

LiSteer: mmWave Beam Acquisition and Steering by Tracking Indicator LEDs on Wireless APs

Muhammad Kumail Haider
Rice University, Houston, TX
kumail.haider@rice.edu

Dimitrios Koutsonikolas
University at Buffalo, SUNY
dimitrio@buffalo.edu

Yasaman Ghasempour
Rice University, Houston, TX
ghasempour@rice.edu

Edward W. Knightly
Rice University, Houston, TX
knightly@rice.edu

ABSTRACT

We present LiSteer, a novel system that steers mmWave beams at mobile devices by repurposing indicator LEDs on wireless Access Points (APs) to passively acquire direction estimates using off-the-shelf light sensors. We demonstrate that LiSteer maintains beam alignment at the narrowest beamwidth level even in case of device mobility, without incurring any training overhead at mobile devices. Our extensive evaluation on a custom dual-band hardware platform comprising highly directional horn antennas as well as practical phased antenna arrays with electronic beam steering shows that LiSteer achieves direction estimates within 2.5 degrees of ground truth on average. Moreover, it achieves beam steering accuracy of more than 97% while in tracking mode, without incurring any client beam training or feedback overhead.

CCS CONCEPTS

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

ACM Reference Format:

Muhammad Kumail Haider, Yasaman Ghasempour, Dimitrios Koutsonikolas, and Edward W. Knightly. 2018. LiSteer: mmWave Beam Acquisition and Steering by Tracking Indicator LEDs on Wireless APs. In *The 24th Annual International Conference on Mobile Computing and Networking (MobiCom '18)*, October 29–November 2, 2018, New Delhi, India. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3241539.3241542>

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.

MobiCom '18, October 29–November 2, 2018, New Delhi, India

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5903-0/18/10...\$15.00

<https://doi.org/10.1145/3241539.3241542>

1 INTRODUCTION

The ever increasing demand for high speed wireless connectivity to support applications such as virtual and augmented reality and uncompressed video streaming is straining the capacity of current WiFi and cellular networks [4]. GHz-scale bandwidth coupled with phased antenna arrays to realize high directionality in the mmWave spectrum (30 GHz to 300 GHz and beyond), can solve this problem by realizing data rates of up to 100 Gb/sec [6]. However, a key challenge in exploiting this expansive bandwidth to realize high data rates is that end nodes need to continually align their beams to establish and maintain directional links.

To this end, commercial products [28, 35] and WLAN standards such as 802.11ad [21] and 802.11ay [6] employ beam-search based training mechanisms, in which one node sends training frames sequentially across all its beams while the other node uses pseudo-omni beams to identify the highest signal strength beam. Although this training, when repeated at both ends, discovers the strongest pair of beams with maximum data rates, the process also requires coordination and setup between nodes and may take 10's of milliseconds. This overhead represents a missed opportunity to transmit 100's of Megabits, severely degrading throughput and disrupting high-rate, low-latency applications. Moreover, the overhead worsens for systems with no pseudo-omni reception, increasing the order of the beam-search space from $2N$ to N^2 for N beams at each end. Mobile devices present an even greater challenge, as beam alignment may be repeatedly lost due to mobility, requiring repeated training to maintain beam alignment and incurring overhead each time [10].

In this paper, we present LiSteer, a system that steers mmWave beams at mobile devices by repurposing indicator LEDs on wireless Access Points (APs) as fixed light anchors and continuously tracking the direction to the AP using light intensity measurements with off-the-shelf light sensors. We demonstrate that LiSteer acquires and maintains beam alignment at the narrowest beamwidth level, even with device mobility, without requiring any beam training at the client

devices. Moreover, our design is scalable, such that the AP can simultaneously align beams with multiple clients by performing a beam sweep only once at its end, with client beams selected via out-of-band light sensing.

Our design is motivated by two key observations. First, most off-the-shelf wireless APs are equipped with light sources such as notification LEDs, which are in close proximity to their mmWave antennas. Second, mmWave channels exhibit pseudo-optical properties due to very short wavelength, i.e., dominant Line of Sight (LOS) propagation, limited scattering and reduced multipath [2]. Therefore, our key idea is that by estimating the Angle of Arrival (AoA) corresponding to the LOS path from the AP’s indicator LED using light intensity measurements, we can approximate it as the AoA in the mmWave band due to the close proximity of the AP’s LED and mmWave band antenna. Hence, we select the client-side beam as the one with the highest directivity gain along the AoA for the LOS path. We show that when the LOS path is available, a LiSteer client can continuously adapt its beams without requiring any beam training by passively tracking the AP’s indicator LED. Moreover, client beam adaptation in LiSteer is completely distributed, such that the AP is oblivious to any beam changes at the client end and hence no feedback is required.

The key challenge in exploiting the AP’s light source for this direction tracking is that, unlike lasers, light intensity from LEDs (or common light bulbs) is *incoherent*, and off-the-shelf light sensors can only measure the *intensity* of the incident light. Therefore, AoA estimation techniques in radio bands via antenna array phase difference (e.g., [38]) cannot be used. Consequently, we devise a novel method for incoherent-light Angle of Arrival (il-AoA) estimation by using an array of light sensors. Our key technique is to approximate the ratio of light intensities at adjacent sensors as a function of their AoA only by exploiting their angular separation on the array. We then estimate the AoA of the LOS path without requiring any calibration or knowledge of the AP’s position or client’s orientation. Moreover, our method estimates il-AoA in both the azimuth and elevation planes, allowing us to steer beams for both 2-D and 3-D beamforming codebooks.

We implement LiSteer on a custom dual-band hardware testbed and perform extensive over-the-air experiments in various environments and under different mobility scenarios to evaluate key components of LiSteer design. Our hardware platforms encompass off-the-shelf light sensors for light sensing, horn antenna based transceivers with 7° beamwidth to achieve extremely directional links, and X60 [27], a 60 GHz wideband platform with an electronically steerable phased array for evaluation with practical antenna arrays which exhibit non-uniform beam patterns and side-lobes. Our key findings are summarized below.

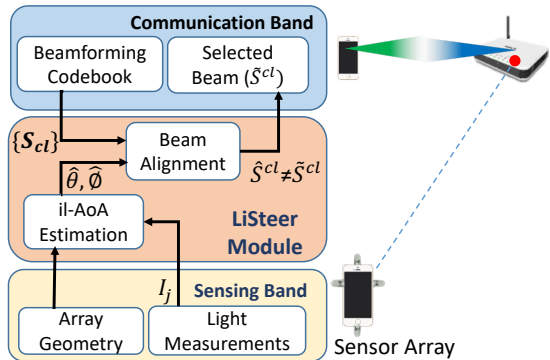


Figure 1: Node architecture for LiSteer.

(i) LiSteer achieves a mean il-AoA estimation accuracy of 2.5° in our experiments encompassing a wide range of AP-client distances and device orientations. (ii) Even with highly directional (7°) beams, LiSteer acquires client beams to within 1 sector of the true highest strength sector in all cases, without any in-band training. (iii) Irregular beam patterns of the phased array can lead to sub-optimal beams (compared to exhaustive beam search) due to side-lobes and reflected paths. However, the difference from exhaustive search is within two beam indexes for more than 80% of instances with an SNR loss of < 1.5 dB. (iv) After initial beam acquisition, LiSteer tracks il-AoA changes due to rotation and translation with even higher accuracy than it does absolute il-AoA due to cancellation of position bias in the error, achieving up to 97% beam steering accuracy. (v) Compared to a baseline 802.11ad scheme, LiSteer achieves up to $2\times$ gain in throughput in dense networks with mobile devices and LOS environments, and its overhead remains almost constant despite an increase in the number of clients due to a single AP-side sweep required for beam training, with client-side beams steered via il-AoA estimates only.

2 LISTEER DESIGN

In this section, we first describe LiSteer’s system architecture and our novel method to estimate the il-AoA using intensity measurements from the AP’s indicator LED. We then describe our beam alignment protocol to steer mmWave beams at the client-side using the il-AoA estimates.

2.1 System Architecture

The LiSteer architecture is divided into two distinct bands; a *Communication Band* and a *Sensing Band*. The former comprises mmWave band radios and phased antenna arrays at the AP and clients, while the Sensing Band includes a light source at the AP and multiple light sensors at the client. In most modern systems, the phases of antenna elements are defined via a 3-D beamforming codebook $\{S\}$, such that by switching between codebook entries, beams can be electronically steered, discretizing the space around the array into

virtual “sectors” [2]. While our design is compatible with any directional antenna design, for the rest of this section we assume a phased array system for both the AP and clients.

Node Architecture: Fig. 1 depicts the LiSteer client node architecture. The client equips an array of J light sensors to measure light intensity (I) from the AP’s LED. The set of intensities $\{I\} = I_j, j = 1, \dots, J$ is input to the LiSteer software module (shown by the middle block), which has two main components: (i) *il-AoA Estimation Block* which uses light measurements to estimate the azimuth and elevation components ($\hat{\theta}_{cl}, \hat{\phi}_{cl}$) of the il-AoA (Sec. 2.3); and (ii) *Beam Alignment Block* which estimates client’s highest strength beam \hat{S}_{cl} using the il-AoA estimates (Sec. 2.4). This estimated beam is then passed on to the Communication Band, and is used as the “selected beam” (\tilde{S}_{cl}) for directional transmission and reception.

2.2 Design Principle

In LiSteer, we exploit the sparsity and dominant LOS propagation of mmWave channels resulting from limited scattering and diffraction due to extremely small wavelength. This is usually characterized using geometric channel models as follows [2]:

$$\mathbf{H} = C \sum_{l=1}^L \alpha_l \mathbf{a}_T(\theta_{T,l}, \phi_{T,l}) \mathbf{a}_R(\theta_{R,l}, \phi_{R,l}), \quad (1)$$

where C is a normalization constant, L is the number of physical paths, α_l is the path gain, \mathbf{a}_T and \mathbf{a}_R are the array response vectors at the transmitter and the receiver, and θ and ϕ denote the azimuth and elevation components of the corresponding Angle of Departure (AoD)/AoA respectively.¹ Due to dominant LOS propagation of mmWave channels shown in prior measurement studies and channel models [1, 39], we expect the LOS channel component to have the maximum gain. Therefore, our key idea is to exploit the AP’s LED to estimate the il-AoA ($\hat{\theta}_{cl}, \hat{\phi}_{cl}$) of the LOS path at the client using light measurements only, and then select the client-side beam with maximum directivity gain (using known beam patterns) along the estimated AoA. As such, we avoid any mmWave in-band training or beam-search at the client in the presence of LOS path. We use the term il-AoA to specify the AoA of the physical LOS path between the AP and the client measured using light intensities. In particular, we use the client’s codebook $\{S_{cl}\}$ to find the beam (or sector for discretized codebooks) \hat{S}_{cl} that has maximum gain along ($\hat{\theta}_{cl}, \hat{\phi}_{cl}$). For codebooks with uniform beam patterns, \hat{S}_{cl} can be computed as follows:

$$\hat{S}_{cl} = \arg \min_{S_{cl,m}; m=1 \dots N} \left| \angle (\Theta_{S_{cl,m}}, \Phi_{S_{cl,m}}) - (\hat{\theta}_{cl}, \hat{\phi}_{cl}) \right| \quad (2)$$

¹Due to channel reciprocity, only the AoD or the AOA needs to be estimated.




Figure 2: Indicator-LED sensing via client’s sensor array.

where $(\Theta_{S_{cl,m}}, \Phi_{S_{cl,m}})$ are the central azimuth and elevation angles of any client beam $S_{cl,m}$. For irregular beam patterns, a simple lookup table can be devised to find \hat{S}_{cl} .

2.3 il-AoA Estimation

Here we describe the light channel model and our method to estimate both azimuth and elevation components of the il-AoA for the LOS path using light measurements only.

Visible Light Channel Model: The intensity (I) of light received at a sensor can be described as:

$$I = T \cdot H \cdot G_r \quad (3)$$

where T is the transmit power of a light source and G_r is the receiver gain, which needs to be calibrated once. H is the channel gain, which is shown extensively in the literature to closely follow the Lambertian radiation pattern for LOS propagation [3] as follows:

$$H(\rho, \gamma, \psi) = A \cdot g(\psi) \cdot \left(\frac{m+1}{2\pi} \right) \cdot \cos(\psi) \cdot \frac{\cos^m(\gamma)}{\rho^2} \quad (4)$$

where A is the sensor area, γ is the irradiance angle defined as the angle between the vector from light source to sensor and the normal vector to the light surface, and ψ is the Angle of Arrival (AoA) at the sensor. Further, ρ is the distance between light source and sensor and g is the optical concentrator, which is a constant factor if AoA lies within the field-of-view of the sensor. m is the Lambertian order, which is equal to one for common LEDs.

Problem Formulation: Fig. 2 depicts an AP at position (x, y, z) with respect to a LiSteer client, where the reference frame is centered at the client’s planar phased array, with the z -axis orthogonal to the array. By geometry, angles θ_{cl} and ϕ_{cl} shown in the figure correspond to the azimuth and elevation components of the AoA from the AP to the client’s array for the LOS path. Our objective is to estimate $(\hat{\theta}_{cl}, \hat{\phi}_{cl})$ as the il-AoA using the incoherent light from the AP’s LED and off-the-shelf light sensors.

Sensor Array Design: The two components of the il-AoA cannot be estimated using a single sensor since the light intensity depends on both the position of and the AoA at

the sensor. Moreover, since the sensor may have an arbitrary orientation, the AoA (ψ) at the sensor may not be the same as the il-AoA, but a projection of it along the sensor's axis. Our key technique is to exploit an array of multiple sensors with known angular separation to estimate θ_{cl} and ϕ_{cl} . When introducing more sensors, the entropy of measurements is maximized by placing sensors at right angles, since this gives maximum angular separation. Therefore, in our sensor-array design, we use at least six sensors arranged mutually orthogonally on the six facets of a mobile device. We discuss the case of a six-sensor array in the rest of this section, but the formulation can easily be extended to larger array sizes.

Estimation Method: A LiSteer client with $J = 6$ light sensors arranged mutually orthogonally is depicted in Fig. 2. In this case, the light intensity from the AP's LED received at the j^{th} sensor of the client is given as:

$$I_j = C \cdot \cos(\psi_j) \cdot \frac{\cos(\gamma_j)}{(\rho_j)^2} \quad (5)$$

where C is a constant parameter for sensors of same type, and ρ_j is the distance between the LED and the j^{th} sensor. If $\vec{P} = [x, y, z]^T$ is the position vector to the AP's LED and \vec{P}_j is that of the j^{th} sensor (with unit normal vector \vec{u}_j), then angles γ_j and ψ_j can be computed as:

$$\cos(\gamma_j) = \frac{\vec{z} \odot (\vec{P}_j - \vec{P})}{\rho_j} \quad (6)$$

$$\cos(\psi_j) = \frac{\vec{u}_j \odot (\vec{P} - \vec{P}_j)}{\rho_j} \quad (7)$$

Since the size of mobile devices is usually much smaller than the AP-client distance, we approximate the irradiance angle and distance from the AP to be the same at all sensors ($\forall j, \gamma_j = \gamma, \rho_j = \rho$). With this, the ratio of intensities at any two adjacent sensors is a function of their AoA only, independent of ρ and γ :

$$\frac{I_1}{I_2} \approx \frac{\cos(\psi_1)}{\cos(\psi_2)} \quad (8)$$

Since the arrangement of sensors is fixed and known at the client, we consider the ratio of intensities at adjacent sensors in three perpendicular planes to estimate the il-AoA component in that plane, without requiring the client's position or orientation. For example, when sensors are arranged mutually orthogonally, this difference in AoA is in fact 90° , such that we can make the substitution $\cos(\psi_2) = \sin(90 - \psi_1)$ in Eq. (8) to estimate ψ_1 as:

$$\hat{\psi}_1 = \tan^{-1} \left(\frac{I_2}{I_1} \right) \quad (9)$$

Note that it is not necessary that the light sensor array and the client's phased array are coplanar and aligned; only

the mapping is required such that angles estimated in the light band can be rotated to find angles with respect to the phased array. However, for simplicity and without loss of generality, here we assume that the two arrays are aligned so that the same reference frame defined in Fig. 2 can be used for the light sensor array as well. With this simplification, we can define \vec{u}_j as unit vectors along $+x, -x, +y, -y, +z, -z$ axes for the six mutually orthogonal sensors.

Moreover, by array geometry, at most three sensors on the array can have a LOS path to the AP, one along each axis (I_x, I_y, I_z). Using the negligible array dimension approximation and solving for $\cos(\psi_j)$ at adjacent sensors in the three perpendicular planes, we estimate θ_{cl} and ϕ_{cl} as follows:

$$\hat{\theta}_{cl} = \tan^{-1} \left(\frac{I_y}{I_x} \right), \quad \hat{\phi}_{cl} = \tan^{-1} \left(\frac{\sqrt{I_x^2 + I_y^2}}{I_z} \right) \quad (10)$$

2.4 Beam Alignment Protocol

Using the aforementioned estimation framework, we design LiSteer to comprise the following two phases.

(i) Beam Acquisition: This is the initial phase where maximal strength beams are not known at the AP or the clients e.g., at association or after link breakage. During this phase, a LiSteer client estimates its maximal strength beam using light measurements as described above, and uses this beam to receive in mmWave band while the AP does a beam sweep at its end. The client then sends feedback about the AP's maximal strength beam to the AP. This may be followed by an optional beam refinement phase, e.g., following the procedure defined in 802.11ad [13], where the client can use the il-AoA estimate to do a local search among the neighboring beams to further improve link strength albeit incurring a small overhead. In any case, exhaustive search is not required at the client if il-AoA estimates are available. In the special case that il-AoA estimates are not available due to blockage of the LOS path, LiSteer falls back to the underlying beam training protocol and performs beam training at the client end as well.

(ii) Beam Steering: After a directional link is established via Beam Acquisition, LiSteer enters the Beam Steering phase, where it passively tracks the il-AoA from the AP's LED and continuously estimates the best client-side beam \hat{S}_{cl} using the il-AoA estimates. Due to client mobility, if this best beam estimate becomes different from the selected beam \tilde{S}_{cl} being used for communication, an interrupt ($\tilde{S}_{cl} \leftarrow \hat{S}_{cl}$) is passed to the MAC layer to adapt the current beam \tilde{S}_{cl} . As such, LiSteer steers client-side beams without incurring any training or feedback overhead. Moreover, the AP is oblivious to any changes in client beams, making client-side beam steering completely distributed.

Multi-Client Beam Training: For AP-side adaptation, we exploit the periodic beacon sweeps which are required by most WLAN standards, including 802.11ad and 802.11ay for 60 GHz WLANs. During a beacon sweep, the AP sends beacon frames across all its beams such that the frames also include training sequences [13]. This process is repeated at regular intervals, usually every 100 ms. Since client-side beams are continuously adapted via il-AoA estimation in LiSteer, we introduce multi-client beam training where the AP can simultaneously train with any number of clients by doing a single beam sweep and getting feedback from the clients following a simple polling procedure (as depicted in Fig. 3), with polling sequence assigned during client association. As such, if the AP has N_{AP} beams and trains with M clients, each with $N_{cl,m}$ beams, then LiSteer requires beam search only over N_{AP} beams during this training phase, compared to $N_{AP} + \sum_{m=1}^M N_{cl,m}$ beam combinations in case of 802.11ad based pseudo-omni training.

If beam alignment is lost due to tracking error or il-AoA estimation error, LiSteer enters Beam Acquisition again.

2.5 Design Considerations

While LiSteer tracks notification LEDs on off-the-shelf APs using an array of off-the-shelf light sensors, there are several challenges in practical system deployment.

Tracking Range: The coverage range of the indicator LED depends on its maximum intensity. Ideally, the light intensity range would be the same as radio coverage range. To realize this, existing techniques such as Pulse Width Modulation (PWM) can be leveraged to increase the coverage area [23, 24]. In particular, the LED driving circuit responsible for the “blinking” of indicator LEDs can be tuned to adjust the duty cycle of the on-period of the LED, such that the LED emits a higher instantaneous light intensity, thus increasing the range, while maintaining the same average intensity. As such, it can still serve the purpose of indicating various system notifications while radiating light to a much longer distance. In fact, recent works have demonstrated designs where the LED even appears switched-off to the human eye for communication in the dark [34]. Thus we can incorporate these techniques in LiSteer if required.

Interference from Luminaries: Another practical challenge is the presence of ambient light and other indoor light sources, which can interfere with the light from the AP’s LED. This problem is also a challenge in Visible Light Communication (VLC); it has been extensively studied in prior literature in this context, and many solutions exist to address this challenge as well. These solutions include using RGB photodiodes for colored LEDs on the AP, encoding a distinct pattern onto the PWM sequence which can identify an AP (similar to BSSID in radio beacons from APs), and detecting




Figure 3: Multi-Client Beam Training for AP-side steering.

the characteristic frequency of the indicator LED [43]. Any of these solutions can be integrated with LiSteer by making simple modifications to the AP’s LED. Nonetheless, we do not require any further knowledge about the LED (e.g., its intensity or position) and there is no data communication required in the light band.

Sensor Array Requirements: As described in Sec. 2.3, we propose an array with at least six light sensors arranged mutually orthogonally on the facets of mobile devices. Since we require off-the-shelf light sensors with sampling rate on the order of 100Hz to 1kHz for tracking at 802.11ad frame transmission times, existing ambient light sensors on mobile devices can be integrated into the array.

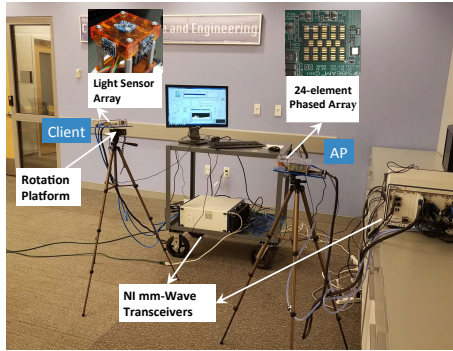
3 IMPLEMENTATION

In this section, we present our implementation of LiSteer first on a dual-band hardware platform, and then on a trace-and model-driven indoor-WLAN simulator. Our evaluation encompasses scenarios where there is always a LOS path between the AP and the client to study the potential gains of LiSteer. When the LOS path is blocked or intensity measurements are not available, LiSteer falls back to the underlying training protocol (e.g., 802.11ad or any other existing in-band training schemes, as discussed in Sec. 6) and shows similar performance as the baseline scheme.


3.1 Dual-band Hardware Testbed

We developed a custom hardware testbed which implements all key components of the LiSteer system design. For the *Sensing Band*, we use off-the-shelf Lumileds LED (900 lm, 100° viewing angle) for the AP, while the client houses a custom $7 \times 7 \times 3$ 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 LiSteer module, implemented in MATLAB.


For the mmWave band, we use two distinct systems to capture different aspects of mmWave channels and system design, encompassing a wide range of radio and antenna technologies possible for current and future mmWave systems.



(a) Wide-band phased array platform setup.



(b) 3D beam pattern.



(c) Azimuth beam pattern.

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

3.1.1 X60 Platform with Phased Antenna Array. First we integrate a configurable Software Defined Radio (SDR) based X60 mmWave platform [27] into our tested, as depicted in Fig. 4a. It is based on National Instruments' mmWave Transceiver System [14] and equipped with a *user-configurable* 24-element phased antenna array from SiBeam. It enables *fully programmable* PHY and MAC layers while still allowing for ultra-wide channels (2 GHz baseband bandwidth) and multi-Gbps data rates (up to 4 Gbps). We made enhancements to the base design to achieve rate selection and adaptation based on il-AoA inferences in LiSteer design, and beam steering in over-the-air experiments.

The in-built phased array has 24 elements; 12 each for TX and RX. The phase of each antenna element can be set to one of four values: $0, \pi/2, \pi, 3\pi/2$. SiBeam's reference codebook defines 25 beams spaced roughly 5° apart (in their main lobe's direction). The beams cover a sector of 120° (in the azimuthal plane) centered around the antenna's broadside direction, with 3 dB beamwidth ranging from 25° to 35° . As such, each beam's main lobe overlaps with several neighboring beams. The beam patterns for the central TX beam are depicted in Fig. 4 as an example (taken from [27]). Thus, the X60 platform allows us to evaluate a realistic mmWave phased array based system with imperfect beam patterns and side-lobes, and their impact on beam steering.

3.1.2 VubiQ 60 GHz Platform with Horn Antennas. While X60 enables experiments over wideband channels and a practical phased-array, the beamwidth is limited to 25° and the array yields irregular beam patterns with significant side-lobes. The phased arrays of other currently available commercial and testbed platforms in the 60 GHz band also exhibit similar beam patterns [19, 30]. However, with more antenna elements beams will have higher directionality. Here we also aim to evaluate whether il-AoA estimates are accurate enough to steer such highly directional beams as well. For




Figure 5: VubiQ Platform with 7° horn antennas.

achieving highly directional beams with minimum side-lobes, we integrate a horn antenna based platform into our testbed.

For this, we develop programmable nodes using the VubiQ [7, 36] transceiver system operating in 57-64 GHz unlicensed frequency band with 1.8 GHz bandwidth (compliant with 802.11ad) and WARP baseband (a software-defined radio platform). A differential I/Q input to the VubiQ transmitter is achieved by feeding WARP's I/Q baseband signal to an evaluation board using a 6 GHz Ultra Dynamic Range Differential Amplifier (ADL5565). To achieve narrow beams, we use horn antennas with 7° beamwidth at both AP and client sides. Moreover, to implement beam-steering, both nodes (depicted in Fig. 5) are mounted on Cine-Moco motion platform which can be rotated to within $1/100^{th}$ of a degree. Using this platform, antennas are rotated in discrete steps to emulate discretized sectors (predefined by a codebook) to achieve sector sweeps during the training phase or to adapt sectors during the tracking phase of LiSteer.

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.

3.2 Trace and Model-Driven Simulator

To explore a broader set of operational conditions beyond the capabilities of the hardware platform including WLAN performance with multiple clients, multi-client beam training, and different mobility patterns and speeds, we also develop a custom MATLAB WLAN simulator. We use channel traces from our hardware experiments to drive the simulator. Since the measurement data points are relatively coarse-grained, we use the 802.11ad [21] and visible light [3] channel models to extrapolate mmWave signal strength and light intensities 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 LiSteer performance.

Furthermore, we use PHY and MAC specifications from 802.11ad in our simulator. Control packets are sent at the lowest bit rate (27.5 Mbps, corresponding to MCS-0), and require minimum receive sensitivity (-78 dBm). Packet headers are considered to be received correctly if signal strength is greater than -78 dBm (threshold for MCS-0, used to encode headers) whereas payloads are received correctly if signal strength is greater than the threshold for the MCS used to encode data. Table 1 lists important simulation parameters.

4 TESTBED EVALUATION

We first describe our experiments with horn antennas and then discuss a subset of results from the X60 platform.

4.1 Experimental Setup

We setup the dual-band testbed within a $4 \times 3 \times 5$ m space bound by walls on two sides and open space on the other two (Fig. 6), such that the AP is fixed in one corner at 1m height, while we consider multiple positions and orientations for the client to study both translational and rotational mobility. Furthermore, we repeat the same experiment with both horn antenna and phased array platforms to compare the performance of different mmWave band systems.

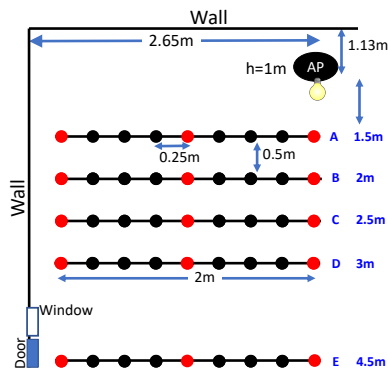


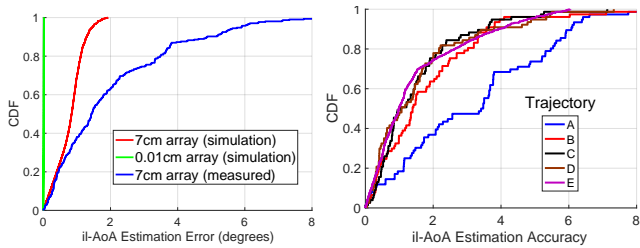
Figure 6: Testbed setup for dual-band experiments.

In particular, we consider five client trajectories (A,B,C,D and E) at longitudinal distances 1.5m, 2m, 2.5m, 3m and 4.5m from the AP for the initial position of the client. The trajectories follow a straight line from respective initial positions and cover 2m lateral distance, while we take measurements every 0.25m. Furthermore, for 15 points (shown as black circles in Fig. 6), we rotate the client by 60° in both clockwise and counter-clockwise directions (in steps of 1° for the Vu-bIQ platform and steps of 5° for the X60 platform) to study rotation. Note that there is always a LOS path in the visible light and 60 GHz bands in these experiments.

4.2 il-AoA Estimation Accuracy

We first evaluate the accuracy of il-AoA estimation using light intensities. To identify various sources of estimation error, we also perform light channel model based simulations to analyze the performance of sensor arrays with different dimensions, and for comparison with over-the-air experiments. Fig. 7a depicts the CDF of il-AoA estimation error. First, for simulation of an array of negligible dimensions (0.01cm), we observe an almost perfect il-AoA estimation accuracy, which validates our key approximation that device size is negligible compared to AP-client distance. Second, simulations for a 7cm array (same dimensions as our testbed array) show up to 2° estimation error, highlighting the impact of collocated sensors approximation on the accuracy of our method. Sensors can be placed closer to device edges to further minimize this error and achieve better accuracy.

Next we analyze the performance in over-the-air (OTA) experiments. The blue curve for measurement results indicates that the estimation error is within 3.5° of the true AoA for more than 90% of measurement instances. The error is also higher than the simulations due to the deviation of intensity measurements from the channel models, which we discover is dependent on the position and the angular separation (irradiance angle) between light source and sensors. To investigate this further, we depict the CDF of il-AoA estimation error across all locations for various trajectories of our



(a) CDF of il-AoA estimation error. (b) Error across various trajectories.

Figure 7: il-AoA estimation accuracy.

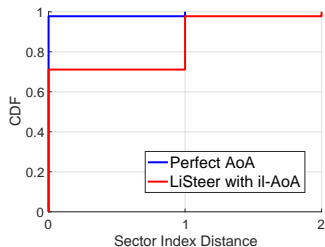


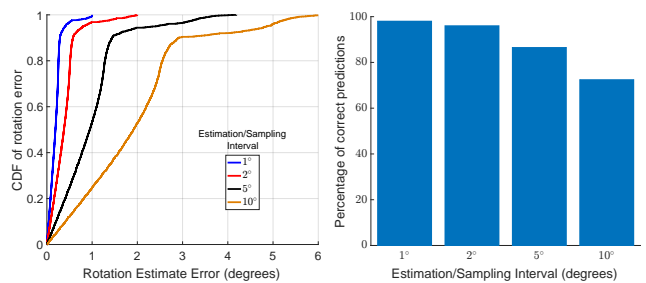
Figure 8: Sector selection accuracy with horn antennas.

experimental setup, such that the average distance to the AP increases across the trajectories. Fig. 7b shows that the error across various positions along the trajectories decreases as the distance of trajectories from the AP increases. This is because of an additional factor which impacts estimation accuracy; the validity of our negligible array dimensions approximation, which improves with an increase in distance.

Findings: LiSteer estimates il-AoA within 3.5° for more than 90% of measurement instances in OTA experiments, demonstrating the viability and high accuracy of our estimation method. Further increase in accuracy can be realized by exploiting device edges to reduce inter-sensor distances.

4.3 Results with Horn Antenna Platform

4.3.1 Beam Acquisition. First we evaluate sector selection accuracy during the Beam Acquisition phase at the client, by comparing sectors selected by LiSteer based on il-AoA estimates to the true optimal sector by performing an exhaustive search over all the sectors at each position. For comparison, we also compute the sectors based on perfect AoA from geometry. Fig. 8 depicts the CDF of client-side sector selection accuracy captured by the metric ‘Sector Index Distance’ defined as the absolute difference between indices of predicted sectors and the optimal sectors from exhaustive search. With perfect AoA information, we always get the optimal sector which is consistent with very narrow and near-perfect beam patterns of horn antennas with no side-lobes. In comparison, LiSteer selects the optimal sectors more than 70% of



(a) CDF of rotation error. (b) Beam steering accuracy.

Figure 9: Rotation with various estimation intervals.

instances with il-AoA estimates, while the error is limited to immediately adjacent sector for all except one instance.

Findings: With highly directional beams of 7° wide horn antennas, LiSteer acquires client beams to within 1 sector of the true highest strength sector without any in-band training. Thus our light based il-AoA estimation predicts the correct client sector with high accuracy and eliminates the need for exhaustive beam search at mobile clients.

4.3.2 Beam Steering. Next we evaluate LiSteer performance under client mobility. For space constraints, we discuss rotation results here. A key factor affecting beam steering accuracy is the frequency at which il-AoA estimates are computed. This is determined by multiple factors, such as sampling frequency of light sensors and computational resources of smart devices. Moreover, rotational speed of the client may also affect steering accuracy; the faster the speed, the harder the tracking since the client may rotate more for the same estimation frequency. Therefore, instead of evaluating all these factors separately, we normalize the estimation frequency to the client’s rotation, such that an il-AoA estimate is computed for every δ degrees of client’s rotation. Here we present results for four δ values: 1° , 2° , 5° and 10° .

First we analyze rotation estimation accuracy by computing the change in estimated il-AoA between the initial and final orientations of the client. Fig. 9a plots the CDF of rotation estimation error for the four δ values. We observe that when il-AoA estimates are computed most frequently (every 1° of client rotation), the estimation error is the lowest since we have the most light measurements to estimate the same rotation compared to the other cases. Further, we observe that the rotation estimates are within 0.5° of the true value for more than 90% of instances, which is much higher than absolute il-AoA estimation accuracy in Fig. 7a. This is because of cancellation of location bias when computing the change in il-AoA to find client rotation. This bias results from deviation of measurements from the theoretical model in Eq. (4), depending on AP-client relative angle and distance. Moreover, the graph also shows that as the estimation interval (δ) increases, rotation estimation error also becomes

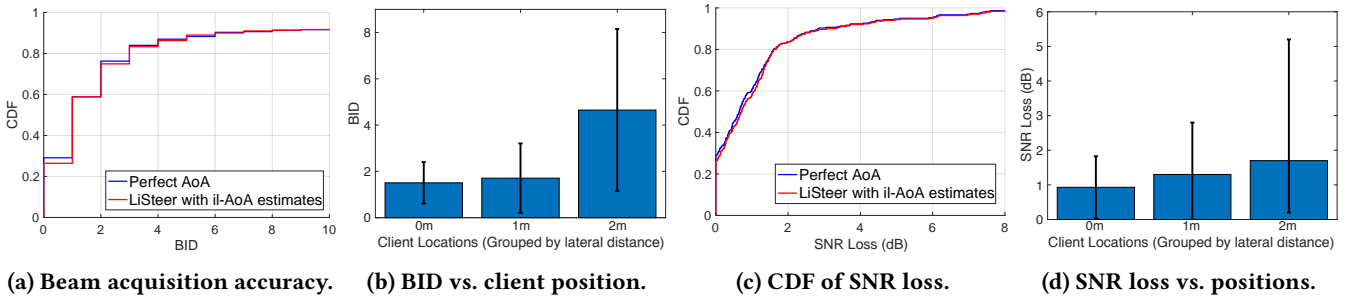


Figure 10: Client-side beam acquisition accuracy using il-AoA estimates in experiments with X60 platform.

large, since there is a greater change in client orientation between two measurements.

Next we analyze the client-side sector steering accuracy (after perfect initial alignment) in Fig. 9b for various values of δ plotted along the x -axis. Consistent with a high rotation estimation accuracy, we observe that LiSteer is able to steer client sectors to the true highest strength sectors more than 97% of the time with $\delta = 1^\circ$. Although steering accuracy decreases with increasing estimation interval, even with a high interval of 10° , which represents very high rotational speeds or conversely very low sampling rate of sensors, LiSteer computes the correct sectors more than 70% of the time.

Findings: By estimating changes in il-AoA from the AP's light source, LiSteer is able to track rotation at even higher accuracy than it does absolute il-AoA, leading to almost 97% steering accuracy when il-AoA estimates are computed at a modest rate of every 1° of client rotation. Consequently, LiSteer maintains beam alignment at the narrowest level despite device mobility solely by passive light sensing.

4.4 Results with X60 Phased Array Platform

4.4.1 Beam Acquisition. We perform Beam Acquisition at each position/orientation combination in the above experimental setup (i.e., 9 positions along each of five trajectories with 0° orientation, and 25 different orientations for 15 positions marked by black circles in Fig.6), such that the AP does a beam sweep at its end, while the client-side beam is selected based on il-AoA estimates. For comparison with optimal beam selection via exhaustive search, we repeat the experiments such that a beam sweep is done at the client side as well. To evaluate beam acquisition accuracy, we devise a metric to capture how far two beams are in terms of codebook search space. Thus we define Beam Index Distance (BID), as the absolute difference between the indices of the exhaustive search based maximal strength beam and the beam inferred by LiSteer. For comparison, we also compute the beam which will be selected if the true LOS path AoA is known (via geometry) at the client.

Beam Selection Accuracy: Fig. 10a depicts the client-side beam selection accuracy plotted as a CDF of BID. First we observe that even with true AoA, there is a non-zero BID for almost 70% experimental instances. This is in contrast to the horn antenna measurements in the previous section, where the geometric AoA always leads to the selection of the highest strength beams for LOS paths. The primary reason for this deviation is the imperfect beam patterns of the SiBeam phased array and presence of side-lobes, which can capture multiple paths in the channel to show slightly higher strength. Nonetheless, for almost 80% instances, the difference is within 2 beams from the beams selected via exhaustive search. The figure also shows that beam selection accuracy for over-the-air il-AoA estimates matches very closely to the beams selected using the geometric AoA and the error is negligible. The reason is that our AoA estimation error is very small compared to the relatively wide (25° to 35°) beams of the phased array platform. This shows that il-AoA estimates in LiSteer are sufficiently accurate for geometric AoA based beam selection in such platforms as well.

We also observe that $BID \geq 6$ for about 10% instances. This is due to the presence of strong reflectors close to some positions (especially close to side-walls), combined with the irregular patterns such that for some orientations there is not sufficient directivity gain along the LOS path and a reflected path is actually stronger. To investigate further, we separately plot BID for the 15 positions where we consider multiple orientations, grouped together as columns with different lateral distances from the AP. Positions in the first column have 0m lateral distance and hence are located right in front of the AP. We expect these client locations to have a dominant LOS path for all orientations. The middle column has positions at lateral distance of 1m from the AP. The third column, with positions at 2m lateral distance from the AP is closest to the side-walls. These positions potentially can have strong non-LOS paths due to reflections from the walls. Fig. 10b plots the average BID across all orientations and positions for the three client-position groups. The figure reveals that the positions in-front of the AP and in the middle in fact show an average BID of 1.5 with relatively small standard

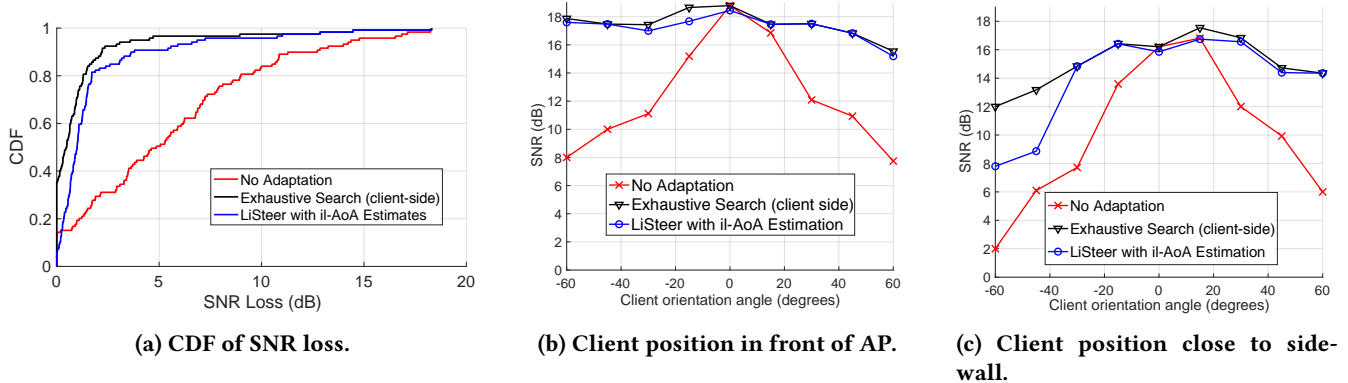


Figure 11: Beam steering accuracy for various adaptation strategies under client rotation with X60 platform.

deviation, whereas the positions closer to the reflectors have significantly large BID with even larger deviation.

SNR Performance: As discussed above, for almost 70% instances even with perfect AoA estimates the predicted beams are sub-optimal as compared to beams selected via an exhaustive search. Next we evaluate the cost of sub-optimal beam selection in LiSteer by analyzing the loss in SNR incurred compared to the maximum SNR achieved via exhaustive search.

Fig. 10c plots the CDF of SNR loss (in dB) incurred by selecting client-side beams using AoA (both il-AoA estimates and perfect geometrical AoA) as compared to optimal selection via exhaustive search. We observe that the SNR loss is within 1.5 dB for more than 80% of instances, and that the performance with il-AoA estimation is very close to that with perfect AoA. This is because of the overlap between adjacent beam patterns of the phased array, which results in a small degradation in link SNR despite selecting beams which are one or two indices away from the optimal beams. Finally, we plot the average SNR loss for the three client-position groups in Fig 10d. Our results show that for the first two groups with expected dominant LOS path, the SNR loss is within 1.2 dB. This is consistent with their lower average BID as discussed above. However, it is interesting to note that for the positions close to the side wall, although we observed a high average BID, the average SNR loss is still below 2 dB. This shows that although the LOS path does not provide the strongest signal component for some orientations at these positions, resulting in a high BID, LiSteer can still establish a strong link within 6 dB of the highest possible SNR.

Our above analysis shows that although il-AoA based beam selection is not always optimal, there is minimal loss in link SNR as compared to overhead-intensive exhaustive search, and LiSteer achieves this beam selection without incurring any overhead. Hence LiSteer successfully eliminates the need for beam training at the client side with a small loss in SNR. If even higher link strength is desired, for more

than 80% of instances, the optimal beam is within two indices of il-AoA based beam, and hence local search in the beam-space around the il-AoA predicted beam can be performed to further improve the accuracy of beam selection. In any case, LiSteer successfully acquires the AP's direction and eliminates the requirement for exhaustive search at the client end.

Findings: Even with phased antenna arrays with non-uniform beam patterns and strong side-lobes, LiSteer is able to acquire client-side beams within 2 indices of the optimal beams for more than 80% instances, with il-AoA estimates almost always selecting the same beam as the true geometric AoA. Moreover, the SNR loss is within 1.5 dB of exhaustive search, showing that there is only a marginal gain for performing exhaustive search at the client-end.

4.4.2 Beam Steering: Client Rotation. Next we evaluate beam steering accuracy for rotational mobility of the client device. For this, the client performs Beam Acquisition at an initial orientation (0°). We then rotate the client by 60° in both clock-wise and counter-clockwise directions in steps of 5° for 15 different client locations and evaluate beam steering accuracy. As per LiSteer design, after beam acquisition, the client adapts its beam in response to any changes in il-AoA estimates. For comparison, we repeat each experiment with no beam adaptation, as well as client-side adaptation with exhaustive search (perfect adaptation). Moreover, AP-side beams remain fixed in this experiment. Fig. 11a depicts the CDF of SNR loss for various adaptation schemes. The SNR loss, plotted along the x -axis, is the link degradation due to sub-maximal beam selection, compared to perfect adaptation via exhaustive beam search at both AP and client ends, which achieves maximum SNR (i.e., 0 dB SNR loss). First, we observe that even with exhaustive search based client beam-adaptation, there is non-zero SNR loss for almost 60% of instances. This is again an artifact of imperfect beam patterns and the reflectors in the environment: even though the position of the client is not changing, the AP-side beam that

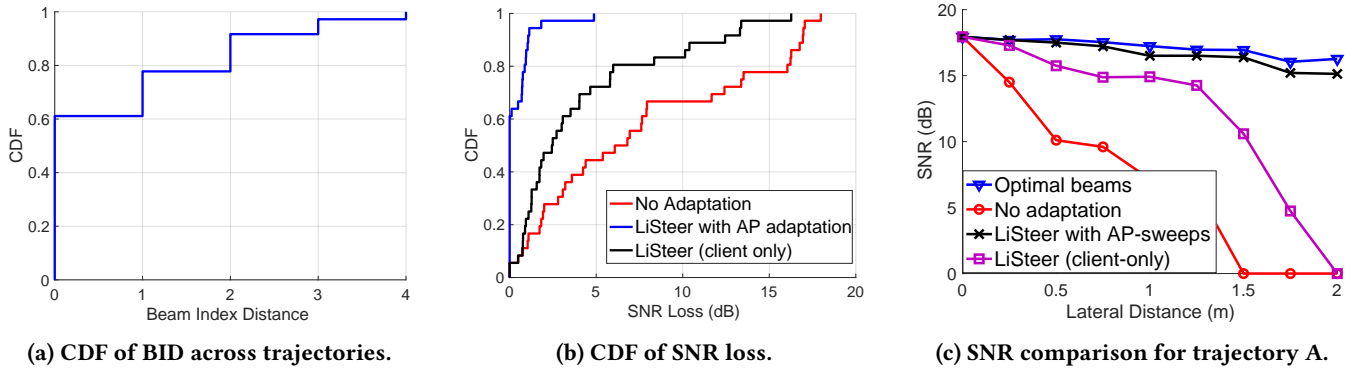


Figure 12: Beam steering accuracy for various adaptation strategies under client translation with X60 platform.

achieves maximal SNR changes and can only be addressed by AP-side adaptation. This is also in contrast to horn antenna experiments with finer beams and near-idealized antenna patterns, where AP-side beam remains fixed in case of client rotation.

LiSteer’s il-AoA based adaptation performs very close to exhaustive search and is able to maintain alignment at the client side despite mobility. Its performance is within 0.5-1 dB of exhaustive search, despite not performing any beam training after the link is setup during beam acquisition. The benefit of LiSteer adaptation is further emphasized by comparison with no-adaptation strategy, which shows high SNR degradation (more than 75% instances have > 3 dB SNR loss) despite wider beam patterns and overlap between multiple beams of the phased array.

The above analysis is based on a CDF of rotation at all 15 positions. Although it captures the overall performance of various beam adaptation schemes, the performance is highly dependent on the angular separation between the AP and client, and the proximity to reflectors in the environment such as side walls. To explore this further, we compare the performance of the aforementioned schemes for two different positions.

First we consider a client position in front of the AP (0m lateral distance, 2m longitudinal distance) such that it is far from the side-walls. Fig. 11b depicts that with no-adaptation, SNR quickly degrades for rotation on either side. However, even with 60° rotation, the link maintains SNR above 7.5 dB, required for 500 Mbps data rate in X60. This is because of the relatively short AP-client distance, for which the initial receiver beam is wide enough to maintain moderate SNR despite rotation. LiSteer, on the other hand, shows up to 10 dB gain in SNR compared to no-adaptation, achieving almost the same SNR as exhaustive search for most rotation angles.

Next we consider a position close to the side-wall in the same row (2m lateral distance, 2m longitudinal distance) in Fig. 11c. Here, the no-adaptation scheme performs even

worse, and SNR is not high enough to support even the base rate for extreme rotation angles. LiSteer again achieves SNR close to exhaustive search for clockwise rotation, with il-AoA estimates only. However, for counter-clockwise rotation, the SNR for -45° and -60° angles is more than 4 dB lower than client-side exhaustive search. This is because for these client orientations, the reflected path from the side wall provides a much stronger signal, which LiSteer is unable to discover since it tracks LOS AoA only. However, even in these few cases, LiSteer still avoids link breakage and achieves significant gain compared to no-adaptation.

4.4.3 Beam Steering: Client Translation. To study LiSteer performance under translational mobility, we consider the five client trajectories described in the experimental setup, such that the client performs beam acquisition at the initial position in front of the AP. It then covers a 2m distance along a straight line laterally from the AP. For each measurement location along a trajectory, we compare the performance of LiSteer with exhaustive search, where beam sweep is performed at both ends.

Beam Steering Accuracy: Fig. 12a depicts the BID of client beams predicted across all points along the five trajectories. We observe that for almost 60% instances, LiSteer predicts the correct client side beam. Beams predicted by LiSteer are almost the same as geometric beams based on true AoA knowledge, similar to the case of rotation experiments. Overall, the geometric beam steering accuracy is higher for translation, since the client orientation remains fixed during translation and for most locations we observe that a dominant LOS path is present. Further, more than 90% of predictions are within two beam indices of exhaustive search.

SNR Performance: Next we evaluate the impact of beam steering accuracy on link SNR. Fig. 12b shows the CDF of SNR achieved along all trajectories for three different adaptation schemes. First, we plot the no-adaptation strategy, where the beams selected at both ends during beam acquisition are

used throughout the client translation motion and there is no beam adaptation. This strategy leads to a significant decrease in link throughput, with more than 50% of instances incurring $> 7\text{dB}$ SNR loss compared to perfect adaptation when optimal sectors are used at both ends. This link degradation corresponds to loss of several MCS levels. This reiterates the importance of beam steering in mmWave networks, such that only slight motion can significantly impact link strength despite wide beam patterns and significant beam overlap of a practical phased array.

Since AP-side adaptation in LiSteer is handled by the beacon sweeps of the underlying protocol (e.g., 802.11ad), AP-side beam selection depends on how often and where along the trajectory the beacon sweep is performed, and also on the translation speed of the client. Therefore, here we present two schemes which serve as upper and lower bounds of performance, depending on the frequency of AP-side adaptation and client speeds. First, we study the case of perfect AP-side adaptation, such that the AP's beams are updated at each measurement location. For this, we observe that il-AoA estimation in LiSteer achieves near-optimal link strength, with $< 1\text{dB}$ SNR loss for more than 90% instances. This may correspond to a single MCS level loss or no loss at all as compared to exhaustive search, depending on the MCS schemes and true SNR. Thus, there is very little gain in repeating beam training at the client end and LiSteer estimates are sufficient to maintain a highly directional link. Second, we consider LiSteer with no AP-side adaptation, such that this strategy is completely training-overhead free. While no AP-side adaptation impacts the link strength, 50% of instances incur $< 3\text{dB}$ SNR loss. Hence, while sub-optimal, there is still significant gain in client-only adaptation as compared to no adaptation. This shows that if AP adaptation is not possible or is delayed due to contention, LiSteer can still delay link breakage and maintain beam alignment for a longer range compared to no adaptation at all.

To analyze the adaptation schemes further, we study the example of SNR changes along client trajectory A in Fig. 12c. First, the curve for *perfect* beam selection via exhaustive search at all points represents the 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. Next we consider the case of no beam adaptation, where the AP and the client keep using the beams initially selected via beam training. For this scenario, the link SNR degrades sharply despite wider beam patterns and significant beam overlap of the phased array, and the 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, LiSteer with AP adaptation maintains near-optimal link SNR for most client positions along the trajectory, and SNR loss is within 1 dB. Finally we consider LiSteer with client-only adaptation such that AP's beam remains fixed throughout the experiment. The figure shows that LiSteer still achieves up to 8 dB gain over no-adaptation scenario, maintaining Gbps data rates up to 1.5m translation. This is especially useful in cases when there is a delay in beacon sweeps at the AP.

Finding: Without beam adaptation, mmWave links can lose multi-Gbps data rates via a mere 0.5m translation despite wider beam patterns of a practical phased array, highlighting their susceptibility to client mobility. With LiSteer, the client maintains a highly directional link with SNR within 1 dB of exhaustive search for most positions along the trajectory via il-AoA estimation. Moreover, if only the client is adaptive, the AP may incorrectly hold on to an older beam too long without necessarily incurring link breakage.

5 SIMULATION RESULTS

5.1 Experimental Setup

Here we evaluate LiSteer performance using our trace-and model-driven simulator. In particular, we study indoor WLAN scenarios under a variety of mobility and contention scenarios. For this, we consider multiple, fully backlogged clients which undergo random waypoint mobility at average human walking speed (1.5 m/s) and different rotational speeds. We adopt the commonly used value of 100 ms for beacon sweep interval. For performance comparison, we also 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 beam training whenever the data rate drops below MCS 1. Hence it is not excessively incurring repeated training overhead, yet maintains data rates above 385 Mbps.

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

5.2 Results

Training Overhead Comparison: First, we compare the training overhead incurred by the two schemes in the aforementioned experiments. Fig. 13 depicts overhead vs. the

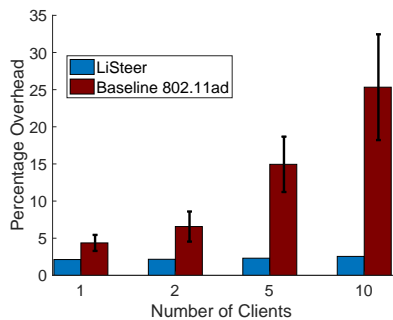


Figure 13: Training overhead comparison.

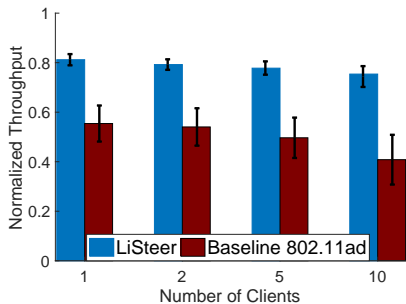


Figure 14: Normalized throughput comparison.

number of clients for 1.5 m/s translational speed. We calculate overhead as the percentage of total time used to adapt 60 GHz beams (for all clients), which comprises beam training overhead for the baseline scheme whereas for LiSteer it includes Beam Acquisition overhead for link establishment and the time spent on multi-client training. The figure shows that for the baseline scheme, even with a single client the overhead is around 5% and includes beacon sweeps and training for beam adaptation. Moreover, as the number of clients increases, the overhead increases steeply due to beam training with multiple clients and contention among clients for training with the AP. In contrast, LiSteer’s overhead comprises entirely of periodic beacon sweeps since client-side steering is driven by il-AoA estimates via passive light sensing and hence is completely distributed. Thus the overhead stays almost constant despite an increase in the number of clients and is limited to around 2%.

Finding: LiSteer incurs negligible overhead for polling after the mandatory beacon sweeps by the AP, with client side beams acquired and steered using il-AoA estimates only. Thus its overhead also remains nearly constant despite an increase in the number of clients, incurring 10× to 15× lower overhead than an 802.11ad based baseline scheme with in-band retraining.

Throughput: Next, we analyze the throughput performance of LiSteer and the baseline scheme. Further, we normalize to the throughput of an omniscient scheme which uses optimal beams and data rates for each transmission. Hence, it achieves maximum data rate, while incurring the same channel access and contention overhead.

Fig. 14 depicts normalized throughput vs. the number of clients for 1.5 m/s translational speed. First we consider the case of a single client, where LiSteer achieves 82% of the maximum throughput on average. The loss in throughput is due to sub-maximal data rates resulting from beam steering inaccuracy and MCS under-selection, or packet losses due to MCS over-selection. In comparison, the 802.11ad baseline scheme achieves only ~ 57% throughput, due to rate adaptation losses, beam misalignment and training overhead as discussed above.

When we increase the number of clients in the network, 802.11ad throughput decreases further due to increased collisions, training overhead and latency since multiple clients are contending to train with a single AP. In contrast, LiSteer normalized throughput remains above 75% for up to 10 clients due to il-AoA driven beam steering which eliminates training overhead at the client-side. Moreover, multiple clients can simultaneously train with the AP with a single sweep during beacon intervals, resulting in a small overhead.

Our simulations also encompass different client speeds and rotational mobility scenarios, which we omit here for want of space. Under rotational mobility, LiSteer exhibits further gains compared to the baseline scheme since rotation requires steering at the client-side only to maintain beam alignment, which has no overhead in LiSteer.

Finding: Due to repeated beam training overhead in 802.11ad, more than 50% of available throughput is lost due to beam misalignment from translational and rotational mobility. By passively tracking AP’s LED using light sensors, LiSteer achieves up to 2× gain in throughput by avoiding training overhead in client-side adaptation, and maintaining AP-side alignment via beacon sweeps only. Moreover, its performance scales much better with rotational speed and the number of client devices.

6 RELATED WORK

Visible Light Signaling: Recent work on Visible Light Communication (VLC), where indoor luminaries are used for data transmission, provides solutions such as controlling the brightness of light sources via pulse width and position modulation to maintain the same average power while still communicating data at reasonable range [23, 24], even in the dark [34]. Similarly, multiple solutions exist to distinguish the target light source from ambient light or interfering sources, such as decoding specific signatures modulated by LEDs [15], periodic beacons from VLC sources [18], frequency hopping [17], and exploiting characteristic frequency of LEDs [43]. Any of these solutions can be integrated in LiSteer to extend the AP’s LED range and distinguish it from other light sources if required. Moreover, if the AP also supports VLC, LiSteer can passively track the VLC source without requiring any special signaling or calibration.

Visible Light Sensing: Recent visible light localization solutions such as model-driven multi-lateration [17, 41] and fingerprinting [42, 43] can also be used to track position at sub-meter accuracy. However, these solutions require multiple luminaries at known locations, whereas multi-lateration approaches also require knowledge of transmit powers as well as fixed and known *orientation* of light sensors. Hence they cannot be used to track client-side beams, which can be affected by both position and orientation changes. In contrast, in LiSteer, we estimate AoA from a *single* light source located at the AP without requiring position or orientation estimation. Prior work on AoA estimation using incoherent light includes model driven AoA calculation for localization from multiple light sources [16] and using two photodiodes of different fields of view such that their angular response patterns are also different and can be modeled as a non-linear function of the AOA [43]. However, [43] requires calibration to discover this differential response function, and is limited to 1-D AoA estimation (insufficient for 3D beam steering). Likewise, [15] uses image sensors to estimate AoA. However, image processing requires higher power and computational resources with a latency of several seconds. In contrast, LiSteer uses off-the-shelf light sensors with significantly lower power consumption, and tracks beam changes at milli-second time scale.

mmWave Beam Training: *In-band* solutions to reduce training overhead include model-driven beam steering and channel profiling [37, 44], hierarchical codebook designs [2, 12], compressive sensing techniques to exploit channel sparsity [20, 30], correlation between beams [33], efficient beam searching [32, 40], sector switching and backup paths [10, 29], and beamwidth adaptation [10]. These solutions reduce training overhead and maintain alignment in certain environments, yet still incur training overhead when constructing channel profiles, searching for backup or redundant paths, or incur SNR loss when switching to wider beams. Moreover, packet level beam tracking solutions to address mobility have also been proposed, e.g., 802.11ad’s beam tracking [13], exploiting multi-lobe beam patterns [19] and beam sounding [10]. 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, in LiSteer, we target to eliminate beam search at mobile devices entirely by obtaining direction estimates from existing LEDs on APs for both beam acquisition and beam adaptation, and without requiring any training or feedback in mmWave band for this purpose. Nonetheless, prior solutions can be integrated to reduce training overhead for AP-side sweeps, to further improve beam steering at packet-level, or when light estimates are not available.

Lastly, prior *out-of-band* solutions also address link adaptation in directional networks, e.g., via session transfer to legacy bands [25, 31], AoA estimation in sub 6-GHz bands to eliminate exhaustive search for beam acquisition [22], and using sensors on mobile devices [5, 8, 26, 40]. In contrast, passive light sensing has significantly less power requirements than mechanical sensors and tracks the AP under both translation and rotation, requires no communication in the sensing band, and is more resilient to multipath due to the dominant LOS propagation of visible light. SearchLight [9] is the only prior work that also uses visible light sensing with a sensor-array to track device mobility (both position and orientation changes). While SearchLight adapts beams at both the AP and client ends, it requires feedback packets for AP-side beam adaptation, AP localization and maintaining LOS with multiple (≥ 3) light sources which presents greater challenge in terms of infrastructure requirements and LOS blockage. In contrast, LiSteer exploits a single LED collocated at the AP to adapt client beams without any feedback and handles AP-side adaptation by repurposing periodic beacon-sweeps.

An earlier version of LiSteer design appears in [11] where we present the key idea of beam adaptation by tracking AP’s indicator LEDs with proof-of-concept implementation on a horn antenna system. In this paper, we have enhanced the protocol design to include multi-client beam steering to address AP-side beam adaptation, extended the hardware implementation to include practical phased arrays with irregular beam patterns which impact the optimality of LOS path based geometric beam steering, and evaluated in WLANs with multiple clients (including impact of contention) and various indoor mobility patterns.

7 CONCLUSION

We present LiSteer to steer mmWave beams at mobile devices by tracking indicator LEDs on wireless APs to passively acquire direction estimates, and demonstrate that LiSteer acquires and maintains beam alignment despite device mobility, without incurring any training overhead at clients. We implement our system on a custom dual-band platform. Our experiments show that LiSteer estimates the incoherent-light AoA to within 2.5° on average and steers beams correctly more than 97% of instances while in tracking mode, without incurring any in-band training or feedback.

8 ACKNOWLEDGMENTS

We thank the reviewers for their valuable feedback and our shepherd for help in improving the paper and getting it ready for the final submission. This research was supported by Cisco, Intel, and by NSF grants CNS-1629929, CNS-1553447, CNS-1642929 and CNS-1514285.

REFERENCES

- [1] Omid Abari, Haitham Hassanieh, Michael Rodriguez, and Dina Katabi. 2016. Millimeter Wave Communications: From Point-to-Point Links to Agile Network Connections. In *Proc. of ACM HotNets*.
- [2] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W Heath. 2014. Channel Estimation and Hybrid Precoding for Millimeter Wave Cellular Systems. *IEEE Journal of Selected Topics in Signal Processing* 8, 5 (2014), 831–846.
- [3] John R Barry. 2012. *Wireless Infrared Communications*. Vol. 280. Springer Science & Business Media.
- [4] Visual Networking Index Cisco. 2016. Global Mobile Data Traffic Forecast Update, 2015–2020 White Paper. *Document ID 958959758* (2016).
- [5] Arjan W Doff, Kishor Chandra, and R Venkatesha Prasad. 2015. Sensor Assisted Movement Identification and Prediction for Beamformed 60 GHz Links. In *Proc. of IEEE CCNC*.
- [6] Yasaman Ghasempour, Claudio RCM da Silva, Carlos Cordeiro, and Edward W Knightly. 2017. IEEE 802.11 ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* 55, 12 (2017), 186–192.
- [7] Yasaman Ghasempour and Edward W Knightly. 2017. Decoupling Beam Steering and User Selection for Scaling Multi-User 60 Ghz WLANs. In *Proc. of ACM MobiHoc*.
- [8] Nuria González-Prelcic, Anum Ali, Vutha Va, and Robert W Heath. 2017. Millimeter-Wave Communication with Out-of-Band Information. *IEEE Communications Magazine* 55, 12 (2017), 140–146.
- [9] Muhammad Kumail Haider, Yasaman Ghasempour, and Edward W Knightly. 2018. SearchLight: Tracking Device Mobility using Indoor Luminaries to Adapt 60 GHz Beams. In *Proc. of ACM MobiHoc*.
- [10] Muhammad Kumail Haider and Edward W. Knightly. 2016. Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation. In *Proc. of ACM MobiHoc*.
- [11] Muhammad Kumail Haider and Edward W Knightly. 2018. iTrack: Tracking Indicator LEDs on APs to Bootstrap mmWave Beam Acquisition and Steering. In *Proc. of ACM HotMobile*.
- [12] Sooyoung Hur, Taejoon Kim, David J Love, James V Krogmeier, Timothy A Thomas, and Amitava Ghosh. 2013. Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks. *IEEE Transactions on Communications* 61, 10 (2013), 4391–4403.
- [13] IEEE 802.11 Working Group. 2012. IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. (2012).
- [14] National Instruments. 2017. Introduction to the NI mmWave Transceiver System Hardware - National Instruments. <http://www.ni.com/white-paper/53095/en/>. (2017). (Accessed on 06/25/2017).
- [15] Ye-Sheng Kuo, Pat Pannuto, Ko-Jen Hsiao, and Prabal Dutta. 2014. Luxapose: Indoor Positioning with Mobile Phones and Visible Light. In *Proc. of ACM MobiCom*.
- [16] Seongsu Lee and Sung-Yoon Jung. 2012. Location Awareness using AoA based Circular-PD-Array for Visible Light Communication. In *Proc. of IEEE APCC*.
- [17] Liqun Li, Pan Hu, Chunyi Peng, Guobin Shen, and Feng Zhao. 2014. Epsilon: A Visible Light Based Positioning System. In *Proc. of USENIX NSDI*.
- [18] Tianxing Li, Chuankai An, Zhao Tian, Andrew T Campbell, and Xia Zhou. 2015. Human Sensing using Visible Light Communication. In *Proc. of ACM MobiCom*.
- [19] Adrian Loch, Hany Assasa, Joan Palacios, Joerg Widmer, Hans Suys, and Bjorn Debaillie. 2017. Zero Overhead Device Tracking in 60 GHz Wireless Networks using Multi-Lobe Beam Patterns. In *Proc. of ACM CoNEXT*.
- [20] Zhinus Marzi, Dinesh Ramasamy, and Upamanyu Madhow. 2016. Compressive Channel Estimation and Tracking for Large Arrays in mm-Wave Picocells. *IEEE Journal of Selected Topics in Signal Processing* 10, 3 (2016), 514–527.
- [21] T. Nitsche, C. Cordeiro, A.B. Flores, E.W. Knightly, E. Perahia, and J.C. Widmer. 2014. IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gigabit-per-second Wi-Fi. *IEEE Communications Magazine* 52, 12 (2014), 132–141.
- [22] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer. 2015. Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement. In *Proc. of IEEE INFOCOM*.
- [23] Mohammad Noshad and Maithe Brandt-Pearce. 2012. Expurgated PPM using Symmetric Balanced Incomplete Block Designs. *IEEE communications letters* 16, 7 (2012), 968–971.
- [24] G Ntogari, T Kamalakis, JW Walewski, and T Spicopoulos. 2011. Combining Illumination Dimming based on Pulse-Width Modulation with Visible-Light Communications based on Discrete Multitone. *Journal of Optical Communications and Networking* 3, 1 (2011), 56–65.
- [25] Eldad Perahia, Carlos Cordeiro, Minyoung Park, and L Lily Yang. 2010. IEEE 802.11 ad: Defining the Next Generation Multi-Gbps Wi-Fi. In *Proc. of IEEE CCNC*.
- [26] Lenin Ravindranath, Calvin Newport, Hari Balakrishnan, and Samuel Madden. 2011. Improving Wireless Network Performance using Sensor Hints. In *Proc. of USENIX NSDI*.
- [27] Swetank Kumar Saha et al. 2017. X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays. In *Proc. of ACM WinTech*.
- [28] Swetank Kumar Saha, Tariq Siddiqui, Dimitrios Koutsonikolas, Adrian Loch, Joerg Widmer, and Ramalingam Sridhar. 2017. A Detailed Look into Power Consumption of Commodity 60 GHz Devices. In *Proc. of IEEE WoWMoM*.
- [29] Sumit Singh, Federico Ziliotto, Upamanyu Madhow, Elizabeth M. Belding, and Mark Rodwell. 2009. Blockage and Directivity in 60 GHz Wireless Personal Area Networks. *IEEE Journal on Selected Areas in Communications* 27, 8 (2009), 1400–1413.
- [30] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. 2017. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11 ad Devices. In *Proc. of ACM CoNEXT*.
- [31] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom*.
- [32] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Proc. of ACM SIGMETRICS*.
- [33] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of USENIX NSDI*.
- [34] Zhao Tian, Kevin Wright, and Xia Zhou. 2016. The DarkLight Rises: Visible Light Communication in the Dark. In *Proc. of ACM MobiCom*.
- [35] TP-Link. 2018. Talon AD7200 Multi-Band Router. <http://www.tp-link.com/en/products/details/AD7200.html>. (2018).
- [36] Vubiq. 2015. V60WGD03 60 GHz Waveguide Development System. <http://vubiq.com/v60wgd03/>. (2015).
- [37] Teng Wei, Anfu Zhou, and Xinyu Zhang. 2017. Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors. In *Proc. of USENIX NSDI*.
- [38] Jie Xiong and Kyle Jamieson. 2013. ArrayTrack: A Fine-Grained Indoor Location System. In *Proc. of USENIX NSDI*.
- [39] Hao Xu, Vikas Kukshya, and Theodore S Rappaport. 2002. Spatial and Temporal Characteristics of 60-GHz Indoor Channels. *IEEE Journal on Selected Areas in Communications* 20, 3 (2002), 620–630.

- [40] Z. Yang, P. H. Pathak, Y. Zeng, and P. Mohapatra. 2015. Sensor-Assisted Codebook-Based Beamforming for Mobility Management in 60 GHz WLANs. In *Proc. of IEEE MASS*.
- [41] Zhice Yang, Zeyu Wang, Jiansong Zhang, Chenyu Huang, and Qian Zhang. 2015. Wearables Can Afford: Light-weight Indoor Positioning with Visible Light. In *Proc. of ACM MobiSys*.
- [42] Chi Zhang and Xinyu Zhang. 2016. LiTell: Robust Indoor Localization using Unmodified Light Fixtures. In *Proc. of ACM MobiCom*.
- [43] Chi Zhang and Xinyu Zhang. 2017. Pulsar: Towards Ubiquitous Visible Light Localization. In *Proc. of ACM MobiCom*.
- [44] Anfu Zhou, Xinyu Zhang, and Huadong Ma. 2017. Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-Driven Beam Steering. In *Proc. of IEEE INFOCOM*.