

LEAP: Location Estimation and Predictive Handover with Consumer–Grade mmWave Devices

Joan Palacios^{‡*}, Paolo Casari[‡], Hany Assasa^{‡*}, Joerg Widmer[‡]

[‡]IMDEA Networks Institute, Madrid, Spain

^{*}Universidad Carlos III de Madrid, Spain

Abstract—Future millimeter-wave networks will support very high densities of devices and access points. This vastly increases the overhead required for access point selection and beam training. Fortunately, the quasi-optical properties of millimeter-wave channels make location-based network optimization a highly promising technique to reduce control overhead in such millimeter-wave WLANs. In this paper, we extract channel state information from off-the-shelf routers, we use it to design a high accuracy location system, and then show how location information enables the optimization of network operations. The resulting scheme, named LEAP, can predict blockage, optimize access point association, and select the most suitable antenna beam patterns while significantly reducing the beam training overhead. We show that compared to standard state-of-the-art 802.11ad systems, LEAP’s location driven management greatly improves network performance and link stability.

I. INTRODUCTION

Millimeter-wave (mmWave) communications have emerged as one of the most promising solutions to deliver multi-Gbit/s data rates to wireless devices. They are being standardized by 3GPP as a component of future mobile networks and have already become a WLAN standard through the introduction of IEEE 802.11ad operating in the 60 GHz band [1]. First commercial off-the-shelf (COTS) mmWave devices started appearing on the market, such as the Netgear Nighthawk X10 [2] and the TP-Link Talon AD7200 [3] routers that implement the 802.11ad standard.

Radio propagation at mmWave frequencies exhibits a quasi-optical behavior, i.e., the useful energy received by a mmWave receiver propagates from the transmitter over a limited number of paths, typically a line-of-sight path and/or several low-order reflections. It is also characterized by a high path loss, and devices use beam forming antennas to compensate this loss with a higher antenna gain. Such antennas require beam training [4] to choose the beam pattern configuration that achieves the best signal quality over a point-to-point link. 802.11ad devices can use hierarchical beam training with different beam widths [5], but faster solutions have been proposed that take advantage of the sparsity of the mmWave multi-path channel for compressive beam training [6], [7].

Most construction and furnishing materials are opaque to mmWaves, which implies that 802.11ad deployments experience blockage events and link outage may be frequent in the presence of obstructions [8], especially for mobile devices. Ubiquitous indoor mmWave coverage thus requires a very dense deployment of access points (APs) [9], so that ideally every location of an indoor space is served by more than one

AP. This way, the probability of outage becomes small and the link quality improves thanks to the lower average distance between an AP and its associated clients. When the density of APs and clients increases, beam training can constitute a significant source of overhead for the network. In principle, each client would have to perform beam training with each AP in its vicinity to determine which is the most suitable one. This makes optimal AP association a combinatorial problem, and hence scaling up the network density poses a significant challenge even with very efficient beam training mechanisms.

For such scenarios, location information is extremely beneficial, since it makes it possible to reduce or ideally avoid the beam training overhead. With a location system that provides device and AP positions with sufficient accuracy, it is possible to directly select the most suitable AP together with the beam pattern providing the highest gain in the known direction to that AP. By learning the environment, location systems also make it possible to prevent blockage by triggering handovers before the link to the associated AP is obstructed.

The idea of exploiting location information in order to improve network performance is envisioned in several works on indoor localization [10]–[13]. However, while these works advocate the advantages offered by location information, they focus mainly on the location system and its performance per se, without exploring the resulting network performance of location-driven beam training and AP association. Other works in the literature combine location information and network optimization, but are not targeting mmWave systems. For example, [14] leverages location data to optimize the load of indoor femto-cells, whereas [15] considers automatic peer node discovery as a possible networking-related application.

We address the above issues by implementing a 3D location system¹ that works solely based on the channel state information (CSI) extracted from the firmware of COTS mmWave devices, specifically Talon AD7200 routers. The CSI includes the amplitude and phase of the most powerful propagation path that links two mmWave devices. By inferring the specific layout and configuration of the Talon antenna arrays, and thus the steering vectors required to point the array towards a given direction, the CSI thus collected makes it possible to estimate the directions from which a signal arrives to each AP. From this angle information, we derive the 3D location both of the APs and of the mmWave clients, and are thus able to

¹2D location systems are suboptimal since mmWave APs are typically deployed below the ceiling to avoid blockage, whereas devices are at table height.

learn the mapping between the location of a device and the quality of the links between the client and the surrounding APs. This makes it possible to develop simple, location-based algorithms to determine not only the best AP a client should associate with, given its current location, but also which beam pattern the client should use to achieve the best link quality and throughput. The resulting system, named Location Estimation And Predictive handover (LEAP), enables location-based, automatic handover and beam training mechanisms that scale to very high device densities and network dynamics.

To the best of our knowledge, this work is the first to employ 3D location information extracted from mmWave communications (without interfaces to external systems) to perform mmWave network optimization for client-AP association, handover, and beam selection. It is also the first work to implement such a system and provide quantitative results for the performance gains of location-based network optimization.

Our evaluation is based on experimental data collected through a testbed of seven Talon routers working as mmWave APs. We deploy our testbed in a large indoor space with furniture and separating walls, located such that they obstruct the line-of-sight (LoS) path from one or more APs as the client moves around the room. Except for the extraction of the CSI to be processed by our location system, the router operation is unmodified and remains standard-compliant. We demonstrate decimeter-level AP localization accuracy and sub-meter client localization accuracy in more than 80% of the cases. We further show that the handover optimization and optimal beam pattern selection process provides a throughput improvement between 8.5% and 57% with respect to the 802.11ad standard, and that links experience 2 to 9 dB higher Signal-to-Noise Ratio (SNR).

Specifically, we provide the following contributions:

- *A mmWave location system* that autonomously localizes both mmWave APs and clients in a given environment via CSI measurements. For this, we efficiently estimate the physical propagation paths that best explain the measured CSI, and translate this information into the direction of departure of the signal from the APs to the client;
- *Location-based handover and beam training* by exploiting location information, such that a client is more likely to avoid blockage and connect to the most suitable AP using the beam pattern providing the best link performance;
- *An implementation of our LEAP scheme using COTS devices*, whose operation is unmodified except for the collection of CSI.

In the following, we present related work (Section II), the design of LEAP's location system (Section III), and of its optimized location-based handover and beam pattern selection schemes (Section IV). Next, we present our hardware platform and implementation (Section V), evaluate our scheme (Section VI), and draw concluding remarks in Section VII.

II. RELATED WORK

Indoor localization is becoming an increasingly important topic for networks, applications and services. Several schemes

and systems are compared in [16] under practical indoor conditions. Among the tested schemes, [17] uses improved likelihood functions that achieve good WiFi fingerprint matching accuracy, whereas [18] relies on arbitrary array processing to extract angle of arrival (AoA) information. Source localization is made possible through the processing of multipath receptions from different AoAs. Among the systems that achieve accurate localization in a WiFi context, Centaur [19] mixes radio and acoustic location fixes through Bayesian inference, in a way that is oblivious to the specific location system employed in each domain. ArrayTrack [20] exploits antenna arrays to suppress multipath, synthesize AoAs from different antenna subsets, and localize a client in a multi-AP scenario. Up to decimeter-level accuracy is achieved by Chronos [15], a system which localizes a client by exploring the phase of the WiFi signal at different frequencies to measure the time of flight (ToF) between each of its antennas and the client.

Fingerprinting-based methods have also been recently revived by crowdsourcing the collection of measurements for the fingerprint database. For example, in [21], where a significant contribution is the use of received signal strength indicator (RSSI) differences, which even out the different devices characteristics while preserving the accuracy of the radio map. Similarly, ARLS [22] exploits crowdsourced RSSI maps to achieve room-level client localization. The theoretical framework in [23] suggests that cooperation represents a good solution to improve localization accuracy, especially if AoA, RSSI and ToF information can be successfully merged with a client's sensors. Practical implementations of this concept include enhancing Bluetooth low-energy (BLE) RSSI measurements with inertial sensor readings [24], or using pedestrian dead reckoning jointly with WiFi and BLE. All of the above systems operate at frequencies below mmWave.

mmWave systems show significant potential to achieve very high positioning accuracy [25], both for the localization of a client, and for the estimation of the environment around it using multipath propagation [12]. Notably, this can be achieved using range-based or AoA-based methods, with no need to resort to special hardware, measurement-intensive fingerprinting or ToF methods typical of ultrawideband systems [26]. Since the small wavelength makes it possible to integrate a very large number of antenna elements in a comparatively small form factor, massive mmWave arrays are feasible, which in turn enable high localization accuracy even with random beamforming [27]. With somewhat smaller arrays, maximum likelihood estimation and the exploiting the sparsity of the MIMO channel matrix in mmWave scenarios have been shown to provide low location errors [28]. In general, the best accuracy and lowest localization errors are enabled by range-based algorithms [10], [29]; however, ranging requires accurate path loss models, whose parameters are environment-specific and may have to be re-trained over time. Multipath propagation can be exploited along with AoA information in order to localize a client [30] even with a single AP [31], although mmWave beamforming upon link establishment may turn out to sparsify the channel and reduce the number of

useful multipath components [8]. Given sufficiently many reflected multi-path components, a mmWave location system can also be used to map the environment [32].

Our work clearly differs from the above literature, as we do not just design a 3D location system, but also exploit it to optimize network-related functions such as handover and initial access by significantly reducing the beam training or feedback between the AP and the client [33]. Unlike previous theoretical approaches [34], we implement our solution on COTS devices, showing that the improvements of our location-based scheme can be achieved with consumer equipment.

III. LOCATION SYSTEM

A. Main idea

The main idea behind LEAP is as follows. We assume the presence of multiple APs in an indoor area. This is realistic given that indoor mmWave deployments are envisioned to be dense [9]. As a client moves, the APs measure the CSI of the link that connects them to the user. Multiple CSI measurements are converted into angle information in order to estimate the 3D location both of the APs and of the client. For each AP, we train a regression tree to learn the mapping between the location of the client and the SNR of its link to that AP. These regression trees help us decide, for any client location, which AP the client should connect to and which beam should be used. No coordination or feedback between the APs and the client is required, and the visibility of just a few APs at the same time (typically 2-3) is sufficient to achieve good results. Additionally, being able to select the correct beam based only on the locations of the APs' and the client significantly reduces the beam training time. The central result of our work is that a location system based only on information readily available to the physical layer of a set of mmWave APs suffices to enable location-aided beam training and handovers, greatly improving the network performance compared to the standard IEEE 802.11ad protocol.

B. Client localization algorithm

The CSI provided by the Talon router's firmware consists of one complex gain per antenna for the strongest multipath component of the received signal [35]. As a first step, our 3D location system converts the CSI into directions of arrival at each AP. Our formulation works in any number of dimensions, and is used to localize both the APs and the client in 3D space. We assume that each AP has a LoS connection to at least two APs, as needed to estimate the AP orientations. The quasi-optical behavior of the mmWave channel implies that the energy collected by a given node has reached that node through a limited number of propagation paths. Call L the number of paths. We can define the uplink channel matrix towards AP a as

$$\mathbf{H} = \sum_{\ell=1}^L \alpha_{\ell} \mathbf{s}_a(\mathbf{v}_{\ell}) \mathbf{s}_C(\boldsymbol{\xi}_{\ell})^H, \quad (1)$$

where the superscript H denotes conjugate-transpose, α_{ℓ} is the complex power gain of path ℓ (which includes path loss and

phase shift), \mathbf{v}_{ℓ} and $\boldsymbol{\xi}_{\ell}$ are the unitary vectors that define the direction of arrival at the AP and the direction of departure at the user of the ℓ th path, whereas $\mathbf{s}_a(\mathbf{v}_{\ell})$ and $\mathbf{s}_C(\boldsymbol{\xi}_{\ell})$ are the steering vectors of the array of AP a and of the client array that point towards directions \mathbf{v}_{ℓ} and $\boldsymbol{\xi}_{\ell}$, respectively. The CSI value measured by the receiver depends on the beam pattern \mathbf{p} employed by the user, as well as on the above quantities. Specifically, the following vector represents the receive channel measured by AP a :

$$\mathbf{h}_a = \sum_{\ell=1}^L \alpha_{\ell} \mathbf{s}_a(\mathbf{v}_{\ell}) \mathbf{s}_C(\boldsymbol{\xi}_{\ell})^H \mathbf{p} = \sum_{\ell=1}^L \beta_{\ell} \mathbf{s}_a(\mathbf{v}_{\ell}), \quad (2)$$

where we define the scalar value $\beta_{\ell} = \alpha_{\ell} \mathbf{s}_C(\boldsymbol{\xi}_{\ell})^H \mathbf{p}$. The vector \mathbf{h}_a is the CSI measured by AP a . Assuming that such CSI is mostly explained by a single main path [35], we can now estimate the vector of arrival $\bar{\mathbf{v}}_a$ (from the user to the receiving AP) that corresponds to such main path, and thus best explains the measured CSI. Assuming that the CSI is affected by circularly complex Gaussian noise, the estimate of $\bar{\mathbf{v}}_a$ is obtained as

$$\bar{\mathbf{v}}_a = \arg \min_{\mathbf{v}} \min_{\beta} \|\mathbf{h}_a - \beta \mathbf{s}_a(\mathbf{v})\|^2. \quad (3)$$

where \mathbf{h}_a is as in (2). In the right-hand side of (3), we have

$$\|\mathbf{h}_a - \beta \mathbf{s}_a(\mathbf{v})\|^2 = \|\mathbf{h}_a\|^2 + |\beta|^2 \|\mathbf{s}_a(\mathbf{v})\|^2 - 2\Re\{\beta \mathbf{h}_a^H \mathbf{s}_a(\mathbf{v})\}, \quad (4)$$

where $\Re\{\cdot\}$ denotes the real part. The last term on the right-hand-side of (4) is the only one affected by the phase of β , and its minimum is equal to $-2|\beta| |\mathbf{h}_a^H \mathbf{s}_a(\mathbf{v})|$. It thus remains to minimize the quadratic expression

$$\|\mathbf{h}_a\|^2 + |\beta|^2 \|\mathbf{s}_a(\mathbf{v})\|^2 - 2|\beta| |\mathbf{h}_a^H \mathbf{s}_a(\mathbf{v})|, \quad (5)$$

whose minimum value is $\|\mathbf{h}_a\|^2 - |\mathbf{h}_a^H \mathbf{s}_a(\mathbf{v})|^2 / \|\mathbf{s}_a(\mathbf{v})\|^2$. Hence (3) simplifies to

$$\bar{\mathbf{v}}_a = \arg \min_{\mathbf{v}} \left(\|\mathbf{h}_a\|^2 - \frac{|\mathbf{h}_a^H \mathbf{s}_a(\mathbf{v})|^2}{\|\mathbf{s}_a(\mathbf{v})\|^2} \right) = \arg \max_{\mathbf{v}} \frac{|\mathbf{h}_a^H \mathbf{s}_a(\mathbf{v})|}{\|\mathbf{s}_a(\mathbf{v})\|}. \quad (6)$$

Eq. (6) can be solved once we know $\mathbf{s}_a(\mathbf{v})$ for a sufficiently large set of directions \mathbf{v} . This requires knowledge of the array topology of all routers, which can be obtained either from antenna modeling or by measuring a sufficiently large set of steering vectors in an anechoic chamber. Assume for the moment that we know the AP locations \mathbf{y}_a , where $1 \leq a \leq N_A$ is the index of the AP, and N_A is the total number of APs deployed in the indoor space. Given $\bar{\mathbf{v}}_a$, we can compute the location \mathbf{x} of the user. Ideally, if we knew the distance d_a between \mathbf{x} and \mathbf{y}_a , we could directly compute the location of the user as $\mathbf{x} = \mathbf{y}_a + d_a \bar{\mathbf{v}}_a$. Since our method is range-free, we have no means of estimating this distance directly. We solve this issue by formulating the location estimation step as a minimum mean-square error (MMSE) problem, where we obtain the estimated user location $\bar{\mathbf{x}}$ as the one that minimizes the difference between the actual location and the estimate $\mathbf{y}_a + d_a \bar{\mathbf{v}}_a$, for minimum distance d_a from any AP. Formally:

$$\bar{\mathbf{x}} = \arg \min_{\mathbf{x}} \min_{d_a, \bar{\mathbf{v}}_a} \sum_{a=1}^{N_A} \|\mathbf{y}_a + d_a \bar{\mathbf{v}}_a - \mathbf{x}\|^2. \quad (7)$$

Minimizing each of the terms $\min_{d_a} \|\mathbf{y}_a + \mathbf{v}_a d_a - \mathbf{x}\|^2$, by definition, means finding the minimum distance between \mathbf{x} and the line passing through \mathbf{y}_a with direction \mathbf{v}_a . By applying the equation for the distance between a point and a line, we get

$$\begin{aligned} \min_{d_a} \|\mathbf{y}_a + \mathbf{v}_a d_a - \mathbf{x}\|^2 &= \|\mathbf{y}_a - \mathbf{x}\|^2 - ((\mathbf{y}_a - \mathbf{x})^\top \mathbf{v}_a)^2 \quad (8) \\ &= \|\mathbf{y}_a - \mathbf{x}\|^2 - (\mathbf{y}_a - \mathbf{x})^\top \mathbf{v}_a \mathbf{v}_a^\top (\mathbf{y}_a - \mathbf{x}) \\ &= (\mathbf{y}_a - \mathbf{x})^\top (\mathbf{I} - \mathbf{v}_a \mathbf{v}_a^\top) (\mathbf{y}_a - \mathbf{x}) \\ &= \mathbf{x}^\top (\mathbf{I} - \mathbf{v}_a \mathbf{v}_a^\top) \mathbf{x} - 2\mathbf{y}_a^\top (\mathbf{I} - \mathbf{v}_a \mathbf{v}_a^\top) \mathbf{x} + \\ &\quad + \mathbf{y}_a^\top (\mathbf{I} - \mathbf{v}_a \mathbf{v}_a^\top) \mathbf{y}_a. \quad (9) \end{aligned}$$

Call $\mathbf{U} = N_A \mathbf{I} - \sum_{a=1}^{N_A} \mathbf{v}_a \mathbf{v}_a^\top$, $\mathbf{r} = \sum_{a=1}^{N_A} (\mathbf{I} - \mathbf{v}_a \mathbf{v}_a^\top) \mathbf{y}_a$ and $z = \sum_{a=1}^{N_A} \mathbf{y}_a^\top (\mathbf{I} - \mathbf{v}_a \mathbf{v}_a^\top) \mathbf{y}_a$. By substituting (9) into (7), we can rewrite (7) as

$$\bar{\mathbf{x}} = \arg \min_{\mathbf{x}} \mathbf{x}^\top \mathbf{U} \mathbf{x} - 2\mathbf{r}^\top \mathbf{x} + z. \quad (10)$$

This MMSE problem has solution $\bar{\mathbf{x}} = \mathbf{U}^{-1} \mathbf{r}$, and error

$$\varepsilon_{\bar{\mathbf{x}}} = z - \mathbf{r}^\top \mathbf{U}^{-1} \mathbf{r}. \quad (11)$$

C. AP localization algorithm

With (11), we can now formulate the AP localization algorithm. We first observe that (11) is in quadratic form. Define \mathbf{y} as the vector containing the vertical concatenation of the coordinates of all access points $1 \leq a \leq N_A$. With a suitable re-arrangement of the terms in (11), we can write

$$\varepsilon_{\bar{\mathbf{x}}} = \mathbf{y}^\top \mathbf{Q} \mathbf{y}, \quad (12)$$

where matrix \mathbf{Q} can be measured, and its terms depend on the CSI collected by each router. Call t the time epoch when a measurement is collected, and call \mathbf{Q}_t the corresponding matrix. We can estimate the vector of the AP locations \mathbf{y} as the one that minimizes (12) by taking into consideration all measurements carried out over time. This leads to the following minimization problem:

$$\bar{\mathbf{y}} = \arg \min_{\mathbf{y}} \mathbf{y}^\top \sum_t \mathbf{Q}_t \mathbf{y} \quad (13a)$$

$$\text{s.t. } \|\mathbf{y}\| = 1 \quad (13b)$$

$$\sum_{a=1}^{N_A} [\mathbf{y}]_{3a-2} = 0 \quad (13c)$$

$$\sum_{a=1}^{N_A} [\mathbf{y}]_{3a-1} = 0 \quad (13d)$$

$$\sum_{a=1}^{N_A} [\mathbf{y}]_{3a} = 0, \quad (13e)$$

where the notation $[\cdot]_k$ denotes the element in position k of the vector in the square brackets. Call \mathbf{T} any orthonormal parametrization of the subspace defined by the linear constraints (13c)–(13e); then \mathbf{T} defines a bijection $\mathbf{y} = \mathbf{T} \tilde{\mathbf{y}}$ between the restricted space of \mathbf{y} and a compressed version in a lower-dimensional space. In this way, the problem in (13) can be expressed as

$$\bar{\mathbf{y}} = \mathbf{T} \arg \min_{\tilde{\mathbf{y}}} (\mathbf{T} \tilde{\mathbf{y}})^\top \sum_t \mathbf{Q}_t (\mathbf{T} \tilde{\mathbf{y}}) \quad (14a)$$

$$\text{s.t. } \|\mathbf{T} \tilde{\mathbf{y}}\| = 1. \quad (14b)$$

Since \mathbf{T} is orthonormal, $\|\mathbf{T} \tilde{\mathbf{y}}\| = \|\tilde{\mathbf{y}}\|$ and (14) becomes

$$\bar{\mathbf{y}} = \mathbf{T} \arg \min_{\tilde{\mathbf{y}}} \tilde{\mathbf{y}}^\top \left(\mathbf{T}^\top \sum_t \mathbf{Q}_t \mathbf{T} \right) \tilde{\mathbf{y}} \quad (15a)$$

$$\text{s.t. } \|\tilde{\mathbf{y}}\| = 1, \quad (15b)$$

which is a minimum eigenvector determination problem. Hence the vector that minimizes (15) is $\mathbf{E}_{\min}[\mathbf{T}^\top \sum_t \mathbf{Q}_t \mathbf{T}]$, where the $\mathbf{E}_{\min}[\cdot]$ operator returns the minimum eigenvector of its argument, and the solution to (13) is

$$\bar{\mathbf{y}} = \mathbf{T} \mathbf{E}_{\min} \left[\mathbf{T}^\top \sum_t \mathbf{Q}_t \mathbf{T} \right]. \quad (16)$$

D. Data smoothing

In order to remove outliers in the measurements and improve the data and location estimation quality, we implement three filtering steps. First, we ensure that the CSI fed to the localization algorithm is being measured from the LoS path. We thus design a filter to determine which measurements correspond to LoS paths and which ones should instead be discarded. A measurement is kept if the measured path satisfies two conditions: *i*) that the path power should exceed a threshold, i.e., $\|\mathbf{h}_a\| > \theta_P$, and *ii*) that the proportion of the measured power explained by the estimated path should also exceed a threshold, i.e., $|\mathbf{h}_a^\top \mathbf{s}_a(\bar{\mathbf{v}}_a)| / (\|\mathbf{h}_a\| \|\mathbf{s}_a(\bar{\mathbf{v}}_a)\|) > \theta_A$, where, e.g., $\theta_A = 0.9$.

Second, spurious CSI measurements provided by the APs should be filtered out to avoid that they contaminate the location estimates. We achieve this by eliminating the measurements that are not in agreement with the AP location estimates. We identify such spurious estimates as those for which the vector of departure $\bar{\mathbf{x}} - \bar{\mathbf{y}}_a$ remains far from the estimated direction of departure $\bar{\mathbf{v}}_a$, i.e., for which the cosine of the angle between the two vectors is small. We therefore check if $\|\bar{\mathbf{x}} - \bar{\mathbf{y}}_a\|^{-1} (\bar{\mathbf{x}} - \bar{\mathbf{y}}_a)^\top \mathbf{v}_a < \cos \epsilon$. If this is the case, we iteratively discard each measurement for which the left-hand-side of the above inequality is smallest (i.e., those leading to the largest discrepancies between $\bar{\mathbf{x}} - \bar{\mathbf{y}}_a$ and $\bar{\mathbf{v}}_a$), until all remaining measurements exceed the threshold.

In the same vein, we filter out client location errors as follows. We check if removing any of the AP coordinates included in vector \mathbf{y} in (12) and the corresponding lines and columns of \mathbf{Q} reduces the error $\varepsilon_{\bar{\mathbf{x}}}$ at least by a factor of 2. We then iteratively eliminate APs until we find no more that can be removed to improve the error, or until we are left with just two APs, which is the minimum number required to estimate the client position.

E. Mobility model and trajectory smoothing

We smooth the trajectory of the client via a rank-deficient Kalman filter. We consider the following mobility model

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \Delta t \mathbf{w}_t + \mathbf{x}_n \quad \mathbf{x}_n \sim \mathcal{N}(\mathbf{0}, \Delta t \sigma_x^2 \mathbf{I}) \quad (17)$$

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \mathbf{w}_n \quad \mathbf{w}_n \sim \mathcal{N}(\mathbf{0}, \Delta t \sigma_v^2 \mathbf{I}), \quad (18)$$

where \mathbf{x}_t is the location of the client at time t , \mathbf{w}_t its movement speed, Δt is the time interval between subsequent

measurements, and σ_x^2 and σ_v^2 are the variance of each component of the position and velocity vector, respectively. The model in (17)–(18) correspond to uniform linear motion with location error \mathbf{x}_n and speed error \mathbf{w}_n . Define the client state as the concatenation of the position and speed, i.e., $\mathbf{z}_t = [\mathbf{x}_t^\top, \mathbf{w}_t^\top]^\top$, and define the matrices

$$\mathbf{E} = \begin{bmatrix} \mathbf{I} & \Delta t \mathbf{I} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad \Sigma_z = \Delta t \begin{bmatrix} \sigma_x^2 \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \sigma_v^2 \mathbf{I} \end{bmatrix}, \quad (19)$$

from which we can compute the state evolution as

$$\mathbf{z}_{t+1} = \mathbf{E}\mathbf{z}_t + \mathbf{z}_n, \quad \mathbf{z}_n \sim \mathcal{N}(\mathbf{0}, \Sigma_z). \quad (20)$$

If $\mathbf{z}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \Sigma_t)$, we obtain the evolution formulas

$$\boldsymbol{\mu}_{t+1} = \mathbf{E}\boldsymbol{\mu}_t, \quad \Sigma_{t+1} = \mathbf{E}\Sigma_t\mathbf{E}^\top + \Sigma_z, \quad (21)$$

where $\boldsymbol{\mu}_0$ and Σ_0 are initialized with the first measurement's value and uncertainty, respectively. By modeling the likelihood of the measurements as a normal distribution with mean $\bar{\mathbf{x}}_t$ and inverse covariance \mathbf{U}_t^{-1} , the posterior distribution of \mathbf{z}_t is

$$\bar{\boldsymbol{\mu}}_t = \bar{\Sigma}_t \left(\Sigma_t^{-1} \boldsymbol{\mu}_t + \begin{bmatrix} \mathbf{U}_t \bar{\mathbf{x}}_t \\ \mathbf{0} \end{bmatrix} \right), \quad \bar{\Sigma}_t = \left(\Sigma_t^{-1} + \begin{bmatrix} \mathbf{U}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \right)^{-1}, \quad (22)$$

so that finally $\bar{\boldsymbol{\mu}}_t$ predicts the client's position and velocity at time t . We remark that only a few measurements from typically 2-3 visible APs are sufficient to achieve a good position estimate. When location estimation should fail, the Kalman filter can be used for *dead reckoning*: namely, when a location cannot be computed from the CSI measurements, the filter automatically fills the gap by assuming that the client maintained the same speed vector.

IV. USING LOCATION INFORMATION FOR HANDOVER AND BEAM SELECTION

The accurate location system developed in Section III opens the way to extremely lightweight schemes to select the best AP. Specifically, we predict how the SNR of the link between the client and each AP will evolve over time, and use this information to rank the link quality with each AP and beam pattern. Considering a specific AP a , we proceed as follows.

For each measurement, after computing the posterior distribution of \mathbf{z}_t in (22), we predict the location of the client over time for a given number N_T of future epochs spaced by a time interval τ . Call these locations $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_{N_T}$. For each location $\hat{\mathbf{x}}_i$, we also compute the inverse square of the distance from AP a , namely $d_i^{-2}(a) = \|\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_a\|^{-2}$. This serves as a measure of link quality. We pass the predicted locations of the client and the inverse squared distance from the AP to a regression tree trained to map these features and the SNR that would be experienced at any location. This yields N_T predicted SNR values for AP a , $\gamma_1, \dots, \gamma_{N_T}$. Finally, we compute the score assigned to AP a as

$$\kappa_a = \sum_{n=1}^{N_T} \exp(-n\tau/\lambda) \gamma_n, \quad (23)$$

where λ is a custom parameter that regulates the exponential weighing of future predicted SNR values. After carrying out the above procedure for all APs, we finally instruct the client

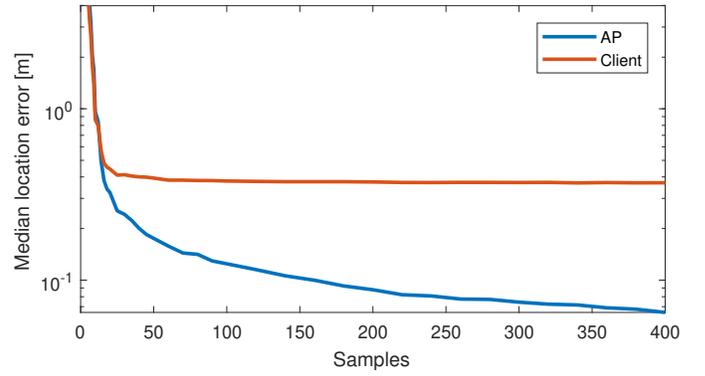


Fig. 1. Median location estimation error in meters as a function of the number of measurements collected. About 15 measurements are sufficient to achieve sub-meter localization accuracy both for the client and for the APs.

to connect to AP $\hat{a} = \arg \max_a \kappa_a$. By knowing the location of the AP and of the client, LEAP can potentially select the right beam without probing any APs. In practice, up to 4 or 9 beams are tested in order to ensure a correct handover in the presence of location errors. This is still a substantial reduction with respect to exhaustive or hierarchical beam training.

The system is bootstrapped by collecting enough measurements to estimate the location of the APs. LEAP achieves sub-meter location errors already with very few measurements. This is demonstrated by Fig. 1, which shows the evolution of the median localization error as a function of the number of measurements collected. We observe that even when no initial information is available to the user about its own location or the location of the APs, about 15 measurements are sufficient to achieve sub-meter median accuracy. Further measurements improve the AP locations even further, achieving an accuracy below 40 cm for about 50 measurements, and converging to sub-decimeter accuracy in the long run. In contrast, due to client mobility, only a limited set of measurements is available for each new client location and thus the median location error for the client converges to approximately 40 cm.

As soon as AP locations become available, we train a regression tree to learn the mapping between the user location and the SNR provided by each AP. Initially, the data set employed to train the regression tree is composed of the collected measurements, enriched with simulated data derived from a channel model. The simulated data is iteratively expunged from the training dataset and substituted with measurements as they become available, until the regression tree can be trained using only measured data.

V. HARDWARE AND METHODOLOGY

We evaluate LEAP using TP-Link Talon AD7200 APs [3]. The routers integrate the Qualcomm QCA9500 chipset that implements the IEEE 802.11ad standard. A 32-element antenna array mounted on one of the eight foldable antenna wings provides beamforming capabilities. We modify the firmware to provide access to CSI measurements. To achieve the required level of control, we flash a compact LEDE linux distribution to the router [36]. LEDE uses the open-source wil6210 driver

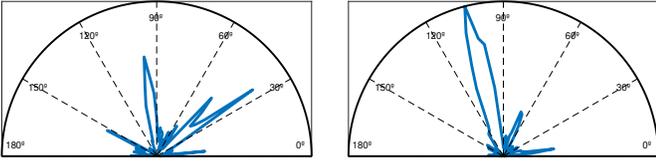


Fig. 2. (Left) Beam pattern from the original Talon router codebook; (Right) directional beam pattern in our improved codebook. Amplitudes are plotted relative to the highest lobe of the beam pattern on the right.

to communicate with the firmware of the Qualcomm chipset. We implement custom vendor commands in the driver to expose CSI measurements to user-space programs. The router measures the CSI by sending 64 training subfields during a beam refinement phase, which make it possible to estimate both phase and amplitude of all 32 elements of the mmWave antenna array. With the Qualcomm firmware, only client-side CSI measurements are enabled. We circumvent this issue by switching the roles of the APs and of the client, such that the node to be located is actually the AP, whereas all other devices work in client mode and can thus measure CSI. Besides CSI, the driver provides the sector ID used for communications and the SNR measured for that sector.

We measure MAC-level throughput using the `iperf3` tool available in LEDE. In addition, the firmware makes it possible to implement custom 3D beam patterns by selecting which antenna elements to enable and specifying the phase shift values for each activated antenna with a 2-bit quantization. We prepare a codebook of 64 beam patterns that have significantly more directional main lobes than the default ones used by the router’s firmware. An example is provided in Fig. 2: the left panel depicts one of the standard beam patterns of the original Talon router codebook, whereas the right panel shows our directional beam pattern. With these functionalities, it is possible to choose at any given time which beam pattern an AP and a device should use to communicate, thus greatly reducing the overhead of the beam training process through a location-aided beam pattern selection.

VI. EXPERIMENTAL RESULTS

A. Experiment setup

Our measurement scenario consists of a laboratory room of size about $17\text{ m} \times 6\text{ m}$. We deploy a total of seven Talon routers along the perimeter of the room, at different heights. As the client to be located, we employ an additional Talon router. Two wood/metal panels of size $2\text{ m} \times 2\text{ m}$ are located around the center of the room. These panels induce a complete signal blockage if they obstruct the LoS path between the client and its AP, e.g., due to client mobility.

We collect measurements by moving the Talon acting as the client along a number of trajectories. At subsequent points spaced between 40 and 60 cm along such trajectories, we collect CSI, SNR and MAC throughput measurements from all APs. Fig. 3 gives an overview of the laboratory setup. Access points are depicted as purple triangles, internal panels are shown in gray, and we depict three example trajectories

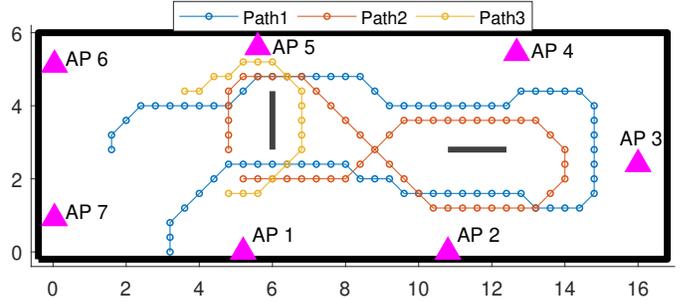


Fig. 3. Experiment scenario: a lab with two internal panels acting as blocking walls. Seven APs are deployed along the boundaries of the room (purple triangles). Three examples of client trajectories are shown. Each marker represents a measurement point. The axes show lengths in meters.

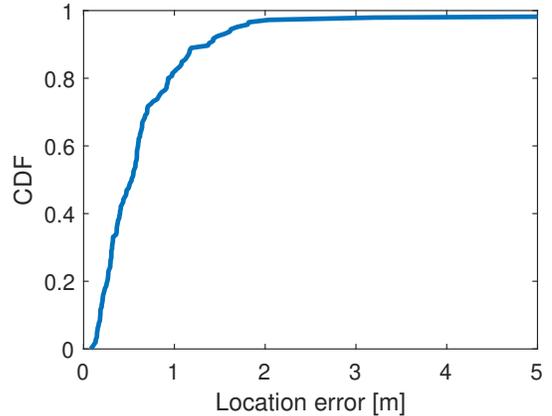


Fig. 4. CDF of LEAP’s location error for all measurements. Sub-meter accuracy is achieved more than 80% of the time.

of our measurement collection process: path 1, that mostly follows the perimeter of the room and revolves around the two panels; path 2, that describes a figure-of-8 around the panels; and path 3, that describes an open loop close to the left side of the room. We remark that it is not necessary to know the location of the APs, which will be estimated by our scheme and fed to the location-based handover algorithm.

For LEAP, we select the best beam pattern among four (LEAP-4) or nine (LEAP-9) tested beam patterns covering the optimum transmit direction predicted through the location system. We compare our performance against that of standard 802.11ad with conventional beam training. For each measurement point, we let the 802.11ad beam training settle upon the best beam pattern. If the connection is lost and throughput drops to zero, we let the device reconnect to a new AP at the next step and perform a full beam training. The performance for 802.11ad thus represents an upper bound on actual 802.11ad protocol performance with the Talon routers, which is characterized by sub-optimal beam fluctuations in case of mobility and noticeable reconnection delays when performing a handover to a different AP.

B. Results

We start by evaluating the accuracy of LEAP’s localization system. Fig. 4 shows the cumulative distribution function (CDF) of the localization error, obtained by jointly considering

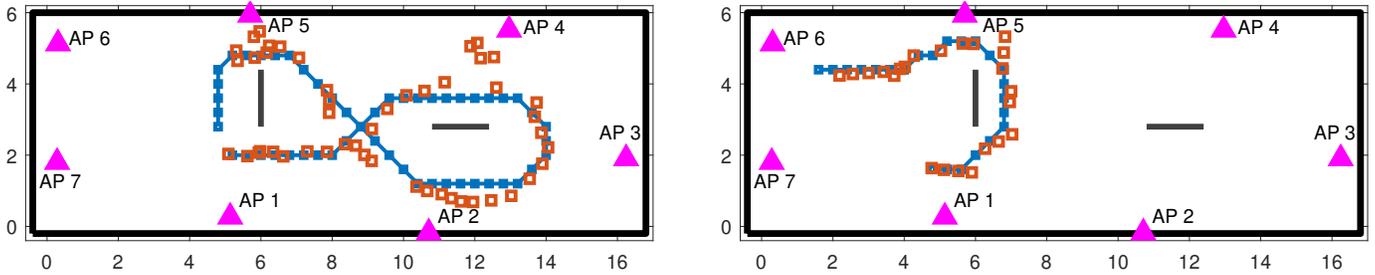


Fig. 5. Path reconstruction for paths 2 (left) and 3 (right) in Fig. 3. The location estimates follow the ground truth reasonably well, and the errors are not critical for the selection of the best AP. The axes show lengths in meters.

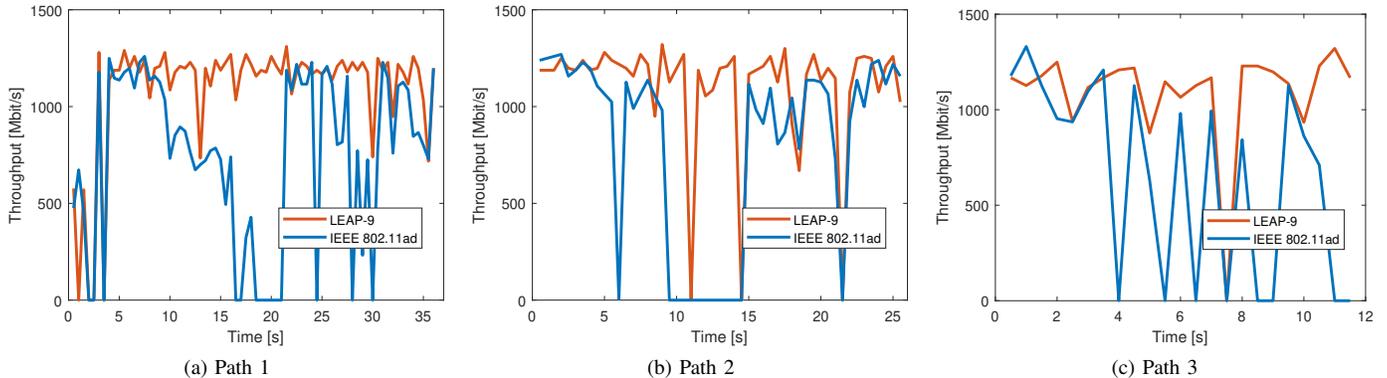


Fig. 6. Measured MAC-level throughput for LEAP against the throughput achieved by IEEE 802.11ad.

all measurement trajectories in the experimental area. Despite the comparatively stretched room and the presence of internal blocking panels, the location system achieves sub-meter error in more than 80% of the cases, with a median location error of about 52 cm. Given that our method only uses angle information (unlike, e.g., [10]) and localizes the client in 3D (unlike, e.g., [12]), these results are very accurate. Moreover, as we will show below, this level of accuracy is more than sufficient to achieve good performance and significant gains in the location-aided handover process.

Fig. 5 provides more details of LEAP’s location errors by showing the reconstruction of two trajectories among those depicted in Fig. 4, namely path 2 (left panel) and path 3 (right panel). Path 2 is a figure-of-8 trajectory that moves around both obstacles and thus experiences several changes in terms of AP visibility. In several cases, the localization error remains around a few decimeters. However, in other cases the collected CSI either leads to some errors (e.g., in the proximity of AP 4) or is insufficient to localize the node and has to be compensated for via the Kalman filter’s dead reckoning capabilities (central section of the trajectory). Path 3 is comparatively simpler and subject to fewer handovers, except when the client disappears behind the internal panel and reappears on the other side. In any event, the trajectory reconstruction is very good, and the closeness of the estimated and actual locations makes it possible to still associate to the correct AP. For example, even for comparatively large errors near the top-right section of path 2, the client correctly connects to AP 4. Similar considerations apply, e.g., to the rightmost section of path 3.

We now show results for a client moving along paths 1 to 3, and plot the MAC-level throughput achieved over time by LEAP against the performance of 802.11ad in Fig. 6. All trajectories are covered with the mmWave antenna array of the device facing forward, which therefore changes the orientation of the client at each curve. Note that firmware and hardware limitations prevent the Talon router from consistently maintaining transmission rates above 1.25 Gbps, and it does not use the fastest modulation and coding schemes. In any event, a conservative rate adaptation strategy makes sense, as the router only provides external connectivity through a 1-Gbps Ethernet port. Starting from path 1, we observe that the movement along its trajectory leads to several blockage events, which often cause reduced performance or even outage to 802.11ad. This is due to the obstruction of the LoS path to AP 1, to which the client connects at the beginning of the path. In the period from 7 to 22 s, 802.11ad maintains the connection to the same AP, which becomes progressively farther and gets blocked by obstacles, which reduces the throughput and causes even prolonged outages. With reference to the shape of path 2 in Figs. 3 and 5, initially both LEAP and 802.11ad connect to AP 1. As the client moves upwards in a clockwise direction, 802.11ad maintains the connection to AP 1, and a long outage occurs when the AP disappears behind the horizontal panel, at about 9 s. Conversely, LEAP performs a handover first to AP 4, and then to APs 3, 2 and 5. At the same time, AP 1 disappears again behind the horizontal panel, causing the 802.11ad outage at about 22 s. By handing the client connection over to the correct AP throughout the trajectory, LEAP successfully avoids most outage events, and in many

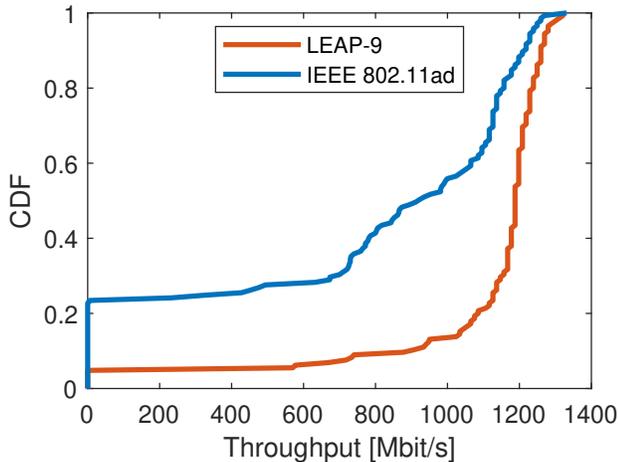


Fig. 7. CDF of the MAC-level throughput for LEAP and 802.11ad. LEAP achieves better throughput at a much lower outage probability than 802.11ad.

cases achieves better throughput than 802.11ad even in the absence of outages. In Fig. 6b, the instances where LEAP’s throughput decreases marginally below 802.11ad’s are due to a location estimate that leads to the choice of a suboptimal AP and beam pattern.

For path 3 (Fig. 6c), LEAP connects to AP 5 at about 4 s (which is when the first outage occurs for 802.11ad), and then to AP 6, and suffers from a single outage at about 7.5 s. Conversely, suboptimal AP selection makes 802.11ad suffer from repeated outages. As a result, the throughput of LEAP is much closer to the maximum throughput achievable by the device throughout the whole trajectory.

The CDF of the MAC-level throughput computed over all tested trajectories is shown in Fig. 7. LEAP shows a much smaller outage probability (about 0.05) than 802.11ad (about 0.23). In fact, in about 85% of the cases, LEAP’s throughput remains above 1 Gbps. For the cases where the 802.11ad is not in outage, LEAP still provides a throughput improvement between 8.5% and 57%. We remark that this improvement is actually limited by the hardware capabilities, not by LEAP’s performance.

Therefore, to show the actual performance gain enabled by our algorithm, in Fig. 8 we compare the SNR achieved

by LEAP against that of 802.11ad, and against the optimum SNR that would be achievable by an oracle-aided selection of the best AP. The three panels refer to the same paths 1 to 3 as in Fig. 6. As a general observation, throughout each trajectory, the SNR achieved by both LEAP-4 and LEAP-9 remains much closer to the optimum, and in the majority of the cases is higher than the SNR achieved by 802.11ad. Also, the performance of LEAP-9 is only marginally better than that of LEAP-4. This indicates that the location-aided AP and beam pattern selection is sufficiently accurate, making it unnecessary to test many beam patterns. At the start of each path, LEAP and 802.11ad tend to connect to the same AP, hence their SNR is similar. However, as time elapses and the client moves, some APs become closer and thus more convenient to associate with, whereas others may disappear behind blockage. While LEAP keeps reacting to mobility and connects to (close to) optimal APs, 802.11ad re-associates to a different AP only upon severe blockage. The new AP is usually the first one from which a beacon is received, rather than the optimal one. As a consequence, 802.11ad’s SNR is often 5 to 15 dB below the optimum. Instead, even in the presence of imperfect location estimates, LEAP still achieves near-optimal SNR, and only in a few cases do such imperfect estimates cause its SNR to drop a few dB below that of 802.11ad.

The CDF of the achieved SNR for the whole ensemble of our measurements is shown in Fig. 9. These results confirm that both LEAP-4 and LEAP-9 remain within 3 dB from the optimum 75% of the times, and improve the SNR by 2 to 9 dB compared to 802.11ad. We believe that these results constitute a very promising demonstration of how mmWave network performance can be improved with the aid of a sufficiently precise location system. This is specifically significant considering that such performance is achieved with only angle information related to the LoS path as provided by consumer-grade mmWave devices, and that we compared against an idealized version of 802.11ad that does not suffer from issues related to mobility.

VII. CONCLUSIONS

We presented LEAP, a system that leverages location information to improve handover decisions and beam pattern selection using CSI and SNR measurements provided by

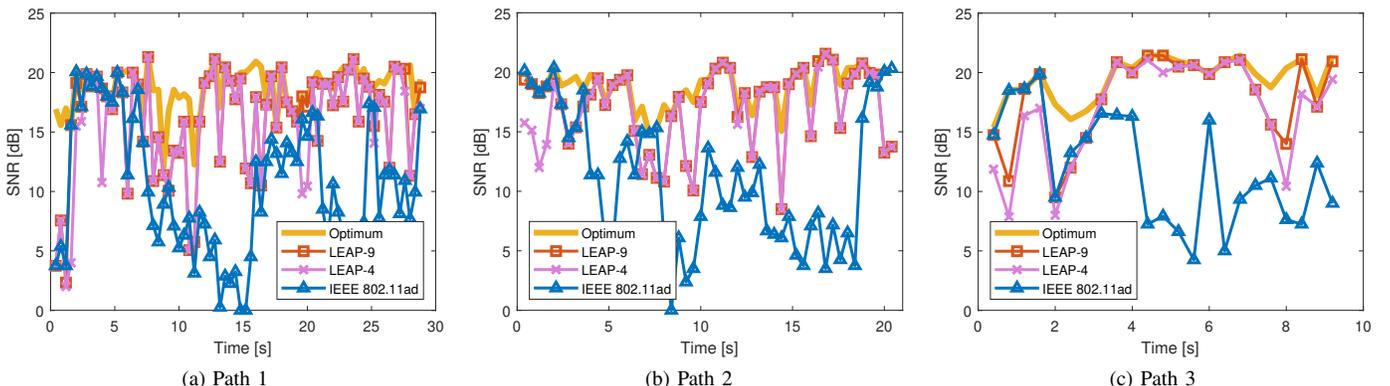


Fig. 8. Measured SNR for LEAP-4, LEAP-9 and IEEE 802.11ad compared to an optimum, oracle-aided AP and beam pattern selection scheme.

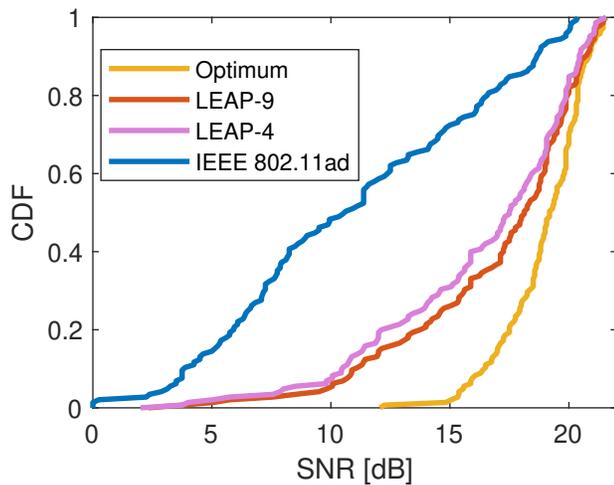


Fig. 9. CDF of the SNR of LEAP-4, LEAP-9 and 802.11ad compared to an oracle-aided AP and beam pattern selection scheme. LEAP's SNR is higher than 802.11ad's by 2 to 9 dB, and is typically within 3 dB of the optimum.

commercial mmWave devices. LEAP localizes a client with sub-meter accuracy in the great majority of the cases, greatly reduces the beam training overhead via location-based beam pattern selection, and avoids blockage-induced outage by using regression trees to choose the best AP to associate with. We implement our approach using Talon AD7200 routers in a way that does not modify 802.11ad's operation. Compared to standard 802.11ad, LEAP achieves a throughput gain between 8.5% and 57% and an SNR typically within 3 dB of the optimum. Most importantly, the beam training overhead reduction enables LEAP to scale to very high device densities and environment dynamics.

ACKNOWLEDGMENT

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), and the Regional Government of Madrid through the TIGRES-CM program (S2013/ICE-2919).

REFERENCES

- [1] C. Cordeiro *et al.*, "IEEE 802.11ad: Introduction and performance evaluation of the first multi-Gbps WiFi technology," in *Proc. ACM mmCom*, Chicago, IL, Sep. 2010.
- [2] "Nighthawk X10 smart WiFi router," <https://www.netgear.com/home/products/networking/wifi-routers/R8900.aspx>, checked Jan. 17, 2019.
- [3] "Talon AD7200 multi-band WiFi router," https://www.tp-link.com/us/products/details/cat-9_AD7200.html, checked Jan. 17, 2019.
- [4] J. Kim and A. F. Molisch, "Fast millimeter-wave beam training with receive beamforming," *J. of Commun. and Networks*, vol. 16, no. 5, pp. 512–522, Oct. 2014.
- [5] T. Nitsche *et al.*, "IEEE 802.11ad: Directional 60 GHz communication for multi-Gbps Wi-Fi," *IEEE Communication Magazine*, vol. 52, no. 12, pp. 132–141, Dec. 2014.
- [6] M. E. Rasekh *et al.*, "Noncoherent mmWave path tracking," in *Proc. HotMobile*, Sonoma, CA, Feb. 2017.
- [7] D. Steinmetzer *et al.*, "Compressive millimeter-wave sector selection in off-the-shelf IEEE 802.11ad devices," in *Proc. ACM CoNEXT*, Seoul, South Korea, Dec. 2017.
- [8] Z. Abu-Shaban *et al.*, "Performance of location and orientation estimation in 5G mmWave systems: Uplink vs downlink," in *Proc. IEEE WCNC*, Barcelona, Spain, Apr. 2018.

- [9] F. Firyaguna *et al.*, "Coverage and spectral efficiency of indoor mmWave networks with ceiling-mounted access points," in *Proc. IEEE GLOBECOM*, Singapore, Dec. 2017.
- [10] J. Chen *et al.*, "Pseudo iteration: Millimeter-wave localization using a single RF chain," in *Proc. IEEE WCNC*, San Francisco, CA, Mar. 2017.
- [11] P. Botsinis *et al.*, "Quantum-assisted indoor localization for uplink mm-wave and downlink visible light communication systems," *IEEE Access*, vol. 5, pp. 23 327–23 351, Jul. 2017.
- [12] J. Palacios *et al.*, "JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems," in *Proc. IEEE INFOCOM*, Atlanta, GA, May 2017.
- [13] F. Lemic *et al.*, "Localization as a feature of mmWave communication," in *Proc. IWCMC*, Paphos, Cyprus, Sep. 2016.
- [14] A. Aguilar-Garcia *et al.*, "Coordinated location-based self-optimization for indoor femtocell networks," *Elsevier Computer Networks*, vol. 106, pp. 1–16, Sep. 2016.
- [15] D. Vasishth *et al.*, "Decimeter-level localization with a single WiFi access point," in *Proc. USENIX NSDI*, Santa Clara, CA, Mar. 2016.
- [16] D. Lymberopoulos *et al.*, "A realistic evaluation and comparison of indoor location technologies: Experiences and lessons learned," in *Proc. IPSN*, Seattle, WA, Apr. 2015.
- [17] 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.
- [18] R. Reimann *et al.*, "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.
- [19] R. Nandakumar *et al.*, "Centaur: Locating devices in an office environment," in *Proc. ACM Mobicom*, Istanbul, Turkey, Aug. 2012.
- [20] J. Xiong and K. Jamieson, "Arraytrack: A fine-grained indoor location system," in *Proc. USENIX*, Lombard, IL, Apr. 2013.
- [21] C. Laoudias *et al.*, "Crowdsourced indoor localization for diverse devices through radiomap fusion," in *Proc. IPIN*, Montbeliard-Belfort, France, Oct. 2013.
- [22] Y. Wang *et al.*, "Adaptive room-level localization system with crowd-sourced WiFi data," in *Proc. SAI IntelliSys*, London, UK, Nov. 2015.
- [23] Y. Shen *et al.*, "Fundamental limits of wideband localization—Part II: Cooperative networks," *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 4981–5000, 2010.
- [24] Y. Guo *et al.*, "Accurate indoor localization with crowd sensing," in *Proc. IEEE ICC*, Kuala Lumpur, Malaysia, May 2016.
- [25] D. Wang *et al.*, "Pursuance of mm-level accuracy: Ranging and positioning in mmwave systems," *IEEE Syst. J.*, 2018, online-first.
- [26] P. Corbalán and G. P. Picco, "Concurrent ranging in ultra-wideband radios: Experimental evidence, challenges, and opportunities," in *Proc. EWSN*, Madrid, Spain, Feb. 2018.
- [27] A. Guerra *et al.*, "On the impact of beamforming strategy on mm-wave localization performance limits," in *Proc. IEEE ICC Workshops*, Paris, France, May 2017.
- [28] 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.
- [29] K. Witrisal *et al.*, "High-accuracy localization for assisted living," *IEEE Signal Process. Mag.*, vol. 33, no. 2, pp. 59–70, Mar. 2016.
- [30] M. Z. Comiter *et al.*, "A structured deep neural network for data-driven localization in high frequency wireless networks," *AIRCC Intl. J. Comp. Networks & Comm.*, vol. 9, no. 3, pp. 21–39, May 2017.
- [31] A. Olivier *et al.*, "Lightweight indoor localization for 60-GHz millimeter wave systems," in *Proc. IEEE SECON*, London, UK, Jun. 2016.
- [32] J. Palacios *et al.*, "Communication-Driven Localization and Mapping for Millimeter Wave networks," in *Proc. IEEE INFOCOM*, Honolulu, HI, May 2018.
- [33] M. E. Rasekh *et al.*, "Noncoherent mmWave path tracking," in *Proc. ACM HotMobile*, Sonoma, CA, Feb. 2017.
- [34] A. Abdelreheem *et al.*, "Location-based millimeter wave multi-level beamforming using compressive sensing," *IEEE Commun. Lett.*, vol. 22, no. 1, pp. 185–188, Jan. 2018.
- [35] "Qualcomm Atheros vendor-specific commands," <http://w1.fi/cgiit/hostap/tree/src/common/qca-vendor.h>, line 767, checked Jan. 17, 2019.
- [36] D. Steinmetzer *et al.* (2017) Talon tools: The framework for practical IEEE 802.11ad research. [Online]. Available: <https://seemoo.de/talon-tools/>