X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays

Swetank Kumar Saha [a,*], Yasaman Ghasempour [b], Muhammad Kumail Haider [b], Tariq Siddiqui [a], Paulo De Melo [a], Neerad Somanchi [a], Luke Zakrajsek [a], Arjun Singh [a], Roshan Shyamsunder [a], Owen Torres [a], Daniel Uvaydov [a], Josep Miquel Jornet [a], Edward Knightly [b], Dimitrios Koutsonikolas [a], Dimitris Pados [c], Zhi Sun [a], Ngwe Thawdar [d]

[a] *University at Buffalo, Buffalo, NY, USA*
[b] *Rice University, Houston, TX, USA*
[c] *Florida Atlantic University, USA*
[d] *Air Force Research Laboratory, Information Directorate, Rome, NY 13441, USA*

ABSTRACT

This paper introduces X60, the first SDR-based testbed for 60 GHz WLANs, featuring fully programmable MAC/PHY/Network layers, multi-Gbps rates, and a user-configurable 12-element phased antenna array. Combined these features provide an unprecedented opportunity to re-examine the most important aspects of signal propagation and performance expected from practical 60 GHz systems. Leveraging the testbed's capabilities, we conduct an extensive measurement study, looking at different aspects of indoor 60 GHz links. We find that the presence of reflective surfaces and imperfect beams generated by practical phased arrays together can result in multiple NLoS paths supporting Gbps rates. Additionally, our comparison of different beam adaptation strategies reveals how beam steering even at one end of the link can often be sufficient to restore link quality. Finally, contrary to a common assumption in recent works, we find that a one-to-one MCS to SNR mapping is hard to obtain in typical indoor environments, and probing of more than one MCS index is often required to discover the optimal MCS for a given SNR.

1. Introduction

The IEEE 802.11ad standard, using 2.16 GHz wide channels in the unlicensed band centered around 60 GHz and directional transmissions, provides data rates of up to 6.7 Gbps in an indoor WLAN setting [1]. In the near future, IEEE 802.11ay [2] promises to provide even higher data rates up to 100 Gbps. Realizing high-speed directional links, however, comes with challenges, sparking off research for the design of efficient link training/adaptation techniques. Nonetheless, most available experimental platforms either offer very limited access to the PHY/MAC layers (commercial devices) or use narrow band transmissions coupled with horn antennas (USRP/WARP combined with a 60 GHz frontend) deviating significantly from 802.11ad's use of ultra-wide channels and phased array antennas. This leaves a vacuum for a testbed that can offer the best of both worlds: a realistic PHY *and* programmability of PHY and MAC layers.

In this work, we introduce **X60** [3], the first highly configurable software defined radio (SDR) 60 GHz testbed, featuring fully *programmable PHY, MAC, and Network* layers while still allowing for ultra-wide channels and multi-Gbps data rates. Based on the National Instruments' (NI) millimeter-wave (mmWave) Transceiver System [4] and equipped with a *user-configurable* 12-element phased array antenna from SiBeam, X60 nodes enable communication over 2 GHz wide channels using realistic TX and RX beams that can be steered in real-time.

X60 offers several key advantages over other existing mmWave experimental platforms. Unlike commercial 802.11ad devices, X60 with its SDR/FPGA based architecture allows access to and complete control over the PHY and MAC layers. This not only enables experimentation that can obtain a full view of the often complex interaction among multiple layers of the networking stack, but also allows for prototyping and testing of new techniques at multiple layers. In contrast to most existing SDR mmWave experimental platforms (based on USRP/WARP), X60 provides high reconfigurability without limiting baseband bandwidth to a few hundred MHz, enabling us to study the impact of ultra-wide channels supported by the 802.11ad standard. Lastly, using SiBeam's phased array, X60 nodes generate beam patterns that are configurable and steerable in real time, overcoming a basic limitation of horn-antenna based platforms where the beam can only be steered using a mechanical

---

rotator and may not always be representative of the often imperfect beams generated by phased arrays.

X60's capabilities provide an opportunity to re-examine the understanding of the most important aspects of 60 GHz WLAN signal propagation and performance. To this end, we undertake an extensive measurement campaign across four characteristic indoor environments (corridor, lab, lobby, and conference room) in a typical academic building. Enabled by the testbed's reference implementation that uses a slotted TDD based MAC and supports multiple modulation schemes (from BPSK up to 16-QAM), we analyze various MAC performance metrics such as goodput while still having access to the underlying PHY parameters. We further study the implications of steering TX/RX beams along different directions. Our measurements encompass a range of propagation environments (dominant LoS, non-LoS only, reflections from multiple obstacles, LoS propagation with side-lobes) and TX/RX orientations.

Our major findings are as follows: *(i)* Contrary to the common belief that only a few TX/RX beam-pairs achieve high SNR in the 60 GHz band, almost 15% of the total 625 possible beam-pair combinations in our setup provide at least 1 Gbps of throughput across all environments. This is primarily due to the shape of imperfect beam patterns, overlap of main and side-lobes between neighboring beams, and richness of NLoS paths in the environment. *(ii)* At short range, the sender and receiver can be together off by several beam pair indices while maintaining high SNR, whereas at longer range, they can be off by only one or two indices, as the relative penalty for imperfect beam selection increases with distance. *(iii)* NLoS paths from strong reflectors can support links with comparable signal strength to LoS paths, and achieve multi-Gbps throughputs. *(iv)* Nodal mobility scenarios such as translation and rotation can severely degrade signal strength for a fixed pair of beams being used at TX and RX ends. Further, our analysis of various beam adaptation schemes reveals that in comparison to exhaustive search, adjusting beams at either the TX or the RX side, although sub-optimal, is adequate in most cases to restore the link. Therefore, if one node (e.g., the AP) is adaptive, the other one (client) may incorrectly hold on to an older beam for a longer time without necessarily incurring link breakage. *(v)* A one-to-one MCS to SNR mapping is hard to obtain, especially for a range of medium SNR values. Consequently, MCS selection based on the measured SNR from a lookup table derived offline, a common practice in the evaluation of the mmWave link adaptation proposals, may not always lead to the best possible performance, and probing of more than one MCS indexes is often required to discover the optimal MCS for a given SNR.

The remainder of the paper is structured as follows: Section 2 discusses related work. Section 3 describes the different components of the X60 testbed. Section 4 describes the experimental methodology of the measurement campaign and Section 5 discusses the results. Finally, Section 6 concludes the paper.

2. Related work

Initial experimental studies of 60 GHz in indoor environments focused on measuring and modeling channel propagation characteristics using dedicated channel sounding hardware (e.g., [5–9]).

The recent work in this domain has largely been driven by observations and models derived from measurements with platforms that implement narrow-band transmissions instead of wide-band [10–13], or/and horn antennas instead of phased arrays [14–17]. While initial studies based on such platforms have provided valuable insights into mmWave propagation, such setups cannot capture the effects associated with wide-band transmissions, e.g., as specified in 802.11ad. For instance, past works relied on SNR measured over few hundred MHz of bandwidth to estimate rates by looking up a receiver sensitivity table. More importantly, the use of horn antennas masks the effects of imperfect beam-patterns, side-lobes, and non-uniform steerability, all typical features of beams realized through practical phased arrays. It is

Fig. 1. Two X60 nodes, each attached to a SiBeam phased array module.

important to understand the impact of these artifacts as they directly affect mmWave link characteristics, interference, or spatial-reuse.

The only SDR testbeds capable of wide-band transmission with phased arrays are OpenMili [18] and the testbed in [19]. OpenMili nodes are based on an off-the-shelf FPGA processor supporting a channel width of 1 GHz. These nodes are equipped with electronically steerable 60 GHz four-element phased arrays, with 2 possible values for each element's weight. The testbed in [19] uses eight-element phased arrays but operates in the 24 GHz band. In contrast, X60 nodes have twelve-element phased arrays, 4 discrete possible phase values per element, and support a 2 GHz channel width, enabling higher rates and higher resolution experiments.

Many works have also explored the performance of WirelessHD, WiGig, or 802.11ad hardware available commercially [20–25,25–31]. These devices offer the chance to understand the often complex interplay between higher layers of the network stack and WirelessHD/WiGig/802.11ad directional PHY. However, they reveal rather limited information about and allow no control over the PHY/MAC parameters or the weights of the antenna elements of the phased array. These limitations, combined with proprietary rate and beam adaptation techniques, often make it hard for researchers to understand the causes of the observed performance. Further, the closed source firmware (which implements most of the lower MAC and PHY functionality) that these devices ship with limits the possibility of prototyping any new protocols.

3. X60 testbed

In this section, we describe the different components of the X60 testbed. All the modules are programmed using NI LabVIEW. Although our description refers to the reference PHY implementation provided by NI, we note that all the MAC and PHY layer components can be customized by reprogramming the FPGAs. Fig. 1 shows two X60 nodes, each attached to a SiBeam phased array module.

3.1. Baseband TX/RX

Each X60 node is based on the NI mmWave Transceiver System. All modules involved in the baseband signal generation are assembled inside a NI PXIe-1085 PXI Express chassis. Most of the inter-module signaling and data transfer happens over the chassis' high-speed backplane using FIFO queues or DMA. The TX/RX chains consist of one or more high-performance FPGAs which handle the majority of the transmit/receive operations including encoding/decoding and modulation/demodulation. The FPGA outputs feed into a wideband DAC/ADC module which generates/samples the baseband signal. In addition, the chassis holds a high-end controller (host machine) running Microsoft®Windows 7. The host generates the source bits for transmission and is the sink for the receive operation. It controls different TX/RX parameters (MCS, uplink/downlink, etc.) and collects information about different parts of the TX/RX chain to allow for user-display and debugging.

3.2. PHY/MAC structure

The current reference PHY implementation allows for the following modulation and coding (Turbo codes) combinations: 1/5 BPSK, 1/4 QPSK, 1/2 QPSK, 3/4 QPSK, 1/2 16QAM, 3/4 16QAM, 7/8 16QAM, resulting in theoretical bit rates from 300 Mbps to 4.75 Gbps. Data transmission takes place in 10 ms frames which are divided into 100 slots of 100 µs each. In each frame (100 slots), the first 2 slots are reserved for achieving time synchronization between the TX and RX, needed for proper frame decoding operation. Both the MCS and operation type (uplink/downlink/sync) can be configured on a *per-slot* basis. A slot is made up of 92 codewords (data bit sequence after encoding), each of which has an attached CRC block. At the RX end, the throughput for a given slot is calculated by counting the number of correctly decoded codewords and multiplying by the number of actual data bits contained in each codeword (which depends on the MCS).

3.3. Antenna array and beam patterns

The SiBeam mmWave module, on the TX path, takes as input the baseband signal (as differential I/Q), up-converts, and transmits over the air a 2 GHz wide waveform centered around one of the 802.11ad channel center-frequencies. The typical transmit power is 30 dBm EIRP at channel 2. The RX side flow is symmetric to the TX path. The in-built phased array has 24 elements; 12 each for TX and RX. The module connects to the baseband chassis over an additional dedicated control path that allows different phase values for the antenna elements through the use of codebooks. Different beams within a codebook can be switched by applying the required index into the module's registers. The phase of each antenna element can be set to one of four values: 0, $\pi/2$, π, $3\pi/2$.

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. The 3 dB beamwidth for the beams ranges from 25° to 30° for TX and from 30° to 35° for RX. As a result, each beam's main lobe overlaps with several neighboring beams. We refer to the beams using index range: −12 (−60°) to +12 (+60°), with index 0 corresponding to the broadside beam.

We first numerically simulated the expected beam patterns using COMSOL Multi-physics [32]. An antenna array integrated by 12 elements with the same size, separation, and spatial distribution was defined and the input at each antenna was set as per the defined codebook.

Further, the generated radiation patterns were experimentally characterized in an anechoic chamber at the Air Force Research Laboratory, Information Directorate, in the City of Rome, New York. The measurement setup in reception consisted of a 60–90 GHz horn antenna mounted on a tripod and connected indirectly to a N9030A PXA spectrum analyzer through a M1970E waveguide harmonic mixer. On the transmission end, the SiBeam array was placed at a distance of 1.5 m from the receiving antenna. The antenna was placed on a tripod on top of a rotating table and aligned so that the center axis of the array was directly above the point of rotation of the table. The main chassis was set behind the rotating table. Received power measurements were taken with the chamber completely sealed. After the platform was turned on and the RF carrier signal was initiated, a codebook value was chosen and peak power measurements were taken from the spectrum generator while rotating the table from −60° to 60° in increments of 1°.

Figs. 2a–2d depict examples of 2D (simulated and measured) and 3D (simulated only) radiation patterns for select beam indices (all 25 beam patterns are included in Appendix). The measured patterns are plotted in blue and the simulated beams are plotted in red. The measured powers were normalized from dBm to normalized dB to provide an accurate comparison to the simulated values. These patterns highlight how, in contrast to beams generated by horn antennas, phased-array generated beams often have strong side-lobes. Such side-lobes on one hand can cause interference to co-located links but on the other hand can also enable communication in scenarios where the main lobe is blocked. Moreover, as beams are steered away from the main lobe, patterns become more imperfect with even stronger side lobes. For instance, comparing beam index +3 (Fig. 2b) and the broadside beam index 0 (Fig. 2c) shows how practical phased-arrays can have *non-uniform steerability* along different directions as opposed to mechanically rotated horn antenna beams. The distorted pattern results from the effect of the microstrip antenna radiation pattern which is strongest at zero degrees, when it is normal to the face of the microstrip. The pattern is strongest and best formed when the beam is pointed towards this direction. We also observe, surprisingly, that beam indices equally apart from the broadside beam (e.g., +3 (Fig. 2b) and −3 (Fig. 2d)) can have radiation patterns that are *not* necessarily mirror images of each other. Finally, a comparison between the simulated and measured patterns reveals that the simulated patterns exhibit side lobes which are much stronger than those present in the measured patterns. This suggests that there may be a method of side lobe suppression implemented in the front end that is unknown and not accounted for in the simulation. With the side lobes discounted, the main lobes of the measured patterns match those of the simulation with a typical difference of 1–2 dB.

The remaining simulated beam patterns also exhibit similar behavior. Further, note that these particular characteristics of the beam patterns result both out of the discretization of the individual antenna element phase weights and the particular geometry in which the elements are arranged in the 2D array [33]. Nitsche et al. [22] and Steinmetzer et al. [29] also found the beam patterns of commercial WiGig and 802.11ad devices to be imperfect with strong side lobes. Also, an inspection of the open source wil6210 driver [34] targeting Qualcomm 802.11ad chipsets suggests 2 bits for phase control of the antenna elements (which allows for 4 possible phase values similar to our platform).

3.4. Enhancements for measurements

We made the following modifications to the reference code to enable logging of all the required PHY/MAC parameters and to allow for more realistic measurements.

Automatic Gain Control (AGC): We implemented an AGC block running on the host machine (every 100 ms) that adjusts the receiver's gain value based on the energy calculated from the raw I/Q samples to achieve an experimentally determined optimal target energy value that ensures best ADC operation. Through a separate set of experiments, we verified that our implementation is throughput optimal (as compared to exhaustive-search manual gain control) for different MCS and channel conditions.

Thin Control Channel: We added an external legacy WiFi radio to all four nodes to implement a reliable control path. This allows us to implement certain features like TX–RX beam selection or MCS selection and to automate parts of our measurements with only few modifications to the existing code base, without the burden of maintaining tight timing requirements of the code running the mmWave channel. The scripts that implement this control path run on the host machine and communicate with the LabVIEW process via IPC over TCP to control parameters like MCS and beam index, and collect link metrics for further processing.

Instrumentation: We instrumented the host side LabVIEW code base to log a whole range of different parameters. Since the host is an active part of the TX/RX flow (e.g., fine synchronization operations) and needs to maintain strict timing guarantees in its generator/consumer loops, we selected different logging frequency for each parameter to minimize overhead. Some parameters (Signal Power Estimation, Noise Power Mean, Throughput, CRC pattern) are logged on a per-frame basis (every 10 ms), while others, generated in already computation-heavy timed-loops (RSSI, SNR, Carrier-to-Noise, Phase, Power Delay Profile), are logged at a lower frequency (every 40 ms).
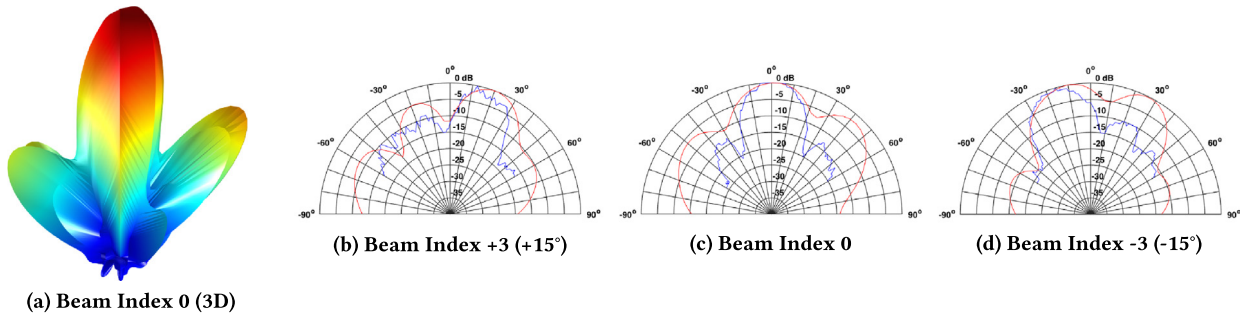
(a) Beam Index 0 (3D)

(b) Beam Index +3 (+15°)   **(c) Beam Index 0**   **(d) Beam Index -3 (-15°)**

Fig. 2. Antenna array beam patterns.

4. Measurement campaign

4.1. Methodology

Our measurement campaign is aimed at collecting key PHY and MAC layer parameters across multiple indoor environments. Measurement locations are selected to characterize static 60 GHz channels, as well as emulate typical mobility patterns like translation and rotation to study the impact of nodal mobility.

At each location, we collect channel measurements in two steps. *(i) Beam Sweep:* The first step encompasses channel estimation for all possible beam pairs in an exhaustive search. The transmitter and receiver co-ordinate their beam switching (over the control channel) to generate all 625 (25 × 25) beam-pair combinations. For each beam pair, 25 frames are transmitted at MCS 0 and SNR is logged for the channel estimation slot in each frame (every 40 ms). This is a crude emulation of 802.11ad's beam-training scheme. However, note that we only evaluate the resultant beams and not the time-efficiency of the process itself. *(ii) MCS sweep:* In the second step, we select a small subset of (TX, RX) beam pairs for which we repeat measurements at all seven achievable MCS levels. We select the three strongest beam pairs out of all 25 × 25 beam pair combinations based on average SNR computed during beam sweep in step (i). Further, to study the impact of selecting neighboring beams, we also include the immediate neighbors of the RX beam in each of the three (TX, RX) pairs, for a total of nine (TX, RX) beam pairs. For each MCS, we log all channel parameters for 500 frame transmissions.

4.2. Environments

We perform measurements in four different indoor environments in a typical office building. A detailed map of three of these environments is depicted in Fig. 3. The arrows at each TX and RX location indicate the physical orientation of the node i.e., the direction along the main lobe of the center beam of the phased array.

Corridor: The TX is fixed at one end of a 1.74 m wide corridor, at 1.23 m height. We consider 10 receiver positions varying the inter-node distance from 2.5 m to 25 m on a straight line in steps of 2.5 m, such that the RX always faces the TX. Apart from characterizing the static environment, this topology emulates receiver translation along a straight line, as it moves way from the transmitter.

Lab: The second environment is an 11.8 × 9.2 × 3.4 m³ lab with four rows of office cubicles, with cardboard partitions and metal cabinets. The TX is fixed close to the center of the right wall at 2.05 m height, a location where 60 GHz WLAN APs would be typically deployed. We choose 17 locations in three rows for taking measurements. As shown in Fig. 3a, row A (locations *1a* through *7a*) and row B (locations *1b* through *7b*) are nearly symmetric about the center partition, while the distance from the right wall increases from 2.9 m to 8.9 m in steps of 1 m. The front row (locations *1b, 0b, C, 0a, 1a*) is selected to study the impact of increased angular separation between TX and RX. The height of the

receiver is 1.26 m at all locations, such that there is always an LoS path between the TX and the RX.

Conference Room: This is a 7 × 4.87 × 3.4 m³ room (Fig. 3b) with a large central table and various metallic/shiny surfaces (e.g., TV, white boards, metal cabinet, chairs, glass windows) which have been shown to be excellent reflectors in 60 GHz band [6,35]. Hence this environment is suitable to study the impact of reflectors, non-LoS paths and sidelobes. The TX is placed in a corner at 2.23 m height, and we consider 10 different RX locations across the room with orientations depicted in Fig. 3b.

Lobby: This is the largest open space in the building, with large glass panels as walls. To study the impact of increasing distance and angular separation, we fix the TX in one corner of the lobby, and select 15 RX locations in four rows. To study the impact of receiver rotation, we vary its orientation between −90° to 90° in steps of 15°, such that 0° corresponds to RX phased array facing the front wall, parallel to side walls. Hence the central beams of both TX and RX are perfectly aligned for 0° orientation at positions 1, 7, 10 and 13 (Fig. 3c).

5. Results

In our measurements, a beam sweep captures the SNR achieved for all 25 × 25 possible beam pair combinations, each resulting in a distinct channel. As such, it can be used to study the distribution of strong SNR beam pairs and their mapping to the physical environment, and how the signal strength changes across different positions and environments. Therefore, we use beam sweeps as the main tool to understand different characteristics of 60 GHz links.

We represent each beam sweep as a heatmap of corresponding SNR values with TX beam indices along the *x*-axis and RX beam indices along the y-axis. Fig. 4 shows the beam-pair heatmap for the center position (C) in the lab with distance of 2.3 m from TX. In this figure, as well as in Figs. 6, 7, 9, 12, yellow colored regions indicate beam pairs with SNR above 10 dB whereas blue regions indicate beam pairs for which the SNR is below the receive threshold (< 0 dB, determined from measurements). The central beam pair (0,0) corresponding to the LoS path between TX and RX achieves the strongest link strength. Due to overlap between neighboring beams (Section 3.3), multiple beams may include the LoS component, albeit with a different directivity gain. Hence we get a cluster of high SNR beam pairs close to the central pair. Besides the LoS central high SNR region, there are smaller clusters of beam pairs with moderate to high SNR, resulting from reflections and side-lobes. According to Fig. 4, for TX beam indices between −2 to 3, the received SNR is above 5 dB regardless of the RX beam choice. The TX and RX are relatively close to each other which makes any RX beam (covering −60° to 60°) to achieve high SNR provided that the TX beam is pointed towards the receiver. Similarly, when RX beam indices between −2 and 4 are used, SNR is above 5 dB for most TX beams.
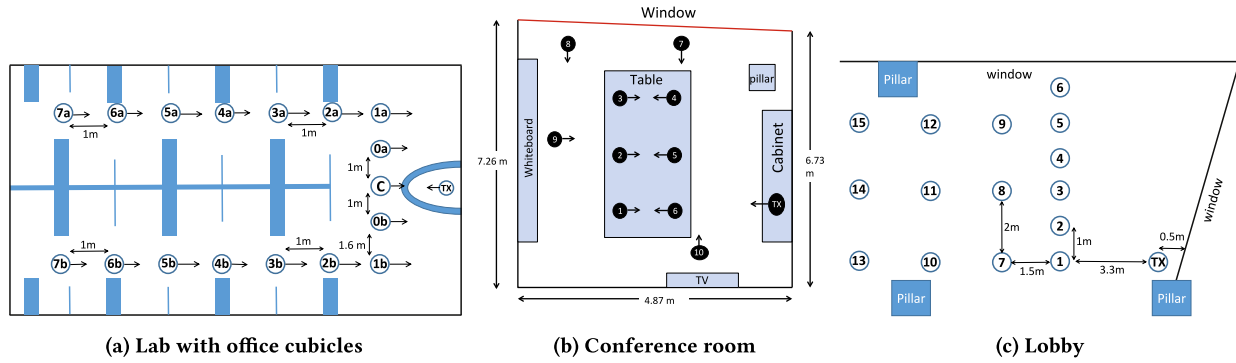
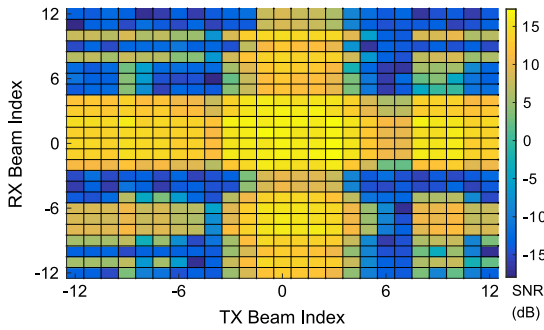**Fig. 3.** Maps of three indoor environments and measurement locations included in our study.

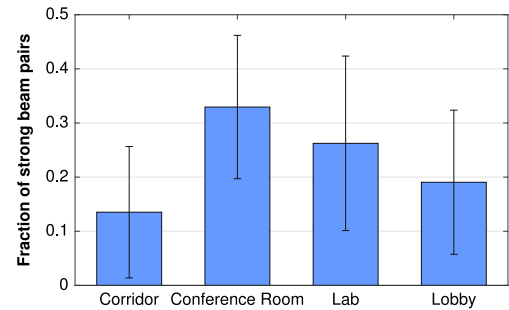**Fig. 4.** Beam-pair heatmap for Lab Position C.

Fig. 5. Richness of strong beam pairs in four environments.

5.1. Richness of strong beam pairs

mmWave channels are expected to be sparse due to the higher path loss and penetration loss [6]. Therefore, we would expect to see only a few physical paths, including LoS and NLoS, between the TX and RX. A particular beam pattern captures a number of these paths and applies different directivity gain. Hence, the provided signal strength by a particular (TX, RX) beam pair depends on the number of captured paths, their link budget, and beam directivity. The isolation of LoS/NLoS paths in order to measure the richness or sparsity of 60 GHz channels is not feasible with our platform; however, in this subsection, we explore the richness of *strong beam pairs*. In particular, we study how many beam pairs can achieve at least 1 Gbps throughput in different environments. From our data set, we found that 10 dB SNR is sufficient for achieving 1 Gbps throughput with at least one MCS. Hence, we define a strong beam pair as a beam pair that provides at least 10 dB SNR. We count the number of such strong beam pairs for each position in the corridor, conference room, lab, and lobby (see Fig. 3). Fig. 5 shows the average (over all measurement positions) ratio of the number of strong beam pairs over the total 625 possible beam combinations for the four environments. E.g., 0.2 in the *y*-axis means that on average 125 beam combinations (out of a total of 625 beam pairs) can provide at least 1 Gbps throughput.

First, Fig. 5 reveals that the fraction of strong beam pairs is highest for the conference room and lowest for the corridor. The materials in the conference room such as metallic cabinet, white board, and glass windows are known to be good reflectors for millimeter waves [6] providing several reflected paths, while there are no strong reflectors in the corridor. Further, the high error bars show that in any given environment, the number of strong beam pairs between two nodes highly depends on the RX position, its relative orientation with the TX, and the distance between them.

Second, the average ratio of strong beam pairs is above 0.13 for all four environments, i.e., more than 80 beam pairs provide at least 1 Gbps of throughput. This shows that, in contrast to the common

belief, there are several beam pairs that are able to provide Gbps throughput for 60 GHz communication. This result is caused by the shape of the imperfect beam patterns in use which have side-lobes as well as overlap (Fig. 2); thus, a physical LoS/NLoS path can be captured by multiple beams. The richness of strong beam pairs implies that beam training/adaptation algorithms might be able to avoid exhaustive time-consuming search through all beam combinations to find the best beam pair. Another implication is that interference between simultaneous transmissions may not be negligible in 60 GHz.

5.2. Relative strength of neighboring beams

In principle, an exhaustive search over all possible beam combinations is required to discover the highest signal strength beam pair. However, the associated training overhead may be prohibitively high, especially in case of mobile links. Hence, it is sometimes desirable to adopt lower overhead strategies which search over a subset of beam pairs. E.g., the 802.11ad standard specifies an initial coarse level search with quasi-omni beams at one end, followed by beam refinement for only a subset of beam pairs. Here, we evaluate the significance of selecting the highest strength beam pairs, and quantify the loss in signal strength if a sub-maximal pair is selected during the training process.

In particular, we are interested in studying the impact of distance and multiple paths (from reflections and side-lobes) on the relative strength of neighboring beams. The corridor data set encompasses both these scenarios, as illustrated by beam-pair heatmaps for 5 m, 15 m, 20 m and 25 m positions (Fig. 6). We make two key observations. *(i)* For the closest position (5 m), beam 0 and its two nearest neighbors at the TX side achieve > 10 dB SNR for all RX-side beams and vice versa. This is due to reflections off of side walls from the narrow corridor and the side lobes, resulting in a strong channel. *(ii)* As the TX–RX distance increases, the high SNR region shrinks and includes only the central beam pairs at 25 m. This is because the impact of reflections from side walls becomes less pronounced with distance. By geometry, the azimuth angle for first-order reflection paths (strongest NLoS components) from
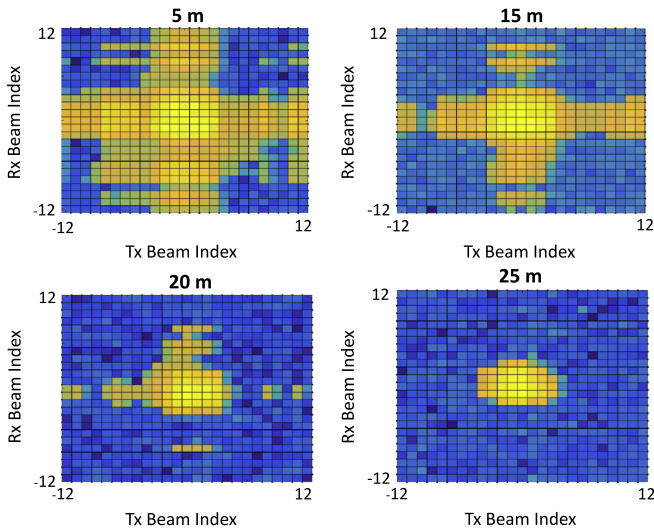
Fig. 6. Beam-pair heatmaps for all 25×25 beam pair combinations for 5 m, 15 m, 20 m and 25 m distances in the corridor.
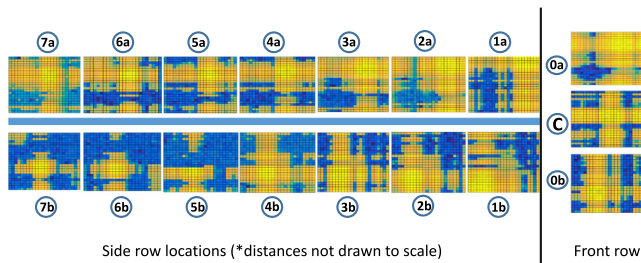

Fig. 7. Beam-pair heatmaps for all 25×25 beam pair combinations and all positions in the lab.
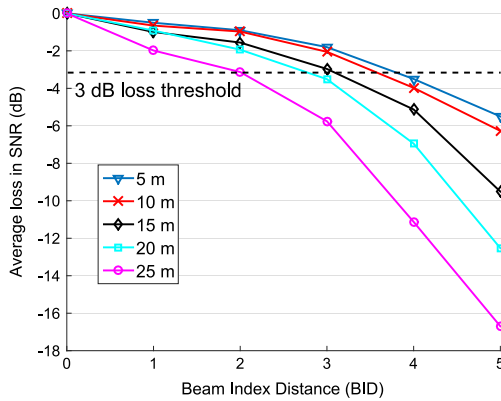

Fig. 8. Average loss in SNR vs. *BID* for corridor positions.

either wall reduces from 19° at 5 m to less than 4° at 25 m. As such, the angular separation between LoS and the strongest NLoS component decreases and only the central beams include these paths. Hence, more distant RX positions will require a larger search space to discover high strength beam pairs.

Fig. 7 shows the beam-pair heatmaps for all the lab positions. We observe the same general trend, i.e., again the high SNR region shrinks as the TX–RX distance increases. Interestingly, we also observe a large amount of asymmetry between positions at the two side row locations, e.g., 7a vs. 7b, due to different object and furniture placement (computers, boxes, a whiteboard, windows, etc.) in the two parts of the lab.

To further quantify the impact of selecting sub-maximal beam pairs, we analyze the loss in SNR as we move away from the highest strength beam pairs. We associate a distance metric with each beam pair *(i,j)*; Beam Index Distance *(BID)*. If (T,R) is the beam pair with highest SNR, we define *BID* as $(|T - i| + |R - j|), \forall\, i, j \in [-12, 12]$. E.g., *BID*=1 indicates a difference of one beam index, either in TX or in RX beam. Fig. 8 shows SNR loss in dB vs. *BID* for five different positions (at distances 5 m, 10 m, 15 m, 20 m, 25 m) in the corridor. Since multiple beam pairs can have the same *BID* for some highest signal strength pair, we plot average SNR loss over all such pairs. Since *BID*=0 indicates the maximal strength pair, SNR loss is 0 in this case for all distances.

The figure reveals that for all positions, SNR drops monotonically with increase in *BID* (i.e., for beam pairs farther and farther away from the maximal pair). However, at 5 m, beam pairs with $BID \leq 2$ are still within 1 dB of the maximal pair, indicating only a small loss in link strength for selecting these sub-maximal beam pairs. Moreover, for $BID \leq 4$, the loss in SNR is still within 3 dB (i.e., 50% of highest achievable SNR). As the TX–RX distance increases, SNR decreases more rapidly with *BID*, indicating a greater degradation in relative strength of neighboring beam pairs. Furthermore, fewer beam pairs on average are within 3 dB of the maximum possible signal strength for greater inter-node distances. Hence, the gain in signal strength is higher if an exhaustive search is performed for longer TX–RX distances, whereas for short distances selecting sub-maximal beam pairs can still yield high signal strength.

5.3. Performance of NLoS links

mmWave signals experience attenuation due to reflection and thus the link budget for NLoS components is expected to be lower compared to the LoS path [6]. Here, we measure and compare the achievable SNR and throughput with and without the presence of the LoS path. In particular, we want to explore the feasibility of Gbps scale throughput via reflections in the absence of LoS path. To this end, we consider the conference room since it has many reflectors such as whiteboard and TV screen (see Fig. 3b). The RX orientation in positions 4, 5, 6 on the table is such that there cannot be a LoS path between the TX and RX (the back of the phased array is blocked and there is no back lobe). On the other side, the RX sees the LoS path when located at positions 1, 2, 3 on the table.

Fig. 9 depicts the beam-pair heatmaps for positions 2 to 5. First, by comparing this figure with Figs. 4 and 6, we observe that more beam pairs provide positive SNR values due the better reflection in the conference room and shorter distance. The cluster of high SNR beam pairs for position 2 and 3 maps to the physical LoS path between two nodes confirming that LoS path was present for these positions. Similarly, the cluster of high SNR beam pairs for position 4 and 5 and the map of the conference room (Fig. 3b) suggest that these beam pairs include a reflected path from the whiteboard.

Next, we measure the highest achievable SNR and throughput through the best beam-pair for each receiver position facing the transmitter (1 to 3) or reverse facing the transmitter (4 to 6). We depict the average SNR and Throughput in Figs. 10a and 10b, respectively. Note that MCS 4 (1/2 16 QAM) was used for modulation since it provides the highest throughput in all positions. Fig. 10a reveals that one can achieve 17–18 dB SNR, even in the absence of a LoS path. Furthermore, the throughout is close to 1.9 Gbps with and without the LoS path. Hence, we conclude that the SNR and throughput values for reflected paths can be as high as for the LoS path in real 60 GHz systems.

5.4. Beam misalignment and nodal mobility

In Section 5.1, we discussed the richness of strong beam pairs across different environments, and saw that any of these beam pairs, if identified by the training procedure, can establish a Gbps 60 GHz link. However, the alignment of the selected beams may subsequently be lost due to nodal mobility, which may lead to a degradation in
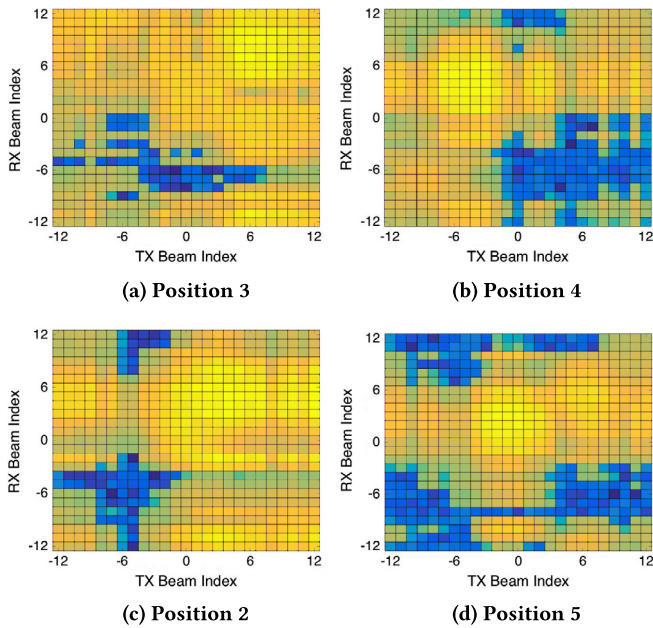
(a) Position 3

(b) Position 4

(c) Position 2

(d) Position 5

Fig. 9. Beam-pair heatmaps for four conference room positions.

(a)

(b)
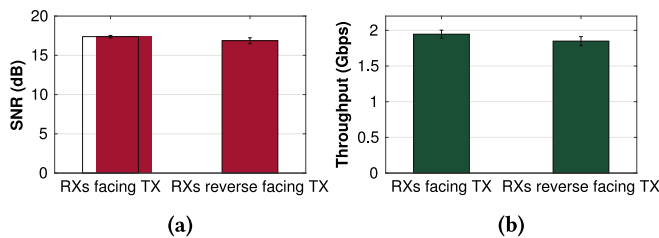
Fig. 10. (a) Average SNR, and (b) Average throughput received in positions 1 to 10 in the conference room.

Fig. 11. Beam adaptation strategies for lateral translation.

Fig. 12. Beam sweep maps for 15 locations in the lobby.

signal strength or may even break the link, depending on the extent of mobility [15,36]. To study the impact of mobility on misalignment of selected beams and the subsequent loss in signal strength, we perform controlled experiments in the lobby isolating two key types of mobility, translation and rotation.

Lateral Translation: First we consider the scenario where the orientation of both TX and RX remains fixed; however, a change in receiver position results in a change in the relative angle between the two nodes. For this, we consider positions 1 through 6 in the lobby for a fixed RX orientation (0°). These positions emulate a path taken by a node as it moves perpendicularly to the TX in steps of 1 m. Further, we compare five different adaptation strategies by the nodes to adjust their beams: (i) *Fixed Beam Pair*, where both nodes keep using the same beams throughout the experiment. For this, we consider beam pair (0,0) which is the strongest at initial position. (ii) *TX and RX* adaptation, where both nodes perform an exhaustive search at each position and re-select the strongest beam pair. (iii) *RX-Only* adaptation, where only the RX locally adapts its beam to maximize link strength, while the TX beam remains fixed (beam 0). (iv) *TX-Only* adaptation, where only the TX locally adapts its beam to maximize link strength, while the RX beam remains fixed (beam 0). (v) *Adjacent Beam RX Only*, where the TX beam remains fixed (beam 0) and the RX performs a limited adaptation and selects the best one out of three beams: the best beam in the last position and one beam immediately adjacent to it in each direction.

Fig. 11 plots SNR vs. TX–RX lateral distance for the three aforementioned strategies across lobby positions 1 through 6 and Fig. 12 plots the beam-pair heatmaps for all the lobby positions. For the case when
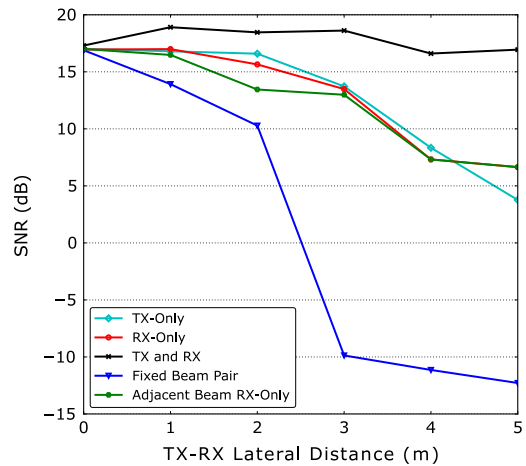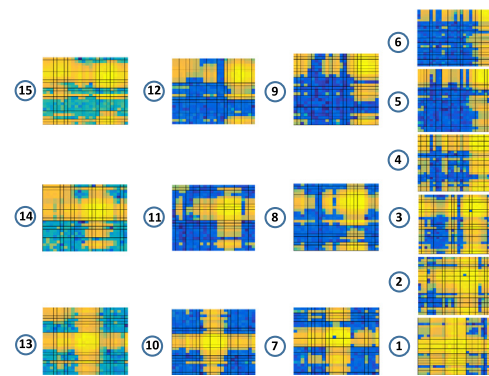
TX and RX beams remain fixed to (0,0), the link strength decreases monotonically from position 1 through 6, as the relative angle between the two nodes increases from 0° at Position 0 to 60° at Position 6. In fact, SNR drops below the receive threshold after Position 3 when the relative angle is 40°. Note that the relative angle at Position 3 is still outside the beamwidth of the main-lobe for beam 0 at the receiver. The high SNR at this position results from a side-lobe of beam 0. However, for positions 4 through 6, the link cannot be sustained for beam pair (0,0) due to higher angular separation between TX and RX, which illustrates the significance of adapting beams for 60 GHz links in response to mobility.

In case of beam adaptation, SNR remains nearly constant across all positions for strategy-*(ii)*, when an exhaustive search is repeated at each position. This is the ideal scenario for beam adaptation, and depicts highest achievable SNR for the mobile receiver. For RX-only adaptation (strategy-*(iii)*), the search space only spans RX side beams and hence this strategy incurs much lower overhead than exhaustive search in the first case. However, in this case, only the RX-side beam has maximum alignment with the TX, while the TX beam remains fixed at 0. This results in SNR degradation as the TX–RX angular separation increases from Position 1 to Position 6. Despite this loss in TX-side alignment, Fig. 11 reveals that SNR achieved with RX-only adaptation is significantly better than that with no-adaptation strategy, and a link is sustained across all positions. Interestingly, strategy-*(v)*, which performs a limited version of the RX-Only adaptation, yields similar performance indicating that simply probing adjacent beams should be the first level of beam adaptation strategy under lateral translation. Lastly, we observe that strategy-*(iv)*, the TX counterpart of the RX-Only strategy can provide similar SNR as strategies *(iii)* and *(v)*. In general, our results strongly suggest that a one-sided local search either at the

(a) Rotation with fixed beam pair (0,0) (b) Beam adaptation for Position 10 (c) Beam adaptation for Position 12
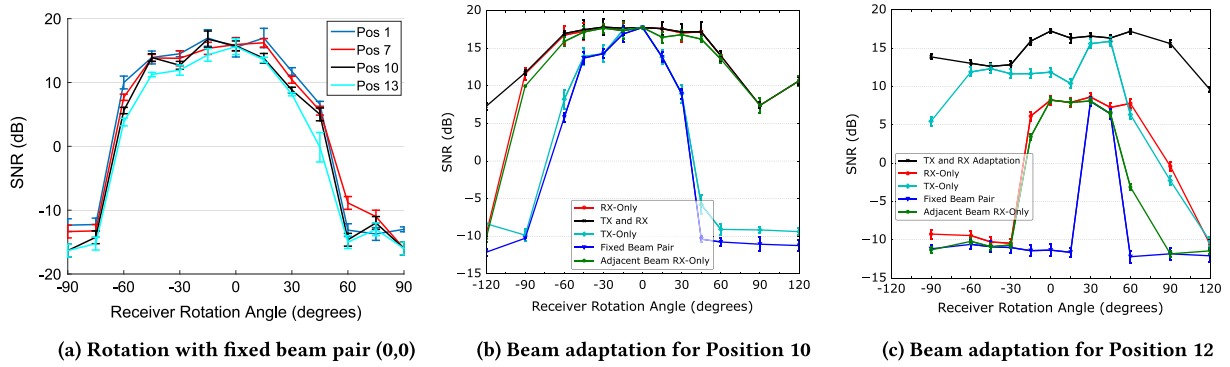
Fig. 13. Impact of receiver rotation on SNR for various lobby positions.

(a) Corridor

(b) Lab

(c) Conference room

(d) Lobby
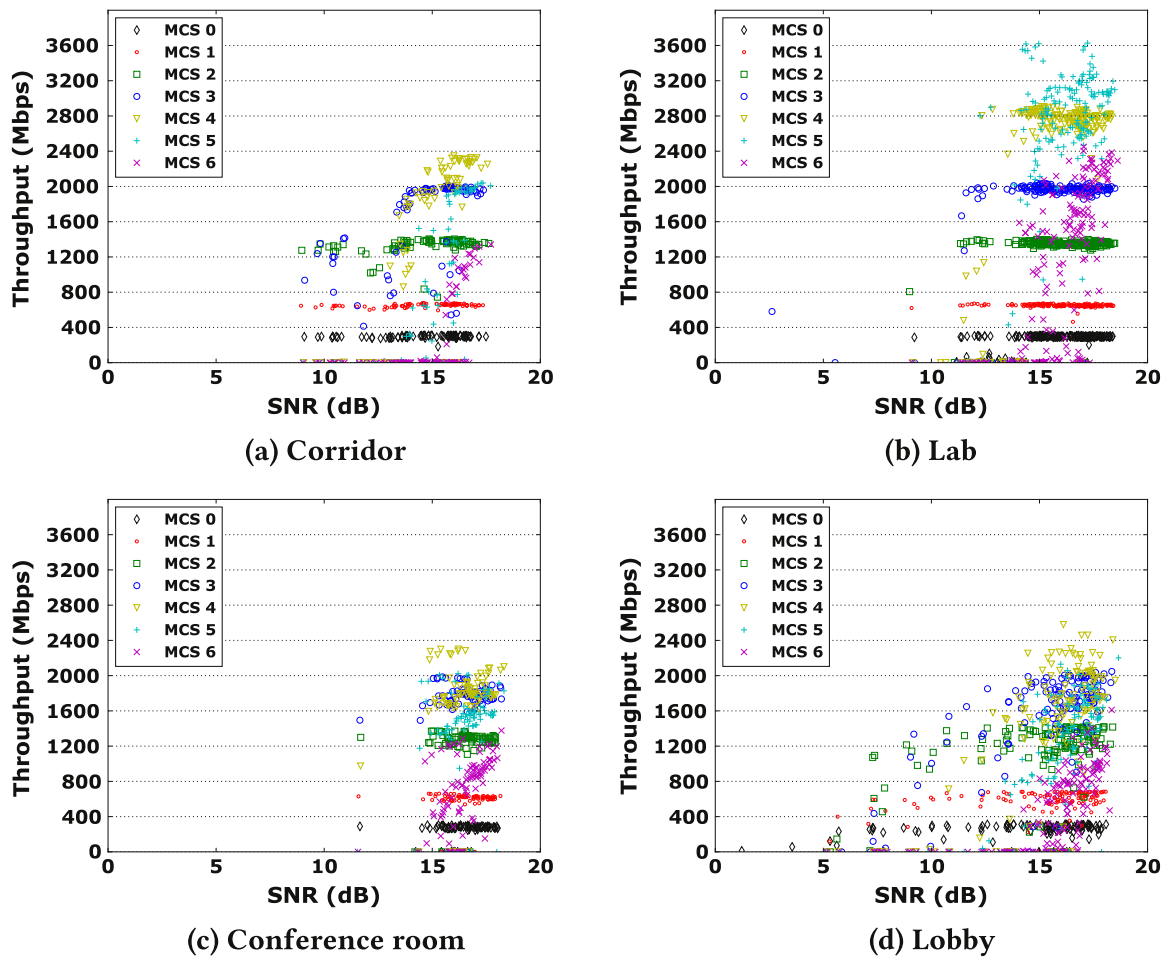
Fig. 14. Measured throughput for different SNR values.

(a) Corridor (b) Lab (c) Conference Room (d) Lobby

Fig. 15. Optimal MCS for different SNR values in each environment.

TX or the RX, although sub-optimal, may be sufficient to maintain a directional 60 GHz link while avoiding exhaustive training.

Rotation: To analyze the impact of receiver rotation, we consider the change in signal strength of the central beam pair (0,0) for positions 1, 7, 10 a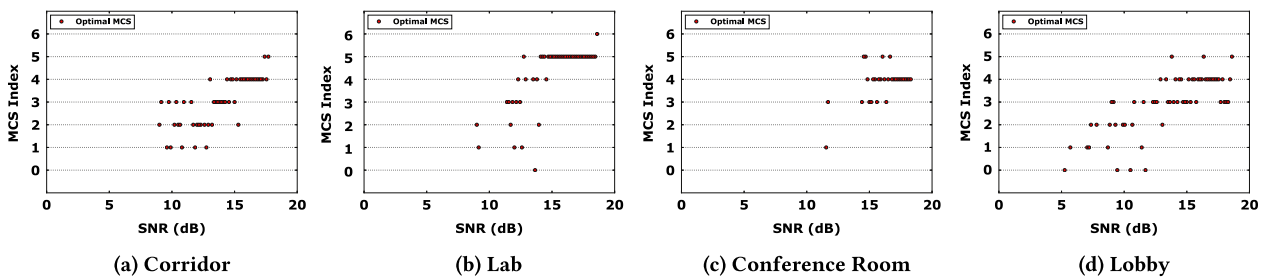nd 13 which are in front of the TX (Fig. 3c). Fig. 13a shows SNR vs. receiver angle such that for all positions, 0° corresponds to a perfect alignment between TX and RX, and hence beam pair (0,0) achieves maximum SNR. As the RX rotates on either side, the SNR decreases sharply due to misalignment of RX beam 0. We also observe that the SNR for counter-clockwise rotation remains steady over a larger range of angles before dropping below 0 dB. This is due to asymmetrical radiation pattern of beam 0, as discussed in Section 3.3, which results in higher gain for counter-clockwise rotation. Moreover, as the inter-node distance increases, the SNR degrades for all orientations. However, the impact of rotation is more pronounced than that of increasing distance. This shows that a 60 GHz link, corresponding to a fixed beam pair (selected during training process), is highly susceptible to misalignment due to nodal rotation, and slight rotation can result in multiple dBs of SNR loss.

To improve link budget, the beams at either the TX or the RX need to be adjusted. Here we compare the same five beam adaptation strategies as we did for the lateral translation case. For this, we consider two different receiver positions; Position 10 which is directly in front of TX (6.3 m apart) and Position 12 which is roughly 30° to the right of TX (7.5 m apart). The latter position captures the impact of both angular separation and receiver rotation. Again, beam 0 is used on the TX side in the case of *RX-Only* and *Adjacent Beam RX Only* adaptation and on the RX side in the case of *TX-Only* adaptation.

Fig. 13b shows that for Position 10, *RX-Only* adaptation achieves similar SNR as TX-and-RX adaptation. This is because the RX is directly in front of TX and is fixed, hence beam 0 is the best TX beam for all orientations of the receiver. However, unlike lateral translation, the maximum achievable SNR diminishes for higher RX angles on either side. This is a consequence of non-uniform angular spread of beam patterns and diminishing directivity gain of beam indices farther from the central beam, a limitation of practical phased array antennas. Further, the increase in SNR for 120° RX angle depicts the impact of side-lobes, since for this orientation the main lobes of all receiver beams are misaligned with the TX. This is an example scenario of side-lobes generating additional paths to provide resilience to receiver mobility, an effect that cannot be observed in the case of horn antenna based systems. The *Adjacent Beam RX Only* adaptation, even with limited beam options, results in SNR values nearly identical to the full RX adaption indicating that under small rotation, typically the adjacent beam index in the direction of the rotation is indeed the SNR-optimal one. This insight can be greatly helpful in the design of efficient beam adaptation mechanisms. The *TX-Only* strategy, on the other hand, fails to provide any significant SNR gains compared to the no-adaptation case (*Fixed Beam Pair*). This is the result of the particular arrangement of TX and RX in this scenario. When RX is performing the adaptation, it can always select a beam index whose main lobe is such that it points directly towards the main lobe of the TX beam, making them essentially aligned. The TX, on the other hand, can at best imperfectly align its main lobe with RX's main lobe (which is pointing at the angle by which the RX was rotated) at a non-zero angle or point towards RX's side lobes to enable communication. In both cases, the SNR is bound to be lower than the one that results out of the direct (zero angular separation) alignment of the main lobes.

For Position 12 (Fig. 13c), the angular separation between TX and RX further degrades signal strength compared to Position 10. For cases with fixed beam (0,0) and *RX-Only* adaptation, the highest SNR is achieved for receiver angle 30°, since for this orientation the RX has maximum alignment with the TX. However, the SNR is low across all angles even with RX-only adaptation, since the TX is still using beam 0, which is misaligned with the receiver due to an angular separation of ∼30°. Consequently, the *Adjacent Beam RX Only* adaptation scheme which only
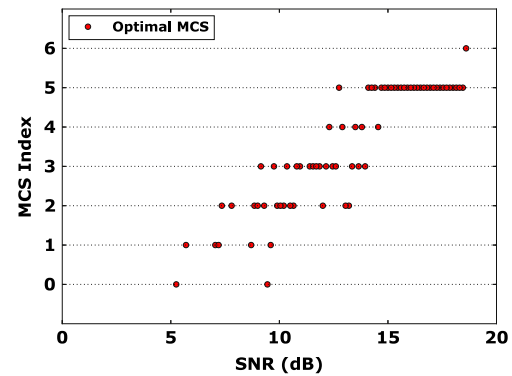

Fig. 16. Optimal MCS vs. SNR for all links in all four environments.

probes a subset of the beam pairs also results in performance equal or worse than its unrestricted *RX-Only* version. When beam adaptation is used at both TX and RX, the TX-side beam also becomes aligned with the RX, adding another 5–7 dB of SNR gain and making the highest signal strength similar to that observed for Position 10. However, in contrast to Position 10, *TX-Only* adaptation is able to outperform RX-side adaptation at almost all RX rotation angles. This difference in behavior (from Position 10) can be accounted for by looking at the ability of the TX and RX to align their beam patterns. When RX is rotated clockwise (positive angles), the TX can align its beam with that of the RX, although imperfectly, by choosing the beam index that counteracts the RX's rotation angle (until ∼45° rotation). However, as the amount of rotation increases further in the clockwise direction, the TX and RX beam alignment gets even worse and results in a sharp drop in SNR. Interestingly, in the counter-clockwise direction (negative angles), *TX-Only* adaptation is able to maintain high SNR even at extreme angles. We believe this happens due to the presence of a metallic wall underneath the north window, which acts a strong reflector and creates a NLOS path between the TX and RX. Note that in the clockwise direction, the environment does not have any strong reflectors and thus only a LOS path is feasible.

These two examples illustrate the importance of identifying different mobility scenarios for protocol design, since adaptation strategies are highly dependent on the type of mobility and the environment. For the first case, RX-side adaptation was sufficient, whereas in the latter case of rotation with angular separation, TX-side adaptation is necessary to realize the highest signal strength link.

5.5. MCS–SNR relationship

The all beam-pairs-SNR data collected during the beam sweeps is an effective way to study the physical characteristics of 60 GHz links under different environments. However, SNR alone does not provide a complete picture of the performance that higher layers and user applications running over mmWave links will ultimately experience. The second step in our experimental methodology (MCS Sweeps) is geared towards providing insights into this area. In this work, we use the *throughput* as the performance metric. For the selected beam pairs (chosen after the exhaustive beam sweep), throughput is measured for all possible MCS indexes to understand how different modulation schemes perform for a link with a given SNR.

Fig. 14 plots the measured throughput under different link SNR values with all MCS indexes available in X60 in each of the four environments. Note that for a given environment, the plot aggregates data from all the links involved in the measurement study for that location. Further, notice that due to the specific characteristics of the measurement environments, each one exhibits different range of SNR values. The bulk of the links experience SNR between 10 and 15 dB.
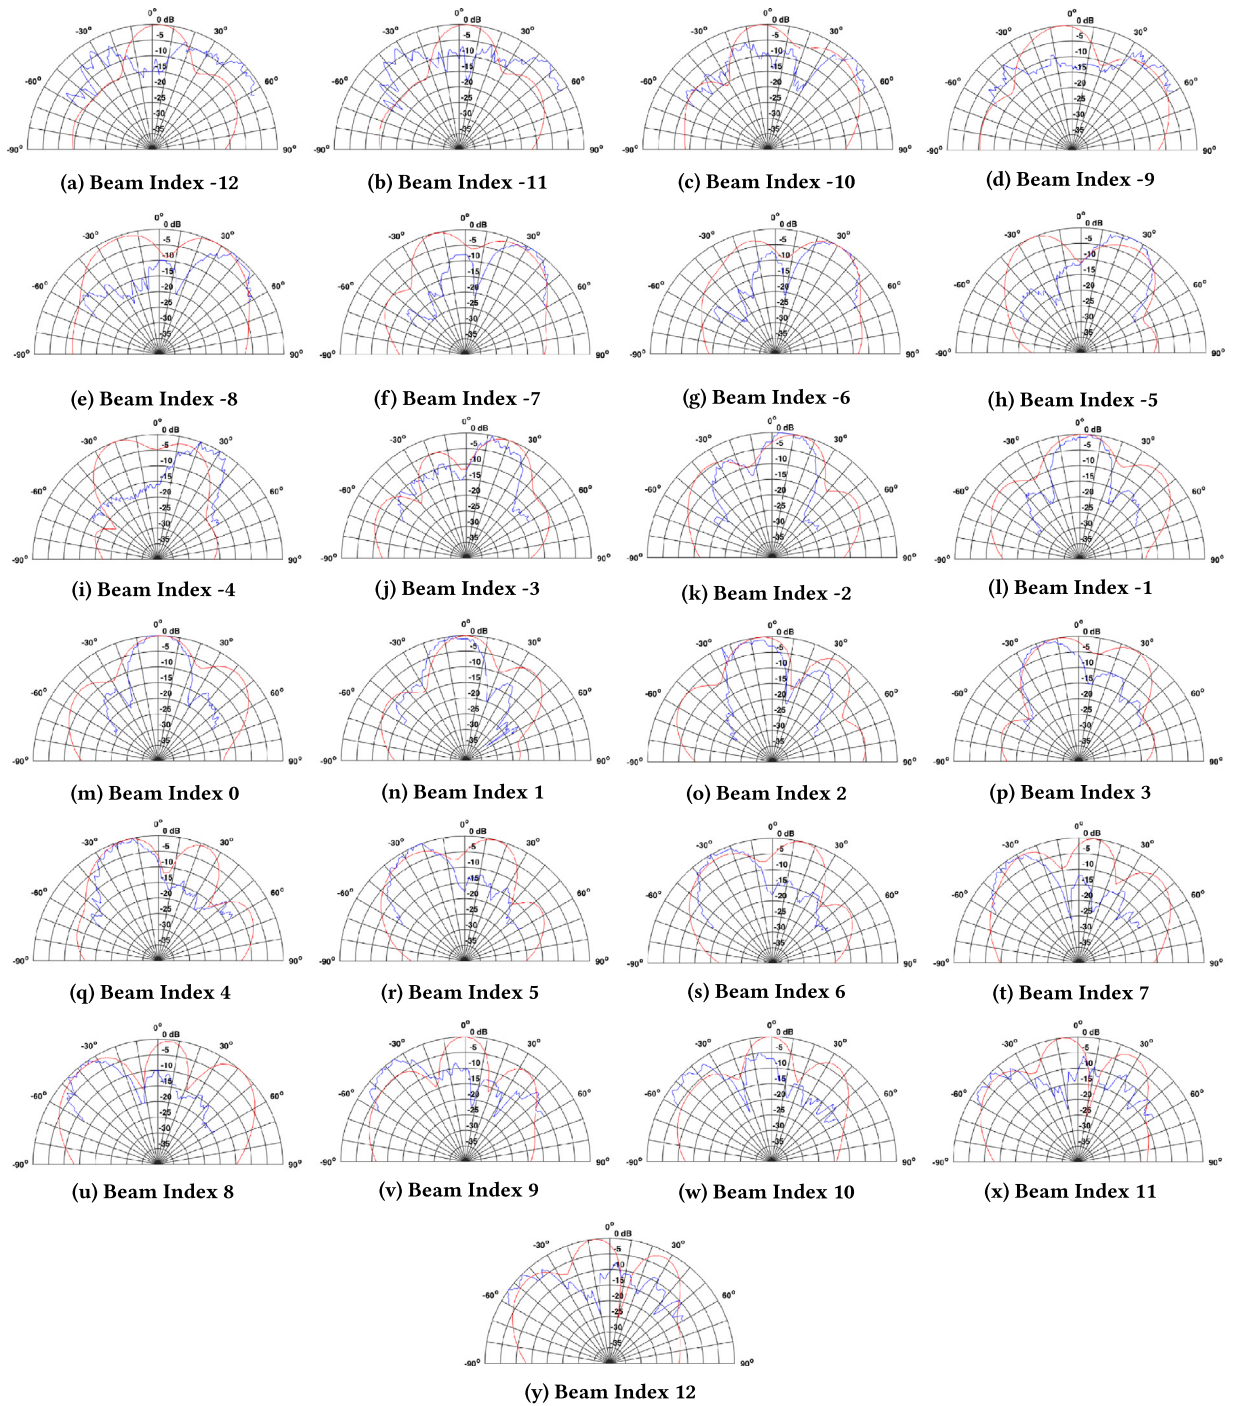
(a) Beam Index -12 (b) Beam Index -11 (c) Beam Index -10 (d) Beam Index -9

(e) Beam Index -8 (f) Beam Index -7 (g) Beam Index -6 (h) Beam Index -5

(i) Beam Index -4 (j) Beam Index -3 (k) Beam Index -2 (l) Beam Index -1

(m) Beam Index 0 (n) Beam Index 1 (o) Beam Index 2 (p) Beam Index 3

(q) Beam Index 4 (r) Beam Index 5 (s) Beam Index 6 (t) Beam Index 7

(u) Beam Index 8 (v) Beam Index 9 (w) Beam Index 10 (x) Beam Index 11

(y) Beam Index 12

Fig. 17. Antenna array beam patterns.

We observe that both MCS 0 and MCS 1 are not only supported at all four environments but mostly yield throughputs (~300 Mbps and ~600 Mbps, respectively) close to their corresponding PHY data rates, for the entire SNR range present in our data set. Similarly, MCS 2 provides ~1.2 Gbps or higher throughput with SNR > 9 dB in three out of four environments (throughput in the lobby exhibits larger variations). Further, as expected, with increasing channel quality (higher SNR), we observe a general improvement in performance of MCS 3–6, until the throughput is close to their corresponding physical data rates in all environments. However, the trend is not strictly increasing; in all four environments, we observe significant dips in throughput with increasing SNR. More importantly, the throughput for a given (MCS, SNR) pair can vary significantly across locations, indicating that sampled link SNR

alone may *not* be sufficient to predict performance for different MCS. For instance, an SNR of 12–13 dB yields a throughput lower than 1200 Mbps in the corridor but as high as 2000 Mbps in the lab. Alternatively, for a given SNR, multiple MCS indexes may provide similar throughputs only differing by a few Mbps. Previous studies [24] with commercial devices have reached similar conclusions.

One of the direct implications of these observations lies on the design of rate adaptation algorithms. A common practice, often employed in the evaluation of the mmWave link adaptation proposals, is the selection of MCS based on the measured SNR from a lookup table derived offline [10,20,37–39]. However, analysis of our data set suggests that such a approach may not always lead to the best possible performance. In particular, we ask the question: *Can an SNR-lookup based rate-selection*

scheme always select the optimal (the one that maximizes throughput) MCS? To answer this question, we look at the optimal MCS across all the links studied in an environment for different SNR values. Fig. 15 plots the Optimal MCS vs. SNR for each of the four environments. Note that in the plots we bin SNR values in increments of 0.15 dB to account for the arbitrary resolution of dB values generated by the system.

Fig. 15 shows that, for low and medium SNR values (<15 dB), a unique optimal MCS cannot be selected in any of the four environments. In other words, links with same or similar SNR, do *not* necessarily lead to the same throughput with a given MCS and hence it is not easy to find a one-to-one mapping between MCS and SNR. On the other hand, for higher SNR values (>15 dB in the corridor and lab, >17 dB in the conference room), it would be possible to select an optimal MCS with high confidence, at least in the corridor, lab, and conference room environments. The lobby environment displays ambiguity (between MCS 3 and 4) even for high SNR values. Further, even under high SNR scenarios, the optimal MCS for a given SNR may not be the same across different environments. For instance, for SNR 16–18 dB, MCS 4 provides best throughput in the corridor but in the lab MCS 5 should be chosen for optimal performance.

Although it may not be possible to directly select the optimal MCS based on SNR, a rate-search scheme that *probes* a fixed limited set of MCS indexes could be designed. Such a scheme however works only if we can limit the search to a small number of MCS indexes for a given SNR, globally across any location. To explore the feasibility of such a scheme, we aggregate data from all links measured in out study across all locations and again look at the SNR to optimal MCS mapping. Fig. 16 plots the global SNR vs. optimal plot with the same SNR bin size as used for per-environment analysis. We observe that under low SNR (5–8 dB), either a unique optimal MCS can be directly selected or at most two MCS indexes need to be probed. Middle SNR range (8–15 dB) exhibits much more variability with certain SNR values requiring probing of up to 4 (out of 7) MCS indices. However, for high SNR links (>15 dB) MCS 5 can be selected as the clear winner. In general, we observe that a direct-mapping or limited probing based rate-adaptation strategy could be employed for low or high SNRs, but such a scheme would end up probing a large set of the MCS indices before converging on the throughput-optimal data rate for medium SNRs. A similar conclusion was made for legacy (802.11b/g) links in [40].

6. Conclusion

In this paper, we introduced X60, the first reprogrammable testbed for 60 GHz supporting wideband transmissions and featuring a user-configurable 12-element phased antenna array. We demonstrated how phased array generated beams can have imperfect radiation patterns, featuring strong side-lobes and non-uniform steerability along different directions; these features result in artifacts not observed by studies conducted on horn antenna based platforms. Our measurement study using X60 allowed us to report new findings with respect to 60 GHz link behavior in different indoor environments. X60's unique capabilities make it an advanced platform for experimentation and prototyping, across layers, towards solving the most important challenges in mmWave research.

Acknowledgments

---

[1] Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of AFRL. Approved for Public Release; Distribution Unlimited: 88ABW-2018-1803.

Appendix. Measured and simulated radiation patterns

Fig. 17 plots all the 2D measured and simulated radiation patterns.

References

[1] IEEE 802.11 Task Group AD. 2012. IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band, 2012.
[2] Y. Ghasempour, C.R.C.M. da Silva, C. Cordeiro, E.W. Knightly, IEEE 802.11ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi, IEEE Commun. Mag. 55 (12) (2017) 186–192, http://dx.doi.org/10.1109/MCOM.2017.1700393.
[3] X60: A Cross-Layer Reconfigurable Multi-Gigabit WLAN Testbed at 60 GHz. https://ubwins.cse.buffalo.edu/projects/x60.html.
[4] National Instruments. 2017. Introduction to the NI mmWave Transceiver System Hardware - National Instruments. http://www.ni.com/white-paper/53095/en/. 2017 (Accessed on 06/25/17).
[5] Peter F.M. Smulders, L.M. Correia, Characterisation of propagation in 60 GHz radio channels, Electron. Commun. Eng. J. 9 (2) (1997) 73–80.
[6] Hao Xu, Vikas Kukshya, Theodore S. Rappaport, Spatial and temporal characteristics of 60-GHz indoor channels, IEEE J. Sel. Areas Commun. (JSAC) 20 (3) (2002) 620–630.
[7] Christopher R. Anderson, Theodore S. Rappaport, In-building wideband partition loss measurements at 2.5 and 60 GHz, IEEE Trans. Wirel. Commun. 3 (3) (2004) 922–928.
[8] Peter F.M. Smulders, Statistical characterization of 60-GHz indoor radio channels, IEEE Trans. Antennas and Propagation 57 (10) (2009) 2820–2829.
[9] Alexander Maltsev, Roman Maslennikov, Alexey Sevastyanov, Alexey Khoryaev, Artyom Lomayev, Experimental investigations of 60 GHz WLAN systems in office environment, IEEE J. Sel. Areas Commun. (JSAC) 27 (8) (2009) 1488–1499.
[10] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, Parameswaran Ramanathan, 60 GHz indoor networking through flexible beams: A link-level profiling, in: Proc. of ACM SIGMETRICS, 2015.
[11] Sanjib Sur, Xinyu Zhang, Parameswaran Ramanathan, Ranveer Chandra, BeamSpy: Enabling robust 60 GHz links under blockage, in: Proc. of USENIX NSDI, 2016.
[12] Anfu Zhou, Xinyu Zhang, Huadong Ma, Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering, in: Proc. of IEEE INFOCOM, 2017.
[13] Teng Wei, Anfu Zhou, Xinyu Zhang, Facilitating robust 60 GHz network deployment by sensing ambient reflectors, in: Proc. of USENIX NSDI, 2017.
[14] Yasaman Ghasempour, Edward W. Knightly, Decoupling beam steering and user selection for scaling multi-user 60 GHz WLANs, in: Proceedings of ACM MobiHoc, 2017.
[15] Muhammad Kumail Haider, Edward W. Knightly, Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs: Protocol design and system implementation, in: Proc. of ACM MobiHoc, 2016.
[16] Sharan Naribole, Edward Knightly, Scalable multicast in highly-directional 60 GHz WLANs, in: Proc. of IEEE SECON, 2016.
[17] Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, Joerg Widmer, Steering with eyes closed: mm-wave beam steering without in-band measurement, in: Proc. of IEEE INFOCOM, 2015.
[18] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, Parameswaran Ramanathan, OpenMili: A 60 GHz software radio platform with a reconfigurable phased-array antenna, in: Proceedings of the ACM MobiCom, 2016.
[19] Omid Abari, Haitham Hassanieh, Michael Rodriguez, Dina Katabi, Millimeter wave communications: from point-to-point links to agile network connections, in: Proc. of ACM HotNets, 2016.
[20] Xiaozheng Tie, Kishore Ramachandran, Rajesh Mahindra, On 60 GHz wireless link performance in indoor environments, in: Proc. of PAM, 2012.
[21] Yibo Zhu, Zengbin Zhang, Zhinus Marzi, Chris Nelson, Upamanyu Madhow, Ben Y. Zhao, Haitao Zheng, Demystifying 60GHz outdoor picocells, in: Proc. of ACM MobiCom, 2014.
[22] Thomas Nitsche, Guillermo Bielsa, Irene Tejado, Adrian Loch, Joerg Widmer, Boon and bane of 60 GHz networks: Practical insights into beamforming, interference, and frame level operation, in: Proc. of ACM CoNEXT, 2015.
[23] Swetank Kumar Saha, Anuj Garg, Dimitrios Koutsonikolas, A first look at TCP performance in indoor IEEE 802.11 ad WLANs, in: 2015 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), IEEE, 2015, pp. 63–64.
[24] Swetank Kumar Saha, Viral Vijay Vira, Anuj Garg, Dimitrios Koutsonikolas, A feasibility study of 60 GHz indoor WLANs, in: Proc. of IEEE ICCCN, 2016.
[25] Swetank Kumar Saha, Viral Vijay Vira, Anuj Garg, Dimitrios Koutsonikolas, Multi-Gigabit indoor WLANs: Looking beyond 2.4/5 GHz, in: Proc. of IEEE ICC, 2016.
[26] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, Kyu-Han Kim, WiFi-assisted 60 GHz networks, in: Proc. of ACM MobiCom, 2017.
[27] Teng Wei, Xinyu Zhang, Pose information assisted 60 GHz networks: Towards seamless coverage and mobility support, in: Proc. of ACM MobiCom, 2017.
[28] Swetank Kumar Saha, Tariq Siddiqui, Dimitrios Koutsonikolas, Adrian Loch, Joerg Widmer, Ramalingam Sridhar, A detailed look into power consumption of commodity 60 GHz devices, in: Proc. of IEEE WoWMoM, 2017.

[29] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, Matthias Hollick, Compressive millimeter-wave sector selection in off-the-shelf IEEE 802.11 ad devices, in: Proc. of the 11th ACM CoNEXT, 2017.

[30] Swetank Kumar Saha, Hany Assasa, Adrian Loch, Naveen Muralidhar Prakash, Roshan Shyamsunder Anantharamakrishna, Shivang Aggarwal, Daniel Steinmetzer, Dimitrios Koutsonikolas, Joerg Widmer, Matthias Hollick, Fast and infuriating: Performance and pitfalls of 60 GHz WLANs based on consumer-grade hardware, in: Proc. of IEEE SECON, 2018.

[31] Hany Assasa, Swetank Kumar Saha, Adrian Loch, Dimitrios Koutsonikolas, Joerg Widmer, Medium access and transport protocol aspects in practical 802.11ad networks, in: Proc. of IEEE WoWMoM, 2018.

[32] COMSOL Multiphysics. https://www.comsol.com.

[33] A. Alkhateeb, G. Leus, R.W. Heath, Limited feedback hybrid precoding for multi-user millimeter wave systems, IEEE Trans. Wirel. Commun. (2015).

[34] wil6210: Linux Wireless. https://github.com/torvalds/linux/blob/master/drivers/net/wireless/ath/wil6210/wmi.h.

[35] George R Maccartney, Theodore S Rappaport, Shu Sun, Sijia Deng, Indoor office wideband millimeter-wave propagation measurements and channel models at 28 and 73 GHz for ultra-dense 5G wireless networks, IEEE Access 3 (2015) 2388–2424.

[36] Muhammad Kumail Haider, Edward W. Knightly, iTrack: Tracking indicator LEDs on APs to bootstrap mmwave beam acquisition and steering, in: Proc. of ACM HotMobile, 2018.

[37] Daniel Halperin, Wenjun Hu, Anmol Sheth, David Wetherall, Predictable 802.11 packet delivery from wireless channel measurements, in: Proc. of ACM SIGCOMM, 2010.

[38] Xia Zhou, Zengbin Zhang, Yibo Zhu, Yubo Li, Saipriya Kumar, Amin Vahdat, Ben Y. Zhao, Haitao Zheng, Mirror mirror on the ceiling: Flexible wireless links for data centers, in: Proc. of ACM SIGCOMM, 2012.

[39] Yibo Zhu, Xia Zhou, Zengbin Zhang, Lin Zhou, Amin Vahdat, Ben Y. Zhao, Haitao Zheng, Cutting the cord: a robust wireless facilities network for data centers, in: Proc. of ACM MobiCom, 2014.

[40] Katrina LaCurts, Hari Balakrishnan, Measurement and analysis of real-world 802.11 mesh networks, in: Proc. of ACM SIGCOMM/USENIX IMC, 2010.