SearchLight: Tracking Device Mobility using Indoor Luminaries to Adapt 60 GHz Beams

Muhammad Kumail Haider, Yasaman Ghasempour and Edward W. Knightly

Rice University, Houston, TX

{kumail.haider,ghasempour,knightly}@rice.edu

ABSTRACT

We present SearchLight, a system that enables adaptive steering of highly directional 60 GHz beams via passive sensing of visible light from existing illumination sources. The key idea is to simultaneously track a mobile device's position and orientation using intensity measurements from lighting infrastructure, and to adapt client and AP beams to maintain beam alignment, without training overhead or outages in the 60 GHz band. Our implementation on custom dual-band hardware with 2 GHz wide channels and 24-element, electronically steerable phased array antennas shows that SearchLight successfully tracks client mobility and achieves up to $3\times$ throughput gains compared to an in-band training strategy, and eliminates millisecond-scale in-band training epochs.

CCS CONCEPTS

• **Networks** → **Network protocol design; Wireless local area networks; Network protocols;**

ACM Reference Format:

Muhammad Kumail Haider, Yasaman Ghasempour and Edward W. Knightly. 2018. SearchLight: Tracking Device Mobility using Indoor Luminaries to Adapt 60 GHz Beams. In *Mobihoc '18: The Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing, June 26–29, 2018, Los Angeles, CA, USA.*, 10 pages. <https://doi.org/10.1145/3209582.3209601>

1 INTRODUCTION

The next 60 GHz Wi-Fi standard promises data rates of 100 Gb/sec [7], more than 10 times faster than today's standards and products [17, 19]. Such data rates are enabled by wide GHz-scale bandwidth coupled with phased array antennas to realize high directionality. Unfortunately, despite such astounding physical-layer bit rates, throughput can be severely degraded by mobility, which can break the highly aligned transmit and receive beams, requiring millisecond-scale delays to re-align. While such a gap might seem quite short, it yields a missed opportunity to transmit 10^6 's of Mb, severely degrading throughput and potentially disrupting high-rate, low-latency applications such as wireless virtual reality.

In this paper, we present the design, implementation and experimental evaluation of SearchLight, a system that re-aligns 60 GHz beams without outages or time-consuming in-band retraining. Our

design is motivated by the observation that despite unconstrained client and environmental mobility, two key indoor WLAN elements do not move: the access point and overhead lighting. Moreover, the ubiquity and dense deployment of indoor luminaries ensures multiple light sources are available in the environment [31], which we repurpose as fixed anchors. Thus, we equip mobile clients with photo diodes that, with near zero energy cost,¹ can view light sources as fixed anchors from which mobility can be inferred. We thereby can track changes in a mobile device's position and orientation by passively sensing light intensity from indoor luminaries, even at *unknown* locations, and continuously infer required changes in 60 GHz beams. In particular, we make the following contributions.

First, we devise a novel method to estimate a mobile device's position and orientation *simultaneously* in 3-D using light intensity measurements from at least three light sources with *unknown* locations. Prior work on visible light positioning uses multi-lateration (e.g., [12]), which requires a *fixed* and *known* orientation of the light sensor, such that light intensity depends solely on position. Since we do not require the location of the luminaries to be known (to simplify deployment) and allow clients to have arbitrary orientation, existing solutions are not directly applicable. In contrast, we employ a light sensor array for mobile devices and fuse measurements from multiple sensors to decouple orientation and position estimation using first order approximations. In particular, we first estimate the Angle of Arrival (AoA) from each light anchor at the client, and then use this knowledge to express light intensities as a function of 3-D coordinates. Although AoA estimation via antenna array phase difference is an effective technique in radio bands, e.g., [27], visible light is incoherent, and light sensors can only measure the magnitude of the incident light. Thus, we exploit the known geometry of the sensor array, such that light intensity at adjacent sensors has a known angle difference and the AoA component in the plane carrying the adjacent sensors can be estimated.

Second, we design an algorithm to steer both client and AP beams by only using light measurements. Our algorithm exploits the dominant Line of Sight (LOS) propagation of the 60 GHz band to track changes in AoA and Angle of Departure (AoD) parameters of the LOS channel component using the aforementioned light-based mobility estimation and steers the 60 GHz beams along the LOS path. In particular, if the SearchLight client determines that a change in client-side steering is required, its beam can be re-steered immediately, without requiring training or feedback to the AP. On the other hand, if the client infers that the AP should re-steer its beam, this inference is fed back to the AP by either piggybacking it to a data transmission, or initiating an on-demand feedback packet. We evaluate this feedback overhead in Sec. 4.

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

Mobihoc'18, June 26–29, 2018, Los Angeles, CA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5770-8/18/06...\$15.00

<https://doi.org/10.1145/3209582.3209601>

¹Indeed, the photo diode's close cousin, the solar cell, provides a net positive energy conversion of light to electricity [1].

Figure 1: SearchLight WLAN scenario.

Finally, we implement the key components of SearchLight on a dual-band hardware testbed comprised of a custom light sensor array integrated with X60 [19], our highly configurable Software Defined Radio (SDR) 60 GHz platform with fully programmable PHY and MAC layers, multi-Gbps rates, 2 GHz wide channels and a user-configurable 24-element phased array antenna which can be electronically steered in real-time. For luminaries, we use bulbs and off-the-shelf LEDs. We conduct extensive over-the-air experiments encompassing a wide range of scenarios with respect to the position and number of light anchors, and mobility patterns. We show that when a LOS path exists, SearchLight correctly steers beams more than 60% of time on average even with unknown light anchor locations, and non-uniform beam patterns and side lobes of the antenna array. We further use traces from extensive channel measurements to drive a custom WLAN simulator implementing both IEEE 802.11ad at 60 GHz and visible light bands with multiple mobile clients and show that SearchLight avoids repeated training overhead under mobility, achieving up to $3\times$ improvement in throughput.

2 SEARCHLIGHT DESIGN

2.1 System Architecture

The SearchLight architecture has three components: (i) three or more luminaries, ceiling mounted and facing downward, which we repurpose as fixed anchors; (ii) a 60 GHz WLAN AP with electronically steerable beams; (iii) a mobile client with 60 GHz radios and off-the-shelf light sensors. An example scenario is depicted in Fig. 1.

In SearchLight, the AP is not required to have any coordination with or control over the luminaries. Moreover, lights are not required to have communication capabilities, i.e., it is not required to support Visible Light Communication (VLC). We only require that at least three lights be turned on and their intensities be distinguishable. That is, the intensity from each light anchor may be measured separately, e.g., via [12] or by exploiting characteristic frequencies of LEDs [31]. Likewise, if the lights transmit signatures, techniques such as [11] can further be employed for distinction.

2.2 Node Architecture

Fig. 2 depicts the SearchLight node architecture. The 60 GHz band, shown at the top of both the AP and the client, uses the “selected beam” from the respective beamforming codebooks for directional transmission and reception. These beams (predefined codewords with phase values for antenna elements) are initially selected via beam training in the 60 GHz band. For subsequent transmissions, these beams are determined by SearchLight (implemented as a software module connected to the 60 GHz MAC) using light measurements at the client.

Figure 2: SearchLight node architecture.

For light measurements, the client equips an array of J light sensors, such that each sensor measures intensity (I) from all available (n) light anchors. The set of intensities $\{I\} = I_j^i$ ($i = 1, \dots, n$, $j = 1, \dots, J$) is input to the client’s SearchLight module which has two main components: (i) *Mobility Estimation Block* which uses light measurements to estimate the position and orientation of the client with respect to the anchors, and (ii) *Beam Steering Block* which translates mobility estimates to changes in AoA and AoD in the 60 GHz band, and infers changes in AP and client beams by computing beams with maximum directivity gain along the estimated angles based on the knowledge of beamforming codebooks at the client and the AP. If a change in client’s beam is inferred, SearchLight triggers the 60 GHz MAC/PHY to switch the codebook entry at the client side, such that the AP is oblivious to any changes in client-side beams. On the other hand, if a change in AP-side beam is inferred, SearchLight transmits feedback to the AP to enable direct selection of the new beam vs. requiring the AP to initiate beam training to search for the new beam.

2.3 Design Overview

We design SearchLight to comprise two phases; an initial *Training Phase* during which the client performs measurements in both 60 GHz and visible light bands and estimates key parameters about the indoor environment, i.e., its initial position, orientation and light anchor positions. After this initial phase, SearchLight enters the *Tracking Phase* which maintains alignment of 60 GHz beams by steering beams solely via passive light sensing. The key steps of both phases are as follows, and are summarized in Algorithm 1.

(i) **Training Phase:** This initial phase is invoked when maximum signal strength beams are not known at the end nodes, e.g., at association or after a link breakage. During the training phase, end nodes perform Beamforming Training (BFT) in the 60 GHz band, e.g., via the 802.11ad standard’s exhaustive search based beam selection procedure [17], during which end nodes discover the pair of beams ($\tilde{S}^{AP}$, $\tilde{S}^{cl}$) that maximize 60 GHz signal strength (R), after steering across all AP and client beams and measuring respective signal strengths $\{R^{AP}\}$ and $\{R^{cl}\}$ (line 1). The client also measures light intensities $\{I\}$ from n light anchors using its sensor array, and uses the method described in Sec. 2.4 to estimate its position $P_o = (x_o, y_o, z_o)$ and orientation angles α_o and β_o in the azimuth and elevation planes, where subscript 0 indicates client’s initial

Training Phase:

- 1 $(\hat{S}^{AP}, \hat{S}^{cl}) = \arg \max_{S^{AP}, S^{cl}} (R) \mid \{R^{AP}\}, \{R^{cl}\}$
- 2 Find $P_o = (x_o, y_o, z_o), \alpha_o, \beta_o \mid \{I\}$
- 3 Find $P^i = (x^i, y^i, z^i) \forall i \mid \{I\}$
- 4 Find $(\theta_o^{AP}, \phi_o^{AP}, \theta_o^{cl}, \phi_o^{cl}) \mid \{I\}, P^{AP}$
- Tracking Phase:**
- 5 Find $P = (x, y, z), \alpha, \beta \mid \{I\}, P^i$
- 6 Find $(\theta^{AP}, \phi^{AP}, \theta^{cl}, \phi^{cl}) \mid P, P_o, \alpha, \alpha_o, \beta, \beta_o, P^i, P^{AP}$
- 7 $\hat{S}^{AP} = \arg \min_{S^{AP}} |\angle(\text{center angles of } S^{AP} - (\theta^{AP}, \phi^{AP})|$
- 8 $\hat{S}^{cl} = \arg \min_{S^{cl}} |\angle(\text{center angles of } S^{cl} - (\theta^{cl}, \phi^{cl})|$
- 9 **if** $(\hat{S}^{cl} = \hat{S}^{cl})$ **then**
 | no change
 else
 | $\hat{S}^{cl} \leftarrow \hat{S}^{cl}$
 end
- 10 **if** $(\hat{S}^{AP} = \hat{S}^{AP})$ **then**
 | no change
 else
 | send feedback to AP ($\hat{S}^{AP} \leftarrow \hat{S}^{AP}$)
 end

Algorithm 1: SearchLight Protocol Overview

parameters during the training phase (line 2). It also estimates the position of light anchors ($P^i, i = 1..n$) during this process (line 3).

Further, to track changes in AoA and AoD with respect to the AP, its position (P^{AP}) is also needed with respect to the light anchors, such that changes in position using light measurements can be translated to changes in position with respect to the AP. There are many recent solutions in 60 GHz band which use Time of Flight (ToF), highly directional beams and reflected paths [2] for localization with a single AP. Conversely, if the AP also houses a light source, the position can be estimated via our method of light anchor positioning (line 3). In any case, this is a one-time computation in SearchLight at the time of association with the AP and any of existing methods can be used to determine (P^{AP}). It is then used in the training phase to estimate AoDs and AoAs $(\theta_o^{AP}, \phi_o^{AP}, \theta_o^{cl}, \phi_o^{cl})^2$ in the 60 GHz band (line 4). Moreover, in our evaluation in Sec.3, we use an existing technique to account for this estimation of (P^{AP}) and the resulting error as well.

(ii) Tracking Phase: After the initial BFT, SearchLight enters the tracking phase which runs as a background process at the client to continuously estimate position (P) and orientation (α, β) using light measurements (line 5). It uses the same estimation method as in line 2, except that anchor positions, which are already estimated in the training phase, can now be used to reduce the number of unknowns in the estimation problem and hence to reduce computation time. Conversely, the entire process can be repeated to obtain further P^i estimates and improve anchor positioning accuracy. Moreover, since light sensors can sample light intensity at multiple kHz frequency consuming very little energy, SearchLight can passively track device mobility at 802.11ad frame transmission times (orders of a few ms).

AoD and AoA in the 60 GHz band are then computed for the LOS path using estimates of the client's current position and orientation, and the AP's position with respect to light anchors (line 6). Using

² θ and ϕ denote the azimuth and elevation components of the respective angles.

knowledge of AP and client codebooks and beam patterns, $\hat{S}^{AP}$ and $\hat{S}^{cl}$ are inferred as highest strength beams, such that they have maximum directivity gain along (θ^{AP}, ϕ^{AP}) and (θ^{cl}, ϕ^{cl}) respectively (lines 7, 8). If a change in the client's beam is required, the local 60 GHz MAC is triggered to steer to $\hat{S}^{cl}$ (line 9). For AP-side beam adaptation, feedback containing $\hat{S}^{AP}$ can be piggybacked to an ongoing transmission, or sent as a separate packet (line 10). We discuss this feedback process in detail in Sec. 4. In any case, there is no 60 GHz in-band training overhead in the tracking phase.

If the link breaks during the tracking phase for any reason (e.g., due to blockage of LOS path or estimation error), SearchLight enters training-phase again. These subsequent training phases only require beam training in the 60 GHz band; the client can learn the indoor topology (anchor and AP positions) within the first training phase at association. However, these processes can be optionally repeated for further improvement of estimates.

2.4 Position and Orientation Estimation

Here we describe the visible light channel model and our method to estimate the position and orientation of a mobile client.

Visible Light Channel Model: The intensity (I) of light received at a sensor is modeled by Lambertian radiation pattern for LOS propagation [1] as follows:

$$I(\rho, \gamma, \psi) = T \cdot A \cdot g(\psi) \cdot \left(\frac{m+1}{2\pi} \right) \cdot \cos^m(\psi) \cdot \frac{\cos(\gamma)}{\rho^2} \quad (1)$$

where T is the transmit power, A is sensor area, γ is the irradiance angle between the vector from light source to sensor and the normal vector to the source, ρ is the source-sensor distance and ψ is the AoA at the sensor. g is optical concentrator, which is a constant if ψ lies within the field-of-view of the sensor. m is the Lambertian order, which is unity for common indoor LEDs. It follows that the light intensity varies inversely to distance, AoA and irradiance angle.

Problem Formulation: Since the mobile device's size is usually much smaller compared to ρ , we consider the client as a cube of edge length $2r$. Then the position of the client can be represented as 3-D coordinates (x, y, z) of the center of this cube with respect to some fixed reference, and its orientation as angles α (angle in the azimuth plane) and β (the elevation angle) about the client's center. Since phased arrays have 3-D beam patterns and clients can have arbitrary orientation, we need to estimate (x, y, z, α, β) to track both the client's position and orientation at any time.

Sensor Array Design: We exploit multiple light sensors with known angular separation to estimate the AoA. When introducing more sensors, the entropy of measurements is maximized by placing sensors at right angles, since it gives maximum separation in AoA. Therefore, in our sensor array design, we use at least six sensors arranged mutually orthogonally on the six facets of a mobile device. For the rest of this section, we discuss this case of six-sensor array, but the formulation can be easily extended to larger array sizes.

Light Measurements: Fig. 3 depicts a SearchLight client and n light anchors. We define the positions of all devices with respect to some light anchor L^0 , the choice of which is arbitrary but should be consistent for a single iteration of the algorithm. Eventually, we localize the AP in this frame (by localizing the AP w.r.t the client using any existing RF techniques), and then redefine all coordinates with respect to the AP using simple translation of coordinates. Nonetheless, the other light anchors are located at

Figure 3: SearchLight client w.r.t. n light anchors and the AP.

$P^i = (x^i, y^i, z^i)$, $i = 1, 2, \dots, (n-1)$. These locations are initially not known at the client.

In this case, the light intensity from the i^{th} anchor received at the j^{th} sensor of the client is given as:

$$I_j^i = C_j^i \cdot \cos(\psi_j^i) \cdot \frac{\cos(\gamma_j^i)}{(\rho_j^i)^2} \quad (2)$$

where C_j^i is a calibration constant, and ρ_j^i is the distance between the i^{th} anchor and the j^{th} sensor. If $\vec{P}^i = [x^i, y^i, z^i]^T$ is the position vector of the i^{th} anchor and $\vec{P}_j$ is that of the j^{th} sensor (with unit normal vector $\vec{u}_j$), then angles γ_j^i and ψ_j^i can be computed as:

$$\cos(\gamma_j^i) = \frac{\vec{z} \odot (\vec{P}_j - \vec{P}^i)}{\rho_j^i}, \quad \cos(\psi_j^i) = \frac{\vec{u}_j \odot (\vec{P}^i - \vec{P}_j)}{\rho_j^i} \quad (3)$$

Since the arrangement of sensors is fixed and known at the client, we use spherical coordinates (r, α, β) to compute sensor positions. In particular, with six mutually orthogonal sensors, position vectors are defined by the following set:

$$\vec{P}_j = \left\{ \begin{pmatrix} x + r \cos \alpha \cos \beta \\ y + r \sin \alpha \cos \beta \\ z - r \sin \beta \end{pmatrix}, \begin{pmatrix} x - r \cos \alpha \cos \beta \\ y - r \sin \alpha \cos \beta \\ z - r \sin \beta \end{pmatrix}, \begin{pmatrix} x - r \sin \alpha \cos \beta \\ y + r \cos \alpha \cos \beta \\ z - r \sin \beta \end{pmatrix}, \right. \\ \left. \begin{pmatrix} x + r \sin \alpha \cos \beta \\ y - r \cos \alpha \cos \beta \\ z - r \sin \beta \end{pmatrix}, \begin{pmatrix} x + r \cos \alpha \sin \beta \\ y + r \sin \alpha \sin \beta \\ z + r \cos \beta \end{pmatrix}, \begin{pmatrix} x - r \cos \alpha \sin \beta \\ y - r \sin \alpha \sin \beta \\ z - r \cos \beta \end{pmatrix} \right\} \quad (4)$$

And the normal vectors are computed as $\vec{u}_j = \frac{\vec{P}_j - [x, y, z]^T}{r}$.

By geometry, at most three sensors on a cube can have LOS path to any single light anchor. Therefore, we consider the three sensors with the highest intensities for each anchor ($I_1^i, I_2^i, I_3^i \in \{I^i\}$). Hence, with a single light anchor, we have at most three measurements and five unknowns (x, y, z, α, β) and as such the system is under determined. We can increase the number of measurements by considering more anchors. With two anchors at known locations, the system becomes solvable. However, we consider the general case where position of light anchors is not known to the client, and in this case adding more anchors to the system also increases the number of unknowns (x^i, y^i, z^i) for each anchor, and the problem remains under-determined. To overcome this under-determinacy, we observe that for typical indoor settings, luminaries are usually installed on ceilings and face downwards, which puts another constraint on anchor positions i.e., the height is the same for all anchors ($z^i = 0 \forall i$). With this constraint, the underdetermination is resolved if we consider at least three anchors, such that we have nine unknowns and nine measurements.

Various numerical methods can be employed to find a solution for this system. However, traditional methods are computationally complex for solving at least nine nonlinear equations simultaneously. Moreover, the solution may be sensitive to initial conditions

due to multiple local extrema of trigonometric functions. Thus, instead of solving the system simultaneously, we present a recursive solution considering a single anchor at a time, which significantly reduces the computational complexity.

Recursive estimation of anchor locations: Our key technique is to neglect the orientation angles α and β , and localize each anchor separately with respect to the client first. That is, we initially assume that the client is aligned with the axes defined in Fig. 3. With this assumption, we can localize each anchor separately since we have three unknowns (x, y, z) and three light measurements. In particular, considering that $r \ll \rho$ in most cases, we can approximate the irradiance angles and distances from the light anchor to be the same at all sensors ($\forall j, \gamma_j = \gamma, \rho_j = \rho$). With this approximation, the ratio of light intensities at any two adjacent sensors is a function of their AoA only, and is independent of γ and ρ , i.e.,

$$\frac{I_{j1}^i}{I_{j2}^i} \approx \frac{\cos(\psi_{j1}^i)}{\cos(\psi_{j2}^i)} \quad (5)$$

Therefore, we consider the ratio of intensities at sensors in three perpendicular planes to estimate the AoA, and substitute it in intensity equations to solve for position coordinates.

We make two key observations: (i) Since the light intensities depend on α and β , the aforementioned localization method projects the anchors to an alternate space where anchor positions are rotated. (ii) Since the light sensors are not collocated, this projection will result in a relative translation between projections as well, which is a function of $\frac{r}{\rho}$. For small device sizes ($r \ll \rho$), this translation is negligible, and the locations of anchors in the alternate space result from a rotation of their true coordinates in the original space (as a function of α and β). Moreover, since the anchors are coplanar in the original space, they will lie in the same plane in this alternate space as well, albeit different from the plane $z=0$ in the original space. Therefore, we can estimate α and β by calculating rotation of this plane in the transformed space around and about the z -axis. Once we have estimates $\hat{\alpha}$ and $\hat{\beta}$, we use intensity measurements from the reference anchor to get position estimates $(\hat{x}, \hat{y}, \hat{z})$ in the fixed reference frame defined in Fig. 3.

3 HARDWARE IMPLEMENTATION

3.1 Dual-band Phased-Array Testbed

We develop a custom hardware testbed which includes all three components of the SearchLight architecture, i.e., luminaries, the AP and the client. For the visible light band, we use multiple off-the-shelf Lumileds LEDs (1200 lm, 33V, 100° viewing angle) for illumination. For light sensing, the client houses a $7 \times 7 \times 3$ cm sensor array (emulating dimensions of a big smartphone or a tablet) with six Lux sensors (Adafruit TSL-2591, 180° field-of-view). The sensors are sampled using an Arduino Mega 2560 board which communicates with the client's SearchLight module, implemented in MATLAB. The AP and client nodes are depicted in Fig. 4a.

For implementing 60 GHz band at the AP and the client, we integrate the world's first highly configurable Software Defined Radio (SDR) based X60 mmWave platform [19] into our testbed. It is based on the National Instruments (NI) mmWave Transceiver System and equips a *user-configurable* 24-element phased array antenna from SiBeam. It enables *fully programmable* PHY and MAC layers while still allowing for ultra-wide channels (2 GHz baseband

(a) SearchLight hardware platform.

(b) 3D beam pattern.

(c) Azimuth beam pattern.

(d) Conference room testbed setup.

Figure 4: Dual-band hardware platform with 24-element phased array antenna and custom light sensor array.

(a) CDF of overall error.

(b) Variation across locations.

Figure 5: Beam steering with perfect LOS path knowledge.

bandwidth) and multi-gigabit data rates (up to 4Gbps). We make further enhancements to achieve beam training and beam steering as per SearchLight design (electronic switching in $< 1\mu s$).

The in-built phased array has 24 elements; 12 each for TX and RX. The module allows four different phase values ($0, \pi/2, \pi, 3\pi/2$) for the antenna elements through the use of codebooks. SiBeam's reference codebook defines 25 beams spaced roughly 5° apart (in their main lobe's direction). The 3 dB beamwidth for the beams ranges from 25° to 35° . As a result, each beam's main lobe overlaps with several neighboring beams. The beam patterns are depicted in Fig.4b and Fig.4c. Thus, the X60 platform allows us to evaluate realistic mmWave phased array based system with imperfect beam patterns and side-lobes, and their impact on beam steering.

3.2 Experimental Setup

Using our dual-band testbed, we conduct extensive over-the-air experiments encompassing multiple indoor environments, light anchor topologies and mobility scenarios including translation and rotation. Due to space constraints, here we present results from one such mobility scenario; client translation, which presents greatest challenge as it may require beam adaptation at both end nodes.

For this, we setup the dual-band testbed in a conference room within $4 \times 3 \times 5m$ space bound by walls on two sides and open space on the other two (Fig.4d), such that the AP is fixed in one corner at 1m height, while we generate multiple trajectories for the client. We also fix client's orientation in these experiments to isolate translational mobility. Moreover, we consider scenarios in which the requirements for SearchLight estimation are satisfied, i.e., the client always has a LOS path to the AP and to three light anchors to evaluate various components of SearchLight design. If these requirements are not satisfied, SearchLight falls back to existing in-band training solutions to recover from link breakages, albeit incurring training overhead, as discussed in Sec.5. However, here we are interested in quantifying possible gains via SearchLight

in presence of LOS paths. Fig. 4d depicts our experimental setup with five light anchor positions at 2.5m height. Although the figure shows five anchors, only three are turned on in a single experiment. Hence we create $\binom{5}{3} - 1 = 9$ distinct topologies by using different combinations of active light anchors in separate experiments. Note that we exclude one topology with collinear light sources, which is under-constrained for 3-D orientation estimation. These topologies incorporate the impact of anchor location, and the results in this section present averages over all topologies.

Here we consider four different client trajectories (A, B, C and D) at radial distances 1.5m, 2m, 2.5m and 3m respectively from the AP and 1m height, as shown in Fig. 4d. For each trajectory, the client is initially placed in front of the AP and covers 2m lateral distance along a straight line. As such, these trajectories create maximum angular separation between the initial and the final client positions for the corresponding radial distances, and require maximum beam adaptations. We take measurements at 25cm distances along each path, resulting in 36 measurement locations in total for each three-anchor topology. In these experiments, any change in AP-side beams inferred by the client is immediately applied at the AP since we consider a single client with no contention. We evaluate the impact of contention and feedback overhead in Sec. 4.

3.3 Validity of LOS Path based Beam Steering

Before evaluating SearchLight's beam steering accuracy, we first investigate whether selecting beams with maximum directivity gain along the LOS path in fact yields the strongest links, given perfect knowledge of the position and orientation. For this, we perform an exhaustive beam search over all 25×25 possible AP-client beam pair combinations at all 36 client locations in the above experimental setup, and find the maximal signal strength beam pair. We then compare it with the "geometric beam pair" with maximum gain along the LOS path for each instance, computed with perfect knowledge of client's position and orientation from our setup. To capture beam separation in codebook space, we define a metric; Beam-Index Distance (BID) as the absolute difference between indices of the geometric beam and the maximal strength beam (from exhaustive search). Hence, $BID=0$ implies that LOS path based beam achieves maximum strength, whereas $BID=1$ indicates that the geometric beam is adjacent to the highest strength beam.

Fig.5a depicts the CDF of BID across all measurements. We find that for more than 60% of instances, the LOS path beams for both AP and client achieve the highest signal strength across all beam

combinations. The slight difference between AP and client beams is due to difference in beam patterns for different arrays and because of unavoidable experimental error in fixing the orientation of the client. Ideally, we expect $BID = 0$ for all instances since the LOS path is available across all client locations, and dominance of LOS path for 60 GHz channels should in theory lead to maximum signal strength for geometrical beam pairs. However, for our practical phased array antenna setup, this deviation results from a combined effect of non-uniform beam patterns with side-lobes and presence of reflected paths in the environment, which make some beams achieve higher signal strength if they include multiple paths, despite having lower gain along the LOS path. To verify this further, we compute the BID of both AP and client side beams across four client locations in the column in front of the AP (0m lateral distance from the AP, as shown in the map in Fig.4d), the four locations in the middle column (1m lateral distance) and those close to the side-wall (2m lateral distance). Fig. 5b illustrates this impact of side-lobes and reflections, where client locations in front of the AP have a dominant LOS path and hence geometric beam pair, selected via perfect knowledge of position and orientation, is almost always the highest strength pair. Whereas for the client locations close to the wall, this ratio is slightly lower due to adjacent beams getting highest strength when they include multiple physical paths. For about 5% instances, the BID is greater than 3 beams, and this happens when a reflected path is much stronger than the LOS path due to imperfect beam patterns.

Finding: Due to wide beams of a practical phased array antenna with imperfect beam patterns and side lobes, and in an environment with availability of reflected paths, the strategy of selecting beams with maximum gain along the LOS path yields the true highest strength beams nearly 60% of instances when the LOS path is available. However, for more than 95% instances, the difference from exhaustive search based highest strength beams is within 2 beams.

3.4 Training Phase Accuracy

Since beam steering in SearchLight during the tracking phase uses anchor position estimates from the training phase, tracking-phase accuracy is coupled with training-phase error. Therefore, we first evaluate the accuracy of anchor position estimation during the training phase using our recursive estimation method. Fig. 6 depicts anchor positioning error averaged over three anchors during the training phase for the four trajectories. Since we consider 9 distinct topologies with different anchor combinations; the error bars capture variance due to different initial positions of the client and anchor-client distances. The graph reveals that the average error increases for trajectories A through C, and then decreases slightly for trajectory D. This results from changes in average anchor-client distances across trajectories, and we find from experiments that positioning error is proportional to the average client-anchor distance. This is because the deviation of measured light intensity from the analytical model increases with distance. Note that for trajectory D, the proximity of two light anchors in the bottom row reduces client-anchor distance for some topologies, resulting in a lower positioning error on average. Overall, we find that SearchLight can localize the light anchors within 40cm of their true position with respect to the client.

In SearchLight, we also require clients to find AP's location at association (i.e., during the first training-phase), and any error in this

Figure 6: Anchor positioning error during training phase.

(a) Beam steering accuracy.

(b) Loss in link SNR.

Figure 7: Tracking-phase estimation accuracy.

estimation will also propagate into tracking-phase estimates. As discussed in Sec.5, decimeter-level localization is achievable using various RF techniques. Although these techniques incur overhead in exchanging several frames to estimate Time of Flight (ToF) or AoA, this represents a one-time overhead in SearchLight and is minimal ($< 10ms$). To capture this error, we simulate an existing ToF based technique [2] to estimate AP position across all client locations in the conference room setup, with an average localization accuracy of 50cm, and inject this error in SearchLight estimation algorithm to account for AP localization error. All the results presented in the rest of this section are based on these estimates of AP and anchor positions to evaluate the scenario when the client has no prior knowledge about the indoor environment. In separate experiments, which we do not discuss here for want of space, we observe that knowledge of these parameters (e.g., via planned deployment) further improves SearchLight beam steering accuracy.

3.5 Tracking Phase Accuracy

Next we evaluate beam steering accuracy during the tracking-phase, when the client only uses light measurements from the three anchors (whose positions are estimated during the training-phase) to predict both AP and client-side beams. Fig. 7a depicts the beam steering accuracy as the CDF of BID of beams predicted by SearchLight along all points on the four trajectories, and across all 9 different light-anchor topologies. Here BID represents absolute difference in indices between SearchLight predicted beams and the true highest strength beams computed via an exhaustive search. The figure shows that SearchLight achieves perfect beam steering accuracy for AP-side beams for more than 50% of prediction instances, and for more than 90% instances the difference is within 2 beam indices. Client-side beam steering accuracy shows a similar trend with a slightly higher error. This increase in error occurs because unlike AP-side beams, client-side beams also depend on orientation estimates. Although the client's orientation is fixed in the above experiments, the client may incorrectly perceive rotation (due to estimation error) and adapt its beams in response. We evaluate this estimation error in detail in later part of this section.

Figure 8: Accuracy of various estimation components during the tracking-phase.

We observe that due to position and orientation estimation error combined with error due to non-uniform beam patterns, side-lobes, and reflections, the predicted beams in SearchLight are not always the highest strength beams. To further investigate how much loss in link's signal strength is incurred when these sub-maximal beams are selected for some instances in SearchLight, Fig.7b depicts the SNR loss in dB for SearchLight's predicted beam pairs along all trajectories and across all anchor topologies, as compared to the maximum possible SNR with beams selected via exhaustive search. We observe that for about 30% instances, SearchLight selects correct beams at both AP and client sides and achieves the same SNR as exhaustive search, without incurring any training overhead in the 60 GHz band. Moreover, for more than 75% instances, the loss in SNR is within 1 dB. This is because for more than 90% prediction instances, SearchLight predicts beams within 2 beam indices of the true highest strength beam, and due to overlap between adjacent beams of the phased array, the loss in SNR is small.

Finding: Even with unknown anchor locations and practical phased array with non-uniform beam patterns, SearchLight can predict beams at both AP and client sides within two indices of the true highest strength beam for more than 90% instances across multiple client trajectories and anchor topologies. Moreover, without performing any further beam training during the tracking-phase, SearchLight achieves link SNR within 1-1.5 dB as compared to exhaustive search based beams at most positions along the trajectories.

3.5.1 Impact of Position/Orientation Estimation Error: Next we do an in-depth analysis of how error in client's position and orientation estimation affects the steering accuracy in the above results. For this, we study the four client trajectories in isolation. Fig. 8a depicts the accuracy of the AP-side beams predicted by the client for the four trajectories (plotted along the x -axis), whereas the corresponding bar graphs show the percentage of the total predicted beams with different BID s over all 9 topologies of different light anchor combinations. The figure reveals that for the closest trajectory A, the client makes correct prediction over 60% of the total instances. Further, prediction accuracy first increases for trajectory B with increase in radial distance, but then decreases for trajectories C and D. This trend results from two opposing factors.

The first factor is the client-position error during the tracking-phase, which is proportional to the distance between the client and the three light anchors and the error in anchor position estimates. As depicted by the CDF of tracking-phase positioning error in Fig. 8b, the trajectory A shows the least error due to its proximity to the light anchors on average, and also because anchor positioning error is minimum along this trajectory. The positioning error worsens across trajectories A through D as the average client-anchor

distance and anchor positioning error increases, which in-turn increases the beam steering error across trajectories.

The second factor is the increase in beam coverage area with distance. As the radial distance increases across trajectories, each AP-side beam covers a larger area along the client's trajectory. Moreover, for the same 2m translation, the change in AP-client angular separation also decreases with radial distance, resulting in a decrease in the number of beams to be adapted. Steering accuracy is also affected by error in the AP's position estimate, and we find that its impact is strongest for closer trajectories. Thus for trajectory B, although positioning error is higher than trajectory A, there is also a greater tolerance to error due to lower adaptation frequency and greater beam-coverage area, and hence we observe an improvement in average steering accuracy. We further observe that beam steering error starts increasing again beyond trajectory B. This is because with higher position-tracking error, the impact of beam coverage area and adaptation frequency becomes dominant.

Client-side beam prediction shows a similar trend, as depicted in Fig. 8c, with average steering error slightly higher compared to AP-side beams. This is because of additional error introduced by client's orientation estimation. As shown by the CDF of client's orientation error in Fig. 8d, there is non-zero error in orientation estimates, which the client perceives as rotation and adapts its beams accordingly. This rotation error is an additional factor affecting client-side beam prediction accuracy. Although for some locations this orientation error may in fact improve beam prediction accuracy, it is slightly worse on average.

Finding: Despite unknown anchor locations, SearchLight simultaneously tracks device position within 40cm and device orientation within 5° for more than 80% instances. Moreover, the beam steering accuracy is impacted by client's position and orientation estimates, beam coverage area and anchor position estimates, with average accuracy decreasing with AP-client distance in our experimental setup.

3.5.2 Beam Adaptation Strategies. Finally we evaluate how various possible adaptation strategies affect link SNR as the client moves along the four trajectories, starting from in-front of the AP. As an example, here we analyze variations in link SNR along the closest trajectory (1.5m radial distance) for one of the anchor topologies from above experiments (with light anchors 1,3 and 4 active). The SearchLight client performs training-phase once in-front of the AP and then moves along the trajectory, without doing any further beam training, and adapting beams solely via light measurements. The AP-client angular separation increases from 0° to 53° along this trajectory, resulting in maximum required beam adaptations to sustain the link. Further, we consider both scenarios for SearchLight; with beam adaptation on both AP and client sides, and client-side adaptation only, which may happen if the AP does not receive

Figure 9: SNR variation for different adaptation strategies.

client’s feedback or if the client decides to adapt its own beams only to avoid any training or feedback overhead. For comparison, we also compute maximum achievable SNR for beams discovered via exhaustive search at all points along the trajectory, and also a no-adaptation strategy where the same beams, discovered during the initial beam training, are used throughout the entire trajectory.

Fig. 9 depicts SNR vs. lateral translation for trajectory A. The curve for *optimal* beam selection via exhaustive search at all points represents maximum achievable SNR, which serves as an upper-bound for all beam adaptation strategies. It mostly stays constant, with slight variations due to imperfect beam patterns. In contrast, *without beam adaptation* at either end, link strength degrades sharply and SNR drops below 10 dB for a mere 0.5m translation. As per 802.11ad PHY sensitivity thresholds and on our platform as well, SNR below 10 dB achieves sub-Gbps rates, severely affecting throughput. After 1m lateral translation, the link is completely broken and cannot support even the base data rate. In comparison, SearchLight client is able to maintain near-maximal link strength for most locations along the trajectory with SNR loss within 1.5 dB. Moreover, with only client-side adaptation, SearchLight still achieves upto 7 dB gain over no-adaptation, and extends the range of beam alignment. This is especially useful in cases when there is a delay in conveying feedback to the AP (e.g., due to contention)

Finding: Without beam adaptation, 60 GHz links can lose multi-Gbps data rates via a mere 0.5m translation, highlighting their susceptibility to client mobility. With SearchLight, the client maintains a highly directional link with SNR within 1.5 dB of the maximum achievable SNR for most locations along the trajectory by adapting beams based on light measurements only. Moreover, if only the client node is adaptive, the AP may incorrectly hold on to an older beam too long without necessarily incurring link breakage.

4 WLAN SIMULATOR IMPLEMENTATION

4.1 Trace and Model-Driven Simulator

To explore a broader set of operational conditions beyond the capabilities of our hardware platform, including multiple clients, different mobility patterns and client speeds, we also develop a custom MATLAB WLAN simulator. To drive the simulator, we use the 802.11ad [16] and visible light channel models to extrapolate 60 GHz signal strength and light intensities from our measurement traces to all possible positions and orientations of the mobile device in the indoor environment. This enables us to study multi-client network performance with mobility models such as random waypoint mobility at different speeds to further evaluate SearchLight performance. Moreover, we use PHY and MAC specifications from 802.11ad, and Table 1 lists important simulation parameters.

Simulation Parameter	Value
Max. transmit slot	2ms
Beacon Interval	100 ms
Preamble Length	1.9 ns
Contention Slot	5 μ s
SIFS	3 μ s
DIFS	10 μ s
Base Rate	27.5 Mbps
Highest Rate	4.62 Gbps

Table 1: List of important simulation parameters.

Figure 10: Protocol overhead comparison.

4.2 Setup

Using our measurement-driven simulator, we study indoor WLAN scenarios with multiple, fully backlogged SearchLight clients. We perform multiple experiments with random waypoint mobility and speeds between 1.5 m/s and 5 m/s, which are average human walking and running speeds respectively. Moreover, we evaluate the maximum impact of feedback overhead by having each AP-side adaptation inference generate a feedback packet at the client which is transmitted via the 60 GHz band.

For performance comparison, we simulate baseline 802.11ad and use the same SNR based rate adaptation for both schemes. Because 802.11ad does not have light assisted beam adaptation, it recovers from link breakages via in-band BFT whenever the data rate drops below MCS 1. Hence it is not excessively incurring repeated BFT overhead, yet maintains data rates above 385 Mbps.

The frequency at which light sensors are sampled and the tracking-phase estimates are computed is an important design factor in SearchLight, as more measurements can improve estimation accuracy, yet require increased power and computational resources. We use 100 Hz estimation rate in the experiments discussed below, which we found is adequate for the indoor mobility scenarios in our analysis, and well within sampling range of light sensors.

4.3 Results

Training Overhead Comparison: First, we compare the training overhead incurred by the two schemes in the aforementioned experiments. Fig.10a depicts overhead vs. the number of clients for 1.5 m/s speed. We calculate overhead as the percentage of total time used to adapt 60 GHz beams (for all clients), which consists of BFT overhead for the baseline scheme, whereas for SearchLight it comprises training-phase overhead and the time spent sending feedback packets. The figure shows that for a single client, overhead is negligible for SearchLight since beam adaptation is achieved either by locally steering the beams at the client without any overhead, or by sending a small ($< 100\mu$ s) feedback packet to change AP-side beams. On the other hand, the baseline scheme performs 802.11ad specified in-band BFT (5 – 10ms) to recover from mobility

Figure 11: Normalized throughput comparison.

induced misalignment which takes almost 5% of the time. Further, while training overhead increases monotonically with the number of clients for both schemes, even with 10 clients SearchLight incurs only 1% overhead, whereas 802.11ad's overhead is almost 25%.

Fig. 10b indicates that training overhead increases for both schemes at 5 m/s speed due to a higher frequency of BFT for the baseline scheme and more feedback packets being generated in SearchLight due to higher mobility. Nonetheless, SearchLight incurs almost negligible in-band training overhead, significantly increasing the time available for data transmissions.

Finding: Even with the maximum feedback strategy, SearchLight achieves beam steering with less than 1% in-band training overhead for most scenarios with moderate to high client mobility, incurring 15× to 30× lower overhead than 802.11ad based in-band re-training. This overhead can be further reduced by piggybacking feedback to ongoing transmissions.

Throughput Performance: Next, we analyze throughput performance and normalize to the throughput of an omniscient scheme that always uses the optimal beams and data rates for each transmission. Hence, to the maximum achievable throughput, while incurring the same channel access and contention overhead.

Fig. 11a depicts normalized throughput per client of the two schemes for 1.5 m/s translational speed vs. the number of clients. With a single client, SearchLight achieves 88% of the maximum throughput on average. The degradation is due to sub-maximal data rates resulting from beam steering inaccuracy and MCS under-selection, or packet losses due to MCS over-selection. The 802.11ad baseline scheme achieves ~55% throughput in comparison, due to rate adaptation losses, beam misalignment and BFT overhead as discussed above.

With increasing client density, 802.11ad throughput degrades further due to increased collisions, BFT overhead and latency since multiple clients are contending to train with a single AP. Note that while the absolute per-client throughput also decreases with an increasing number of clients since they share a single medium, due to normalization, our metric only represents losses due to BFT or feedback overhead, and rate and beam mis-selection, our parameters of interest in this experiment.

In comparison, SearchLight normalized throughput remains above 80% for up to 10 clients due to out-of-band beam adaptation which eliminates BFT overhead in most cases. In fact, SearchLight throughput slightly improves for two clients compared to a single client. This is due to a slight increase in channel access delay, resulting in more light measurements in between two transmissions, which slightly improves beam prediction accuracy. However, as the number of clients increases further, rate estimation from the previous transmission also starts to get stale due to increased inter-packet transmission time. This results in lower rate selection accuracy, which starts dominating the steering accuracy factor and throughput starts degrading beyond two client networks.

Finally, at a higher speed (5 m/s), both schemes suffer throughput degradation as depicted in Fig. 11b, due to increased BFT frequency and overhead for the baseline scheme and reduced beam steering accuracy for SearchLight. This reduction in steering accuracy is due to fewer light measurements, or conversely, the client covers more distance between two tracking-phase estimation cycles, requiring faster adaptation. This is also coupled with increased errors in rate selection. Overall, SearchLight maintains more than 65% throughput for up to 10 clients.

Finding: With 801.11ad beam adaptation, more than 50% of available throughput is lost due to beam misalignment for nodes moving at human walking speeds, with even greater impact for faster speeds. With light assisted beam steering, SearchLight achieves between 2× and 3× improvement in throughput by avoiding BFT overhead in most cases, and its performance scales much better with speed and client density. Further, channel access delay due to increased contention benefits SearchLight beam steering due to an increase in the number of light measurements between transmissions.

5 RELATED WORK

Sensing Multiple Luminaries: Existing approaches to distinctly detect different luminaries include decoding specific signatures modulated by LEDs [11], periodic beacons from VLC sources [13], frequency hopping [12], and exploiting characteristic frequency of fluorescent lights and LEDs [30, 31]. Any of these solutions can be incorporated in SearchLight.

Indoor Localization: RF solutions include fingerprinting [3], RSSI mapping [4], Time of Flight [15, 28] and AoA [24, 27] based techniques, achieving up to decimeter-scale accuracy. Likewise, in the 60 GHz band, device tracking with repeated beam sweeps [25] and localization using distance bounding and reflected paths with a single AP [2] have been proposed. These existing techniques can be used to estimate AP position in SearchLight at association, however, they do not estimate orientation and incur significant overhead. Hence they are not suitable for repeatedly estimating position during tracking-phase. We introduce passive light sensing instead for mobility tracking.

Prior work on *visible light localization* uses multi-lateration to achieve sub-meter accuracy [12, 29], but requires known luminary locations as well as fixed and known *orientation* of light sensors, and hence cannot be used to track position and orientation simultaneously. In [31], non-linear intensity differences between two sensors of different fields of view was employed to estimate AoA and localize using fingerprinting with 3+ light sources of known location, or 4+ sources to estimate orientation as well. However,

orientation is limited to the azimuth plane due to 1-D AoA estimation. Finally, camera based fingerprinting and image processing solutions also achieve sub-meter localization with known anchor locations [6, 30], whereas [11] also estimates orientation with at least four anchors at known locations. However, cameras require much higher power and computational resources than light sensors, with [11] requiring 8-9 seconds for image processing.

60 GHz Link Adaptation: *In-band* solutions to reduce training overhead include model-driven beam steering and channel profiling [26, 32], compressive sensing techniques to exploit channel sparsity [18, 20], correlation between beams [23], efficient beam searching [22], sector switching and backup paths [8], and beamwidth adaptation [9]. These solutions help reduce steering overhead and maintain alignment in certain environments, however, they still incur training overhead when constructing channel profiles, searching for backup paths, or SNR degradation when switching to wider beams. Moreover, in-band beam tracking solutions to address mobility have also been proposed, e.g., 802.11ad's beam tracking [16], exploiting multi-lobe beam patterns [14] and beam sounding [9]. While these solutions help refine beam alignment with small-scale mobility, they also incur in-band overhead and do not work if alignment is lost in-between transmissions. In contrast, we target to eliminate in-band beam re-training while maintaining alignment at the narrowest beamwidth. Nonetheless, when visible light hints are not available, e.g., due to blockage or insufficient anchors, prior solutions can be integrated into SearchLight to further reduce training overhead.

Lastly, prior *out of band* solutions also address mobile 60 GHz clients, e.g., via session transfer to legacy bands [21], AoA estimation in legacy bands to eliminate exhaustive search [17], and using sensors on mobile devices [5]. In contrast, SearchLight uses passive light sensing which has much less power requirements than mechanical sensors, requires no communication in the sensing band, and is more resilient to multipath due to dominant LoS propagation of visible light. iTrack [10] proposes to exploit indicator LEDs on APs to track LOS path's AoA at mobile devices and adapt beams. However, beam adaptation is limited to client-side only. In SearchLight, we address beam steering at both AP and client sides by tracking both position and orientation of mobile devices. AP-side steering is critical since APs usually have larger antenna arrays and hence larger beam-search space.

6 CONCLUSION

We present SearchLight, a system that replaces in-band re-training of 60 GHz links in response to device mobility with beam steering based on passive light sensing, by using indoor luminaries as fixed anchors to track changes in position and orientation. Our implementation on custom dual-band hardware platform and WLAN simulator shows that SearchLight successfully tracks device mobility even with light anchors at unknown locations and achieves up to 3× throughput gains over an in-band training approach.

7 ACKNOWLEDGEMENTS

The authors thank Dimitrios Kousonikolas and Swetank K. Saha from University at Buffalo, SUNY, for help with experiments. This research was supported by Cisco, Intel, the Keck Foundation, and by NSF grants CNS-1642929 and CNS-1514285.

REFERENCES

- [1] BARRY, J. R. *Wireless Infrared Communications*, vol. 280. Springer Science & Business Media, 2012.
- [2] CHEN, J., STEINMETZER, D., CLASSEN, J., KNIGHTLY, E., AND HOLLICK, M. Pseudo Lateralization: Millimeter-Wave Localization Using a Single RF Chain. In *Proc. of IEEE WCNC* (2017).
- [3] CHEN, Y., LYMBERPOULOS, D., LIU, J., AND PRIYANTHA, B. FM-based Indoor Localization. In *Proc. of ACM MobiSys* (2012).
- [4] CHINTALAPUDI, K., PADMANABHA IYER, A., AND PADMANABHAN, V. N. Indoor Localization Without the Pain. In *Proc. of ACM MobiCom* (2010).
- [5] DOFF, A. W., AND ET AL. Sensor Assisted Movement Identification and Prediction for Beamformed 60 GHz Links. In *Proc. of IEEE CCNC* (2015).
- [6] GAO, R., TIAN, Y., YE, F., LUO, G., BIAN, K., WANG, Y., WANG, T., AND LI, X. Sextant: Towards Ubiquitous Indoor Localization Service by Photo-taking of the Environment. *IEEE Transactions on Mobile Computing* 15, 2 (2016), 460–474.
- [7] GHASEMPOUR, Y., DA SILVA, C. R., CORDEIRO, C., AND KNIGHTLY, E. W. IEEE 802.11 ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* 55, 12 (2017), 186–192.
- [8] HAIDER, M. K. Overhead Constrained Joint Adaptation of MCS, Beamwidth and Antenna Sectors for 60 GHz WLANs. Master's thesis, Rice University, 2015.
- [9] HAIDER, M. K., AND KNIGHTLY, E. W. Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation. In *Proc. of ACM MobiHoc* (2016).
- [10] HAIDER, M. K., AND KNIGHTLY, E. W. iTrack: Tracking Indicator LEDs on APs to Bootstrap mmWave Beam Acquisition and Steering. In *Proc. of ACM HotMobile* (2018).
- [11] KUO, Y.-S., PANNUTO, P., HSIAO, K.-J., AND DUTTA, P. Luxapose: Indoor Positioning with Mobile Phones and Visible Light. In *Proc. of ACM MobiCom* (2014).
- [12] LI, L., HU, P., PENG, C., SHEN, G., AND ZHAO, F. Epsilon: A Visible Light Based Positioning System. In *Proc. of USENIX NSDI* (2014).
- [13] LI, T., AN, C., TIAN, Z., CAMPBELL, A. T., AND ZHOU, X. Human Sensing using Visible Light Communication. In *Proc. of ACM MobiCom* (2015).
- [14] LOCH, A., ASSASA, H., PALACIOS, J., WIDMER, J., SUYS, H., AND DEBAILLIE, B. Zero Overhead Device Tracking in 60 GHz Wireless Networks using Multi-Lobe Beam Patterns. In *Proc. of ACM CoNEXT* (2017).
- [15] MARCALETTI, A., REA, M., GIUSTINIANO, D., LENDERS, V., AND FAKHREDDINE, A. Filtering Noisy 802.11 Time-of-Flight Ranging Measurements. In *Proc. of ACM CoNEXT* (2014).
- [16] NITSCHKE, T., CORDEIRO, C., FLORES, A., KNIGHTLY, E., PERAHIA, E., AND WIDMER, J. IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gigabit-per-second Wi-Fi. *IEEE Communications Magazine* 52, 12 (2014), 132–141.
- [17] NITSCHKE, T., FLORES, A. B., KNIGHTLY, E. W., AND WIDMER, J. Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement. In *Proc. of IEEE INFOCOM* (2015).
- [18] RASEKH, M. E., MARZI, Z., ZHU, Y., MADHOW, U., AND ZHENG, H. Noncoherent mmWave Path Tracking. In *Proc. of ACM HotMobile* (2017).
- [19] SAHA, S. K., AND ET AL. X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays. In *Proc. of ACM WinTech* (2017).
- [20] STEINMETZER, D., WEGEMER, D., SCHULZ, M., WIDMER, J., AND HOLLICK, M. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11 ad Devices. In *Proc. of ACM CoNEXT* (2017).
- [21] SUR, S., PEFKIANAKIS, I., ZHANG, X., AND KIM, K.-H. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom* (2017).
- [22] SUR, S., VENKATESWARAN, V., ZHANG, X., AND RAMANATHAN, P. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Proc. of ACM SIGMETRICS* (2015).
- [23] SUR, S., ZHANG, X., RAMANATHAN, P., AND CHANDRA, R. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of USENIX NSDI* (2016).
- [24] VASISHT, D., KUMAR, S., AND KATABI, D. Decimeter-Level Localization with a Single WiFi Access Point. In *Proc. of USENIX NSDI* (2016).
- [25] WEI, T., AND ZHANG, X. mTrack: High-Precision Passive Tracking using Millimeter Wave Radios. In *Proc. of ACM MobiCom* (2015).
- [26] WEI, T., ZHOU, A., AND ZHANG, X. Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors. In *Proc. of USENIX NSDI* (2017).
- [27] XIONG, J., AND JAMIESON, K. ArrayTrack: A Fine-Grained Indoor Location System. In *Proc. of USENIX NSDI* (2013).
- [28] XIONG, J., SUNDARESAN, K., AND JAMIESON, K. ToneTrack: Leveraging Frequency-Agile Radios for Time-Based Indoor Wireless Localization. In *Proc. of ACM MobiCom* (2015).
- [29] YANG, Z., AND ET AL. Wearables Can Afford: Light-weight Indoor Positioning with Visible Light. In *Proc. of ACM MobiSys* (2015).
- [30] ZHANG, C., AND ZHANG, X. LiTell: Robust Indoor Localization using Unmodified Light Fixtures. In *Proc. of ACM MobiCom* (2016).
- [31] ZHANG, C., AND ZHANG, X. Pulsar: Towards Ubiquitous Visible Light Localization. In *Proc. of ACM MobiCom* (2017).
- [32] ZHOU, A., ZHANG, X., AND MA, H. Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-Driven Beam Steering. In *Proc. of IEEE INFOCOM* (2017).