performance of our system in the auditorium when less than eight APs are available. To assess the performance degradation, we consider four sparse arrangements of four to six APs labeled “Scenario 1” to “Scenario 4” in Figure 16. For the setups with five and six APs, we achieve a median error meter between $1.5m$ and $1.7m$ for the single antenna case. Four antennas improves the median error to around $1.2m$. However, when reducing the number of APs to four in our last setup, the median error increases to $2m$ and $1.8m$, respectively. We also observe occasional large outliers due to the limited information available to the location system, which are typically in parts of the map with sparse AP coverage.

Figures 14 and 15 depict an equivalent analysis for our office environment. In this case, the performance of the different sparse AP deployments is more similar. Surprisingly, the performance is the same or even slightly better compared to the more benign auditorium case. Even the scenario with

APs aligned in a single line which we expected to perform badly shows results close to those of the other scenarios. The system can locate a single antenna device with a median error of $1.2m$ to $1.6m$ and with sub-meter accuracy 35% of the time. Using four antennas improves these values to $1.1m$ to $1.4m$ and 40% of sub-meter accuracy. From the error maps (not shown due to space constraints), we conclude that the areas performing worst are the ones close to the walls, which are the areas that experience the worst coverage.

VI. CONCLUSIONS

We design and implement a mmWave indoor localization system that operates on commercial off-the-shelf IEEE 802.11ad hardware. This is particularly challenging due to the strong limitations of consumer-grade hardware in terms of SNR accuracy, beam pattern shapes, and the lack of receiver beam training. Our system uses particle filtering along with Fourier analysis to tackle these practical challenges. In contrast to earlier work, our system does not change the operation of IEEE 802.11ad, and thus allows for simultaneous communications and localization. We modify the firmware of a commercial IEEE 802.11ad router to implement and validate our approach. The system operates in real-time and achieves sub-meter accuracy in 70% of the cases, which is a remarkable results given the limited capabilities of the hardware.

ACKNOWLEDGMENTS

This work has been supported in part by the ERC project SEARCHLIGHT grant no. 617721, the Ramon y Cajal grant RYC-2012-10788, the grant TEC2014-55713-R (Hyperadapt),




Fig. 12. CDFs: Auditorium, single Antenna




Fig. 13. CDFs: Auditorium, 4 antennas




Fig. 14. CDFs: Office, single antenna




Fig. 15. CDFs: Office, 4 antennas




Fig. 16. Auditorium scenarios

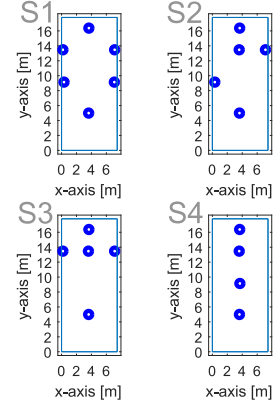


Fig. 17. Office scenarios

the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919), the German Research Foundation (DFG) within the Collaborative Research Center (CRC) 1119 “CROSSING — Cryptography-Based Security Solutions”, the German Federal Ministry of Education and Research (BMBF), the State of Hesse within CRISP-DA, and the Hessian LOEWE excellence initiative within NICER.

REFERENCES

- [1] T. Nitsche *et al.*, “Boon and bane of 60 GHz networks: Practical insights into beamforming, interference, and frame level operation,” in *Proc. ACM CoNEXT*, Heidelberg, Germany, Dec. 2015.
- [2] T. Wei and X. Zhang, “mTrack: High-precision passive tracking using millimeter wave radios,” in *Proc. ACM MobiCom*, Paris, France, 2015.
- [3] A. Olivier *et al.*, “Lightweight indoor localization for 60-GHz millimeter wave systems,” in *Proc. IEEE SECON*, London, UK, Jun. 2016.
- [4] 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.
- [5] D. Steinmetzer *et al.*, “Compressive millimeter-wave sector selection in o-the-shelf IEEE 802.11 ad devices,” in *Proc. ACM CoNEXT*, 2017.
- [6] H. Liu *et al.*, “Survey of wireless indoor positioning techniques and systems,” *IEEE Trans. Syst., Man, Cybern. C*, vol. 37, no. 6, pp. 1067–1080, Nov. 2007.
- [7] F. Lemic *et al.*, “Localization as a feature of mmWave communication,” in *Proc. IWCMC*, Paphos, Cyprus, Sep. 2016.
- [8] 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.
- [9] B. Cook *et al.*, “Indoor location using trilateration characteristics,” in *Proc. London Communications Symposium*, London, UK, Sep. 2005.
- [10] A. Zanella and A. Bardella, “RSS-based ranging by multichannel RSS averaging,” *IEEE Wireless Commun. Lett.*, vol. 3, no. 1, pp. 10–13, 2014.

- [11] M. Rea *et al.*, “Filtering noisy 802.11 time-of-flight ranging measurements from commoditized wifi radios,” *IEEE/ACM Transactions on Networking*, vol. PP, no. 99, pp. 1–14, 2017.
- [12] J. Xiong and K. Jamieson, “Arraytrack: A fine-grained indoor location system,” in *Proc. USENIX*, Lombard, IL, Apr. 2013.
- [13] R. Nandakumar, K. K. Chintalapudi, and V. Padmanabhan, “Centaur: Locating devices in an office environment,” in *Proc. ACM Mobicom*, Istanbul, Turkey, Aug. 2012.
- [14] C. Zhang *et al.*, “iLocScan: Harnessing multipath for simultaneous indoor source localization and space scanning,” in *Proc. ACM SenSys*, Memphis, TN, Nov. 2014.
- [15] C. Gentner and T. Jost, “Indoor positioning using time difference of arrival between multipath components,” in *Proc. IPIN*, Montbeliard-Belfort, France, Oct. 2013.
- [16] Alex Mariakakis *et al.*, “SAIL: Single access point-based indoor localization,” in *Proc. ACM MobiSys*, Bretton Woods, NH, Jun. 2014.
- [17] 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.
- [18] A. Shahmansoori *et al.*, “5G position and orientation estimation through millimeter wave MIMO,” in *Proc. IEEE GlobeCom*, Dec. 2015.
- [19] A. Jafari *et al.*, “NLOS influence on 60 GHz indoor localization based on a new TDOA extraction approach,” in *Proc. EuMC*, Nuremberg, Germany, Oct. 2013.
- [20] J. Chen *et al.*, “Pseudo lateration: Millimeter-wave localization using a single RF chain,” in *Proc. IEEE WCNC*, San Francisco, CA, Mar. 2017.
- [21] M. E. Rasekh *et al.*, “Noncoherent mmWave path tracking,” in *Proc. Hotmobile*, Sonoma, CA, Feb. 2017.
- [22] D. Steinmetzer, D. Wegemer, and M. Hollick. (2017) Talon tools: The framework for practical IEEE 802.11ad research. [Online]. Available: <https://seemoo.de/talon-tools/>
- [23] M. Schulz, D. Wegemer, and M. Hollick. (2017) Nexmon: The c-based firmware patching framework. [Online]. Available: <https://nexmon.org>
- [24] “Lede project.” [Online]. Available: <https://lede-project.org/>