

Multi-Stream Beam-Training for mmWave MIMO Networks

Yasaman Ghasempour
Rice University
ghasempour@rice.edu

Muhammad K. Haider
Rice University
kumail.haider@rice.edu

Carlos Cordeiro
Intel Corporation
carlos.cordeiro@intel.com

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

Edward W. Knightly
Rice University
knightly@rice.edu

ABSTRACT

Multi-stream 60 GHz communication has the potential to achieve data rates up to 100 Gbps via multiplexing multiple data streams. Unfortunately, establishing multi-stream directional links can be a high overhead procedure as the search space increases with the number of spatial streams and the product of AP-client beam resolution. In this paper, we present *MULTI-stream beam-Training for mm-wave networks (MUTE)* a novel system that leverages channel sparsity, GHz-scale sampling rate, and the knowledge of mm-Wave RF codebook beam patterns to construct a set of candidate beams for multi-stream beam steering. In 60 GHz WLANs, the AP establishes and maintains a directional link with every client through periodic beam training. MUTE repurposes these beam acquisition sweeps to estimate the Power Delay Profile (PDP) of each beam with zero additional overhead. Coupling PDP estimates with beam pattern knowledge, MUTE selects a set of candidate beams that capture diverse or ideally orthogonal paths to obtain maximum stream separability. Our experiments demonstrate that MUTE achieves 90% of the maximum achievable aggregate rate while incurring only 0.04% of exhaustive search's training overhead.

CCS CONCEPTS

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

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

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

© 2018 Association for Computing Machinery.

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

<https://doi.org/10.1145/3241539.3241556>

KEYWORDS

MIMO; mmWave; Beam Training; IEEE 802.11ay

ACM Reference Format:

Yasaman Ghasempour, Muhammad K. Haider, Carlos Cordeiro, Dimitrios Koutsonikolas, and Edward W. Knightly. 2018. Multi-Stream Beam-Training for mmWave MIMO Networks. In *The 24th Annual International Conference on Mobile Computing and Networking (MobiCom '18), October 29-November 2, 2018, New Delhi, India*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3241539.3241556>

1 INTRODUCTION

Today, the GHz-scale unlicensed spectrum available at 60 GHz band coupled with phased array antennas is able to support directional transmissions with Gbps data rates. Next-generation devices aim to realize rates up to 100 Gbps via simultaneous transmission of up to eight independent data streams, i.e., downlink single-user and multi-user MIMO¹ [5]. However, the key challenge to realize such rates via mmWave MIMO is to efficiently discover the analog beams at the Access Point (AP) and clients that support *concurrent* directional transmission/reception of multiple data streams and providing the maximum multiplexing gain. Establishing such multi-stream directional links can be a high overhead procedure as the search space increases with the number of spatial streams and the product of AP-client beam resolution. In particular, for simultaneous transmission of m data streams, there are N^{2m} possible beam combinations, where N is the total number of beams in the AP's and clients' RF codebook. While exhaustive search to test all combinations might be acceptable for static use cases such as wireless backhaul, it would incur prohibitively large training overhead in mobile use cases such as virtual reality due to the constant need for beam training.

¹We adopt the IEEE 802.11ay nomenclature and refer to multi-stream communication (via multiple RF chains) as MIMO so that IEEE 802.11ad is referred to as SISO, despite its use of antenna arrays.

We present **MULTI-stream beam-Training for mm-wave networks (MUTE)**, a novel system that identifies dominant paths between the AP and each client in order to efficiently steer 60 GHz beams over diverse or ideally orthogonal paths, such that undesired channel correlations are minimized. Our design is motivated by two key observations: (i) Unlike legacy MIMO in sub 6 GHz bands that are privileged with high multiplexing gain as a result of rich scattering, mmWave channels are sparse, i.e., only a few dominant Line of Sight (LOS) and reflected Non-LOS (NLOS) paths characterize the channel between any two nodes [17, 28]. Furthermore, exploiting an analog beam may impact the mmWave channel as it amplifies certain paths and weakens others. Hence, multiplexing independent streams should avoid common paths as it will otherwise incur throughput degradation due to channel correlation and inter-stream interference. (ii) While a “perfect” non-overlapping set of beam patterns would ensure that use of different beam codebooks would yield non-overlapping paths, practical 60 GHz beams generated via phased array antennas have irregular beam patterns [18, 19]. Nonetheless, despite their irregularity, the directivity gain is known *a priori* in each direction as it is a deterministic function of the codebook and antenna spacing. We exploit these two properties, combined with GHz-scale sampling rate to design MUTE. In particular, we make the following contributions:

First, MUTE repurposes beam acquisition sweeps, which occur periodically in 60 GHz WLANs to establish and maintain a directional link between the AP and each client, to estimate the Power Delay Profile (PDP) of each beam with *zero* additional overhead. In contrast to sub-6 GHz bands, the GHz-scale sampling rate and sparsity of 60 GHz channels provide the unique capability to obtain the high-resolution PDP. While PDP reveals the presence of multiple paths as well as their relative timing, it does not convey any direction information and cannot solely identify orthogonal paths across different clients. Thus, MUTE next couples the known radiation patterns over the suite of irregular beam patterns with PDP measurements for each pattern in order to infer the direction of each path. Lastly, MUTE leverages these direction inferences to construct a candidate set of transmit and receive beams over diverse or ideally orthogonal paths for a multi-stream transmission which will be further trained.

Leveraging the design of MUTE, we present the *first experimental exploration of MIMO beam steering* in mmWave networks. We implement key components of MUTE on X60, a programmable testbed for wide-band 60 GHz WLANs with electronically-steerable phased arrays [18], and modify it to access link-level statistics such as SNR, channel, and PDP in fine-resolution. We collect channel samples (in time and frequency domains) from over-the-air measurements in

an indoor setting and subsequently perform trace-driven emulations. Our key findings are as follows:

It may seem that high SNR beams are always good candidates for multi-stream beam steering as they focus the signal energy towards the intended receiver and that digital precoding can mitigate or ideally cancel any residual interference. However, our results reveal that this is not the case with practical beam patterns since high SNR beams might share a dominant path causing high channel correlation as a result of sparse scattering. This is due to the irregularity of beam patterns generated by phased arrays including the presence of partial overlap among different beams in the RF codebook as well as strong side lobes that cause a particular path to be captured via multiple beams, yet with different directivity gains. We show that in such cases digital precoding methods such as zero-forcing are of little help. That is, digital precoding cannot compensate for a *bad* choice of analog beams that obtain low stream separability in the analog domain.

In contrast, MUTE achieves 90% of the maximum aggregate rate for both single-user and multi-user MIMO, with only 0.04% of the training overhead compared to exhaustive search. In particular, MUTE maximizes stream separability in the analog domain by selecting a candidate subset of beams that capture diverse or ideally orthogonal paths. Although this candidate beam selection itself does not require any additional signaling overhead, discovering the final choice of beams out of this candidate set requires further training with overhead proportional to the candidate set size. MUTE targets that the candidate set size be in the order of the number of LOS/NLOS paths, which is small due to sparse scattering in 60 GHz band. Hence, MUTE approximates the PHY throughput of exhaustive search, while searching over only a few beams with diverse paths.

The rest of this paper is organized as follows. Section 2 illustrates the 60 GHz MIMO node architecture and beam training framework. Section 3 presents the design of MUTE. Section 4 introduces the experimental platform. Our benchmarking algorithms are introduced in Section 5. Section 6 describes the experimental evaluation of our design. Section 7 reviews the related work and Section 8 concludes this paper.

2 60 GHZ MIMO ARCHITECTURE

In this section, we describe the node architecture for multi-stream transmissions in 60 GHz band and our beam training framework.

2.1 Node Architecture

60 GHz radios realize analog beamforming by applying different phase delays to the various antenna elements of a




Figure 1: Node architecture.

phased array antenna. The IEEE 802.11ad standard supports such beam steering but limits the AP to transmit a single stream at a time [9]. Consequently, both the AP and client require only a single RF chain connected to the phased array for digital processing. In contrast, the IEEE 802.11ay standard allows for simultaneous downlink transmission of up to 8 spatial streams. This requires multiple RF chains at the AP (at least one per stream) whereas clients might have a single RF chain (connected to a phased array) or multiple RF chains. The latter is capable of both single-user MIMO and multi-user MIMO reception.

Fig. 1 depicts a node architecture that supports such multi-stream transmission/reception. As shown, each RF chain drives a separate set of phase shifters, each controlling the phase of one antenna element, to be able to independently steer each stream. In this paper, we interchangeably use the terms beam steering, RF beamforming, and analog beamforming to refer to the application of different phase delays to different antenna elements in the RF domain. Moreover, the AP’s plurality of RF chains can also be used for digital pre-coding at baseband to complement analog beam steering (i.e., hybrid analog/digital beamforming). However, the baseband channel is impacted by the phase of antenna elements in the RF domain. Hence, the performance of digital precoding depends on the choice of analog beams [2, 17].

2.2 Beam Training Framework

Multi-stream analog beam steering determines transmit and receive antenna weight vectors for simultaneous transmission of multiple spatial streams. The IEEE 802.11ay standard introduces a two-layer beam training framework for multi-stream beam steering [5]. The first layer involves AP and client beam sweeps to discover the best analog configuration for a directional single-stream transmission (see Sec. 3.2 for details). Next, right before a multi-stream data transmission, IEEE 802.11ay performs a local search over a subset of analog beams from the RF codebook. The

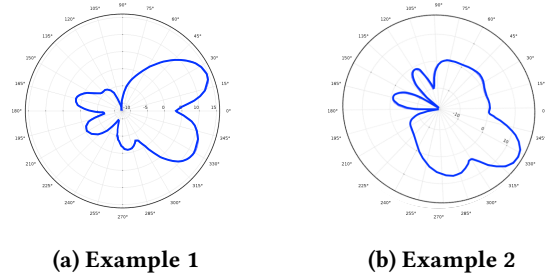


Figure 2: Irregular beam pattern examples from X60 platform [18].

mechanism of candidate beam selection is not specified in the standard and is left to vendors; however, the rationale is that searching over all possible analog configurations is unnecessary as there might be several beams that cannot provide the link budget for a single-stream transmission, let alone multi-stream communication with inter-stream interference. Clearly, there is a tradeoff between the achievable sum rate and the number of candidate beams to be trained in the second layer. We adopt the IEEE 802.11ay two-layer beam training framework but strategically select the minimum number of candidate beams that provide maximum stream separability for the MIMO transmission.

3 MUTE DESIGN

In this section, we describe the design of MUTE that aims to provide the best analog configurations for downlink MIMO transmissions.

3.1 Design Overview

60 GHz channels lack rich scattering, i.e., a few dominant paths fully characterize the channel [28]. An analog beam acts as an amplifier, boosting the strength of certain paths within its main lobe (and side lobes) and weakening the others. As a result, beams that cover the same physical paths have highly correlated channels; transmitting multiple data streams using such beams hinders multiplexing gains. Hence, MUTE targets selection of analog beams that capture diverse or ideally orthogonal paths to the intended receivers. With idealized analog beam patterns, i.e., non-overlapping pencil-shaped beams without side lobes, knowing the received SNR via each beam corresponds to knowledge of path direction. In other words, if a particular beam provides high SNR, we can infer that it captures a path whose direction is within its main lobe. However, beam patterns generated by phased arrays are highly irregular and may even have multiple equally strong lobes [18] as depicted in Fig. 2a. If using such beams at the AP achieves high SNR at a client, it is hard to conclude the number of paths captured




Figure 3: MUTE system architecture.

by each lobe. Hence, selection of this beam pattern (i.e., codebook entry) to exploit one of these lobes might preclude use of another codebook entry that directs energy along the other lobe (e.g., the beam pattern shown in Fig. 2b), as it causes interference in multi-stream transmissions. Note that irregularity in the beam pattern is a byproduct of small array size and limited phase levels at the phase shifters. We observe the same order of irregularity with commodity off-the-shelf 802.11ad devices with 32 antenna elements [19]. Although next-generation 60 GHz devices tend to have even more antennas, it seems that the beam patterns are still far from the “perfect” pencil-shaped patterns due to complexity, size, and power consumption constraints.

We propose MUTE, a MIMO beam steering protocol to select a set of analog beams with diverse or ideally orthogonal paths. Fig. 3 depicts the MUTE system architecture. MUTE runs a background process on the training frames received during the initial beam acquisition phase, which occurs periodically to establish and maintain a directional link between the AP and every client. Thanks to the high sampling rate at 60 GHz band, we can estimate the high-resolution PDP corresponding to each beam. Unfortunately, solely adding PDP information does not solve the MIMO beam steering problem. Namely, while PDP does reveal the presence of multiple paths as well as their relative timing, it does not provide any direction information.

While practical beam patterns are highly irregular, they nonetheless have beamforming gain that is known a priori in each direction. MUTE couples radiation pattern knowledge over a suite of irregular patterns with PDP estimates for each pattern to infer the direction of each path. In other words, by weighting each PDP to the known directional gain for that beam pattern, we can narrow down the direction interval that each path may fall into. For example, in the beam pattern examples of Fig. 2, by receiving the same path (inferred via time delays in the PDP) with both beams, we can correspondingly weight the likelihood that the path direction is within the overlapping area of their main lobes.

MUTE leverages this information to construct candidate sets of transmit and receive beams at the AP and the target user group, respectively. This candidate set selection is

obtained with *zero* additional overhead, only by repurposing the link establishment beam sweeps. As shown in Fig. 3, the final analog configuration is discovered by a local search over the candidate set of beams at the AP and clients.

Lastly, we note that MUTE requires directional beam patterns so that if beam patterns are perfectly omni-directional (or are otherwise identical), we cannot infer the path angles (directions). Furthermore, if the beam patterns are “perfect” and divide 360 degrees into non-overlapping regions, then the solution is quite trivial: one only needs to select distinct high SNR beams for multi-stream transmission, as distinct non-overlapping beams cannot share a common path. Thus, our approach is applicable not only to irregular beams (e.g., Fig. 2), but also symmetric side lobes and any other deviation from strictly non-overlapping beams.

3.2 Primer on Initial Beam Acquisition

Commercial products [19] and WLAN standards such as IEEE 802.11ad [9] and IEEE 802.11ay [5] establish directional links through a training mechanism, in which one node sends training frames sequentially across all beams in the predetermined RF codebook while the other node employs a quasi-omni antenna pattern to find the beam providing the highest signal strength. Repeating this procedure at both ends achieves a beam-pair configuration that can support single-stream (i.e., SISO) directional communication. Although prior efforts have attempted to reduce the frequency of such beam training procedures in mobile 60 GHz WLANs via in-band and out-of-band solutions [1, 11, 16, 21], the AP is bound to periodically repeat these beam sweeps in order to retain directional link connectivity. For example, link failure due to blockage or mobility or the presence of a new user triggers the beam acquisition sweeps.

In MUTE, we assume that the channel between the AP and each client is reciprocal, i.e., under a fixed analog configuration at the AP and the client, the uplink channel is the same as the downlink channel. Thus, the AP sending a training frame via a directional beam to a client in quasi-omni reception, is equivalent to the client sending the training frame via its quasi-omni pattern and the AP receiving it via the same directional beam. Hence, we run MUTE at the AP and process the received training frames in two stages: (i) the AP sweeps across its directional beams while the quasi-omni client sends training frames, (ii) the client sweeps and sends training frames while the AP is in quasi-omni reception. The underlining rationale is that the AP usually has more computation resources than mobile clients and thus MUTE can easily process the received training frames as discussed below.

3.3 Stream Separability Inference

3.3.1 Beam-specific Delay Profile. MUTE specifies the dominant paths by repurposing the received training frames during the initial beam acquisition phase. The GHz-scale sampling rate and sparsity of 60 GHz channels provides the unique opportunity to obtain the high-resolution PDP in contrast to sub-6 GHz bands [28]. PDP gives the distribution of signal power received over a multipath channel as a function of propagation delays and is specified as the spatial average of the complex baseband channel impulse. PDP is typically computed via transmission of a known pilot block [12, 13]. The auto-correlation of the received signal with a local copy of the ideal pilot block provides an estimation of the channel impulse response and leads to PDP estimation.

Unlike sub-6 GHz bands, we cannot represent a 60 GHz channel via a single power delay profile. The reason is behind the fact that omnidirectional transmission is not feasible with an omnidirectional reception in mmWave networks [9]. In other words, due to higher path loss, at least one side needs to be directional. Employing directional beam biases a number of paths by amplifying their strength over others and thus makes the PDP change based on the analog beams in use.

Let $\mathbf{H} \in \mathbb{C}^{N_{cl} \times N_{AP}}$ be the wireless channel between a client (with N_{cl} antennas) and the AP (with N_{AP} antennas). The uplink signal received by the AP can be written as

$$\mathbf{x}[m] = \mathbf{H}[m]\mathbf{w}_{cl}s[m] + \mathbf{n}[m], \quad m = 0, 1, \dots, M-1, \quad (1)$$

where $s[m]$ is the m th symbol of the transmitted training frame (M is total number of symbols), \mathbf{w}_{cl} represents the client's analog beamforming vector and \mathbf{n} is additive white Gaussian noise. With the phased-array antenna, we can only access the signal at the RF chain. The received signal at the RF chain, denoted by $\bar{\mathbf{y}}$, is expressed as:

$$y[m] = \mathbf{w}_{AP}^H \mathbf{x}[m], \quad m = 0, 1, \dots, M-1, \quad (2)$$

where \mathbf{w}_{AP} is the analog beamforming vector at the AP. Both \mathbf{w}_{AP} and \mathbf{w}_{cl} are implemented using analog phase shifters; hence, the modulus of all the elements is 1. Furthermore, for an omnidirectional case, the phase delays are equal to zero for all antenna elements (i.e., $\mathbf{w}_{AP} = \mathbf{1}_{N_{AP} \times 1}$ or $\mathbf{w}_{cl} = \mathbf{1}_{N_{cl} \times 1}$). From Equation (1) and (2), we can write:

$$y[m] = h_{\text{eff}}[m]s[m] + \mathbf{w}_{AP}^H \mathbf{n}[m], \quad m = 0, 1, \dots, M-1, \quad (3)$$

where \mathbf{h}_{eff} is the effective channel between the client and the AP when they employ \mathbf{w}_{cl} and \mathbf{w}_{AP} , respectively and can be written as follows:

$$h_{\text{eff}}[m] = \mathbf{w}_{AP}^H \mathbf{H}[m] \mathbf{w}_{cl} \quad (4)$$

From Equation (4), the effective channel clearly depends on the choice of analog beams. During the initial beam acquisition phase, the AP and each client exchange a series

of training frames while sweeping through their codebook entries while the other side is in omnidirectional mode. These training frames are typically fixed across codebook entries and are known at both the AP and client. Hence, the autocorrelation of the received training signal with the known training frame obtains the PDP as follows :

$$(\bar{\mathbf{y}} \star \bar{\mathbf{s}})[n] = \tilde{h}_{\text{eff}}[n] \quad (5)$$

where $\bar{\mathbf{y}} = [y[0], y[1], \dots, y[M-1]]^T$ is the received signal and $\bar{\mathbf{s}} = [s[0], s[1], \dots, s[M-1]]^T$ is the ideal known training frame.

Peak Detection. A physical path between the AP and the client is reflected as a peak in $|\tilde{h}_{\text{eff}}|^2$. MUTE accounts for noise and hardware impairments by setting thresholds for peak detection (e.g., the peak value should be at least 10 times greater than the noise level). To emphasize that the identified multi-path components are beam dependent, we denote $P_b(\tau)$ the delay profile when the AP employs beam b (with quasi-omni client). $P_b(\tau)$ includes the intensity of detected peaks in $|\tilde{h}_{\text{eff}}|^2$ as a function of time delay τ .

Hardware Imperfections. The estimated power delay profile is subject to error due to carrier frequency offset or packet detection delay. Fortunately, prior work presented several solutions to deal with such hardware imperfections in practice [15, 29]. In this paper, we obtain the PDP measurements directly from our platform (see Sec. 4) which deals with carrier frequency offset and synchronization issues. Today's WiFi drivers already provide channel impulse response (CIR), RSSI, and SNR information and we expect future drivers for the 60 GHz devices to continue this trend and allow access to PDP.

Next, without loss of generality, we explain MUTE's design for selecting transmit beams (at the AP) to enable downlink MIMO transmission. At the end, we briefly explain how candidate beam selection is performed for the MIMO reception at clients in a similar manner.

3.3.2 Delay Profile Aggregation. MUTE obtains the PDP corresponding to all codebook entries. Fig. 4 depicts an example scenario in which a LOS path and two reflected paths exist between the AP and the client. Two analog beam patterns are shown. Beam m (blue) has high directivity along LOS path 1 and NLOS path 3 while having low directivity gain along NLOS path 2. Hence, $P_m(\tau)$ includes two peaks at τ_1 and τ_3 , corresponding to path 1 and path 3, respectively. Likewise, beam n (red) captures two dominate paths 1 and 2. The common LOS path between is reflected in both profiles by a peak at time τ_1 . MUTE aggregates all beam-specific delay profiles to identify all dominant paths between the AP and the client (denoted by $P_{\text{agg}}(\tau)$ in Fig. 4). We avoid double-counting the peaks if their time delay difference is below a configurable threshold to account for noise and




Figure 4: An example scenario with 3 dominant physical paths between the AP and client.

hardware impairments that can cause a slight shift in the delay profile. Note that such aggregation over beam-specific PDPs do not provide any information about the relative strength of different paths as the strength of a path is always biased by the analog beams in use. To bold this fact, we depict same-length peaks for all paths in $P_{agg}(\tau)$ in Fig. 4.

3.3.3 Ambiguity Across Different Clients. By comparing the delay profile of any two beams, we can observe whether they capture similar paths, diverse paths, or completely orthogonal paths. Hence, beam-specific delay profiles can be directly leveraged to infer stream separability in the analog domain for *single user* MIMO transmissions. However, there is ambiguity in stream separability inference for multi-user transmissions. The reason is that $P_{agg}(\tau)$ represents the relative delay of dominant paths with respect to the shortest path (with smallest time delay); comparing the delay profile of different users is problematic as there is no common time reference point. Even if there exists a global time reference point in PDP, we would still not be able to perfectly identify orthogonal paths to multiple users without having a sense of path direction. For instance, two users each having one path with the same path angle (see U_1 and U_2 in Fig. 6a) may experience different time delays as the time of flight or distance to the AP might be different. Hence, solely adding PDP information does not solve the multi-stream beam selection problem as it does not contain any direction information. MUTE tackles this issue by coupling PDP estimates with the *known* radiation patterns.

3.4 Radiation-Weighted Direction Inference

MUTE integrates the obtained beam-specific PDP with the knowledge of beam patterns to estimate the direction of all dominant paths from the AP to each client. Since

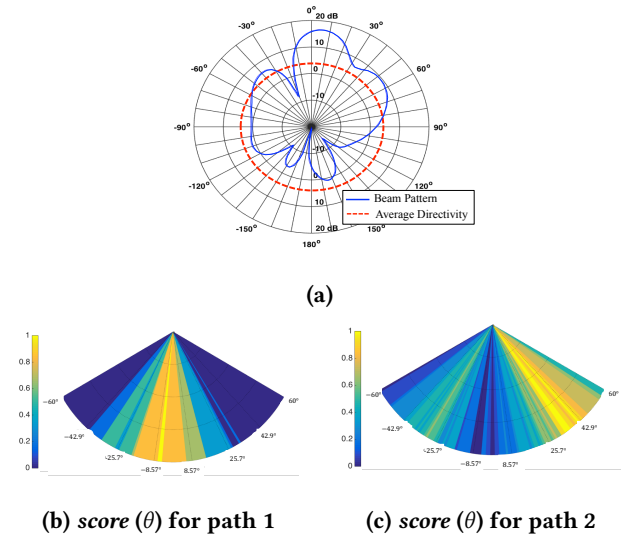


Figure 5: (a) An example irregular beam pattern and its average directivity, (b) $score(\theta)$ for path 1, and (c) $score(\theta)$ for path 2 in Fig. 4.

the AP is the common reference point in space, MUTE is able to infer stream separability for any choice of analog beams in both multi-user and single-user multi-stream transmissions. We denote C_{AP} the AP's RF codebook. In particular, $C_{AP} = [c_1, c_2, \dots, c_N]$ where c_b denotes the radiation pattern of codebook entry b and $c_b(\theta)$ denotes its directivity gain along azimuth angle θ . Furthermore, we define \bar{c}_b as the average directivity of beam b and compute it as $\bar{c}_b = \frac{1}{|c_b|} \sum_{\phi=0}^{2\pi} c_b(\phi)$. Fig. 5a depicts an example beam pattern and its average directivity. For a pencil-shaped beam pattern, only the directivity of the main lobe is above the average directivity; however, for an irregular beam pattern (shown in Fig. 5a), multiple lobes might provide higher than average directivity.

Our key idea is based on the following two insights:

- The PDP of beam b (i.e., $P_b(\tau)$) **containing** peak p (corresponding to a path) implies that the angle of that path is *likely* to be among those directions satisfying $c_b(\theta) > \bar{c}_b$.
- The PDP of beam b (i.e., $P_b(\tau)$) **not containing** peak p (corresponding to a path) implies that that the angle of that path is *unlikely* to be among those directions satisfying $c_b(\theta) > \bar{c}_b$.

Therefore, by weighting each PDP by the known directional gain for that radiation pattern, we can narrow the direction interval that each path may fall into. Algorithm 1 presents the details of MUTE's direction inference mechanism. This algorithm takes beam-specific delay profiles $P_b(\tau)$, the aggregated delay profile $P_{agg}(\tau)$,

Algorithm 1 MUTE's Radiation-Weighted Direction Inference

Input: $\{P_b(\tau)\}, P_{agg}(\tau), C_{AP}$
Output: $A(\Theta)$

- 1: **for** each peak p in $P_{agg}(\tau)$ **do**
 - 2: **for** $c_b \in C_{AP}$ **do**
 - 3: $\bar{c}_b = \frac{1}{|c_b|} \sum_{\phi=0}^{2\pi} c_b(\phi)$
 - 4: **if** $P_b(\tau)$ contains peak p **then**
 - 5: $I(b,p) = 1$
 - 6: **else**
 - 7: $I(b,p) = -1$
 - 8: **end if**
 - 9: **end for**
 - 10: $score(\theta) = \sum_{c_b \in C_{AP}} I(b,p) \left(\frac{|c_b(\theta) - \bar{c}_b| + c_b(\theta) - \bar{c}_b}{2} \right)$
 - 11: $\theta_p^* = \arg \max_{\theta} score(\theta)$
 - 12: **end for**
-

and the pre-known AP's beam patterns C_{AP} . MUTE examines peaks in $P_{agg}(\tau)$ (corresponding to physical paths) one by one. The angle of each path is estimated after examining all beam patterns and their collected delay profiles. For each beam pattern, the algorithm first computes the average directivity gain (line 3) and then indicates whether the delay profile of that beam contains the under-examined path or not. If beam b captures peak p , the algorithm sets the indicator $I(b,p)$ to 1 and -1 otherwise (line 4-7). Then, we weight the likelihood of directions with higher than average directivity via a score-based mechanism (line 10):

$$score(\theta) = \sum_{c_b \in C_{AP}} I(b,p) \left(\frac{|c_b(\theta) - \bar{c}_b| + c_b(\theta) - \bar{c}_b}{2} \right) \quad (6)$$

If beam b captures the peak (i.e., $I(b,p) = 1$), we increase $score(\theta)$ by $c_b(\theta) - \bar{c}_b$ for those angles with higher than average directivity gain (note that $\frac{|x|+x}{2} = 1$ for $x > 0$ and is zero otherwise). Otherwise (i.e., $I(b,p) = -1$), we decrease $score(\theta)$ for those angles with higher than average directivity gain to show the likelihood reduction. Finally, the algorithm returns the direction with the highest score for each path and creates an angular profile $A(\Theta)$ containing the direction of all paths for each user (line 11).

Fig. 5b and Fig. 5c demonstrate $score(\theta)$ for path 1 and path 2 in the example scenario of Fig. 4. We observe that the high score (yellow region) matches the true physical angle. These directions were inferred using imperfect beam patterns generated by practical phased array antennas (see Sec. 4 for details). In general, the accuracy of MUTE's direction inference algorithm depends on the shape of beam patterns, their beamwidth, and the overlap of different codebook entries' beam patterns.




Figure 6: Candidate selection example.

3.5 Candidate Beam Selection

So far, we have described how MUTE applies initial direction acquisition beam sweeps to estimate the delay profile and angular profile of each user. Here, we discuss how MUTE leverages this information to select a subset of candidate beams. MUTE aims to select the analog beams with diverse paths or ideally orthogonal paths. The rationale is that two analog beams sharing a dominant path would have high channel correlation which hinders the multiplexing gain of multi-stream transmission. In particular, [14] showed that in 2×2 single-user MIMO, both streams cannot operate under LOS condition unless they are separated with orthogonal polarizations.

Let G be the target user group. For any user u in G , MUTE attempts to maximize the received signal strength at u while minimizing the interference at all unintended users in G . To this end, for any dominant path i between the AP and user u , MUTE includes a beam in the candidate set that provides maximum directivity along path angle $\theta_{u,i}$ (to increase signal strength at the intended user), while inducing minimum directivity along the path angles of all other unintended users in G (to reduce interference). The general optimization for candidate beam selection is therefore:

$$B_u(G) = \left\{ \arg \max_b \frac{c_b(\theta_{u,i})}{\sum_{\substack{v \in G \\ v \neq u}} \sum_{\substack{\theta_{v,x} \in A_v \\ \theta_{v,x} \neq \theta_{u,i}}} c_b(\theta_{v,x})}, \forall \theta_{u,i} \in A_u(\Theta) \right\} \quad (7)$$

where $c_b(\theta)$ is the directivity of b^{th} entry in the AP's codebook along azimuth angle θ and $A_u(\Theta)$ represents the angular profile of user u . We further elaborate on MUTE's candidate selection mechanism via a simple scenario with three users U_1 , U_2 and U_3 as depicted in Fig. 6a. Fig. 6b shows the angular profiles of these users; U_1 and U_2 hold two dominant physical paths with the AP, i.e., $A_1(\Theta) = \{\theta_{1,1}, \theta_{1,2}\}$, and $A_2(\Theta) = \{\theta_{2,1}, \theta_{2,2}\}$ while $\theta_{1,1} = \theta_{2,1}$ (they share a LOS path). Note that in practice, we flag such

shared path if the path angle distance (here, $|\theta_{1,1} - \theta_{2,1}|$) is a below a configurable threshold. Furthermore, U_3 has only a LOS path with the AP, i.e., $A_3(\Theta) = \{\theta_{3,1}\}$. We refer to U_1 and U_2 as *path-sharing users* because their angular profiles share a common path. We also call U_1 and U_3 *distinct users*. Next, we discuss two examples of candidate beam selection:

3.5.1 Example 1: Distinct Users. First, we target selection of candidate transmit beams for a two-stream MIMO transmission to U_1 and U_3 . To maximize the stream separability in the analog domain, we choose analog beams such that they provide high directivity gain along path angles of U_1 while inducing minimum interference at U_3 , and vice versa. According to Equation (7), MUTE constructs the candidate set of beams, denoted by $B(U_1, U_3)$, as follows:

$$B(U_1, U_3) = \left\{ \arg \max_b \frac{c_b(\theta_{1,1})}{c_b(\theta_{3,1})}, \arg \max_b \frac{c_b(\theta_{1,2})}{c_b(\theta_{3,1})}, \arg \max_b \frac{c_b(\theta_{3,1})}{c_b(\theta_{1,1}) + c_b(\theta_{1,2})} \right\} \quad (8)$$

This candidate set contains: (i) The beam that has maximum directivity along $\theta_{1,1}$ while having minimum directivity along $\theta_{3,1}$; (ii) The beam with maximum directivity along $\theta_{1,2}$ and minimum directivity along $\theta_{3,1}$; and (iii) The beam with maximum directivity along $\theta_{3,1}$ while having minimum directivity along $\theta_{1,1}$ and $\theta_{1,2}$. The candidate set includes three beams in which two of them are eventually selected for 2×2 MU-MIMO transmission. In this toy example, it is not hard to predict that the third candidate beam in the $B(U_1, U_3)$ should be one of the final selected beams as it is the only candidate beam that provides high directivity gain (and SNR) at U_3 .

3.5.2 Example 2: Path-Sharing Users. Second, we target a two-stream MIMO transmission to U_1 and U_2 in Fig. 6b. Here U_1 and U_2 share a common LOS path (i.e., $\theta_{1,1} = \theta_{2,1} = \theta'$) and the transmission of two streams under LOS condition would result in significant channel correlation. Therefore, the analog beam along LOS path should be employed for either user 1 or user 2, and not both. In the former case, its directivity gain along $\theta_{2,2}$ should be minimized while in the later case, its directivity gain along $\theta_{1,2}$ should be minimum. Thus, the set of candidate beams for MU-MIMO to U_1 and U_2 can be directly derived from Equation (7) as follows:

$$B(U_1, U_2) = \left\{ \arg \max_b \frac{c_b(\theta')}{c_b(\theta_{2,2})}, \arg \max_b \frac{c_b(\theta')}{c_b(\theta_{1,2})}, \arg \max_b \frac{c_b(\theta_{1,2})}{c_b(\theta_{2,2}) + c_b(\theta')}, \arg \max_b \frac{c_b(\theta_{2,2})}{c_b(\theta_{1,2}) + c_b(\theta')} \right\} \quad (9)$$

Note that single-user MIMO beam steering is a special case in Equation (7) in which the angular profiles of all users in G are identical.




Figure 7: The X60 platform for 60 GHz band.

Discussion. We have described how MUTE selects a set of candidate transmit beams by processing the received training frames at the AP. There are two possible ways to replicate the same procedure for client-side beam selection: (i) Each client processes the received training frames during the initial beam acquisition sweeps, taking into account its known RF beam patterns; or (ii) clients send beam training frames (while sweeping over their directional beams) and the AP receiving the training frame takes care of processing. In the latter case, the AP needs to know the clients' RF beam patterns and announce the candidate beams to the target user group before the final local search.

3.6 Local Search

Once the sets of candidate transmit and receive beams are constructed, local training is triggered to discover the optimum analog configuration. This training involves testing all possible beam combinations in the candidate sets for a multi-stream transmission to the target user group and estimating the achievable aggregate rate under each configuration. While MUTE's candidate beam selection does not entail any additional overhead, this local search requires active signaling and thus incurs overhead that is proportional to the number of candidate beams at the AP and clients. We evaluate the candidate set size in Sec. 6.3.

4 EXPERIMENTAL PLATFORM

We conduct over-the-air experiments utilizing X60, a programmable testbed for wideband 60 GHz WLANs [18]. X60 provides signal level accessibility, and is engineered to provide CSI, SNR, and PDP. Fig. 7 depicts X60 in which each node is equipped with National Instruments' mm-Wave transceiver systems and a *user-configurable* 24-element phased array antenna from SiBeam. It enables communication over 2 GHz channels via fixed codebook based beam patterns that can be steered in real-time (electronic switching in $< 1\mu s$). X60 enables fully programmable PHY, MAC, and Network layers. The

reference PHY implementation allows for modulation and coding combinations from 1/5 BPSK to 7/8 16 QAM, resulting in 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.

The built-in phased array has 24 antenna elements; 12 for transmission and 12 for reception. The phase of each antenna element can take one of the four values: 0, π/2, π, and 3π/2. SiBeam’s reference codebook consists of 25 3D beam patterns spaced roughly 5° apart (in their main lobe direction) and covering a sector of -60° (corresponding to beam index -12) to 60° (corresponding to beam index +12) around the antenna’s broadside direction. The half power beamwidth of beams are 25° – 35°. X60’s beam patterns (two of them are shown in Fig. 2) are similar to the patterns of commodity off-the-shelf 802.11ad devices which have main lobe overlaps and strong side-lobes [19]. Thus, X60 allows us to evaluate realistic, imperfect and irregular beam patterns, and their impact on multi-stream beam steering.

Due to hardware limitations (availability of only one RF chain at each node), over-the-air MIMO transmission is not feasible. Our key experimental methodology is to collect channel samples (in time and frequency domain) from over-the-air measurements and subsequently perform trace-driven emulation to study MIMO beam steering in millimeter-wave networks. Sec. 6 elaborates more on our measurement setup.

5 BENCHMARKING ALGORITHMS

In this section, we describe two benchmarking algorithms for evaluation purposes.

5.1 SNR-based Beam Selection (Baseline)

We introduce a baseline scheme that down selects a subset of beams from the fixed RF codebook based on their achievable SNR in SISO communication. The underlying rationale is that higher received signal strength provides greater margin for interference tolerance. Therefore, this scheme selects a subset of beams (for multi-stream transmission) entirely based on their SNR values and exploits zero-forcing to mitigate or ideally cancel any residual interference. In particular, the baseline scheme picks the top n beams at the AP and the top m beams at each client, according to their achievable SNR in the initial AP-side and client-side beam sweeps. The final analog configuration is realized via a local search among all possible combinations of candidate beams.

Similar to MUTE, this baseline approach does not introduce additional overhead for candidate selection as the received SNR associated with each beam is already available after initial beam sweeps as discussed in Sec. 3.2. However, unlike MUTE which infers stream separability




Figure 8: Experimental floorplan. Square boxes represent client positions.

by building an angular profile, the baseline scheme select beams entirely based on SNR. The final analog setting is found by further training whose overhead depends on the number of transmit and receive beams in candidates sets (i.e., n and m). Throughout the paper, we refer to this method as baseline scheme or SNR-based beam selection strategy.

5.2 Exhaustive Search

Exhaustive search assesses all beam combinations to find the optimal configuration. For simultaneous transmission of m data streams, exhaustive search tests total $O(|C_{AP}|^m \times |C_{cl}|^m)$ analog combinations where $|C_{AP}|$ is the transmit codebook size at the AP and $|C_{cl}|$ is clients’ codebook size. Implementation of exhaustive search may not be practical in real scenarios due to the prohibitively large beam training overhead. Nonetheless, we study this algorithm for comparison purposes as it provides an upper-bound for the achievable sum-rate of multi-stream transmission.

6 EVALUATION

In this section, we conduct over-the air measurements to evaluate and compare the performance of MUTE against the benchmarking algorithms.

Setup. We deploy X60 nodes and conduct an extensive set of experiments in multiple indoor environments and many AP-client settings. In this paper, as the first attempt to explore MIMO beam steering in 60 GHz WLANs with over-the-air channel traces, we zoom into one experimental setup depicted in Fig. 8 including the AP and 12 client locations (represented by square boxes). The AP is mounted on a tripod at 0.9 m height from the floor and pointing North. The client is at the same height but pointing South at all locations. The presence of windows and metal coating beneath them (not shown) create reflections.

Methodology. For each AP-client setting, we collect channel statistics for all possible 625 (25 × 25) beam-pair combinations. For each beam-pair, 100 frames are transmitted at Modulation and Coding Scheme (MCS) 0 and

SNR, channel magnitude and phase, and PDP are logged every four frames (every 40ms). We assume that the multiple virtual RF chains are co-located at the AP. This is because we find out that changing the location of the AP by $\lambda/2$ or 2.5mm does not change the received PDP, composite channel, and SNR. In other words, due to wide ($\approx 25^\circ$) beamwidth of codebook entries, the physical paths being captured by a beam are not sensitive to small movement of the AP or client. To emulate hybrid analog/digital beamforming, we process the channel traces under a fixed choice of analog beams at the AP and clients and compute zero-forcing weights. We then map the SINR to the corresponding data rate using the protocol-specific minimum SNR tables [4].

6.1 Performance Analysis of MUTE

First, we explore the performance of MUTE in selecting the best analog beams at the AP and client(s) to be used for multi-stream simultaneous transmission and reception, respectively. For simplicity, we focus on a two-stream case (2×2 multi-user and single-user MIMO configurations) and later, in Sec. 6.4, we increase the number of spatial streams.

6.1.1 Multi-User MIMO. We consider a downlink two-user MIMO case in which one client, R_1 , is placed at position 1, while we consider all other 11 positions for the second client (R_2) by sequentially repeating experiments at these positions. For each position of R_2 , MUTE generates a subset of candidate beams at the AP and each client by post processing the measured beam-specific PDPs. Then, to discover the final multi-stream configuration, MUTE performs a local search and computes the zero-forcing weights for any combination of analog beams in the candidate sets. Applying the zero-forcing weights, we compute the expected SINR at R_1 and R_2 and infer the per-user data rate by employing the protocol-specific minimum SNR tables [4]. In particular, for each user and potential multi-stream analog configuration, MUTE selects the MCS index whose corresponding SNR is less than or equal to the calculated SINR. The corresponding number of data bits per symbol is the per-user data rate (each stream can use a different MCS but the coding rate for two streams are the same). The final MIMO beam configuration is the one providing highest aggregate data rate. For comparison, we also implement Exhaustive Search and the baseline SNR-based Beam Selection. To have a fair comparison, we ensure that the candidate set size for the baseline scheme is the same as MUTE, e.g., if MUTE provides k candidate beams, the baseline scheme would down select the top k beams (based on SNR).

Fig. 9 shows the aggregate PHY rate of the baseline, MUTE, and Exhaustive Search for a two-user simultaneous




Figure 9: Aggregate PHY rate of a two-user MIMO transmission to R_1 (fixed at position index 1) and R_2 when placed at other 11 positions.

transmission to R_1 and R_2 . First, we observe that MUTE is able to achieve more than 90% of Exhaustive Search’s rate across all locations. This implies that MUTE’s multi-path inference successfully discovers candidate beams with maximum stream separability. Second, Fig. 9 shows that the baseline strategy provides around 60% of the Exhaustive Search’s aggregate rate, except for when R_2 is placed at position indices 4, 7 and 10. In those cases, R_1 and R_2 are approximately along the same LOS path with the AP. Hence, the baseline scheme selects analog beams along the shared LOS path and induces inter-user interference. In other words, a candidate beam (at the AP) that is intended for R_1 has high directivity along the LOS path and will incur high interference at R_2 . Transmission of two streams along the LOS path reduces the aggregate PHY rate. In contrast, MUTE is able to create “separate beam” transmissions via an NLOS path, even if users are along the same LOS path, and still obtain 90% of maximum aggregate rate.

Finding: MUTE achieves more than 90% of Exhaustive Search’s aggregate rate in two-user MIMO case: If the LOS paths from the AP to two users have enough spatial separation, MUTE selects beams along the LOS path for both users; otherwise, MUTE is able to create “separate beam” transmissions via NLOS paths.

6.1.2 Single-User MIMO. Next, we perform a similar experiment for 2×2 single-user MIMO. In contrast to downlink multi-user MIMO in which users are spatially separated, single-user MIMO requires multiple independent data streams to be successfully decoded at one spatial location. This makes the problem of beam selection even more challenging. To study single-user MIMO, we employ the same node deployment as in Fig. 8, and explore the performance of MUTE across all client positions and




Figure 10: Multiplexing gain of 2×2 single-user MIMO as a function of client position.

compare it against Exhaustive Search and the baseline scheme. Note that again, to ensure a fair comparison, the number of candidate beams in the baseline is the same as in MUTE.

Fig. 10 depicts the multiplexing gain of 2×2 single-user MIMO across different client positions. The multiplexing gain is computed by dividing the aggregate PHY rate of two-stream transmission over the maximum data rate of the corresponding SISO transmission being realized by employing the (TX, RX) beam-pair with the highest SNR. Theoretically, this setup should achieve close to $2 \times$ gain; however, this does not hold true for every client position even under Exhaustive Search as shown in Fig. 10. This is because we find in measurements that the signal strength along the NLOS path is typically lower than the LOS one. Hence, even if the inter-stream interference is negligible, the aggregate PHY rate of a two-stream transmission (via a LOS and an NLOS path or two distinct NLOS paths) might not obtain $2 \times$ gain over the SISO transmission.

Although single-user MIMO beam training is inherently more challenging, MUTE is still able to obtain 90% of the Exhaustive Search’s multiplexing gain across all receive positions. The reason is that MUTE, by design, identifies all dominant paths and includes analog beams capturing diverse or ideally orthogonal paths.

In contrast, Fig. 10 reveals that the baseline scheme does not support transmission of two streams as the multiplexing gain is lower than unity for several client positions. This is because the baseline scheme selects analog beams that capture the LOS path as they provide higher SNR. Hence, to due lack of rich scattering, the vector channel of first stream is highly correlated with the second stream and the client cannot successfully decode both streams. Sec. 6.2 elaborates more on the limitations of the baseline (SNR-based Beam Selection) scheme.

Finding: In single-user MIMO, transmission of multiple streams across the LOS path hinders multiplexing gain; however, MUTE successfully creates separate beam transmissions via diverse paths and achieves 90% of optimal performance.

6.2 Limitations of SNR-based Beam Selection

6.2.1 Sparsity of Channel, Richness of Strong Beams. We hypothesize that only a few beams can provide sufficient SNR for multi-Gbps communications. This implies that there might be a few beams that can potentially support MIMO as most beams would not even provide the sufficient link budget for a SISO transmission. To investigate this premise, we conduct over-the air experiments and explore the distribution of “strong” beams (i.e., beams that support at least 1 Gbps data rate in the SISO configuration). We use the same node placement as depicted in Fig. 8 and measure the received SNR for all 25×25 beam-pair combinations under two scenarios: (i) the AP has a LOS path to the client; (ii) the LOS path is blocked with a wooden table as shown in Fig. 12. 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. 11 presents the SNR heatmaps for client position indices 3, 6, 9 and 12 under LOS connectivity (top row) and blockage (bottom row) (due to space limit, we do not show the SNR heatmap for all positions). The SNR range in the heatmaps is between -20 dB to 15 dB with yellow colored regions indicating beam-pairs with SNR above 10 dB whereas blue regions indicate beam-pairs with negative SNR.

Surprisingly, we observe several beam-pairs providing above 10 dB SNR that corresponds to 1 Gbps data rate in our platform. The received SNR corresponding to an analog configuration is a function of its captured physical paths and the directivity gain along them. Imperfect beam patterns cause a physical LOS/NLOS path to be captured by multiple beams, albeit with different directivity gains. We confirm that the beam-pair with the highest SNR corresponds to the physical LOS path in Fig. 11(a)-(d). For instance, the



Figure 12: LOS blockage with a wooden table.

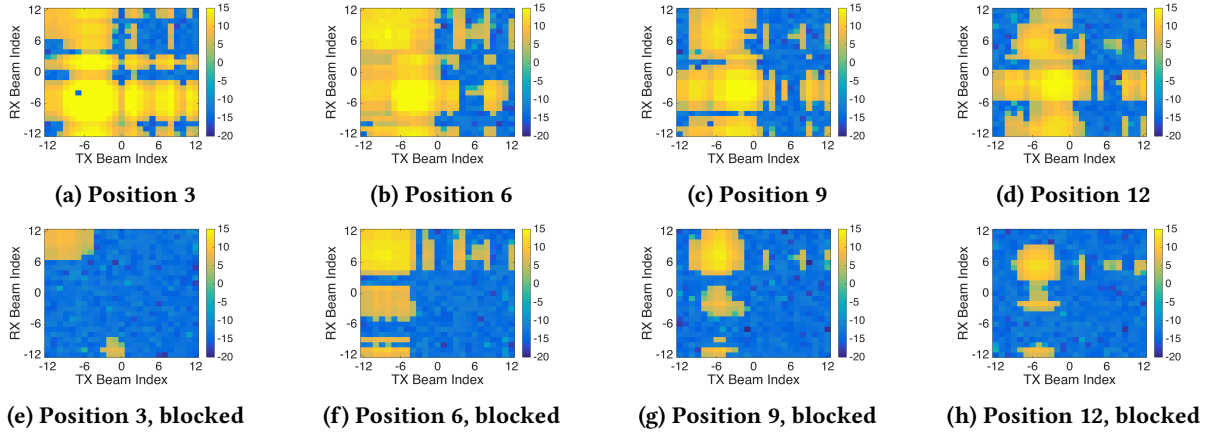


Figure 11: The SNR heatmaps for all 25×25 beam-pair combinations for client position indices 3, 6, 9 and 12 under LOS connectivity (top row) and blockage (bottom row).

direction of the geometrical LOS path between TX and RX at position 3 is $\tan^{-1}\left(\frac{-1.5}{3.3}\right) \approx -25^\circ$. Since the main lobes of X60 beam patterns are spaced roughly 5° apart, beam index -5 at TX and -5 at RX should provide highest directivity gain along the LOS path. Fig. 11a confirms that beam-pair (-5,-5) is within the high SNR region; however, due to overlap between neighboring beams, multiple beams include the LOS path and we see a cluster of high SNR beam-pairs around (-5,-5).

Under LOS blockage, the yellow region corresponding to the LOS component disappears in the bottom row plots of Fig. 11. This confirms that the signal strengths of neighboring beams are highly correlated as they capture one common path. Surprisingly, we observe that other high SNR beam-pairs with *large codebook distance* from the LOS region also experience significant SNR reduction after blockage. This implies that whether two beams capture the same path cannot simply be inferred from their RF codebook distance due to the irregularity and imperfections of beam patterns. Lastly, the highest SNR region after blockage achieves similar SNR under LOS conditions, i.e., LOS blockage has not degraded their SNR; thus, these beam-pairs must be capturing a reflected path. While MUTE discovers the LOS path as well as reflected paths, the baseline scheme likely selects LOS beams as they provide higher SNR.

Finding: While there exist a few physical paths between any two nodes, several beams may capture at least one path and thus provide high SNR. Whether two beams capture the same path or not cannot simply be inferred from their codebook distance due to the irregularity and imperfections of phased array generated beam patterns.

6.2.2 Zero-Forcing to the Rescue? Traditionally, digital precoding schemes such as zero-forcing are employed to mitigate or ideally cancel inter-stream interference. Here, we




Figure 13: Aggregate PHY rate of baseline scheme as a function of candidate set size in 2×2 single-user MIMO.

explore the impact of such digital precoding techniques on the performance of SNR-based beam selection. In particular, we explore whether digital precoding can compensate for a *bad* choice of analog beams. To this end, we repeat the same experiment as in Sec. 6.1.2 and consider a special case that the candidate set size at the AP and client are equal (i.e., $m = n$). In particular, we vary m from 2 to 25, which is the total number of available beams in X60 (note that for $m = 25$, the baseline scheme turns into the Exhaustive Search).

Fig. 13 shows the normalized aggregate PHY rate with and without zero-forcing for the baseline scheme as the candidate set size varies between 2 to 25. The normalization is computed based on the achievable aggregate rate of the Exhaustive Search (i.e., $m = n = 25$). As expected, applying zero-forcing mitigates interference and boosts the achievable aggregate rate. Interestingly, to provide 90% of Exhaustive Search's aggregate rate, the baseline scheme requires the candidate set size to be 16 (which is more than half of the RF codebook size) for a two-stream analog-only beam steering. Applying zero-forcing mitigates interference and enables the baseline approach to achieve 90% of Exhaustive




Figure 14: The number of candidate beams by MUTE in each client’s location.

Search’s aggregate rate with 10 candidate beams. This result implies that although zero-forcing reduces the inter-stream interference, it cannot completely compensate for a *bad* analog beam selection as the performance gap with the optimal solution remains significant.

The reason is behind how zero-forcing performs: Zero-forcing cancels inter-user interference by projecting the channel vector of a user on a precoding vector that is orthogonal to the channel vector of the other user so that the precoded (i.e., projected) channels become orthogonal to each other. The penalty one pays for such interference cancelation depends on the mutual channel correlation between users, i.e., channel projection incurs signal energy loss if the original channel vectors are not orthogonal. In sub-6 GHz bands, the rich scattering propagation causes semi-orthogonal channels and thus zero-forcing can successfully cancel interference (with low penalty). In contrast, 60 GHz channels are sparse and, more importantly, the effective channel vector of each user depends on the choice of analog beams that amplifies certain paths and weakens others. We have demonstrated that the high SNR beams typically share a common LOS path and thus incur high channel correlation. Hence, exploiting zero-forcing in the digital domain cannot compensate for low stream separability in the analog domain.

Finding: Although zero-forcing mitigates inter-stream interference, it cannot compensate for a bad choice of beams with high channel correlation and low analog domain separability.

6.3 Training Overhead

As discussed in Sec. 3, MUTE constructs a candidate set of beams and the final analog configuration is found by a local search among all combinations of beams in the candidate sets. Although MUTE obtains the candidate beams by passively overhearing training frames from initial beam acquisition phase with zero additional overhead, local search incurs additional time overhead that is proportional to the candidate beam set size.




Figure 15: The aggregate PHY rate as a function of the number of spatial streams.

Fig. 14 presents the size of candidate sets selected by MUTE for single-user MIMO configuration (the number of candidate beams for a multi-user transmission depends on the choice of target user group). We observe the candidate set size can be as low as 2 and is at most 5. MUTE, by design, adds exactly one candidate beam corresponding to every physical path in the aggregate delay profile; hence, the sparsity of 60 GHz channels results in selection of only a few beams (on average four).

Combining this result with the one in Sec. 6.1, we conclude that MUTE is able to achieve 90% of optimal aggregate rate by performing local search over only a few (on average four) candidate beams at the AP and client. In contrast, Exhaustive Search has to test all $\binom{25}{2} \times \binom{25}{2}$ beam combinations. Therefore, MUTE can achieve 90% of optimal aggregate rate while incurring only $\frac{\binom{4}{2}\binom{4}{2}}{\binom{25}{2}\binom{25}{2}} = 0.04\%$ of Exhaustive Search’s training overhead. In other words, MUTE reduces the training overhead by 99.6% with only 10% throughput loss. Furthermore, we observed in Sec. 6.2.2 that the baseline scheme provides 90% of the optimal aggregate rate when the candidate set size is 10. Consequently, under similar aggregate rate, MUTE requires only 1.8% of the baseline schemes’ training overhead ($\frac{\binom{4}{2}\binom{4}{2}}{\binom{10}{2}\binom{10}{2}} \approx 1.8\%$).

Finding: MUTE achieves 90% of optimal aggregate rate while inducing only 0.04% of Exhaustive Search’s training overhead or 1.8% of baseline scheme’s overhead.

6.4 Scaling the Number of Spatial Streams

So far, we have evaluated the performance of MUTE for two-stream transmission. Here, we increase the number of spatial streams in a multi-user MIMO setting. To this end, we use the same node deployment as in Fig. 8 and for a k stream transmission, we consider all possible user groups consisting of k users out of 12 (i.e., total of $\binom{12}{k}$ different user groups). For each user group, we find the beam steering configuration under baseline, MUTE, and Exhaustive Search and report the achievable aggregate PHY rate after applying zero-forcing.

Fig. 15 depicts the achievable aggregate PHY rate as a function of the number of spatial streams. First, we observe that exhaustive beam steering is able to achieve on average about 2× and 3× throughput gain via simultaneous transmission of two, and three data streams, respectively. However, by further increasing the number of spatial streams from 4 to 6, the system’s aggregate rate deviates from the *ideal* case which would linearly scale with number of spatial streams. This implies that the MIMO multiplexing gain does not endlessly increase proportionally with the number of streams because of undesired channel correlations. However, the saturation point depends on the AP-clients setting.

Second, while the baseline strategy provides 69% of the maximum sum-rate for two-stream transmission, its performance gap with Exhaustive Search increases with more streams such that it provides only 50% of optimal rate with 6 spatial streams. This is because the baseline scheme attempts to choose beams based on their achievable signal strength at the intended client and relies on zero-forcing to cancel inter-stream interference. However, the undesired channel correlation among a larger set of users cause inefficient interference cancellation and degradation in the relative performance of the baseline scheme. In contrast, MUTE accounts for undesired channel correlations by selecting beams over diverse or orthogonal paths. As a result, with increasing the number of spatial streams, MUTE’s gain remains close to the performance of Exhaustive Search (i.e., with marginal loss).

Finding: While multi-stream beam streaming becomes more challenging with an increasing number of spatial streams, MUTE is able to provide 87% of the maximum aggregate rate (realized by Exhaustive Search) with six streams.

7 RELATED WORK

Prior work on 60 GHz WLANs mainly focuses on efficient beam training and tracking for *single-stream* transmission. To the best of our knowledge, MUTE is the first work on multi-stream beam steering.

Single-Stream Beam Training. Prior efforts reduced the frequency of single-stream beam training procedures in mobile 60 GHz WLANs via a variety of in-band and out-of-band solutions [1, 10, 11, 16, 21]. Such work is complimentary to MUTE as our method can be employed whenever the AP invokes beam sweeps, even if their frequency has been optimized. For example, link failure under blockage and mobility and the presence of a new user triggers the direction acquisition beam sweeps. MUTE passively overhears these beam sweeps and infers the beam separability by estimating PDPs.

Channel Profiling. Reverse-engineering 60 GHz channels has been explored in prior work with a different purpose of improving network connectivity in mobility and blockage [23, 24, 30]. Although such mechanisms obtain an aggregate channel profile (including LOS direction and location of reflectors), they do not provide stream separability inference, primarily due to lack of beam-specific multi-path profile knowledge, especially with imperfect beam patterns generated by practical phased arrays.

Direction Estimation. Direction inference techniques in sub-6 GHz bands have been studied in prior work [25–27]. These techniques employ the phase difference at multiple antennas for Angle of Arrival (AOA) estimation. However, due to a different node architecture at 60 GHz band (lacking one RF chain per antenna), we can only acquire a composite channel at the RF chain, where signals from multiple antenna elements are mixed, thereby thwarting AOA estimation. Likewise, direction estimation for sub 6 GHz MU-MIMO with hybrid beamforming was explored in [3]. However, their approach is limited to LOS detection and requires the analog weight vectors to be orthogonal, which does not hold true for practical phased arrays with limited phased levels. In contrast, MUTE takes advantage of GHz-scale sampling rate and sparsity of 60 GHz channel to obtain high resolution power delay profiles yielding to stream separability inference.

60 GHz MIMO. Prior work studied the potential of beam steering and spatial multiplexing in 60 GHz WLANs and showed that its directional nature motivates spatial reuse [7, 8, 20, 22]. Other works have explored hybrid beamforming [2] and user selection for multi-user 60 GHz WLANs [6]. In contrast, this paper aims to find the best analog configuration for multi-stream beam steering to a given target user-group.

8 CONCLUSION

In this paper, we present MUTE, a novel system that enables multi-stream transmission in 60 GHz WLANs via diverse steering. MUTE repurposes initial beam acquisition sweeps to estimate the beam-specific PDP without incurring additional overhead. Combining PDP measurements with knowledge of beam patterns, MUTE suggests a candidate set of beams that capture diverse or ideally orthogonal paths. Over-the-air measurements with practical phased array antennas show that MUTE provides about 90% of the maximum achievable aggregate rate while incurring only 0.04% of exhaustive search’s training overhead.

9 ACKNOWLEDGMENTS

We appreciate the valuable comments and feedback from the anonymous reviewers and our shepherd. This research was supported by Intel, and by NSF grants CNS-1642929, CNS-1514285, CNS-1629929, and CNS-1553447.

REFERENCES

- [1] Omid Abari, Haitham Hassanieh, Michael Rodriguez, and Dina Katabi. 2016. Millimeter Wave Communications: From Point-to-Point Links to Agile Network Connections. In *Proc. of ACM HotNets*.
- [2] Ahmed Alkhateeb, Geert Leus, and Robert W. Heath. 2015. Limited Feedback Hybrid Precoding for Multi-User Millimeter Wave Systems. *IEEE Transactions on Wireless Communications* 14, 11 (2015), 6481–6494.
- [3] Zhe Chen, Xu Zhang, Sulei Wang, Yuedong Xu, Jie Xiong, and Xin Wang. 2017. BUSH: Empowering Large-Scale MU-MIMO in WLANs With Hybrid Beamforming. In *Proc. of IEEE INFOCOM*.
- [4] Claudio R. C. M. da Silva, Artyom Lomayev, Cheng Chen, and Carlos Cordeiro. 2018. Analysis and Simulation of the IEEE 802.11ay Single-Carrier PHY. In *Proc. of IEEE ICC*.
- [5] Yasaman Ghasempour, Claudio da Silva, Carlos Cordeiro, and Edward W. Knightly. 2017. IEEE 802.11ay: Next-generation 60 GHz Communication for 100 Gbps Wi-Fi. *IEEE Communications Magazine* 55, 12 (2017), 186–192.
- [6] Yasaman Ghasempour and Edward W. Knightly. 2017. Decoupling Beam Steering and User Selection for Scaling Multi-User 60 GHz WLANs. In *Proc. of ACM MobiHoc*.
- [7] Yasaman Ghasempour, Narayan Prasad, Mohammad Khojastepour, and Sampath Rangarajan. 2017. Link packing in mmWave networks. In *Proc. of IEEE ICC*.
- [8] Yasaman Ghasempour, Narayan Prasad, Mohammad Khojastepour, and Sampath Rangarajan. 2017. Managing Analog Beams in mmWave Networks. In *Proc. of Asilomar Conference on Signals, Systems and Computers*.
- [9] IEEE 802.11 Working Group. 2012. IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. (2012).
- [10] Muhammad Kumail Haider, Yasaman Ghasempour, and Edward W. Knightly. 2018. Search Light: Tracking Device Mobility Using Indoor Luminaries to Adapt 60 GHz Beams. In *Proc. ACM MobiHoc*.
- [11] 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*.
- [12] Kun C. Hung and David W. Lin. 2010. Pilot-Based LMMSE Channel Estimation for OFDM Systems With Power Delay Profile Approximation. *IEEE Transactions on Vehicular Technology* 59, 1 (2010), 150–159.
- [13] Young J. Kim and Gi H. Im. 2012. Pilot-Symbol Assisted Power Delay Profile Estimation for MIMO-OFDM Systems. *IEEE Communications Letters* 16, 1 (2012), 68–71.
- [14] Alexander Maltsev, Ali Sadri, Carlos Cordeiro, and Andrey Puduev. 2015. Practical LOS MIMO Technique for Short-Range Millimeter-Wave Systems. In *Proc. of IEEE ICUWB*.
- [15] Apurva N. Mody and Gordon L. Stuber. 2001. Synchronization for MIMO OFDM Systems. In *Proc. of IEEE GLOBECOM*.
- [16] 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*.
- [17] Wonil Roh, Ji Y. Seol, JeongHo Park, Byunghwan Lee, Jaekon Lee, Yungsoo Kim, Jaewon Cho, Kyungwhoon Cheun, and Farshid Aryanfar. 2014. Millimeter-Wave Beamforming as an Enabling Technology for 5G Cellular Communications: Theoretical Feasibility and Prototype Results. *IEEE Communications Magazine* 52, 2 (2014), 106–113.
- [18] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Owen Torres, Daniel Uvaydov, Josep Miquel Jornet, Edward Knightly, Dimitrios Koutsonikolas, Dimitris Pados, and Zhi Sun. 2017. X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays. In *Proc. of ACM WiNTECH*.
- [19] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. 2017. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices. In *Proc. of ACM CoNEXT*.
- [20] Shu Sun, Theodore S. Rappaport, Robert W. Heath, Andrew Nix, and Sundeeep Rangan. 2014. MIMO for Millimeter-Wave Wireless Communications: Beamforming, Spatial Multiplexing, or Both? *IEEE Communications Magazine* 52, 12 (2014), 110–121.
- [21] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom*.
- [22] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A link-Level Profiling. In *Proc. of ACM SIGMETRICS*.
- [23] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of NSDI*.
- [24] Teng Wei, Anfu Zhou, and Xinyu Zhang. 2017. Facilitating Robust 60 GHz Network Deployment by Sensing Ambient Reflectors. In *Proc. of NSDI*.
- [25] Carl Wong, R. Klukas, and G. Messier. 2008. Using WLAN Infrastructure for Angle-of-Arrival Indoor User Location. In *Proc. of IEEE VTC*.
- [26] Jie Xiong and Kyle Jamieson. 2012. Towards Fine-grained Radio-based Indoor Location. In *Proc. of ACM HotMobile*.
- [27] Jie Xiong and Kyle Jamieson. 2013. ArrayTrack: A Fine-grained Indoor Location System. In *Proc. of NSDI*.
- [28] 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 (April 2002), 620–630.
- [29] Jian Zhang, Yahong R. Zheng, Chengshan Xiao, and Khaled B. Letaief. 2010. Channel Equalization and Symbol Detection for Single-Carrier MIMO Systems in the Presence of Multiple Carrier Frequency Offsets. *IEEE Transactions on Vehicular Technology* 59, 4 (2010), 2021–2030.
- [30] 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*.