

m^3 MIMO: An 8×8 mmWave Multi-User MIMO Testbed for Wireless Research

Khandaker Foysal Haque[†], Francesca Meneghelli^{*}, K M Rumman[†] and Francesco Restuccia[†]

[†] Institute for the Wireless Internet of Things, Northeastern University, United States

^{*} Department of Information Engineering, University of Padova, Italy

ABSTRACT

In this paper, we present m^3 MIMO a mmWave fully-digital multi-user multi-input multi-output (MU-MIMO) testbed for advanced wireless research. m^3 MIMO operates in the 57-64 GHz frequency range and supports up to 1 GHz of bandwidth enabling large data multiplexing in the frequency domain through orthogonal frequency-division multiplexing (OFDM). The testbed features three custom-designed Zynq UltraScale+ RFSoC-based Software Defined Radios (SDRs) empowered with the Pi-Radio fully digital transceivers. Two of these SDRs support eight transmit and receive streams each (8 × 8 MIMO), while the third SDR supports up to four channels. m^3 MIMO supports three different communication modes: (i) point-to-point (P2P) transmissions; (ii) single-user multi-input multi-output (SU-MIMO), where multiple streams are transmitted to a single end-device; and (iii) MU-MIMO, where two devices are simultaneously served by a single transmitter. To showcase the m^3 MIMO's versatility, we present two research use cases: tracking-based beamforming and mmWave-based sensing. **We will open-source the m^3 MIMO code along with the relevant use-case datasets, facilitating further analysis¹.**

CCS CONCEPTS

• Networks → Network experimentation; Network measurement; • Hardware → Networking hardware.

KEYWORDS

MU-MIMO, mmWave MU-MIMO, MU-MIMO testbed, mmWave testbed

¹<https://github.com/kfoysalhaque/m3MIMO>

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM *MobiCom '24*, November 18–22, 2024, Washington D.C., DC, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0489-5/24/11.

<https://doi.org/10.1145/3636534.3697321>

ACM Reference Format:

Khandaker Foysal Haque[†], Francesca Meneghelli^{*}, K M Rumman[†] and Francesco Restuccia[†]. 2024. m^3 MIMO: An 8×8 mmWave Multi-User MIMO Testbed for Wireless Research. In *The 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '24)*, November 18–22, 2024, Washington D.C., DC, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3636534.3697321>

1 INTRODUCTION

Emerging technologies based on augmented and virtual reality (AR/VR), such as the Metaverse, will provide new entertainment applications [1], providing ultra-realistic online learning experiences [2], and transforming healthcare through remote surgery opportunities [3]. A key issue currently stymieing the provisioning of such applications is that commercial VR headsets do not deliver adequate performance to the end user. Experts believe that video frames should have at least 120 Hz frame rate with 8K resolution to avoid pixelation and motion sickness [4–6]. However, sending frames at 120 Hz with 8K resolution would require about 40 Gbps of data rate for each AR/VR device, which is currently unachievable through wireless networks. For example, Wi-Fi supports a maximum of 1.2 Gbps network-wide [7, 8]. Laha et al. show that considering a system bandwidth of 100 MHz, fifth generation (5G) new radio can serve only a single user per base station while guaranteeing the stringent AR/VR requirements [9]. The limited available bandwidth together with the small practically achievable multiple-input multiple-output (MIMO) dimensionality is bounding wireless network performance. This has led the research community to further delve into MIMO research to improve the system performance. Indeed, current MIMO implementations are limited to four spatial streams even if higher order MIMO is envisioned (e.g., up to eight streams in the IEEE 802.11ax standard). As one of the research directions to improve network performance, researchers are exploring the use of wider bandwidths available in the higher bands of the radio spectrum (mmWave and Terahertz). While MIMO leverage spatial diversity to multiplex multiple streams in the same spectrum band, the availability of a wider bandwidth allows multiplexing in the frequency domain, through, e.g., orthogonal frequency-division multiplexing (OFDM).

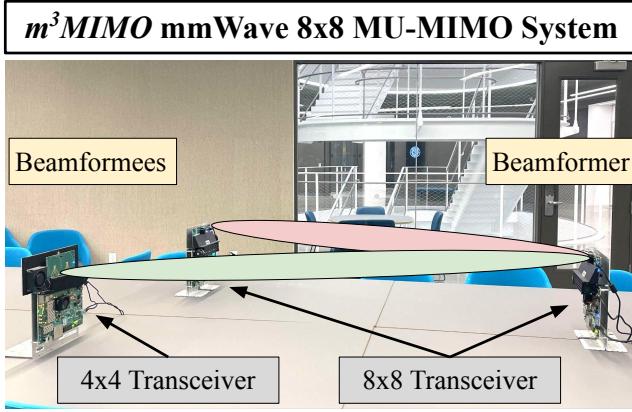


Figure 1: m^3 MIMO overview.

These research efforts highly depend on experimental analysis with appropriate testbeds for evaluating the systems in real-world situations. To this end, in this paper, we present m^3 MIMO, an 8×8 mmWave multi-user multi-input multi-output (MU-MIMO) testbed for next-generation wireless research. The testbed is depicted in Figure 1 and enables both MIMO and mmWave research under the same umbrella making it a flexible and versatile tool for research. The testbed features three Zynq Ultra-Scale+ RFSoC Software Defined Radios (SDRs) empowered with Pi-radio radio transceivers [10]. Two of the SDRs support up to 8×8 MIMO having 8 independent transmitter and receiver RF chains and antennas whereas the third SDR is limited to 4×4 . m^3 MIMO operates in the mmWave band of the radio spectrum, specifically within the 57-64 GHz frequency range, supporting up to 1 GHz of bandwidth. Unlike existing mmWave testbeds, m^3 MIMO operates in a fully digital fashion, enabling 8×8 MIMO and supports OFDM transmissions. This facilitates research on mmWave band and OFDM fully digital MIMO (both single-user multi-input multi-output (SU-MIMO) and MU-MIMO) to improve spectrum efficiency for next-generation applications.

Summary of Contributions.

- We present m^3 MIMO, an 8×8 fully digital MU-MIMO testbed operating in the mmWave spectrum, capable of OFDM transmissions with support for bandwidths up to 1 GHz.
- We detail the implementation of the channel sounding and precoding procedure integrated in the m^3 MIMO testbed to enable SU-MIMO and MU-MIMO transmission modes. We demonstrate through experiments implemented with m^3 MIMO that digital beamforming in SU-MIMO improves the network performance in terms of bit error rate (BER) and signal-to-noise ratio (SNR) by up to 13.3x and 2.05x times respectively.
- We showcase two research use cases of the m^3 MIMO testbed – tracking-based beamforming and mmWave-based gesture

recognition. The experimental evaluations show that tracking-based beamforming improves the BER and SNR performances by 4x and 2.8x times respectively when compared to unbeamformed transmissions. Whereas the mmWave-based gesture recognition achieves an accuracy of up to 99.52% when trained and tested in the same environment and up to 91.33% when trained in one environment and tested in others.

2 EXISTING MMWAVE MIMO TESTBEDS

A few testbeds were developed by the research community to carry out experimental evaluations of algorithms for mmWave MIMO. Among them, Lacruz et al. propose MIMORPH [11], a multi-band platform working on the mmWave and sub-6 GHz portions of the radio spectrum. The system supports up to 4×4 MIMO when operating in the mmWave spectrum while 8×8 can be enabled in the sub-6 GHz frequencies. For the waveform generation in the mmWave and sub-6 GHz bands, MIMORPH follows the IEEE 802.11ay and 802.11ac/ax standards respectively. However, this system works only in single-carrier mode at mmWave. Moreover, following the 802.11ay standard, MIMORPH currently implements analog beamforming strategies, i.e., MIMO is used for directionality gain while the multiplexing gain is not exploited. Instead, m^3 MIMO also enables multiplexing up to 8 different streams to single or multiple users through the SU-MIMO and MU-MIMO procedures we developed. In [12], the authors propose a single-carrier mmWave (60 GHz) MIMO system featuring a 16-antenna transmitter and two 4-antenna receivers. However, the use of only 2 digital-to-analog converters (DACs) and 2 analog-to-digital converters (ADCs) limits the system to 2×3 MIMO. Zhao et al. propose M-cube [13], an OFDM MU-MIMO system that uses beamforming to multiplex up to 4 streams working on the 60 GHz band and following the IEEE 802.11ay standard. In addition to the lower number of streams that can be multiplexed, the fast fourier transform (FFT) size for OFDM transmission is limited to 64 thus reaching a lower level of frequency multiplexing than our m^3 MIMO system – where the FFT size is up to 1024. The recent work in [14] proposes Agora-UHD, an OFDM MIMO system working on the 28 GHz band. However, also for this testbed, the maximum MIMO dimensionality is 2×2 . Finally, an OFDM-based 4×2 MIMO system working on the 24 GHz band is proposed by Ozkaptan et al. together with a radar-assisted precoding mechanism for SU-MIMO [15].

3 M^3 MIMO DATA TRANSMISSION PIPELINE

In this section, we detail the m^3 MIMO transmission pipeline. The system uses OFDM to multiplex data in the frequency domain by dividing the available bandwidth into small orthogonal sub-channels. In addition to the frequency domain,

multiplexing is also applied in the space domain. To achieve this, we tailored the channel sounding and precoding procedures from the IEEE 802.11 standard, customizing them to fit the m³MIMO setup for optimal performance.

3.1 OFDM Operation

m³MIMO supports OFDM transmissions with up to 1 GHz of bandwidth. However, to avoid inter-symbol interference, it is recommended to confine the transmission to the central 800 MHz bandwidth. For the implementation of OFDM, this 800 MHz bandwidth is divided into $K = 800$ sub-channels with a sub-channel bandwidth of at least 1 MHz. Hence, each OFDM symbol consists of up to $X = 800$ modulated IQ samples. Each OFDM symbol $x[n]$ is then processed through an inverse fast fourier transform (IFFT) block to generate the corresponding time domain signal as

$$x[n] = \frac{1}{K} \sum_{l=0}^{K-1} X_l e^{j2\pi ln/K}, \quad n = 0, 1, \dots, K-1. \quad (1)$$

This transformation ensures that the data symbols are orthogonal in the frequency domain. At the receiver, the system captures the incoming signal and converts it back to the frequency domain using an FFT:

$$X_l = \sum_{n=0}^{K-1} x[n] e^{-j2\pi ln/K}, \quad l = 0, 1, \dots, K-1. \quad (2)$$

3.2 MIMO for Directionality at the Receiver

The m³MIMO fully-digital MIMO system can be used to implement analog-like receive beamforming to capture the transmitted beams from specific directions, thus increasing the SNR at the receiver antennas. This involves applying a beamforming vector to the received signals, which includes phase shifts designed to align the phases of the incoming signals with each possible angle of arrivals (AoAs) with a pre-defined granularity. Two parameters that allow evaluating the performance of the implemented analog-like receive beamforming are the directional response pattern and the received power spectrum. Directional response pattern illustrates the variation of power received by the antennas as a function of AoAs, whereas the received power spectrum, refers to the distribution of received signal power over different frequencies. For example, in Figure 2 we present the directional response patterns and the receive power spectrum for different AoAs (emulated by rotating the receiver at different angles along its axis without any horizontal or vertical displacement).

The directional response pattern exhibits moderate angular sensitivity with distinct peaks and nulls, indicating variations in reception quality across different AoAs. The receive power spectrum shows consistent performance across antennas, with slight roll-offs as frequencies deviate from

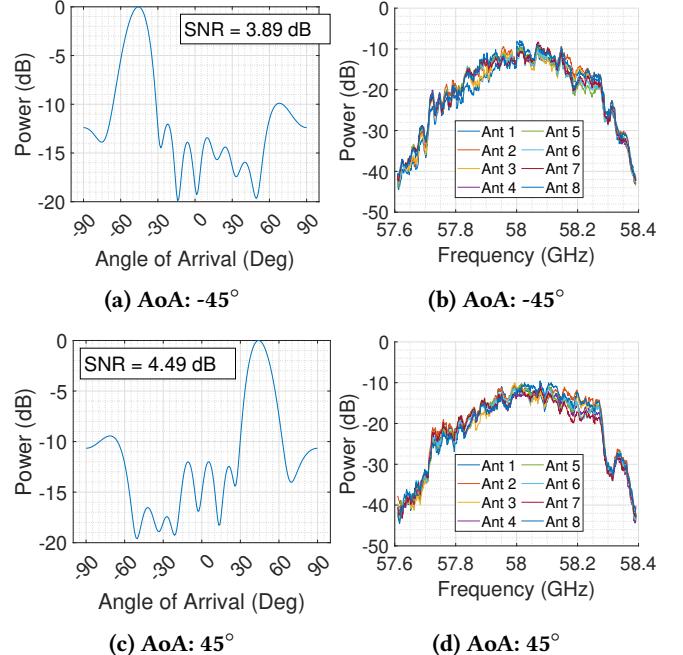


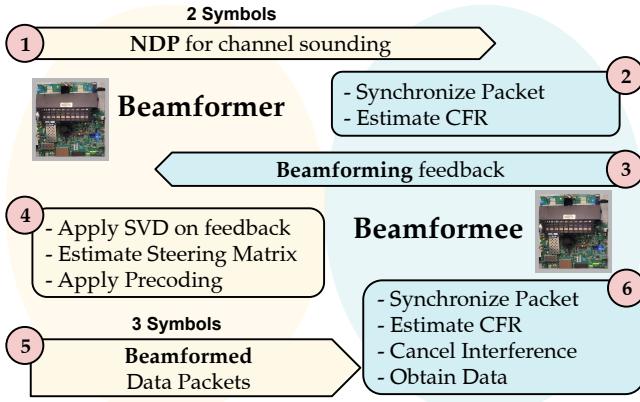
Figure 2: Directional response pattern (a and c) and corresponding received power spectrum (b and d) for different AoAs at a transmission distance of 3 m.

the center frequency of 58 GHz, reflecting effective antenna calibration and system uniformity in the operational band.

3.3 MIMO for Multiplexing and Diversity

To enable proper decoding and equalization at the receiver device, we transmit two additional OFDM symbols containing known samples before transmitting the actual data. The first is used as a synchronization symbol to remove the time offset. The offset is obtained by computing the cross-correlation between the received synchronization symbol and the known transmitted one. The estimation obtained through this procedure is then used to compensate for the timing offset of all the subsequent transmitted symbols. The second control symbol is used to estimate the impairments introduced by the wireless channel, i.e., the channel frequency response (CFR), and properly equalize the actual data. Note that, before using it for channel estimation, the second symbol is also compensated for the timing offset. The symbols containing the actual user data follow the control symbols. The data symbols are first synchronized and equalized using the estimates performed on the first two known symbols. Hence, they are demodulated following the adopted modulation procedure and the transmitted bits are retrieved.

In our specific MIMO implementation, we are constrained to the Zynq RFSoC's default buffer size for DACs/ADCs. Specifically, the transmission is performed by repeatedly sending the data stored in a circular buffer that has a maximum capacity of 2048 samples. As we need to transmit two

Figure 3: m^3 MIMO Beamforming.

mandatory symbols for synchronization and equalization, we are currently constrained to three symbols in total (i.e., two control and one data symbol) with 512 OFDM samples each.

SU-MIMO and MU-MIMO require the transmitter device (beamformer) to combine the different streams at the available antennas through the so-called *precoding* procedure. This strategy allows to achieve multiplexing and diversity gain. Indeed, the different streams are combined at the different antennas using orthogonal weighting vectors that enable simultaneous transmission while reducing the interference among them. Such precoding weights are obtained by combining the channel estimates performed by all the devices (beamformees) connected to the beamformer. The channel estimation procedure is usually referred to as *channel sounding*. Once the precoded data is received, the beamformees retrieve the streams directed to it by applying an *interference cancellation* matrix. The complete m^3 MIMO MIMO transmission procedure is depicted in Figure 3 and detailed next.

3.3.1 Channel Sounding. At first, the beamformer transmits a null data packet (NDP) packet consisting of only two symbols, i.e., the synchronization symbol and the channel sounding symbol (step 1 in Figure 3). The NDP does not contain the data symbol as its only objective is to provide the beamformees with known samples to estimate the channel. For this packet, the beamformer transmits the samples through all eight different antennas such that the channel between each pair of the N transmitter and M receiver antennas over the entire operational bandwidth can be estimated. Specifically, each beamformee i obtains an estimate of the CFR $\mathbf{H}_{k,i}$ on each OFDM sub-channel K by comparing the received samples with the known transmitted ones (step 2 in Figure 3). This corresponds to dividing the received signal in the frequency domain $\mathbf{Y}_{k,i}$ by the transmitted one $\mathbf{X}_{k,i}$, i.e., $\mathbf{H}_{k,i} = \mathbf{Y}_{k,i}/\mathbf{X}_{k,i}$. The CFR for each OFDM sub-channel $\mathbf{H}_{k,i}$ is $N \times M$ dimensional. For example, Figure 4 presents the CFR

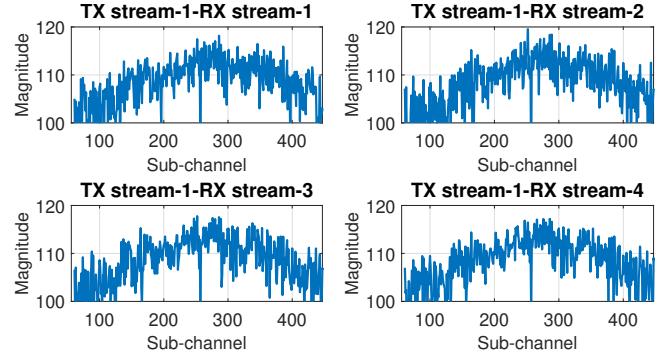


Figure 4: CFR of received streams 1-4, for transmit antenna 1 using beamformed transmission at a distance of 6 m.

of the received streams 1-4 for the first transmit antenna obtained from real-world experiments with m^3 MIMO. Next, the CFR is compressed through singular value decomposition (SVD) as done in the IEEE 802.11ac/ax standard procedure: $\mathbf{H}_{k,i} = \mathbf{U}_{k,i} \mathbf{S}_{k,i} \mathbf{V}_{k,i}$ and only the right singular matrix $\mathbf{V}_{k,i}$ is transmitted to the beamformer (step 3 in Figure 3).

3.3.2 Precoding. Once the compressed CFR $\mathbf{V}_{k,i}$ from all the beamformees is received, the beamformer obtains the precoding weight matrix \mathbf{W}_k through zero-forcing (ZF) as (step 4 in Figure 3)

$$\mathbf{W}_k = \mathbf{V}_k \left(\mathbf{V}_k^\dagger \mathbf{V}_k \right)^{-1}, \quad (3)$$

where \mathbf{V}_k is obtained by combining the $\mathbf{V}_{k,i}$ for all the different beamformees in the second dimension. The beamformed signal to be transmitted is hence obtained as $\mathbf{W}_k \mathbf{X}_k$ maintaining a shape of $K \times S \times M \times N$ where S represents the total number of symbols. Specifically, the packet is composed of three symbols, i.e., the synchronization symbol, the channel estimation symbol, and the data symbol. All of the symbols in the packet are precoded (step 4 in Figure 3) and transmitted to the beamformees (step 5).

3.3.3 Interference Cancellation. The beamformed signal received by beamformee i can be expressed as

$$\mathbf{Y}_{k,i} = \mathbf{H}_{k,i} \mathbf{X}_k + \mathbf{n}_{k,i}. \quad (4)$$

To effectively retrieve the streams directed to user i and separate their contributions, an *interference cancellation matrix* $\mathbf{G}_{k,i}$ is applied to the received signal $\mathbf{Y}_{k,i}$ (step 6 in Figure 3). Specifically, $\mathbf{G}_{k,i}$ is obtained as

$$\mathbf{G}_k = \mathbb{I}_{N_{ss,i} \times N_{ss}} \mathbf{W}_k^\dagger \mathbf{H}_k^\dagger (\mathbf{W}_k \mathbf{H}_k \mathbf{W}_k^\dagger \mathbf{H}_k^\dagger + \mathbf{F}_k)^{-1}, \quad (5)$$

where \mathbf{F}_k is the noise covariance matrix, and $\mathbf{W}_k \mathbf{H}_k$ is the estimate of the beamformed channel obtained from the beamformed channel estimation symbol.

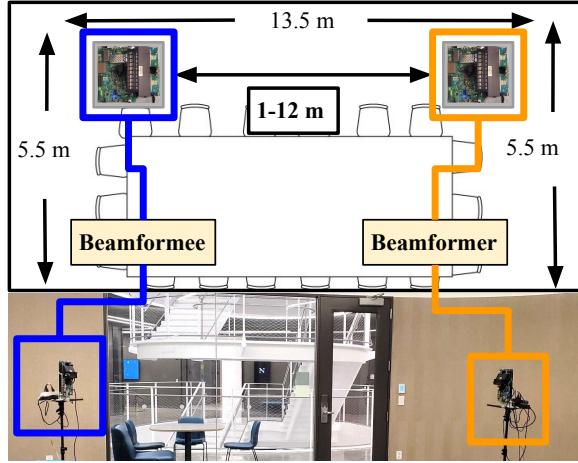


Figure 5: Experimental setup for beamforming test with SU-MIMO.

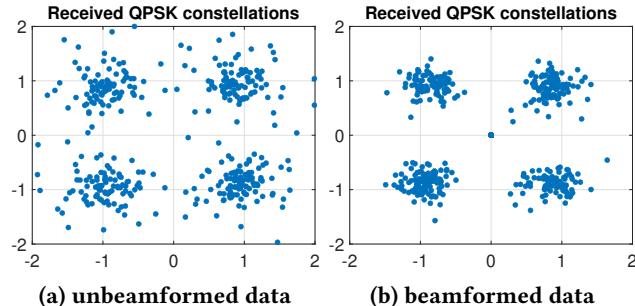


Figure 6: Received constellations with SU-MIMO at 6m of transmission distance.

3.3.4 Performance of m³MIMO MIMO: A SU-MIMO Test. We analyze the performance of SU-MIMO transmission in terms of BER and SNR. We perform the test in a conference room as presented in Figure 5, with both beamformed and unbeamformed transmissions at multiple distances ranging from 1 to 12 m. Figure 6 presents the comparison of the received QPSK constellations of the unbeamformed and beamformed data considering a transmission distance of 6m. This clearly shows that the constellations of the beamformed signal are better mapped in comparison to the unbeamformed one indicating better signal quality with beamforming. To have a clearer idea of the beamforming performance we also compare the communication performance, i.e., BER and SNR of the beamformed and unbeamformed signal. The results are presented in Figure 7 and show that beamforming in SU-MIMO implemented through m³MIMO improves the BER and SNR performance by 13.3x and 2.05x times respectively.

4 M³MIMO RESEARCH USE CASES

To demonstrate the versatility and flexibility of the m³MIMO we demonstrate two different research use cases: (i) tracking-based beamforming, which demonstrates the versatility of

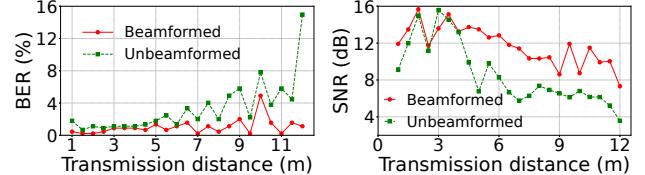


Figure 7: BER and SNR of beamforming test with SU-MIMO at a transmission distance of 1-12 m.

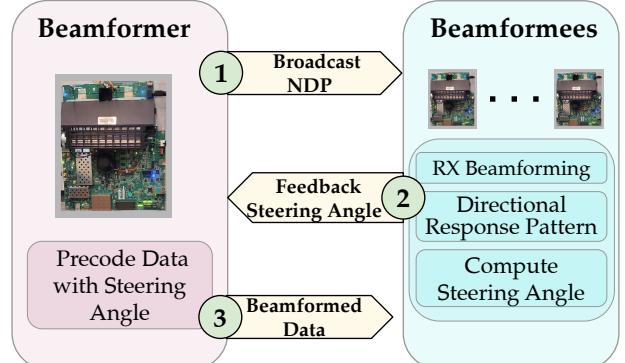


Figure 8: Tracking-based beamforming.

m³MIMO in different beamforming research, and (ii) gesture recognition, demonstrating the prospect of leveraging m³MIMO in the emerging trend of integrated sensing and communication in mmWave band.

4.1 Tracking-based Beamforming

Since the last decade, tracking-based beamforming [16, 17] has attracted the attention of the research community due to its extensive and versatile applications, especially with mmWave [18, 19] and THz [20] technologies. However, most of these works are based on simulations and emulations due to the unavailability of suitable testbeds. In the following, we show how to use m³MIMO for such use cases.

4.1.1 System for Beam-tracking and Beamforming. We have designed a tracking-based beamforming system based on feedback from beamformees estimated through the received directional response as presented in Figure 8.

The system operates by having the beamformer broadcast the NDP (as shown in step 1 of Figure 8). The beamformees in the vicinity receive the NDP and perform analog-like receive beamforming at the receiver end to measure the directional response pattern. Figure 9 illustrates the directional response pattern estimated from the received NDP for different receiver locations, i.e., different AoAs. Each beamformee calculates the steering angle based on the peak of its directional response and sends this information back to the beamformer (step 2 of Figure 8). Subsequently, the beamformer uses the number of beamformees and the received steering angles to precode the data. The beamformer performs precoding by applying the necessary phase shifts at each antenna element

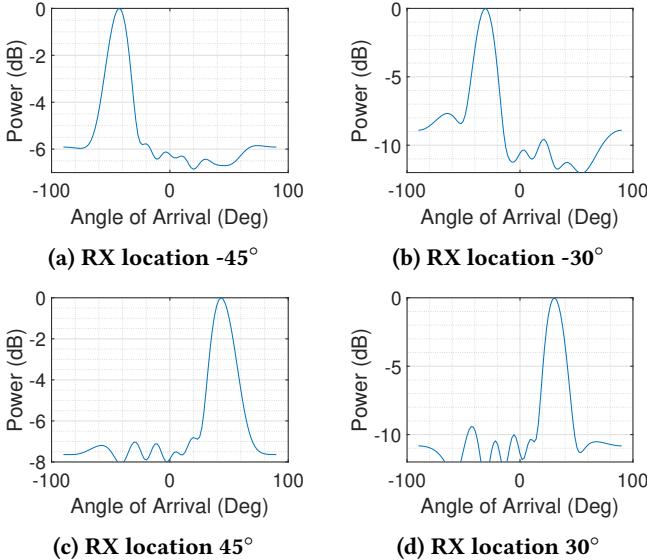


Figure 9: Computed steering angles for beam-tracking at different receiver locations.

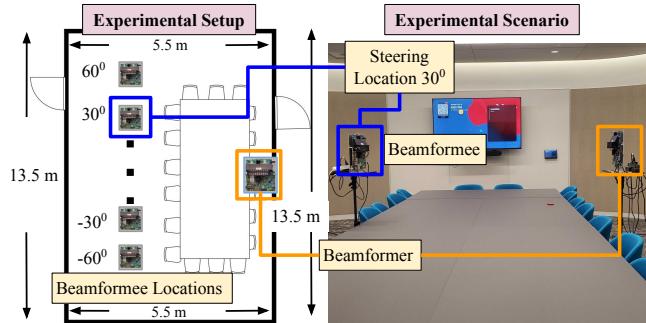


Figure 10: Experimental setup and scenario for beam-tracking.

to focus the signal energy in the direction of the steering angle. Finally, the beamformer transmits the beamformed data to the connected beamformees (step 3 of Figure 8).

Tracking-based beamforming is an essential technique in modern wireless communication systems, especially for high-frequency bands like mmWave. By dynamically adjusting the beam's direction to follow the beamformees, this technique ensures optimal communication performance.

4.1.2 Evaluation of Tracking-based Beamforming System. We evaluated the designed tracking-based beamforming system by implementing it in m^3 MIMO testbed. The system is tested through real-world over-the-air experiments involving one beamformer and one beamformee, as shown in Figure 10. We kept the beamformer's location fixed while moving the beamformee horizontally, maintaining a constant distance from the beamformer. This setup allowed us to form different AoAs for evaluating the performance of the communication in terms of BER and SNR when leveraging the

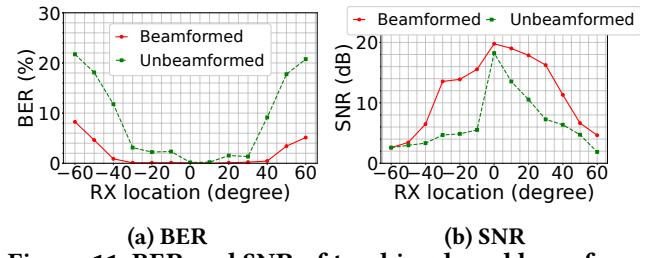


Figure 11: BER and SNR of tracking-based beamforming test with different AoAs ranging from -60° to 60° .

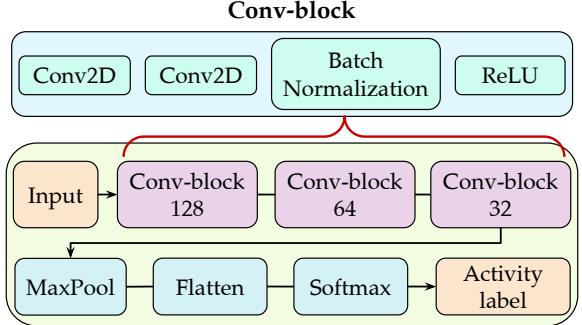


Figure 12: DNN architecture for gesture recognition.

designed tracking-based-beamforming system with m^3 MIMO. Figure 11 presents the comparative analysis of BER and SNR for unbeamformed transmission and beamformed transmission performed through the designed tracking-based beamforming system. The results show that the tracking-based beamforming improves the BER and SNR performance by up to 4x and 2.8x times, respectively, compared to unbeamformed transmissions.

4.2 mmWave Sensing: Gesture Recognition

Beyond enhancing connectivity, the unprecedented rise in the number of smartphones, laptops, and various other wirelessly connected devices [21] is driving the emergence of wireless sensing applications. These include activity recognition [22], radio fingerprinting [23], and gesture recognition [24], among others [25].

The vast majority of such works are based on mmWave radar systems [24, 26, 27]. Despite leading to good performance, radar-based sensing techniques require dedicated infrastructure which hampers the scalability. In this section, we provide the implementation details of *gesture recognition* on m^3 MIMO. The mmWave CFR is processed through a deep neural networks (DNN)-based learning approach for recognition of the gestures representing digits 0-9. Specifically, the DNN classifier relies on the CFR estimated by the beamformees in the network, captured during the depiction of the gestures.

4.2.1 Data Processing and Learning Approach. Firstly, the estimated CFR (see step 2 of Figure 3) corresponding to the

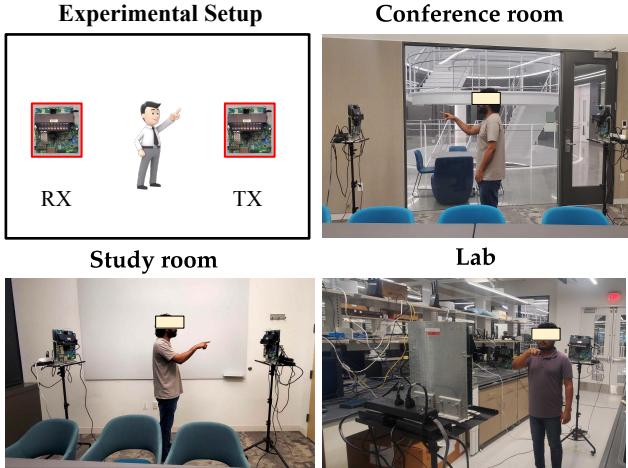


Figure 13: m³MIMO setup and data collection campaigns for gesture recognition.

depiction of different digits are stacked individually. As the duration of the depiction of different gestures is different, the number of CFR frames also varies across different gestures. However, to reduce the complexity of the DNN model, we need to have the same input dimension for different gestures. Thus, for every gesture, only the first $S = F \times T = 600$ number of CFR samples are considered. Here, $F = 100\text{Hz}$ and $T = 6\text{s}$ represent the sampling frequency of CFR and the considered time duration for individual gestures respectively. Thus, our input size of the DNN learning model becomes $G \times S \times K \times M \times N$ where, $G = 10$ represents the total number of gestures, and $K \times M \times N = 1024 \times 8 \times 8$ is the dimension of every CFR sample. To further reduce the computational complexity, we reshaped the input dimension and performed SVD to have a reduced dimension of $G \times S \times (K \cdot M \cdot N)$, i.e., $10 \times 600 \times 2048$ in our case.

In the last decade, convolutional neural networks (CNNs) have achieved tremendous success in addressing computer vision tasks [28, 29]. The convolution layer, at the basis of CNNs, can efficiently extract relevant features from the input by performing convolution operations on the data through learnable kernels. As we aim to investigate the effectiveness of mmWave OFDM CFR as a sensing primitive, we use a widely adopted VGG-based CNN architecture as the classifier for gesture recognition. The network is depicted in Figure 12 and entails stacking three convolutional blocks (conv-block) and a max-pooling (MaxPool) layer. The Softmax activation function is applied to the flattened output to obtain the probability distribution over the gestures. The three conv-blocks have 128, 64, and 32 filters, respectively. We choose a descending order of the number of activation maps to reduce the model size. Moreover, features in lower layers are usually sparser and thus require extracting more activation maps to be properly captured.

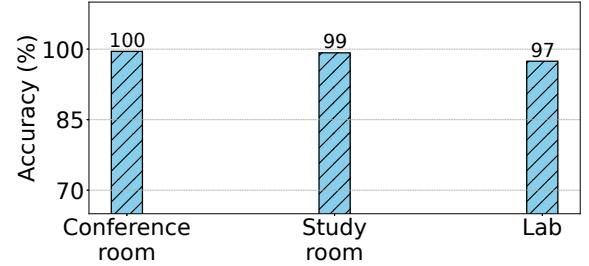
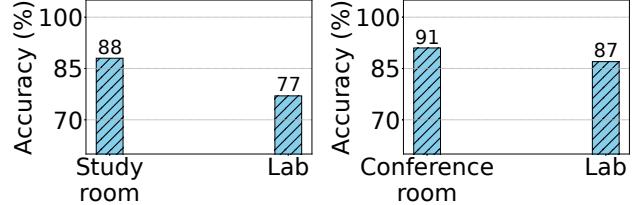


Figure 14: Accuracy of gesture recognition in three different environments.



(a) Trained in Conf. room (b) Trained in Study room
Figure 15: Accuracy of gesture recognition when trained with data from one environment and tested with others.

4.2.2 Experimental Setup and Data Collection. We deploy the m³MIMO testbed with one transmitter and one receiver for an extensive data collection campaign for ten different gestures showing 0 to 9 at three different environments – conference room, study room, and lab. The experimental setup and the data collection campaign in different environments are depicted in Figure 13. Data collection in each environment involves two subjects performing each of the gestures, 50 times each. The CFR corresponding to each gesture is captured and labeled synchronously to create the training dataset. The captured dataset is then processed and fed to the learning model as presented in Section 4.2.1.

4.2.3 Performance Evaluation. Figure 14 shows that the designed gesture recognition system with mmWave CFR achieves an accuracy of up to 99% whereas the average accuracy across three different environments is 98.66%. Moreover, generalization across different domains, i.e., different environments being one of the open challenges in integrated sensing and communication in general, we also present the domain generalization performance. Figure 15 shows that the designed system can achieve an accuracy of up to 91% with an average of 86.21 % when the model is trained with the data from one environment and tested with the others.

5 CONCLUSIONS AND REMARKS

We have presented m³MIMO, an 8 × 8 mmWave MU-MIMO testbed, designed to facilitate advanced research in areas such as MIMO, OFDM, and mmWave integrated sensing. m³MIMO is equipped with three Zynq UltraScale+ RFSoC-based SDR

and supports OFDM transmission and beamforming over a 1 GHz bandwidth. m^3 MIMO's unique capabilities include an integrated channel sounding and precoding procedure, enabling both SU-MIMO and MU-MIMO transmission modes. This integration allows m^3 MIMO to support and validate emerging research algorithms in the mmWave spectrum, thus contributing significantly to the advancement of communication technologies. To demonstrate its versatility, we explored two research use cases with m^3 MIMO: tracking-based beamforming and mmWave sensing through gesture recognition. These examples underscore m^3 MIMO's potential as a comprehensive tool for exploring and testing innovative solutions in OFDM, MU-MIMO, mmWave communications, and integrated sensing applications. Overall, m^3 MIMO represents a significant step forward in providing researchers with a robust and flexible platform for developing and validating next-generation communication systems.

ACKNOWLEDGMENTS

This work is funded in part by the National Science Foundation (NSF) grant CNS-2134973, CNS-2120447, ECCS-2229472, and ECCS-2329013, by the Air Force Office of Scientific Research under contract number FA9550-23-1-0261, by the Office of Naval Research under award number N00014-23-1-2221.

REFERENCES

- [1] Declan McGlynn, Rolling Stones. Music and the Metaverse: Are we on the brink of a virtual artist revolution? <https://tinyurl.com/39ed34dd>, 2022.
- [2] Bokyoung Kye, Nara Han, Eunji Kim, Yeonjeong Park, and Soyoung Jo. Educational Applications of Metaverse: Possibilities and Limitations. *Journal of Educational Evaluation for Health Professions*, 18, 2021.
- [3] VentureBeat.com. Medicine and the Metaverse: New Tech Allows Doctors to Travel Inside of Your Body. <https://tinyurl.com/243a4jnm>, 2022.
- [4] Andrej Somrak, Iztok Humar, M Shamim Hossain, Mohammed F Alhamid, M Anwar Hossain, and Jože Guna. Estimating VR Sickness and User Experience Using Different HMD Technologies: An Evaluation Study. *Future Generation Computer Systems*, 94:302–316, 2019.
- [5] Maria Gallagher and Elisa Raffaella Ferrè. Cybersickness: a Multisensory Integration Perspective. *Multisensory research*, 31(7):645–674, 2018.
- [6] Kyle Orland (Ars Technica). Virtual Perfection: Why 8K resolution per eye isn't enough for perfect VR. <https://tinyurl.com/vckvpbt>, 2013.
- [7] Institute of Electrical and Electronics Engineers (IEEE). Ieee standard for information technology– part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications. *IEEE Std 802.11ax-2021 (Amendment to IEEE Std 802.11-2020)*, pages 1–767, 2021.
- [8] David Lopez-Perez, Adrian Garcia-Rodriguez, Lorenzo Galati-Giordano, Mika Kasslin, and Klaus Doppler. IEEE 802.11be Extremely High Throughput: The Next Generation of Wi-Fi Technology Beyond 802.11ax. *IEEE Communications Magazine*, 57(9):113–119, 2019.
- [9] Moyukh Laha, Dibbendu Roy, Sourav Dutta, and Goutam Das. AI-assisted Improved Service Provisioning for Low-latency XR over 5G NR. *IEEE Networking Letters*, 2023.
- [10] Aditya Dhananjay, Kai Zheng, Marco Mezzavilla, Lorenzo Iotti, Dennis Shasha, and Sundeep Rangan. Pi-Radio v1: Calibration techniques to enable fully-digital beamforming at 60 GHz. *Computer Networks*, 196:108220, 2021.
- [11] Jesus O Lacruz, Rafael Ruiz Ortiz, and Joerg Widmer. A real-time experimentation platform for sub-6 GHz and millimeter-wave MIMO systems. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 427–439, 2021.
- [12] Steve Blandino, Giovanni Mangavitti, Claude Desset, Andre Bourdoux, Piet Wambacq, and Sofie Pollin. Multi-user hybrid MIMO at 60 GHz using 16-antenna transmitters. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(2):848–858, 2018.
- [13] Renjie Zhao, Timothy Woodford, Teng Wei, Kun Qian, and Xinyu Zhang. M-cube: A millimeter-wave massive MIMO software radio. In *Proceedings of the 26th MobiCom*, pages 1–14, 2020.
- [14] Zhenzhou Qi, Zhihui Gao, Chung-Hsuan Tung, and Tingjun Chen. Programmable millimeter-wave MIMO radios with real-time baseband processing. In *Proc. of the 17th ACM WiTECH*, pages 17–24, 2023.
- [15] Ceyhun D Ozkaptan, Haocheng Zhu, Eylem Ekici, and Onur Altintas. A mmWave MIMO Joint Radar-Communication Testbed with Radar-assisted Precoding. *IEEE TWC*, 2023.
- [16] Tsubasa Ochiai, Marc Delcroix, Tomohiro Nakatani, and Shoko Araki. Mask-based Neural Beamforming for Moving Speakers with Self-Attention-Based Tracking. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:835–848, 2023.
- [17] Ha-Lim Song, Young-chai Ko, Jungil Cho, and Chanho Hwang. Beam Tracking Algorithm for UAV Communications Using Kalman Filter. In *2020 ITC*, pages 1101–1104. IEEE, 2020.
- [18] Wang Yi, Wei Zhiqing, and Feng Zhiyong. Beam training and tracking in mmwave communication: A survey. *China Communications*, 2024.
- [19] Sun Hong Lim, Sunwoo Kim, Byonghyo Shim, and Jun Won Choi. Deep Learning-based Beam Tracking for Millimeter-wave Communications Under Mobility. *IEEE TCOM*, 69(11):7458–7469, 2021.
- [20] Yasemin Karacora, Christina Chaccour, Aydin Sezgin, and Walid Saad. Event-based Beam Tracking with Dynamic Beamwidth Adaptation in Terahertz (THz) Communications. *IEEE TCOM*, 2023.
- [21] IoTForAll.com. The Role of WiFi in the IoT. <https://www.iotforall.com/wifi-role-iot>, 2020.
- [22] Khandaker Foysal Haque, Francesca Meneghelli, and Francesco Restuccia. Bfa-sense: Learning beamforming feedback angles for wi-fi sensing. In *2024 IEEE PerCom Workshops*, pages 575–580. IEEE, 2024.
- [23] Francesca Meneghelli, Michele Rossi, and Francesco Restuccia. DeepCSI: Rethinking Wi-Fi Radio Fingerprinting Through MU-MIMO CSI Feedback Deep Learning. In *Proc. of IEEE ICDCS*, pages 1062–1072. IEEE, 2022.
- [24] Baiju Yan, Peng Wang, Lidong Du, Xianxiang Chen, Zhen Fang, and Yirong Wu. mmGesture: Semi-Supervised Gesture Recognition System using mmWave Radar. *Expert Systems with Applications*, 213:119042, 2023.
- [25] Yongsen Ma, Gang Zhou, and Shuangquan Wang. Wifi sensing with channel state information: A survey. *ACM Comput. Surv.*, 52(3), jun 2019.
- [26] Jih-Tsun Yu, Li Yen, and Po-Hsuan Tseng. mmWave Radar-based Hand Gesture Recognition using Range-angle Image. In *2020 VTC-Spring*, pages 1–5. IEEE, 2020.
- [27] Yanhua Zhao, Vladica Sark, Milos Krstic, and Eckhard Grass. Novel Approach for Gesture Recognition using mmwave Fmcw Radar. In *2022 VTC-Spring*, pages 1–6. IEEE, 2022.
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [29] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2014.