

Scaling Multi-User mmWave WLANs: The Case for Concurrent Uplink Transmissions on a Single RF Chain

Keerthi Priya Dasala¹, Josep Miquel Jornet², *Senior Member, IEEE*, and Edward W. Knightly³, *Fellow, IEEE*

Abstract—Today’s mmWave WLANs can realize simultaneous multi-user multi-stream transmission solely on the downlink. In this paper, we present *Uplink Multi-user Beamforming on single RF chain AP (UMBRA)*, a novel framework for supporting multi-stream multi-user uplink transmissions via a single RF chain. We design multi-user overlaid constellations and multi-user receiver mechanisms to enable concurrent time-triggered uplink multi-user transmissions received on a single RF chain AP. We devise exemplary beam selection policies to jointly adapt beams at users and the AP for targeting aggregate rate maximization without increasing training requirements compared to single-user systems. We implement the key components of *UMBRA* using a programmable WLAN testbed using software-defined radios and commercial 60-GHz transceivers and collect over-the-air measurements using phased-array antennas and horn antennas with varying beamwidth. We find that in comparison to single-user transmissions, *UMBRA* achieves more than $1.45\times$ improvement in aggregate rate regardless of the choice of the user group, geometric separation, receiver beamwidth, and also under LOS blockage.

Index Terms—MU-MIMO, 60 GHz, beam selection, analog and digital beamforming, IEEE 802.11ad, IEEE 802.11ay.

I. INTRODUCTION

MILLIMETER wave WLANs can realize *downlink* multi-user transmission by exploiting directional transmission and physical separation of clients [1]. In contrast, simultaneous *uplink* transmission must address the inevitable interference from clients directing their transmissions towards a common point in space, namely, the receiving Access Point (AP). In this paper, we design and experimentally evaluate *UMBRA*, *Uplink Multi-User Beamforming via a Single RF chain AP*, the first system for multi-user mmWave uplink. In particular, we make the following contributions.

First, we propose a 60 GHz WLAN architecture for a multi-user uplink using only a single Radio Frequency (RF)

Manuscript received 16 January 2022; revised 26 October 2022; accepted 29 December 2022; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor K. Chowdhury. This work was supported in part by Cisco; in part by Intel; and in part by NSF under Grant CNS-1955075, Grant CNS-1923782, Grant CNS-1824529, and Grant CNS-1801857. (Corresponding author: Keerthi Priya Dasala.)

Keerthi Priya Dasala and Edward W. Knightly are with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005 USA (e-mail: keerthi.dasala@rice.edu).

José Miquel Jornet is with the Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115 USA.

Digital Object Identifier 10.1109/TNET.2023.3239438

chain at the AP. That is, the AP has a phased array for receive and transmit beamforming, but does not have Multiple Input Multiple Output (MIMO).¹ Transmission is initiated by the AP with a downlink trigger frame as employed by standards such as IEEE 802.11ax [2]. After the trigger frame, we stagger uplink client PHY preambles so that the AP can obtain a “clean” (interference-free) channel measurement for each user to be used during decoding. Subsequently, the triggered clients transmit their uplink data frames in parallel. While these frames are temporally aligned by the trigger, they arrive at the AP offset by the clients’ different propagation delays. Consequently, we design *UMBRA* to enable asynchronous decoding, i.e., we do not require symbol-level synchronization. To realize this feature, we design *Scalable Multi-User Overlaid Constellations*. In particular, we overlay Amplitude and Phase Shift Keying (APSK) constellations such that each user is assigned one or more consecutive rings and groups of rings are assigned to users such that the highest SNR user has the outermost ring. With sufficient SNR spread among the rings, the AP can then successively decode one user at a time starting with the highest SNR user, i.e., we enable the use of Successive Interference Cancellation (SIC) decoding [3]. We show that with this multi-user overlay strategy, at each stage of stream separation, the current symbol being decoded on a particular nearly-constant amplitude constellation ring is resilient to the detrimental impact of phase noise impairment caused by interference from other streams which are being received at significantly different amplitudes. Moreover, we design a Carrier Frequency Offset (CFO) compensation method comprised of pre-compensation and iterative correction. This allows the AP to apply the offset of each user to the composite stream at each interference cancellation iteration, while treating the rest of the signals as noise. When decoding the signal from one user, the AP employs an interference alleviation filter specifically designed from the training preamble of that user to cancel the interference and recover the signal.

Next, we show how to use beam selection to attain the desired ring separation and hence, Signal to Noise Ratio (SNR) separation, at the access point in order to realize a high aggregate rate. We show how both AP and client beams can

¹While *UMBRA* can be extended to the MIMO case with multiple RF chains at the AP, for ease of exposition, we focus on a system with a single RF chain.

be re-steered to maximize the aggregate multi-user rate using the outcome of single user training, i.e., without transmission of additional training frames.

For *UMBRA*, beam selection is constrained in multiple ways: (i) the AP must use a single receive beam from its codebook to receive the superposition of the data streams without MIMO processing due to having only a single RF chain. Nonetheless, both the AP and all clients can steer their beams in directions that yield maximum aggregate rate. Unfortunately, (ii) the search space for considering all such combinations may render joint optimization impractical, as it in principle requires exhaustive testing of all AP-user beam combinations. Therefore, we present three complementary beam selection policies with different computational requirements, spanning from testing all AP and user beam combinations, to only letting the AP re-steer its beam. We first introduce a policy with an objective of maximizing the aggregate rate for a target set of users without regard to computational overhead (*All-Steer*). In particular, with *All-Steer*, all users and the AP jointly steer their beams by considering all possible beam directions (codebook entries) to find the combination that yields the maximum aggregate rate. At the other extreme, we propose *AP-Steers* which allows only the AP to re-steer its beam and, to reduce overhead, requires that clients use the beams that they have previously selected for single-user transmission. Finally, we study the *Freeze-Subset* policy which establishes a balance between *All-Steer* and *AP-Steers* by exploiting the key source of rate gains for *UMBRA*, namely, SNR spread. Thus, unlike *All-Steer* which allows all users to re-steer, *Freeze-Subset* freezes some users' beams and allows other users to re-steer if more SNR spread is needed.

Finally, we implement the key components of *UMBRA* using *X60*, a programmable testbed for wide-band 60 GHz WLANs with electronically-steerable phased arrays [4]. Moreover, we also deploy a *WARP-60* testbed using a steerable 60 GHz RF-frontend combined with the software defined radio platform *WARP* [5]. This platform utilizes mechanically steerable horn antenna with configurable beamwidths. Using these two testbeds, we perform over 67,000 over-the-air measurements and subsequently perform trace-driven emulations to study *UMBRA*. Our experiments demonstrate that with beam re-steering at the AP and at least one of the grouped users, *UMBRA* yields aggregate rate gains of up to $1.45\times$ over *Single User* irrespective of the choice of the user group and the geometric separation between them. We study the critical role of SNR spread among concurrently transmitting users as determined by multi-user beam selection, and show how it helps in limiting inter-user interference and leads to increased SINR for each user and increased gains of *UMBRA*. Second, we study the case that a Line-Of-Sight (LOS) path is unavailable due to blockage, and hence the *UMBRA* AP must connect via a reflected path. While extensive studies [1], [6], [7] prove that Non-LOS (NLOS) paths offer reduced SNR due to a lack of LOS path and could be detrimental to multi-user performance, we have shown that *UMBRA* exploits these NLOS paths to the advantage of increasing performance, with our multi-user policies achieving more than $1.26\times$ respective multi-user gains over *Single User* under blockage.

Next, we explore scaling the number of simultaneous users from 2 to 5 clients. We experimentally show that when four simultaneous users perform beamforming using *UMBRA* results in $1.36\times$ aggregate rate improvement over *Single User*. Lastly, we explore how increasing receive beamwidth at the AP can increase aggregate rate for *UMBRA* via increased SNR diversity and beam selection efficiency.

The remainder of this article is organized as follows: Section II provides *UMBRA*'s multi-user architecture and the timeline for the data transmission, scalable multi-user overlaid constellations and subsequent stream demultiplexing. Section III presents a suite of beam selection policies that achieve the objective of *UMBRA* in aggregate rate maximization yet with systems constraints of single receive beam and computational overhead. Section IV presents our implementation setup which includes the testbeds and measurement methodologies used for data collection. Section V describes experimental investigations of the performance of *UMBRA* and factors including geometric separation of grouped receivers, receiver beamwidth adaptation, and user group size, all compared to single-user transmission as a baseline. Finally, we discuss related work in Section VI and conclude the paper in Section VII.

II. UMBRA FRAMEWORK

In this section, we describe the key components of *UMBRA* for realizing uplink multi-user multi-stream transmission on a single RF chain AP.

A. System Architecture and Timeline

Here, we describe a 60 GHz WLAN architecture that supports simultaneous uplink users and streams exceeding the number of RF chains at the receiver. We describe a special case of a single RF chain system with simultaneous reception from more than one user at a time.

Figure 1 shows the system model that coordinates and supports multi-stream transmissions on the uplink. As shown, each user requires only a single RF chain driving a set of phase shifters, each controlling the phase of one antenna element, to be able to independently beam steer a single data stream. The application of different phase delays to the different antenna elements generates a directional beam. The set of possible beams (or equivalently phase delays) is fixed and is typically chosen from a predefined codebook. The beam width and beam direction are selected from one of these codebook entries for each stream transmission. The AP has a single RF chain and is constrained to receive data streams from multiple users on a single receive beam. The AP must also select its beam from a pre-defined codebook.

UMBRA's high-level timeline is depicted in Fig. 2. As shown, prior to transmission we consider that beam training has occurred. The precise test transmissions and feedback are described in Section III. Because training need not immediately precede the transmission, we depict a potential discontinuity in the timeline. The target user group selection precedes data transmission in *UMBRA*, and user and beam selection framework provided in [8] can serve as a foundation for selection of target user group with certain design

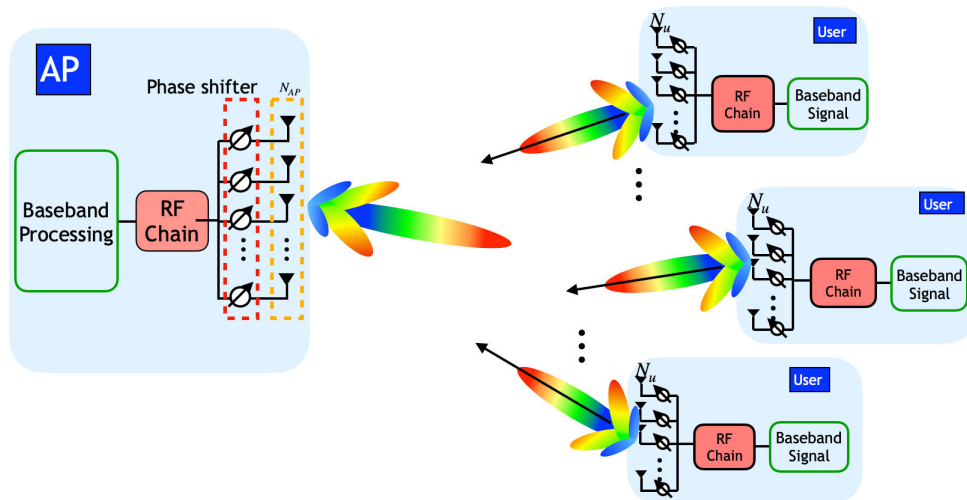


Fig. 1. AP with single RF chain system supporting multiple users on the uplink.

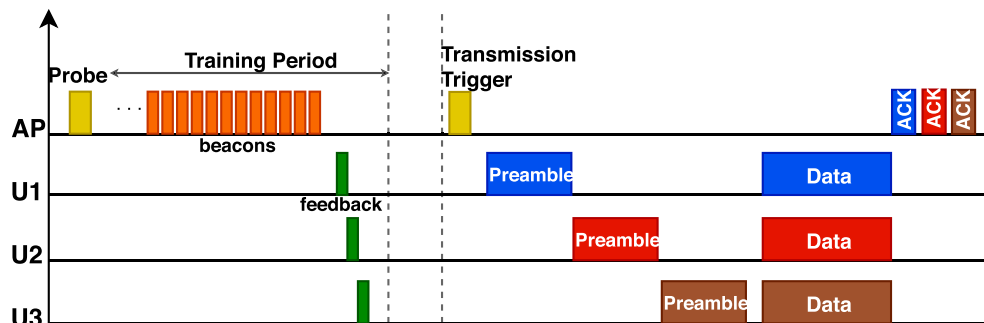


Fig. 2. Different stages of UMBRA timeline model.

trade-offs in data rate, training overhead, and computational overhead.

Uplink transmission in *UMBRA* begins with a group announcement trigger when the AP wins contention to serve a target set of users. The receipt of this trigger serves as a coarse initial time synchronization. The trigger identifies the users to be served, the beams that they should use, and the order for preamble staggering. As shown (not to scale), the users transmit staggered preambles. This allows the AP to receive each preamble corresponding to each stream without interferences which in turn enables estimation of CFO, symbol timing and other channel parameters necessary to decode different streams. Finally, the clients transmit the data in parallel followed by ACKs.

The AP uses its single RF chain to receive a superposition of data streams transmitted from multiple users. In order for the AP to decode these concurrent data streams from the composite signal, we design a SIC framework at the AP to train its receiver to perform stream separation. Typically, SIC consists of an iterative receiver, i.e., it decodes one stream at a time, whereas the remaining streams are considered interference. Then, after decoding a particular stream, its contribution to the overall signal is reconstructed and removed from it. This procedure continues until all streams are decoded.

B. Scalable Multi-User Overlaid Constellations

We present *Multi-User Overlaid APSK Constellations* for *UMBRA* in which an APSK constellation with a varied parameter configuration represents the data stream corresponding to each transmitting user and the AP receives an *overlay* of these constellations from multiple users. In particular, we design an overlaid constellation which defines the configuration of the APSK constellation (and hence MCS level) for each user. The AP will assign one or more APSK rings, each with a specified number of symbols per ring, to each client. At the receiver, this will yield an overlaid constellation yielding comprising all rings. Thus, to enable efficient decoding, the AP must assign rings such that different user's rings are sufficiently separated. While this could in principle be achieved with client power control, we instead use beam steering and user selection to ensure that sufficient SNR spread is available. Moreover, our design allows the phase orientation of each ring to be different, reducing synchronization requirements, and enabling asynchronous decoding.

More specifically, we assign different users to different sets of APSK constellations with an SNR ordering such that increasing SNR clients are assigned constellations with increasing radius. This is illustrated in Fig. 3(a) where the users are arranged in increasing order of SNRs. In the example, user U_1 has lowest SNR and is assigned BPSK (here viewed as

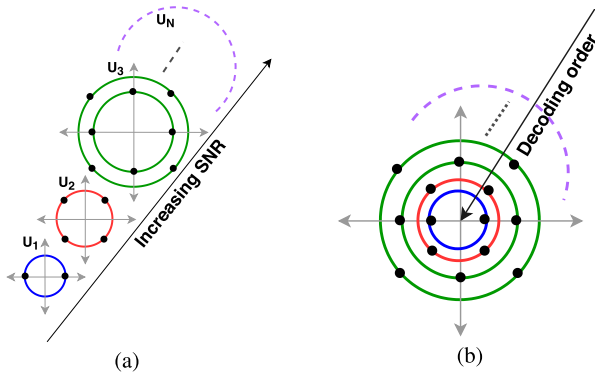


Fig. 3. (a) Multi-user Overlaid APSK constellations showing users with constellation rings of increasing radius. (b) SNR based SIC decoding order.

a one-ring constellation or 2-APSK). Likewise, U_3 has higher SNR than U_1 and is assigned an 8-APSK constellation via two 4-APSK rings. This design of increasing SNR ordering of the users as shown in Fig. 3(b) presents *UMBRA* the ability to better separate the user streams (rings) using ordered SIC decoding. Namely, the AP will first decode the stronger and more robust user with higher SNR, therefore reducing the number of errors propagated from one stage to another. Furthermore, *UMBRA* will never underperform *Single User* in cases where there is no multi-user gain. For example, if users have negligible SNR spread, then *UMBRA* resorts to transmitting to those users separately using the maximum SNR beam pair found from the initial beam training. This is equivalent to the *Single User* solution and corresponds to the worst-case scenario for *UMBRA*.

The performance of stream demultiplexing of *UMBRA* using SIC is highly dependent on two factors: (i) The relative signal strength (SINR) of current user stream, i.e., difference between the power of the signal being decoded and the interference plus noise components of the composite signal. (ii) As the users and the AP do not share a common clock, the transmission trigger may not always guarantee a fully synchronized transmission thus leading to asynchronous reception of multiple user data streams. The impact of both these factors would manifest as phase noise and non-linear distortion and causes rotation of the current symbol being decoded on each ring. However, in our design, the distortion of the signal of one user constellation ring due to inter-stream interference from other user constellation rings tends to have less effect on the Euclidean distance between the symbols compared to an alternate assignment such as QAM. Namely, if we overlaid QAM, the outer constellation points would be more sensitive to phase variations and nonlinear distortions associated with inter-stream interference and various sources of RF impairments. In contrast, we assign different users to concentric rings of constant amplitude to better tolerate the same amount of phase variations. This is illustrated with an example in Fig. 4 which shows the maximum angle error tolerance for a symbol '0011' transmitted using 16-QAM and 16-APSK would be $\pm 16.9^\circ$ and $\pm 22.5^\circ$ respectively.

C. Multi-User Stream Demultiplexing

While the AP's trigger coordinates the multi-user uplink transmissions and provides coarse-grained synchronization,

there are additional offsets due to different propagation delays and CFOs. Here, we present how *UMBRA* can separate the multi-user overlaid APSK constellations with the SIC receiver enhanced with the following capabilities (i) multi-user CFO correction comprised of a combination of pre-compensation and iterative correction, and (ii) interference filtered stream separation.

More specifically, after CFO estimation for all streams and once the AP determines which stream to decode (in SNR order), the AP applies CFO correction of the current stream to the entire composite signal. Next the SIC receiver decodes the current stream using an interference alleviation filter which is constructed from user-specific preambles and aims to recover this signal and cancel the unknown interference from other users. After removing the decoded signal component from the original composite signal, the applied CFO correction is removed from the composite signal. This process is repeated until all streams are decoded. More formally, in what follows we focus on decoding the signal from user i .

Multi-User CFO correction: Let $s_i(n)$ be the time domain symbol n transmitted by the i th user, and $h_i(n)$ be the channel impulse response of the i th channel. The signal of the i th user after passing through channel is $x_i(n) = s_i(n) * h_i(n)$. The composite received baseband signal is given as

$$y(n) = \sum_{i=1}^N x_i(n) e^{j2\pi\Delta f_i n} + w(n) \quad (1)$$

where N is the number of users transmitting simultaneously to the AP, and Δf_i denotes the i th user's CFO normalized by symbol period. Assuming that the AP decodes user i , it will correct CFO in time domain by multiplying $y(n)$ with the term $e^{-j2\pi\Delta f_i n}$. After an iteration, the composite signal after correcting the i th user is given by

$$\begin{aligned} c(n) &= y(n) e^{-j2\pi\Delta f_i n} \\ &= x_i(n)(1) + x_k(n) e^{j2\pi(\Delta f_k - \Delta f_i)n} + \dots \\ &\quad + x_N(n) e^{j2\pi(\Delta f_N - \Delta f_i)n} + w'(n) \end{aligned} \quad (2)$$

The AP can decode the i th user's stream from the composite signal using the filter described next.

Signal Decoding using Interference Alleviating Filter: We model the received baseband signal vector from Equation (2) as

$$\mathbf{C} = \mathbf{H}_i \mathbf{S}_i + \sum_{k \neq i} \mathbf{H}_k \mathbf{I}_k + \mathbf{W} \quad (3)$$

where \mathbf{S}_i is the signal from user i and \mathbf{I}_k is the interfering signal from user k ($k \neq i$).

An interference alleviating filter \mathbf{P}_i exclusively designed from preambles of user i can be employed to decode \mathbf{S}_i . The estimated signal from user i can be denoted as

$$\hat{\mathbf{S}}_i = \mathbf{P}_i^H \mathbf{C}. \quad (4)$$

The optimal filter can be derived by solving the mean squared error (MSE) optimization and is given as

$$\mathbf{P}_i = \mathbb{E}[\mathbf{C}\mathbf{C}^H]^{-1} \mathbb{E}[\mathbf{C}\mathbf{S}_i^H]. \quad (5)$$

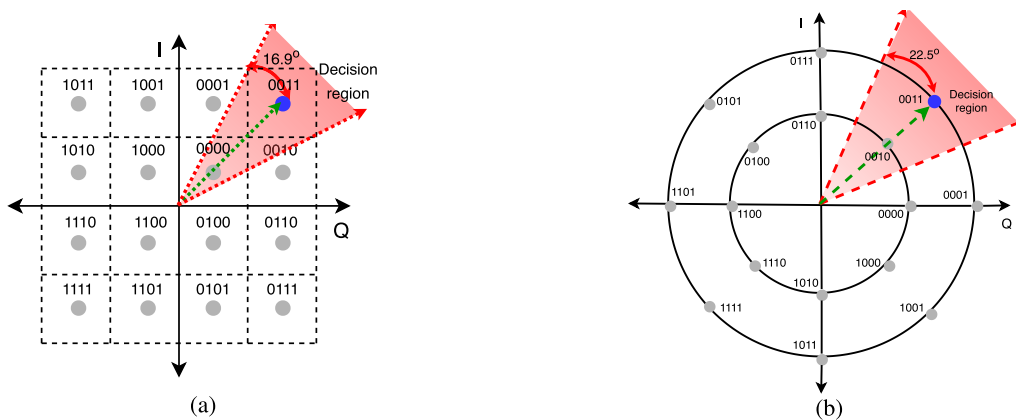


Fig. 4. Comparison of maximum angle error tolerance for symbol '0011' sent using (a) 16-QAM. (b) 16-APSK.

We take advantage of the preamble symbols sent by each user i and estimate $\mathbb{E}[\mathbf{C}\mathbf{C}^H]$ and $\mathbb{E}[\mathbf{C}\mathbf{S}_i^H]$ using statistical averaging operation over the preamble symbols $[\tilde{S}_i(1), \tilde{S}_i(2), \dots, \tilde{S}_i(L)]$ of user i and the received preamble symbols at the AP $[\tilde{C}(1), \tilde{C}(2), \dots, \tilde{C}(L)]$ which also include the interfering signals from other users. This is given as

$$\mathbb{E}[\mathbf{C}\mathbf{C}^H] \leftarrow \frac{1}{L} \sum_{l=1}^L \tilde{\mathbf{C}}(l)\tilde{\mathbf{C}}(l)^H \quad (6)$$

$$\mathbb{E}[\mathbf{C}\mathbf{S}_i^H] \leftarrow \frac{1}{L} \sum_{l=1}^L \tilde{\mathbf{C}}(l)\tilde{\mathbf{S}}_i(l)^H \quad (7)$$

In order to decode the signal from user i at the AP, the filter \mathbf{P}_i is constructed using Equation (5) and then estimate the signal using Equation (4). After the first stream has been decoded, the composite signal is multiplied by $e^{j2\pi\Delta f_i n}$ to remove the CFO component of the i th user's stream. This process is continued for the rest of the users until all streams are decoded.

Lastly, we have found with symbol level simulations that symbol offset due to propagation delay differences has a relatively minor impact on the decoding reliability of *UMBRA* provided user streams are decoded with a sufficient SNR spread. For example, with SNR difference of at least 9 dB, the AP achieves a BER of 10^{-5} to decode the overlay constellations of 4-APSK and 2-APSK streams received with a maximum propagation delay of up to 1 μ s and CFO of 400 Hz. Hence, we do not compensate for this effect.

III. CONSTRAINED BEAM ADAPTATION

Because *UMBRA* decoding is improved with high SNR spread among users, we employ beam steering at the clients and AP to ensure sufficient SNR spread. In our system architecture, the AP is constrained to use a single receive beam from its predefined codebook to receive a superposition of data streams. Moreover, while the AP and all clients have previously been trained for single-user transmission, these beams that maximize the SNR for single user transmission may not be the best beams for multi-user transmission. Thus, some beams may need to be re-steered to improve the aggregate multi-user rate. In this section, we describe *UMBRA*'s AP and client beam selection with a joint focus on aggregate rate and overhead. We describe policies in which some or all users are prohibited from re-steering their beams in order to not incur computation overhead.

A. Training and Computational Overhead

In single-user 802.11ad beam training, the sender sequentially transmits on each of its sectors (codebook entries) and the receiver subsequently identifies and feeds back the ID of the sector that yields the highest SNR. For example, the AP first sends training frames sequentially on all C_{AP} of its codebook entries (beams) while the user employs quasi-omni reception to find the highest SNR transmit beam from the AP. The AP's highest SNR sector is identified by the user and fed back in a control message. Conversely, user u sweeps through its C_u beams while the AP is in quasi-omni receive mode in order to find the user's highest SNR beam. The AP similarly feeds this information back to the user. This training yields the best bi-directional AP-user beam pair via a total of $C_{AP} + C_u$ test transmissions and two feedback messages.

In contrast to solely feeding back the ID of the maximum SNR sector, *UMBRA* requires clients to feedback the SNR of all measured sectors. This enables the AP to re-steer its receive beam in a way that yields the best SNR spread and aggregate throughput. In our design, the AP does *not* need to feed back SNR-sector measurements to the clients: In *UMBRA*, if a client should re-steer to further increase throughput, the AP will notify the client which sector to use in the trigger message. Thus, *UMBRA* does not require additional training compared to 802.11ad, but does require a more rich feedback message from clients to report the results of the training.

We define computational cost as the number of beam combinations that must be compared at the AP. If only the AP re-steers, it must check all of its receive beams and select the beam with the maximum rate, a computation that is $\mathcal{O}(C_{AP})$. If one or more clients are allowed to re-steer, then there are additional computations required at the AP. Below, we introduce policies which vary in their computation cost, but all require the same feedback as above. The computations are required to be performed per-frame (millisecond time scale) vs. per-symbol (nanosecond time scale). Thus, the total computations to find the final beam configuration and MCS assignment for the policies within our framework are easily implementable with negligible computational overhead.




Fig. 5. Illustrative scenario of constrained beam adaptation using (a) *All-Steer*. (b) *AP-Steers*. (c) *Freeze-Subset*.

B. All-Node Re-Steering for Rate Maximization

Here, we present *All-Steer* as a *UMBRA* beam selection policy that specifies for each uplink transmission, the transmit and receive beams (codebook entry) to be used and the modulation and coding scheme for each data stream at the user. This policy targets to maximize the aggregate rate of a user group without any constraints on computational overhead.

We consider N users selected for uplink transmission. Let $b_j \in C_{AP}$ denote the beam index in the AP codebook and $b_k \in C_u$ denote the beam index in the user's u codebook. The input to the beam selection policy is the training information for each user that comprises the measured signal-to-noise ratio (SNR) for each beam pair $\text{SNR}_u(b_j, b_k)$, computed as the sum of the respective SNRs measured with one node in pseudo-omni reception. The achievable data rate $R_u(b_j, b_k)$ by user u on each beam pair can be expressed as

$$R_u(b_k, b_j) = \text{MCS}(\text{SNR}_u(b_k, b_j)) \quad (8)$$

where $\text{MCS}(\cdot)$ gives the data rate achievable for a particular SNR via the single-user minimum SNR tables.

In *Single User* transmission, the AP maximizes the single-user rate by choosing the beam pair having maximum SNR. We denote this *Single User* beam pair for user u as $(b_{k,u}^{\max}, b_{j,u}^{\max})$ and the corresponding *Single User* rate as R_u^{\max} and it is given by

$$(b_{k,u}^{\max}, b_{j,u}^{\max}) = \arg \max_{(b_k, b_j)} \text{SNR}_u(b_k, b_j) \quad (9)$$

$$R_u^{\max} = \text{MCS}(\text{SNR}_u((b_{k,u}^{\max}, b_{j,u}^{\max}))). \quad (10)$$

Let G be the user group to be triggered by the AP. *All-Steer* must determine the best user beams $\{b_{k,u}\}_{u \in G}$ to transmit to a shared receive beam b_j at the AP. Such beams will be those that result in the highest SNR spread and thereby maximum achievable aggregate rate for the receive beam b_j and these beams may not be the best *Single User* beam pair $(b_{k,u}^{\max}, b_{j,u}^{\max})$. The objective of *All-Steer* is as follows

$$(b_j^*, \{b_{k,u}^*\}_{u \in G}) = \arg \max \sum_u R_u(b_k, b_j) \quad (11a)$$

$$\text{s.t. } b_{k,u} \in C_u, u \in G \quad (11b)$$

$$b_j \in C_{AP}. \quad (11c)$$

Equation (11a) optimizes the beam selection to maximize the sum rate by finding the best receive beam at the AP b_j^* that could be shared by all the users in the group G while each user $u \in G$ will be using their best transmit beam $\{b_{k,u}^*\}$ for the selected receive beam. Furthermore, the

maximum aggregate rate resulting from this beam selection is achieved only when sufficient SNR spread is available as this enables efficient decoding with maximal separation among the inter-user rings as discussed in Sec. II-B. The two constraints ensure that the transmit and receive beams are selected from the predefined user and AP codebooks respectively. Fig. 5(a) depicts an example scenario where *All-Steer* jointly optimizes the beam selection by enabling the AP and the two users to simultaneously beam steer to create sufficient SNR difference for successful stream decoding and targets to increase the sum rate if possible. Computationally, with *All-Steer*, the optimal solution of Equation (11) yields an exhaustive search over all possible AP-user beam tuples combinations. Hence, the AP finds the final beam configuration by checking a total of $C_{AP}^G \cdot \prod_{u=1}^G C_u$ distinct beam combinations and then feeds back the final beam IDs to the re-steering users.

C. AP Only Re-Steering

To obtain the maximum aggregate rate using *All-Steer*, we need to exhaustively search every combination of the beams at AP and the target set of users. Unfortunately, implementation of this exhaustive search may not be practical in real scenarios due to the high computational overhead.

Here, we introduce *AP-Steers* as a policy on the other end of the design spectrum. Namely, in order to limit the search overhead, only the AP re-steers while all the users in the group G freeze their best transmit beams $\{b_{k,u}^{\max}\}_{u \in G}$ found during the initial beam training. Namely, the AP picks the top candidate set of receive beams that are optimal in potentially maximizing aggregate rate for multi-user transmission while the users use their best TX beam under all AP receive beams. Subsequently, the AP finds the final receive beam b_j^* by performing a search among all the possible combinations of candidate receive beams at its end. Note that this policy does not introduce additional overhead for user beam selection as the SNR associated with each beam is already available at the AP after initial beam sweeps. Hence, the maximum computational cost at the AP to find the best analog configuration is $\mathcal{O}(C_{AP})$. Fig. 5(b) depicts the *AP-Steers* beam selection mechanism in which the two users use their best transmit beams as the AP re-steers at its end to find the best receive beam with maximum SNR spread and thereby the aggregate rate.

D. Freezing a Subset of Users

Here, we present a final *UMBRA* strategy that represents a balance between *All-Steer* and *AP-Steers*. Namely, we present




Fig. 6. (Left) X60 testbed with mmWave transceiver and phased array antenna. (Middle) X60 Irregular azimuth beam pattern. (Right) WARP-60 testbed with horn antennas.

Freeze-Subset as a policy that significantly reduces the search space of *All-Steer* by exploiting the fact that grouped users with maximum aggregate rate are typically composed of streams having high SNR spread. Furthermore, if all users have the same SNR, *UMBRA* cannot realize a gain over *Single User* as there is no SNR margin between the users to counteract the inter-user interference and improve decoding reliability. Hence, *Freeze-Subset* avoids computing all possible beam combinations by choosing a subset of users from the target user group to re-steer jointly with the AP to find the optimal multi-user beam configuration that targets to maximize the SNR spread and potentially increase aggregate rate. The remaining non re-steering users in the group freeze their transmit beams to share the common receive beam at the AP.

More formally, we define *Freeze-Subset* as follows. For a target set of grouped users G , *Freeze-Subset* first sorts users in decreasing order of their maximum single-user SNR corresponding to their best beam pair $(b_{k,u}^{\max}, b_{j,u}^{\max})$ with the first sorted user having the highest SNR. *Freeze-Subset* begins with an initial “prime user” to re-steer with the AP, while the other users are held to their best transmit beams for the chosen AP receive beam. While any user can be a prime user, we select the user having the highest SNR as the prime user as this user will be decoded first using SIC to have control on the inter-user interference and reduce decoding propagation errors. *Freeze-Subset* iterates the same procedure by searching for other users in the group to re-steer jointly with the AP and can form a higher aggregate rate multi-user transmission with the existing users. Finally, at the end of beam selection process, the AP notifies the re-steering users with their final beam IDs. Let $g \subset |G|$ denote the final outcome of number of users in *Freeze-Subset* that freeze their beams and not re-steering with AP. Namely, $g = |G|$ corresponds to the case of *AP-Steers* where all users have frozen their beams and only AP re-steers and $g = 0$ corresponds to the case of *All-Steer* where none of the users freeze their beams and all re-steer along with the AP. Therefore, computationally, *Freeze-Subset* performs up to $\binom{|G|}{|G|-g} (C_{AP} \times C_u)^{|G|-g}$ tests of aggregate rate to find the final beam configuration among the users in G . Fig. 5(c) depicts an example scenario using *Freeze-Subset* when one of the users i.e., $|G| - g = 1$ perform joint beam steering with the AP to target to increase the aggregate rate. The other user froze its best transmit beam for the chosen receive beam at the AP.

IV. EVALUATION SETUP: TESTBEDS AND OVER-THE-AIR EXPERIMENTS

In this section, we present our implementation of key components of *UMBRA* and collect over-the-air data to evaluate its performance. As our evaluation encompasses scenarios to study the impact of various parameters such as beamwidth, antenna array beam patterns and multi-user capability, we employ two different platforms, X60 and WARP-60 enhanced with a 60 GHz front end.

A. X60 Phased Array Platform

We perform over-the-air experiments with X60, a configurable Software Defined Radio based mmWave platform [4]. X60 features a fully programmable cross-layer architecture for PHY, MAC and Network layers. Fig. 6 (left) shows the X60 platform where each X60 node is built with National Instruments’ (NI) millimeter-wave transceiver system and employs a user configurable SiBeam phased array antenna module with 24 elements, 12 for TX and 12 for RX. The transmit power at each X60 node is 30 dBm EIRP. Communication is established over wide-band 2 GHz channels that can reach multi-gigabit data rates using real-time electronically steerable (switching time of $1\mu s$) TX and RX beams from a predetermined codebook consisting of 25 beams which are spaced roughly 5° apart along the mainlobe direction, thereby covering a sector of -60° to 60° in the azimuthal plane centered around the antenna’s broadside direction. Each beam has a 3 dB bandwidth of 25° to 35° causing each main-lobe to overlap with other neighboring beams. Therefore, it is evident as in Fig. 6 (middle) that X60 provides irregular and imperfect beam patterns with predominant main lobes overlapping and strong side-lobes. We collect channel samples from over-the-air measurements and subsequently perform trace-driven emulation to study *UMBRA*. More details on this methodology are presented in Sec V.

B. WARP-60 Horn Antenna Platform

While X60 enables experiments over wideband channels and a practical phased-array, the beamwidth is limited to 25° . Hence, to evaluate *UMBRA* under changing beamwidth, we integrate a WARP horn antenna platform with a 60 GHz front end into our testbed form from [8] and [9]. We used the MI-WAVE’s WR-15 60 GHz gain horns with a WR-15

waveguide connector which allows us to connect horn antennas with different levels of directivity. The horn antennas' regularly shaped beam patterns can emulate beam forming of a many-antenna phased array. In particular, we use the testbed setup in Fig. 6 (right) which consists of commercial mm-wave transmitter and receiver modules from the VubIQ 60 GHz development system, WARP v1 boards and daughter boards for signal adjustment. These modules can communicate in the 57-64 GHz unlicensed band with up to 1.8 GHz signal bandwidth and can accept and output I/Q baseband signals. To achieve directional beams of varying beamwidth, we use 7° , 20° and 80° horn antennas. For our measurements, we use a 25 dBi gain horn antenna for all the users.. To implement beam-steering, TX-RX nodes (depicted in Fig. 6 (right)) are mounted on commercial motion control setup which enables rotation to sub-degree accuracy. Using this WARP-60 system, we measure the SNR of a point to point transmission and perform numerous measurements varying the user location, antenna beamwidth, and use this data to study the performance of *UMBRA*.

V. EXPERIMENTAL EVALUATION

In this section, we perform over-the-air measurements to evaluate the performance of *UMBRA* and compare to baseline schemes.

A. Impact of Geometry Between Grouped Users

The spread in signal strengths among concurrently transmitting users affects the successful decoding of the composite stream at the AP. More specifically, the relative signal strength of all users as determined by the beam selection can be influenced by geometry of users and this affects the SIC user decoding order by allowing the AP to first decode the stronger and more robust users, therefore, reducing the number of errors propagated from one stage to the other. To demonstrate this, we consider a simplified setting of two users transmitting in the uplink and study both the impact of geometry on the ability of the AP to decode the different user streams. We conduct over-the-air experiments using the X-60 testbed and experimentally explore the multi-user gains of *UMBRA* in comparison to *Single User* transmission scheme.

Setup: We deploy X60 nodes as depicted in the scenario in Fig. 7 which includes the AP and 15 different user locations. The AP is fixed at one corner of the lobby at a height of 1.23 m facing North and all users are at the same height facing South. The presence of windows and metal coating beneath them (not shown) create reflections. For each AP-user position, we collect SNR for all possible 625 (25×25) beam-pair combinations.

Two User uplink with UMBRA: In this experiment, we consider a two-user group for simultaneous uplink transmission and can share a common receive beam at the AP via *UMBRA*. The first user U_1 is fixed, 3.3 m apart from the AP and the second user U_2 can be placed in any of the marked positions 2-15. To study the impact of geometry in terms of distance and angular separation between U_1 and U_2 , the user groups have been divided into three categories as highlighted

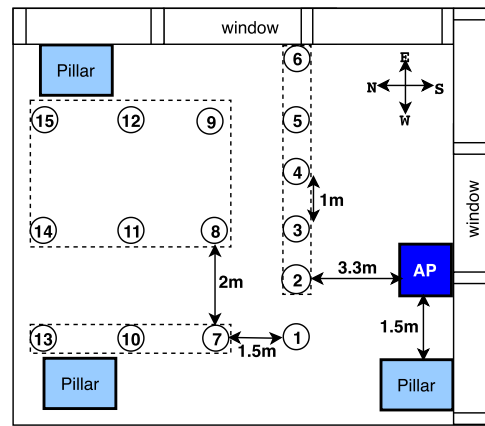


Fig. 7. Experimental floor plan used for measurement of data using X-60 testbed.

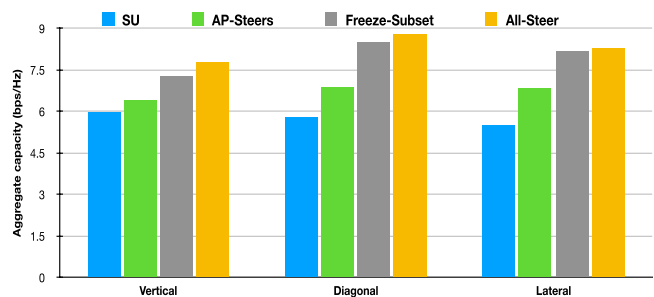


Fig. 8. Average aggregate capacity for different user groups.

in Fig. 7: (i) *vertical* user groups consist of U_2 (7,10,13) vertically separated from U_1 , (ii) *lateral* user groups consist of U_2 (2,3,4,5,6) that has lateral separation with U_1 , and (iii) *diagonal* user groups as U_2 (8,9,11,12,14,15) has diagonal angular separation and increasing distance with U_1 .

Aggregate rate calculation: For all two-user combinations in each category, we find the PHY capacity of a two-user uplink communication as follows: For each user group and potential multi-stream analog beam configuration, we compute the expected SNR at U_1 and U_2 , then the AP selects the MCS index for each user whose corresponding SNR \leq SNR threshold as computed in the APSK two-user minimum SNR decoding tables. The corresponding number of data bits per symbol is the per-user capacity (each stream can use a different MCS but we maintain the same coding rate for the two streams). For comparison, we also implement the *Single User* transmission scheme. *Single User* aggregate capacity is measured by considering that the two users U_1 and U_2 each get half of the air time. The AP uses its best receive beam for U_1 and U_1 uses its best transmit beam. The best TX-RX beam pair is likewise used for U_2 . Unlike the multi-user case, the AP can select different receive beams for U_1 and U_2 . The AP selects the MCS for each user whose corresponding SNR \leq SNR threshold as computed in the APSK minimum SNR decoding table presented in Table I.

Aggregate rate: Figure 8 shows the aggregate capacity obtained by the three user group categories under different beam selection policies. First, the *Single User* scheme achieves an average capacity of 5.7 bps/Hz across all user

TABLE I
SINGLE USER MCS TABLE

Modulation Index	Modulation	Coding Rate	No. of Rings	No. of Symbols per Ring	Minimum SNR Threshold(dB)
1	2-APSK	1/2	1	2	1
2	2-APSK	5/8	1	2	2.5
3	2-APSK	3/4	1	2	3.5
4	4-APSK	1/2	1	4	3.65
5	2-APSK	13/16	1	2	4.5
6	4-APSK	5/8	1	4	4.8
7	4-APSK	3/4	1	4	5.8
8	4-APSK	13/16	1	4	7
9	8-APSK	1/2	2	(4,4)	9.4
10	8-APSK	5/8	2	(4,4)	11
11	8-APSK	3/4	2	(4,4)	12.2
12	8-APSK	13/16	2	(4,4)	13.1
13	16-APSK	1/2	2	(12,4)	13.5
14	16-APSK	5/8	2	(12,4)	12.8
15	16-APSK	3/4	2	(12,4)	13.4
16	16-APSK	13/16	2	(12,4)	14.7
17	32-APSK	1/2	3	(16,12,4)	16.6
18	32-APSK	5/8	3	(16,12,4)	17
19	32-APSK	3/4	3	(16,12,4)	18.2
20	32-APSK	13/16	3	(16,12,4)	19.8
21	64-APSK	1/2	4	(28,20,12,4)	21.1
22	64-APSK	5/8	4	(28,20,12,4)	23.4
23	64-APSK	3/4	4	(28,20,12,4)	24.8
24	64-APSK	13/16	4	(28,20,12,4)	26.6

group categories with the lowest capacity obtained by *lateral* groups due to the impact of both angular separation and distance leading to SNR loss at U_2 and hence decrease in contribution to the aggregate capacity; while the highest capacity is obtained by *vertical* groups in which both the grouped users have high SNR (best MCS) links. Thus the performance of *Single User* scheme depends on the geometry of transmitting users.

Second, *AP-Steers* has marginal gains over *Single User* scheme for *vertical* user group. In particular, as the inter-user distance increases, due to the presence of strong reflection from sidewalls and overlapping beams, the high SNR at both the users remains steady over a larger range of beam steering directions before dropping below 0 dB. The users use the beam 0 which is the best TX beam for both the users. It was observed that for this beam, the maximum achievable SNR diminishes only for higher beam steering angles on either side of central beam at the AP. This is a consequence of the non-uniform angular spread of beams and diminishing directivity gain of beam indices only farther from central beam, a limitation of practical phased array antennas. Therefore the optimal beam selected by the AP results in high SNRs at these users and thereby excessive inter-user interference. The rates improve for *diagonal* and *lateral* user groups in which with increasing inter-node distance, U_2 experiences SNR degradation for almost all beam steering directions chosen at the AP. Moreover, angular separation has a pronounced impact than that of distance and thus results in SNR loss. For example, user 12 uses the beam index 0 for transmission, the SNR is low across all beam steering angles and the highest SNR is achieved for receiver angle 30° . From decoding perspective, this SNR loss translates to the effect of having low inter-user interference on U_1 and thereby increasing the aggregate rate in these user groups.

Third, *Freeze-Subset* achieves $1.22\times$, $1.47\times$ and $1.49\times$ multi-user rate gains over *Single User* scheme for *vertical*,

diagonal and *lateral* user groups respectively and has performance close to 96% of the *All-Steer* policy across all user groups despite its significantly reduced search space. This is because the AP and one of the users re-steer jointly to discover a multi-user beam. This policy exploits the combined effect of the non line of sight (NLOS) paths due to usage of wide beamwidths, non-uniform angular spread of beam patterns and diminishing directivity gain of beams for all angles of the user to the advantage of increasing the SNR spread and therefore results in increased sum capacity. Moreover, the rate gains also highlight that at-least one of the users should re-steer with the AP to find the beams that efficiently tackle the inter-user interference and potentially increase the sum capacity.

Finally, *All-Steer* achieves approximately $1.3\times$, $1.52\times$ and $1.5\times$ multi-user rate gains over *Single User* scheme for *vertical*, *diagonal* and *lateral* user groups respectively. This is attributed to the fact that with large ($\sim 30^\circ$) beamwidth and presence of strong side-lobes giving rise to more flexibility of choosing beams and with AP and the two users jointly re-steering, this policy creates huge number of beam-sharing opportunities such that majority of the beams have highest potential to increase the sum capacity. Theoretically, *All-Steer* should achieve close to $2\times$ gains over *Single User*; however, this does not hold true for every user group in the setup. This is because we find in the measurements that the SNR along the beams of NLOS paths is typically lower than the LOS paths. Hence, even if the inter-user interference is efficiently managed by choosing the best beams (mix of both NLOS and LOS) at AP and users, the aggregate capacity of two user group might not obtain $2\times$ gain over *Single User*. Although this policy requires exhaustive search to find the optimal multi-user beam, nonetheless it offers aggregate capacity ≤ 8.8 bps/Hz and achieves more than $1.35\times$ gain over *Single User* across all user groups indicating that sum capacity increases with increasing number of simultaneous beam steering nodes.

Findings: For any choice of a two-user group and irrespective of user separation, the rate gains over Single User can be controlled through optimal beam steering at AP and at least one of the users. Our Freeze-Subset policy achieves 96% gains in comparison to the optimal sum-rate achieved by All-Steer strategy indicating that at least one out of two users along with the AP should be beam steered in multi-user optimal sense for increasing multiplexing gains.

B. Role of SNR Difference Between Grouped Users

Here, we explore how with the advantage of performing beam selection using *UMBRA*, we can create sufficient SNR diversity among the users that can enable the AP to successfully separate the user streams and manage inter-stream interference. We use the same node deployment as in Fig. 7.

Figure 9 shows the SNR difference obtained by all the multi-user beam selection policies under study. As shown, the beams selected with *AP-Steers* for *vertical* groups resulted in a negligible SNR difference which may be insufficient to separate and successfully decode the composite user stream at the AP. Moreover, a very slight improvement in SNR difference arises from *diagonal* and *lateral* user groups having angular separation and distance. However, the improvement is

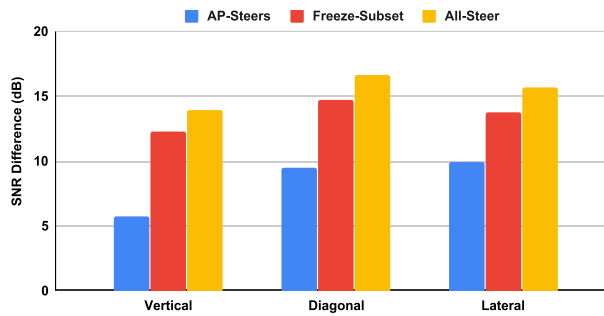


Fig. 9. Received SNR difference for different user groups.

not high enough that could lead to a significant rate increase as seen in Fig. 8. The reason is that sometimes the receive beam selected by the AP has a powerful sidelobe pointing towards U_2 which results in interference power from the U_2 being equal to the signal power and thus the SINR at U_1 is low even with high angular separation at U_2 .

With *Freeze-Subset* there is an increase in the SNR difference and thus higher sum rate over *AP-Steers* and *Single User* schemes. Here the joint beam steering at the AP and one of the users offers more flexibility to choose beams that are optimal in multi-user setting. We observe that as the achievable SNR difference increases, the SINR at U_1 increases with increasing separation between U_1 and U_2 thus causing lower inter-user interference.

All-Steer checks all the possible beam combinations and efficiently exploits the SNR disparity among the users. Allowing all the nodes to re-steer resulted in reduced inter-user interference as these nodes jointly discover the beam that provides greater SINR boost and consequently capacity boost even in case of high SNR links between the users.

Findings: In order to efficiently deal with inter-user interference which becomes a rate-limiting factor, selecting beams optimal in multi-user sense becomes necessary as this controls the increase in the SNR difference (increase in SINR at each user) and therefore the potential increase in sum rate.

C. LOS Vs. NLOS Connectivity

Here we consider scenario in which LOS path is unavailable due to blockage and therefore the AP must connect solely via a reflected path.

Setup: We adopt the same node deployment setup as in Fig. 7 except that we block the LOS path from the AP to each user. Consequently, the AP finds a beam from the reflected paths (e.g., off of a table or wall) available in the environment. We measure the received SNR for all 25×25 beam-pair combinations under two scenarios: (i) the AP has LOS path to the client; (ii) LOS path is blocked with a table as shown in Fig. 10. Note that we study the gains of *UMBRA* under a *worst-case* scenario when all the users have LOS blockage due to environment or nodal mobility. 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.

SNR Heatmaps: Figure 11 presents the SNR heatmaps for user location indices 1, 2, 3 and 4 under LOS connectivity (first



Fig. 10. Experimental setup showing LOS blockage with a wooden table.

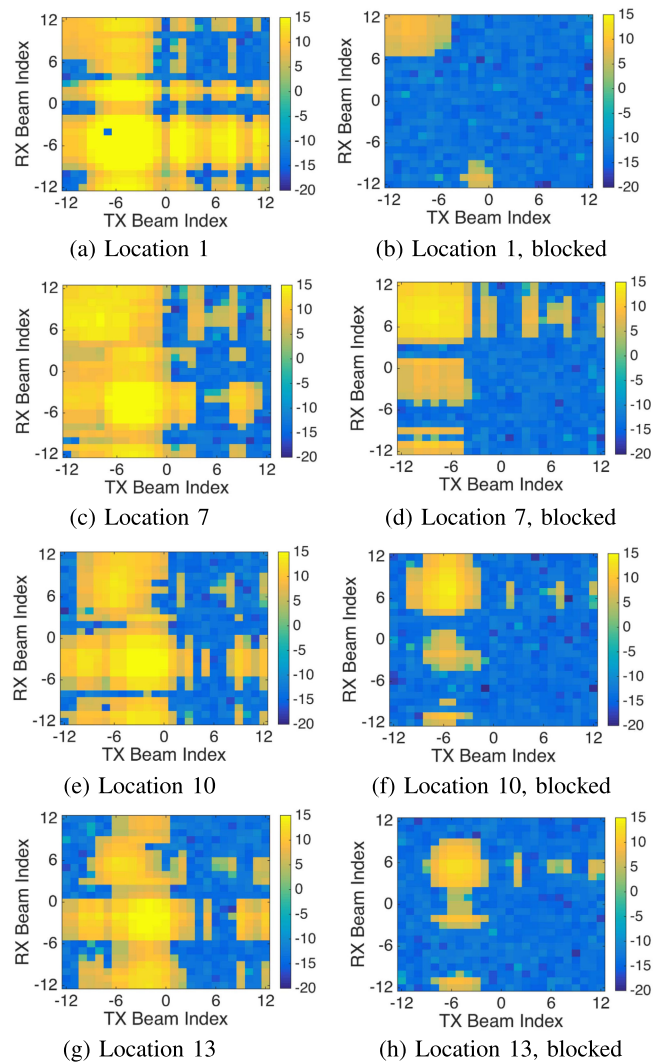


Fig. 11. SNR heatmaps for all beam pair combinations for user location indices 1, 7, 10 and 13 in the lobby under LOS connectivity (left column) and blockage (right column).

column) and blockage (second column). We do not show the SNR heatmap for all locations due to space limit. The SNR range as seen in the heatmaps is between -20 dB to 15 dB with yellow colored regions indicating beam pairs above 10 dB whereas blue regions indicate beam pairs with negative SNR.

We observe several beam-pairs providing above 10 dB SNR that corresponds to 1 Gbps data rate in our platform. The received SNR for each beam-pair is dependent on the physical

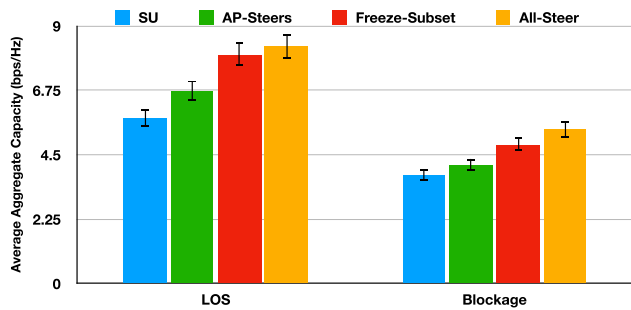


Fig. 12. Achievable Aggregate Rate under LOS Connectivity and Blockage.

paths and the directivity gain along them. Imperfect beam patterns can cause the LOS/NLOS paths to be captured by multiple beams albeit with different directivity gains. The beam pair with the highest SNR corresponds to the physical LOS path in Fig. 11(a), 11(c), 11(e) and 11(h). For instance, the beam index -5 at TX and -5 at RX provides the highest directivity gain along the LOS path at position 1. Fig. 11(a) confirms that while beam-pair $(-5, -5)$ is within the high SNR region but due to overlap between neighboring beams, multiple beams include the LOS paths and we see a cluster of high SNR beam pairs around $(-5, -5)$. Under LOS blockage, the yellow region corresponding to LOS component disappears as seen in the second column of Fig. 11. This conforms that the high SNR region has experienced significant SNR reduction after blockage. Lastly, the highest SNR region after the blockage achieves similar SNR under LOS conditions i.e., LOS blockage has not degraded their SNR; indicating that the irregularities and imperfections as well as the side lobes present in the beam pattern must be capturing a reflected path at these beam pairs.

Aggregate Rate: Figure 12 depicts the achievable aggregate rate of all multi-user schemes of *UMBRA* in comparison with *Single User* under scenarios where LOS path is available and unavailable due to blockage. For each *UMBRA* policy, we average the performance over all two user groups in the setup.

First, as expected the aggregate rate achieved by *Single User* is higher in LOS case in comparison to the case with blockage. As discussed earlier, while some of the users in the setup as seen in Fig. 11(b), 11(d), 11(f) and 11(g) could be co-located and may see strong reflected paths (i.e., off pillars or windows) from the AP, not all users have the benefit of exploiting the reflected paths and therefore achieve degraded SNR under blockage. Consequently, these users get reduced per-user rate leading to nearly 34% reduction in overall aggregate rate due to absence of LOS connectivity. This would not be an issue in LOS scenario, because the received signal strength is relatively higher than any NLOS paths such that the beam selection performed by *Single User* under LOS connectivity are mostly LOS paths with significantly higher received SNR and therefore results in higher achievable aggregate rate.

Second, as observed in Fig. 12, with blockage, *Freeze-Subset*, and *All-Steer* achieve about 28% and 43% multi-user gains over *Single User* scheme. By design, these both schemes select analog beam configuration that results

in at least one transmit beam at each of the user that could share a receive beam at the AP while maximizing the SINR spread and potentially increase the aggregate rate. In the NLOS scenario, because the received signal strength is lower, there is less residual tolerance for inter-user interference. Therefore, as evident from the performance of these both schemes, it is important to apply a multi-user strategy like *AP-Steers* or *All-Steer* which selects beams by allowing users to perform simultaneous beam re-steering to minimize inter-user interference and boost the achievable aggregate rate in such scenarios.

Unfortunately *AP-Steers* policy yields only a marginal improvement over *Single User* under blockage. Considering the lower SNR of NLOS users and decrease in overlap of SNR regions across most of the users, it becomes increasingly unlikely to find a common rate maximizing receive beam when only AP re steers and forgoes user side re steering by the virtue of its design and hence the performance of *AP-Steers* is close to *Single User*.

Finding. With blockages in the environment, beam selection by *UMBRA* effectively exploits the resulting lower SNR inherent to the NLOS paths to the advantage of providing sufficient SINR spread and increasing multi-user gains, whereas transmission to the same ungrouped users using beam separation based policies like *Single User* perform poorly in NLOS environments in which the dominant LOS path is blocked.

D. Impact of Increased Beamwidth

In this experiment, we illustrate the impact of increasing beamwidth on the performance of *UMBRA*. Narrow beamwidth achieves maximum signal strength due its higher directivity gain. However, in case of mobility, there is significant degradation in link strength and leads to increased overhead due to high frequency of beam sounding. In contrast, wider beams offer lower link budget and can reduce SNR and data rate but provide greater resilience to mobility and sufficient signal strength across a larger spread of TX-RX relative angles. This implies that the presence of reflected paths in addition to the LOS paths, makes the wider beamwidth much resilient to blockage or misalignment while the signal spread of narrower beamwidths is not sufficient to exploit these additional paths. Therefore, wider beamwidth can also exploit multiple paths in addition to a much wider reception signal along the LOS paths, thereby providing better beam selection possibilities and yielding more opportunities to achieve maximum aggregate rate via *UMBRA*. To explore this beamwidth-signal coverage tradeoffs, we employ the WARP-60 testbed which generates directional beams of varying beamwidth using horn antennas.

Setup: We consider the experimental floor plan as shown in Fig. 13. We fix the AP location at one end of the conference table. We place the users in 10 different positions. We select the 20° horn for the user's receive antenna to meet the size and power constraints in the mobile user. For each user position, we perform a 360° sweep of the AP in steps of 5° and at each point of AP's sweep, we take RMS baseband measurements to estimate the SNR. We conduct the AP's sweep using 7° ,




Fig. 13. Experimental floor plan used for measurement of data using WARP-60 testbed.

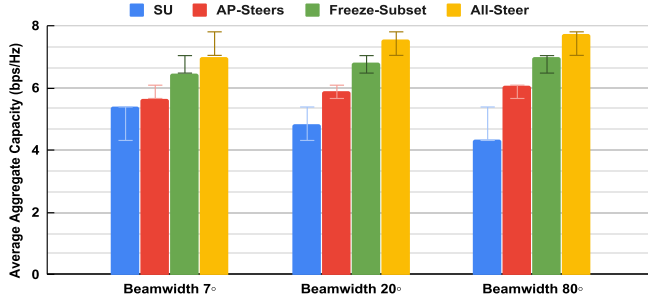


Fig. 14. Average of the aggregate capacity for all the users in the setup.

20° and 80° horns. In order to study the multi-user capacity gains of *UMBRA*, we assume a two stream transmission and consider all possible user groups consisting of 2 users out of 10 (i.e., a total of $\binom{10}{2}$ different user groups). In all topologies, both the users always have LOS connectivity with the AP.

Beamwidth and Aggregate Rate: Figure 14 shows the aggregate rate of each beam selection policy averaged over all two-user groups for different receive beamwidths at the AP. First, as expected, for the narrowest beamwidth of 7°, beam steering provides higher antenna gain due to AP's more focused beams and consequently yielding high *Single User* SNRs that most of the users are served with their best possible MCS. While in principle, an extremely narrow beamwidth would not be useful for multi-user grouping, this was not the case in our experimental setup. *UMBRA* provides modest multi-user gains even in this case of 7° beams with *AP-Steers*, *Freeze-Subset* and *All-Steer* policies providing an aggregate rate improvement over *Single User* by 5%, 22% and 40% respectively.

Second, as the beamwidth increases to 20°, there is a significant drop in the *Single User* average aggregate capacity due to the reduction in directivity gain owing to the inherent tradeoff between selected beamwidth and rate. However, increasing the beamwidth at the AP increases the beam sharing possibilities as *UMBRA* exploits the multiple paths in addition to the LOS paths. This in effect increases the SNR spread among the users and reduces the inter-user interference making opportunity for the policies to boost the SINR at each user. Therefore, the achievable aggregate rate of *AP-Steers*, *Freeze-Subset* and *All-Steer* strategies is approximately 19%, 41% and 62% higher than *Single User* respectively.

Third, the multi-user gains of *UMBRA* are more pronounced in case of 80° beamwidth despite the highest drop in

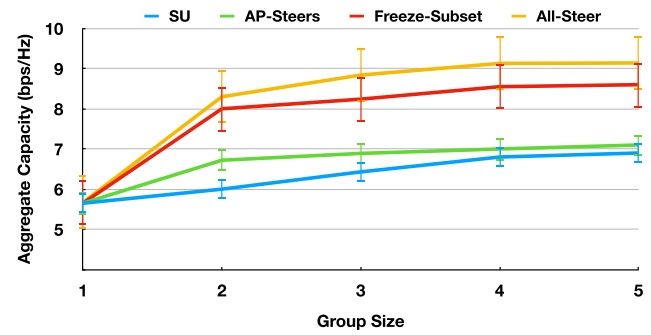


Fig. 15. Achievable Aggregate Capacity vs. Group Size.

Single User aggregate capacity. Fig. 14 reveals that *AP-Steers*, *Freeze-Subset* and *All-Steer* strategies reflect more than 1.3×, 1.5× and 1.7× multi-user capacity gains over *Single User*. This is attributed to the fact that increased beamwidth increases the SNR diversity among the users and leads to a huge number of beam sharing opportunities and the number of such potential beams that could result in maximum achievable aggregate rate outweigh the former beams from usage of narrower beamwidth.

Findings: Increasing receive beamwidth at the AP leads to larger SNR spread among the users thus providing more beam sharing possibilities and hence increased aggregate rate for multi-user uplink transmission. *AP-Steers*, *Freeze-Subset* and *All-Steer* strategies outperform *Single User* by 30%, 56% and 78% with 80° beamwidth.

E. Scaling the Number of Users/Streams

So far, we have evaluated the performance of *UMBRA* for a two-user transmission. Here, we increase number of grouped users beyond two to study the viability of further multi-user rate gains.

Setup: We use the same node deployment as in Fig. 7 and for a n stream transmission, we consider all possible user groups consisting of n users out of 15 (i.e., a total of $\binom{15}{n}$ different user groups). For each user group and each *UMBRA* policy, we find the average performance over all user groups and compare it with the baseline scheme *Single User*.

Aggregate Rate: Fig. 15 depicts the achievable aggregate PHY rate as a function of the user group sizes. First, we observe that *All-Steer* is able to achieve on average about 1.35× rate gains over *Single User* via simultaneous transmission of 3 and 4 users respectively. However, the improvement in aggregate rate has diminishing returns when increasing group size from 4 to 5. When five users are grouped, one user may not contribute enough to the aggregate rate due to lack of sufficient SINR spread as this would have yielded excessive interference making it difficult for reliable user stream decoding at the AP. The reason is behind how the SIC performs: interference-alleviation filtering cancels the inter-user interference by projecting the received signal onto the subspace orthogonal to the receive spatial signatures of all the other streams. By doing so, it optimally trades-off capturing the energy of the data stream of interest and nulling the inter-stream interference. The penalty one pays for decoding

current stream with successful interference cancellation incurs energy loss especially if there is no SNR margin to tolerate the excessive inter-stream interference. Consequently, *UMBRA* suffers from *error propagation*: if one layer is decoded incorrectly, all subsequent layers are affected due to inefficient interference cancellation and hence causes degradation in the overall multi-user performance. This puts a practical limit on the number of users which can transmit concurrently via *UMBRA* and therefore, the number of simultaneous users that can be successfully decoded by the AP while yielding a gain saturates for a group size of 4 in our setup.

Second, although *AP-Steers* achieves 12% aggregate capacity gain over *Single User* for a simultaneous transmission of two users, it has only marginal aggregate rate increase as the group size increases beyond 2. The reason is that the AP-only re-steering can tackle the inter-user interference for a group size of 2, unfortunately, it fails to efficiently cancel the interference as the group size increases. As *AP-Steers* freezes the transmit beam of users to share the common AP's receive beam, finding an optimal receive beam that often results in sufficient SINR spread with minimal inter-user interference becomes difficult for larger user groups, thus resulting in degraded relative stream decoding performance under *AP-Steers* scheme.

Third, *Freeze-Subset's* performance is quite close to *All-Steer*, despite its significantly reduced search space. This is because *Freeze-Subset* always selects the receive beam in favor of the user having the highest single-user SNR as this user will be decoded first using SIC thus having control over the amount of inter-user interference and reduces error propagation from one stream to another. Fig. 15 also shows that *Freeze-Subset* never loses sum-rate due to an additional transmitting user and the reason is that *Freeze-Subset* allows users to re-steer jointly with the AP in a sequential manner and evaluates the net effect of joint beam re-steering of the newly chosen user before data transmission. However, by increasing the number of users from 4 to 5, *Freeze-Subset* shows saturation point similar to *All-Steer* and as discussed above, this depends on the SINR spread resulting from AP-user re-steering beam configuration and successful stream separation of SIC.

Finding: While multi-user beam steering becomes challenging with increased inter-user interference arising from increasing number of concurrent user streams, UMBRA with its All-Steer scheme is successfully able to provide $1.36\times$ rate gains over Single User for simultaneous transmission of 4 users.

VI. RELATED WORK

Multi-User Uplink 60 GHz Networks. Prior work in mmWave WLANs target downlink multi-user multi-stream transmissions with the AP using at least one RF chain per stream [1], [8], [10], [11], [12]. However, all these systems are limited to downlink transmissions and not applicable to uplink. Although a significant amount of research has focused on theoretical capacity analysis for mmWave uplink [13], [14], [15], [16], little is known about the performance of such systems in practice. *UMBRA* in contrast, realizes the first

mmWave WLAN system design and experimentation in which the number of users exceeds the number of RF chains.

Multi-User Uplink Sub 6 GHz Networks. Prior work on uplink MU-MIMO below 6 GHz focuses on information-theoretic capacity exploration [17], [18], multi-user transmission via successive interference cancellation, interference alignment or orthogonal preambles [19], [20], [21], [22], [23]. Unfortunately, techniques used in these works can't be applied to our scenario due to a different node architecture at 60 GHz (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.

Multi-User Uplink with Single RF Chain: Extensive prior work in uplink has realized MU-MIMO while requiring a single RF chain. References [24], [25], [26]. More recently, NOMA has been used in mmWave Uplink which has the same philosophy of overlaid constellations of using same time, frequency and space resources [27], [28], [29], [30], [31], [32], [33], [34], [35]. Furthermore, user streams are separated via power allocation policies and employ SIC at receiver to remove the multi-user interference. Instead, we address the stream separation with design of multi-user overlaid APSK constellations and multi-user stream decoder and rely on beam steering and user locations to enable SNR spread. In contrast to all prior work, our focus is not only the design of mmWave uplink WLAN system, but we also experimentally evaluate the functionality under practical system constraints and identified the key factors that affect the multi-user gains.

Lastly, our recent conference paper in [36] studies the problem of multi-user beamforming on a Single RF Chain in 60 GHz MIMO WLANs. This article is the extended version of [36] and has the same core idea. However, here, we have included more over-the-air experiments and analyses.

VII. CONCLUSION

In this paper, we present the design and experimental evaluation of *UMBRA*, a novel framework that supports multi-user uplink transmissions on a single RF chain AP. We introduced Multi-User Overlaid APSK constellations for enabling simultaneous uplink transmissions of APSK signals with a design feature that allows multi-stream separation via SIC. We designed the multi-user receiver with the mechanism to perform multi-user CFO correction and compensation and interference cancellation filter based stream separation. We proposed constrained beam adaptation that enables *UMBRA* to maximize aggregate rate under multiple system constraints. Our experiments demonstrate that *UMBRA* achieves about $1.45\times$ improvement in multi-user gains over single-user systems irrespective of choice of users grouped, geometric separation and receive beamwidth.

REFERENCES

- [1] Y. Ghasempour, M. K. Haider, C. Cordeiro, D. Koutsonikolas, and E. Knightly, "Multi-stream beam-training for mmWave MIMO networks," in *Proc. 24th Annu. Int. Conf. Mobile Comput. Netw.*, Oct. 2018, pp. 225–239.
- [2] IEEE P802.11 TASK GROUP AX, *IEEE P802.11 Wireless LANs: Specification Framework for TGax*, Standard P802.11ax, IEEE Standard in Progress, 2017.

- [3] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [4] S. K. Saha et al., "X60: A programmable testbed for wideband 60 GHz WLANs with phased arrays," *Comput. Commun.*, vol. 133, pp. 77–88, Jan. 2019.
- [5] N. Anand, E. Aryafar, and E. W. Knightly, "WARPlab: A flexible framework for rapid physical layer design," in *Proc. ACM Workshop Wireless Students, Students, Students*, 2010, pp. 53–56.
- [6] Y. Ghasempour and E. W. Knightly, "Decoupling beam steering and user selection for scaling multi-user 60 GHz WLANs," in *Proc. 18th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, Jul. 2017, pp. 1–10.
- [7] M. K. Haider and E. W. Knightly, "Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs: Protocol design and system implementation," in *Proc. ACM MobiHoc*, 2016, pp. 61–70.
- [8] K. P. Dasala, J. M. Jornet, and E. W. Knightly, "SIMBA: Single RF chain multi-user beamforming in 60 GHz WLANs," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Jul. 2020, pp. 1499–1508.
- [9] S. Naribole and E. Knightly, "Scalable multicast in highly-directional 60-GHz WLANs," *IEEE/ACM Trans. Netw.*, vol. 25, no. 5, pp. 2844–2857, Oct. 2017.
- [10] K. P. Dasala, J. M. Jornet, and E. W. Knightly, "Scaling mmWave WLANs with single RF chain multiuser beamforming," *IEEE/ACM Trans. Netw.*, vol. 30, no. 6, pp. 2630–2643, Dec. 2022.
- [11] *IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band*, IEEE Standard 802.11ad-2012 (Amendment to IEEE Std 802.11-2012, as amended by IEEE Std 802.11ae-2012 and IEEE Std 802.11aa-2012), 2012, pp. 1–628, doi: [10.1109/IEEESTD.2012.6392842](https://doi.org/10.1109/IEEESTD.2012.6392842).
- [12] Y. Ghasempour, C. R. da Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi," *IEEE Commun. Mag.*, vol. 55, no. 12, pp. 186–192, Dec. 2017.
- [13] L. Zhu, J. Zhang, Z. Xiao, X. Cao, D. O. Wu, and X.-G. Xia, "Joint power control and beamforming for uplink non-orthogonal multiple access in 5G millimeter-wave communications," *IEEE Trans. Wireless Commun.*, vol. 17, no. 9, pp. 6177–6189, Sep. 2018.
- [14] P. Raviteja, Y. Hong, and E. Viterbo, "Millimeter wave analog beamforming with low resolution phase shifters for multiuser uplink," *IEEE Trans. Veh. Technol.*, vol. 67, no. 4, pp. 3205–3215, Apr. 2017.
- [15] J. Li, L. Xiao, X. Xu, and S. Zhou, "Robust and low complexity hybrid beamforming for uplink multiuser mmWave MIMO systems," *IEEE Commun. Lett.*, vol. 20, no. 6, pp. 1140–1143, Jun. 2016.
- [16] O. Onireti, A. Imran, and M. A. Imran, "Coverage, capacity, and energy efficiency analysis in the uplink of mmWave cellular networks," *IEEE Trans. Veh. Technol.*, vol. 67, no. 5, pp. 3982–3997, May 2018.
- [17] A. Goldsmith, S. A. Jafar, N. Jindal, and S. Vishwanath, "Capacity limits of MIMO channels," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 5, pp. 684–702, Jun. 2003.
- [18] Y. Kim, S. Cho, and D. K. Kim, "Low complexity antenna selection based MIMO scheduling algorithms for uplink multiuser MIMO/FDD system," in *Proc. IEEE 65th Veh. Technol. Conf. (VTC-Spring)*, Apr. 2007, pp. 1663–1667.
- [19] K. Tan et al., "SAM: Enabling practical spatial multiple access in wireless LAN," in *Proc. 15th Annu. Int. Conf. Mobile Comput. Netw.*, 2009, pp. 49–60.
- [20] A. Zhou, T. Wei, X. Zhang, M. Liu, and Z. Li, "Signpost: Scalable MU-MIMO signaling with zero CSI feedback," in *Proc. 16th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, 2015, pp. 327–336.
- [21] A. B. Flores, S. Quadri, and E. W. Knightly, "A scalable multi-user uplink for Wi-Fi," in *Proc. USENIX NSDI*, 2016, pp. 1–14.
- [22] K. C. Lin, S. Gollakota, and D. Katabi, "Random access heterogeneous MIMO networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 146–157, Aug. 2011.
- [23] T. Wei and X. Zhang, "Random access signaling for network MIMO uplink," in *Proc. IEEE 35th Annu. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, Apr. 2016, pp. 1–9.
- [24] N. I. Miridakis and D. D. Vergados, "A survey on the successive interference cancellation performance for single-antenna and multiple-antenna OFDM systems," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 1, pp. 312–335, 1st Quart., 2013.
- [25] M. Zeng, W. Hao, A. Yadav, N.-P. Nguyen, O. A. Dobre, and H. V. Poor, "Energy-efficient joint power control and receiver design for uplink mmWave-NOMA," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, Jun. 2020, pp. 1–5.
- [26] Z. Xiao, L. Zhu, Z. Gao, D. O. Wu, and X. Xia, "User fairness non-orthogonal multiple access (NOMA) for millimeter-wave communications with analog beamforming," *IEEE Trans. Wireless Commun.*, vol. 18, no. 7, pp. 3411–3423, Jul. 2019.
- [27] H. Wang, R. Zhang, R. Song, and S.-H. Leung, "A novel power minimization precoding scheme for MIMO-NOMA uplink systems," *IEEE Commun. Lett.*, vol. 22, no. 5, pp. 1106–1109, May 2018.
- [28] Z. Ding, R. Schober, and H. V. Poor, "A general MIMO framework for NOMA downlink and uplink transmission based on signal alignment," *IEEE Trans. Wireless Commun.*, vol. 15, no. 6, pp. 4438–4454, Jun. 2016.
- [29] Z. Yang, Z. Ding, P. Fan, and N. Al-Dhahir, "A general power allocation scheme to guarantee quality of service in downlink and uplink NOMA systems," *IEEE Trans. Wireless Commun.*, vol. 15, no. 11, pp. 7244–7257, Nov. 2016.
- [30] M. S. Ali, H. Tabassum, and E. Hossain, "Dynamic user clustering and power allocation for uplink and downlink non-orthogonal multiple access (NOMA) systems," *IEEE Access*, vol. 4, pp. 6325–6343, 2016.
- [31] M. A. Sedaghat and R. R. Müller, "On user pairing in uplink NOMA," *IEEE Trans. Wireless Commun.*, vol. 17, no. 5, pp. 3474–3486, May 2018.
- [32] X. Chen, A. Benjebbour, A. Li, and A. Harada, "Multi-user proportional fair scheduling for uplink non-orthogonal multiple access (NOMA)," in *Proc. IEEE 79th Veh. Technol. Conf. (VTC Spring)*, May 2014, pp. 1–5.
- [33] B. Kim et al., "Uplink NOMA with multi-antenna," in *Proc. IEEE 81st Veh. Technol. Conf. (VTC Spring)*, May 2015, pp. 1–5.
- [34] Y. Gao, B. Xia, K. Xiao, Z. Chen, X. Li, and S. Zhang, "Theoretical analysis of the dynamic decode ordering SIC receiver for uplink NOMA systems," *IEEE Commun. Lett.*, vol. 21, no. 10, pp. 2246–2249, Oct. 2017.
- [35] Z. Wei, L. Yang, D. W. K. Ng, J. Yuan, and L. Hanzo, "On the performance gain of NOMA over OMA in uplink communication systems," *IEEE Trans. Commun.*, vol. 68, no. 1, pp. 536–568, Jan. 2020.
- [36] K. P. Dasala, J. M. Jornet, and E. W. Knightly, "Uplink multi-user beamforming on single RF chain mmWave WLANs," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, May 2021, pp. 1–10.