

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

Swetank Kumar Saha^{1*}, Yasaman Ghasempour^{2*}, Muhammad Kumail Haider^{2*}, Tariq Siddiqui¹, Paulo De Melo¹, Neerad Somanchi¹, Luke Zakrajsek¹, Arjun Singh¹, Owen Torres¹, Daniel Uvaydov¹, Josep Miquel Jornet¹, Edward Knightly², Dimitrios Koutsonikolas¹, Dimitris Pados¹, Zhi Sun¹

¹University at Buffalo, SUNY, Buffalo NY, USA, ²Rice University, Houston, TX, USA

{swetankk,tariqsid,paulodem,neeradso,lukezakr,asingh29,owentorr,danieluv,jmjornet,dimitrio,pados,zhisun}@buffalo.edu, {ghasempour,kumail.haider,knightly}@rice.edu

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.

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 [5]. 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-gigabit data rates. Based on the National Instrument's (NI) millimeter-wave (mmWave) Transceiver System [10] 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 extra wide channels supported by the 802.11ad standard. Lastly, using SiBeam's phased array, X60 generates 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.

*Primary co-author

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.

WiNTECH'17, October 20, 2017, Snowbird, UT, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5147-8/17/10...\$15.00

<https://doi.org/10.1145/3131473.3131479>

Our major findings are as follows: (i) Unlike the common belief of only a few TX/RX beam-pairs achieving 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 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 rates. (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 (client) may incorrectly hold on to an older beam too long without necessarily incurring link breakage.

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., [7, 12, 19, 20, 25]).

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 [21, 22, 24, 28], or/and horn antennas instead of phased arrays [8, 9, 13, 15]. 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 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 [27] and the testbed in [4]. 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 [4] 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 or WiGig hardware available commercially [14, 16–18, 23, 29]. These devices offer the chance to understand the often complex

inter-play between higher layers of the network stack and WirelessHD/WiGig 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.

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. 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

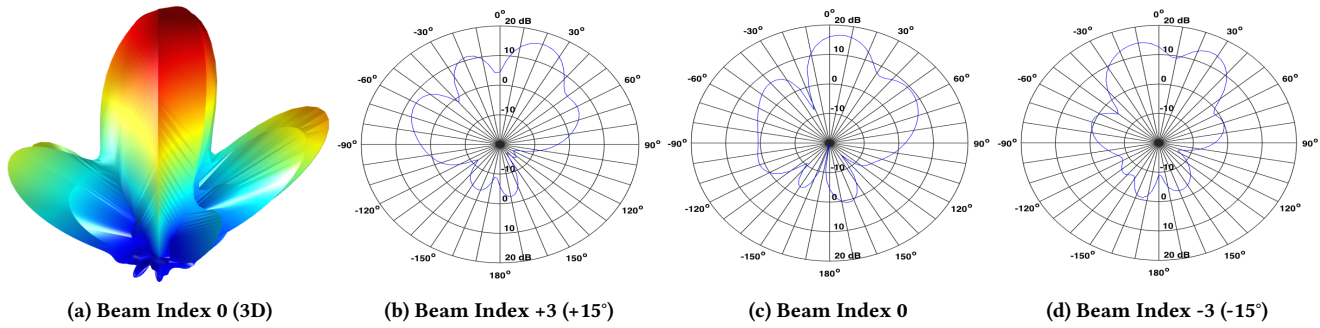


Figure 1: Antenna array beam patterns.

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: $0, \pi/2, \pi, 3\pi/2$.

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

We first computed the idealized beam patterns using COMSOL Multi-physics [1]. 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. Figures 1a-1d depict examples of 2D and 3D radiation patterns for select beam indices. These patterns highlight how, in contrast to beams generated by horn antennas, phased-array generated beams often have strong side-lobes. Moreover, as beams are steered away from the main lobe, patterns become more imperfect with even stronger side lobes and considerably weaker main lobe. For instance, comparing beam index 3 (Fig. 1b) and 12 (Fig. 1c) shows how practical phased-arrays can have *non-uniform steerability* along different directions as opposed to mechanically rotated horn antenna beams. Surprisingly, beam indices equally apart from the broadside beam (e.g., +3 (Fig. 1b) and -3 (Fig. 1d) can have radiation patterns that are *not* necessarily mirror images of each other.

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 that the elements are arranged in the 2D array [6]. Nitsche et al. [14] also found the beam patterns of commercial WiGig devices to be imperfect with strong side lobes. Also, an inspection of the open source wil6210 driver [2], 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).

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 (25x25) 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 40ms). This is

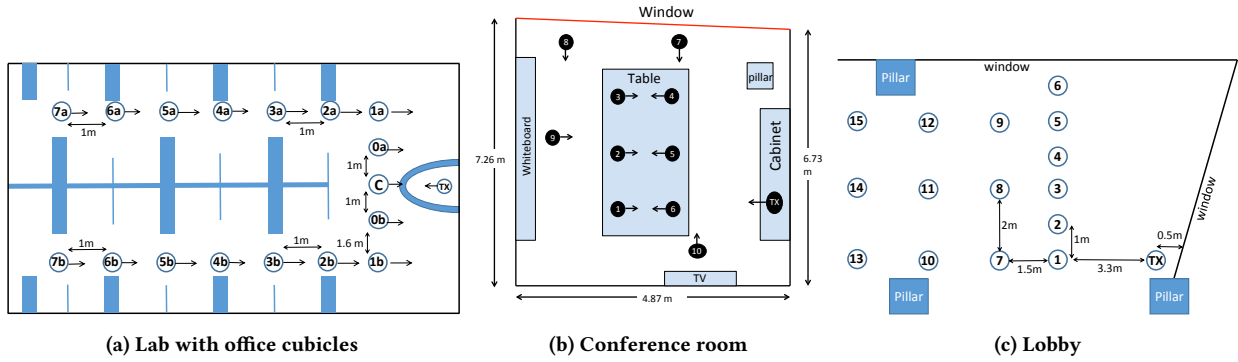


Figure 2: Maps of four indoor environments and measurement locations included in our study.

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 25x25 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. 2. 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 away from the transmitter.

Lab: The second environment is an 11.8x9.2x3.4 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. 2a, row A (locations 1a through 7a) and row B (locations 1b through 7b) are nearly symmetrical 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 7x4.87x3.4 m³ room (Fig. 2b) 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 [11, 26]. Hence this environment is suitable to study the impact of reflectors, non-LoS




Figure 3: Beam-pair heatmap for Lab Position C

paths and side-lobes. The TX is placed in a corner at 2.23 m height, and we consider 10 different RX locations across the room.

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. 2c).

5 RESULTS

5.1 Beam Sweeps

In our measurements, a beam sweep captures the SNR achieved for all 25 x 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. 3 shows the beam-pair heatmap for the center position (C) in the lab with distance of 2.3 m from TX, with yellow colored regions indicating beam pairs with the SNR of 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 (Sec. 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. 3, for TX beam indices between -2 to 4, 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 5 are used, SNR is above 5 dB for most TX beams.

5.2 Richness of Strong Beam Pairs

mmWave channels are expected to be sparse due to the higher path loss and penetration loss [25]. 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 data rate in different environments. From our data set, we found that 10 dB SNR is sufficient for achieving 1 Gbps data rate. 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. 2). Fig. 4 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 data rate.

First, Fig. 4 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 [25] 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 data rates 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. 1); 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

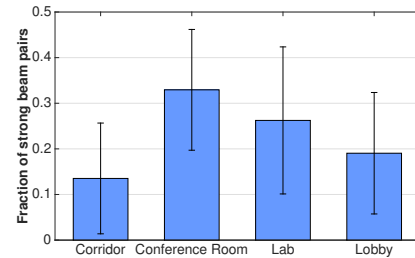


Figure 4: Richness of strong beam pairs in four environments.

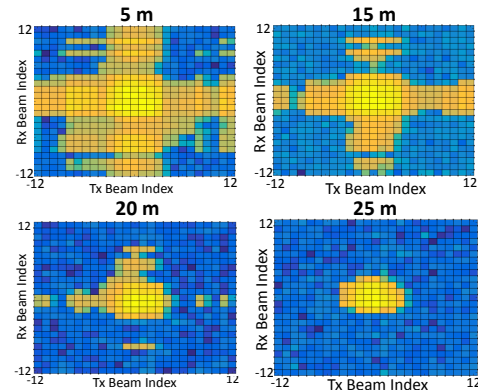


Figure 5: Beam-pair heatmaps for all 25×25 beam pair combinations for 5m, 15m, 20m and 25m distances in the corridor.

that interference between simultaneous transmissions may not be negligible in 60 GHz.

5.3 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 5m, 15m, 20m and 25m positions (Fig. 5). We make two key observations. (i) For the closest position (5m), 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 25m. 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

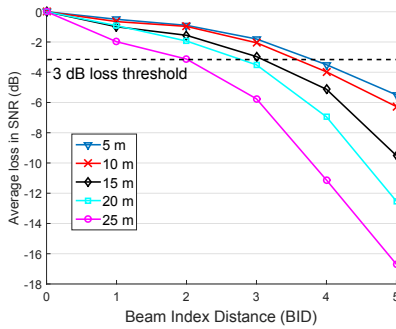


Figure 6: Average loss in SNR vs. BID for corridor positions.

either wall reduces from 19° at 5m to less than 4° at 25m. 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.

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. 6 shows SNR loss in dB vs. BID for five different positions (at distances 5m, 10m, 15m, 20m, 25m) 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 5m, beam pairs with $BID \leq 2$ are still within 1dB 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 3dB (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 3dB 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.

5.4 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 [26]. 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. 2b). 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. 7 depicts the beam-pair heatmaps for positions 2 to 5. First, by comparing this figure with Fig. 3 and Fig. 5, 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 2b) suggest that these beam pairs include a reflected path from the whiteboard.

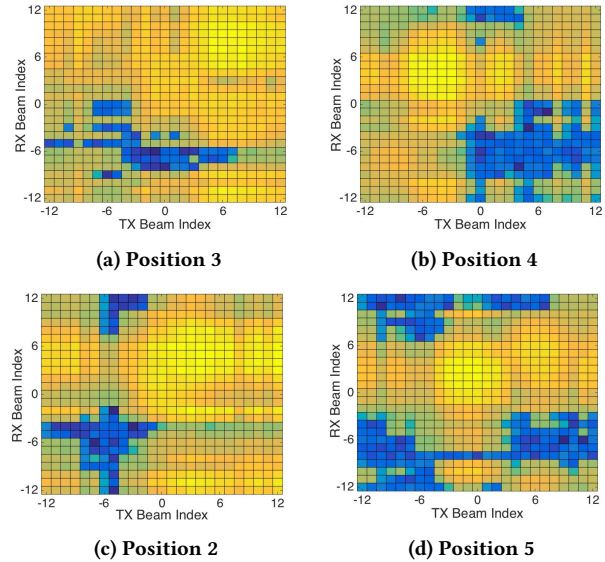


Figure 7: Beam-pair heatmaps for four conference room positions.

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 Fig. 8a and Fig. 8b, respectively. Note that MCS 4 (1/2 16 QAM) was used for modulation since it provides the highest throughput in all positions. Fig. 8a reveals that one can achieve 17-18 dB SNR, even in the absence of a LoS path. Furthermore, the throughput 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.

5.5 Beam Misalignment and Nodal Mobility

In Sec. 5.2, 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 signal strength or may even break the link, depending on the extent of mobility [9]. 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.

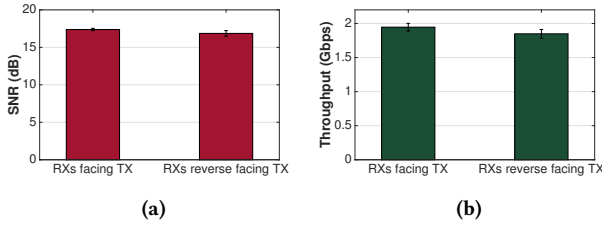


Figure 8: (a) Average SNR, and (b) Average Throughput received in positions 1 to 10 in the conference room.

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 1m. Further, we consider three possible adaptation strategies by both nodes to adjust their beams. (i) Fixed beams i.e., 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, when both nodes perform an exhaustive search at each position and re-select the strongest beam pair. (iii) RX-only adaptation, when only the RX locally adapts its beam to maximize link strength, while the TX beam remains fixed.

Fig. 9 plots SNR vs. TX-RX lateral distance for the three aforementioned strategies across lobby positions 1 through 6. 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. 9 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. This illustrates that a local search at the receiver, 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 and 13 which are in front of the TX (Fig. 2c). Fig. 10a shows SNR vs. receiver angle such that for all positions,

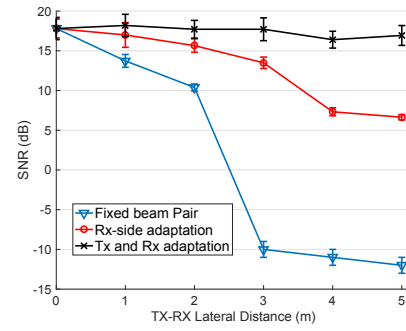


Figure 9: Beam adaptation strategies for lateral translation.

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 Sec. 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 three beam adaptation strategies, i.e., *fixed beam pairs*, *TX and RX adaptation*, and *RX adaptation only*, 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 (6.3m apart) and Position 12 which is roughly 30° to the right of TX (7.5m apart). Thus the latter position captures impact of both angular separation and receiver rotation.

Fig. 10b shows that for Position 10, RX-side 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.

For Position 12 (Fig. 10c), 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° . When beam adaptation is used




Figure 10: Impact of receiver rotation on SNR for various lobby positions.

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. 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.

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 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.

7 ACKNOWLEDGMENTS

This research was supported by NSF grants CNS-1629929, CNS-1553447, CNS-1642929, CNS-1514285, and CNS-1444056, and by Cisco, Intel, and the Keck Foundation.

REFERENCES

- [1] COMSOL Multiphysics. <https://www.comsol.com>
- [2] *wil6210: Linux Wireless*. <https://github.com/torvalds/linux/blob/master/drivers/net/wireless/ath/wil6210/wmi.h>
- [3] *X60: A Cross-Layer Reconfigurable Multi-Gigabit WLAN Testbed at 60 GHz*. <https://ubwins.cse.buffalo.edu/projects/x60.html>
- [4] 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*.
- [5] IEEE 802.11 Task Group AD. 2012. IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. (2012).
- [6] A. Alkhateeb, G. Leus, and R. W. Heath. 2015. Limited Feedback Hybrid Precoding for Multi-User Millimeter Wave Systems. *IEEE Transactions on Wireless Communications* (2015).
- [7] Christopher R. Anderson and Theodore S. Rappaport. 2004. In-building Wideband Partition Loss Measurements at 2.5 and 60 GHz. *IEEE Transactions on Wireless Communications* 3, 3 (2004), 922–928.
- [8] Yasaman Ghasempour and Edward W. Knightly. 2017. Decoupling Beam Steering and User Selection for Scaling Multi-User 60 GHz WLANs. In *Proceedings of ACM MobiHoc*.
- [9] 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*.
- [10] 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).
- [11] George R Maccartney, Theodore S Rappaport, Shu Sun, and Sijia Deng. 2015. 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.
- [12] Alexander Maltsev, Roman Maslennikov, Alexey Sevastyanov, Alexey Khoryaev, and Artyom Lomayev. 2009. Experimental investigations of 60 GHz WLAN systems in office environment. *IEEE Journal on Selected Areas in Communications (JSAC)* 27, 8 (October 2009), 1488–1499.
- [13] Sharan Naribole and Edward Knightly. 2016. Scalable Multicast in Highly-Directional 60 GHz WLANs. In *Proc. of IEEE SECON*.
- [14] Thomas Nitsche, Guillermo Bielsa, Irene Tejado, Adrian Loch, and Joerg Widmer. 2015. Boon and Bane of 60 GHz Networks: Practical Insights into Beamforming, Interference, and Frame Level Operation. In *Proc. of ACM CoNEXT*.
- [15] Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer. 2015. Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement. In *Proc. of IEEE INFOCOM*.
- [16] Swetank Kumar Saha, Anuj Garg, and Dimitrios Koutsonikolas. 2015. A first look at TCP performance in indoor IEEE 802.11 ad WLANs. In *2015 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 63–64.
- [17] Swetank Kumar Saha, Viral Vijay Vira, Anuj Garg, and Dimitrios Koutsonikolas. 2016. A Feasibility Study of 60 GHz Indoor WLANs. In *Proc. of IEEE ICCCN*.
- [18] Swetank Kumar Saha, Viral Vijay Vira, Anuj Garg, and Dimitrios Koutsonikolas. 2016. Multi-Gigabit indoor WLANs: Looking beyond 2.4/5 GHz. In *Proc. of IEEE ICC*.
- [19] Peter F. M. Smulders. 2009. Statistical Characterization of 60-GHz Indoor Radio Channels. *IEEE Transactions on Antennas and Propagation* 57, 10 (October 2009), 2820–2829.
- [20] Peter F. M. Smulders and L. M. Correia. 1997. Characterisation of propagation in 60 GHz radio channels. *Electronics & Communication Engineering Journal* 9, 2 (April 1997), 73–80.
- [21] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parameswaran Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Proc. of ACM SIGMETRICS*.
- [22] Sanjib Sur, Xinyu Zhang, Parameswaran Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of USENIX NSDI*.
- [23] Xiaozheng Tie, Kishore Ramachandran, and Rajesh Mahindra. 2012. On 60 GHz wireless link performance in indoor environments. In *Proc. of PAM*.
- [24] Teng Wei, Anfu Zhou, and Xinyu Zhang. 2017. Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors. In *Proc. of USENIX NSDI*.
- [25] 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 (JSAC)* 20, 3 (April 2002), 620–630.
- [26] 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.
- [27] Jiali Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. 2016. OpenMili: A 60 GHz Software Radio Platform With a Reconfigurable Phased-Array Antenna. In *Proceedings of the ACM MobiCom*.
- [28] 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*.
- [29] Yibo Zhu, Zengbin Zhang, Zhinus Marzi, Chris Nelson, Upamanyu Madhoo, Ben Y. Zhao, and Haitao Zheng. 2014. Demystifying 60GHz Outdoor Picocells. In *Proc. of ACM MobiCom*.