



CommRad: Context-Aware Sensing-Driven mmWave Networks

Ish Kumar Jain, Suriyaa MM, Dinesh Bharadia
University of California San Diego
La Jolla, CA, USA
{ikjain,smm233,dineshb}@ucsd.edu

Abstract

Millimeter-wave (mmWave) technology is pivotal for next-generation wireless networks, enabling high-data-rate and low-latency applications such as autonomous vehicles and XR streaming. However, maintaining directional mmWave links in dynamic mobile environments is challenging due to mobility-induced disruptions and blockage. While effective, the current 5G NR beam training methods incur significant overhead and scalability issues in multi-user scenarios. To address this, we introduce CommRad, a sensing-driven solution incorporating a radar sensor at the base station to track mobile users and maintain directional beams even under blockages. While radar provides high-resolution object tracking, it suffers from a fundamental challenge of lack of context, i.e., it cannot discern which objects in the environment represent active users, reflectors, or blockers. To obtain this contextual awareness, CommRad unites wireless sensing capabilities of bi-static radio communication with the mono-static radar sensor, allowing radios to provide initial context to radar sensors. Subsequently, the radar aids in user tracking and sustains mobile links even in obstructed scenarios, resulting in robust and high-throughput directional connections for all mobile users at all times. We evaluate this collaborative radar-radio framework using a 28 GHz mmWave testbed integrated with a radar sensor in various indoor and outdoor scenarios, demonstrating a 2.5x improvement in median throughput and an 8x improvement in 20th percentile throughput compared to a non-collaborative baseline.

CCS Concepts

• **Hardware** → **Wireless devices**; • **Networks** → **Physical links**; **Wireless access points, base stations and infrastructure.**

Keywords

Millimeter-wave, Radar, Sensing-Driven Communication, 5G NR, Tracking, Blockage, Mobility.

ACM Reference Format:

Ish Kumar Jain, Suriyaa MM, Dinesh Bharadia. 2024. CommRad: Context-Aware Sensing-Driven mmWave Networks. In *ACM Conference on Embedded Networked Sensor Systems (SenSys '24)*, November 4–7, 2024, Hangzhou, China. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3666025.3699363>



This work is licensed under a Creative Commons Attribution International 4.0 License. *SenSys '24*, November 4–7, 2024, Hangzhou, China
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0697-4/24/11...\$15.00
<https://doi.org/10.1145/3666025.3699363>

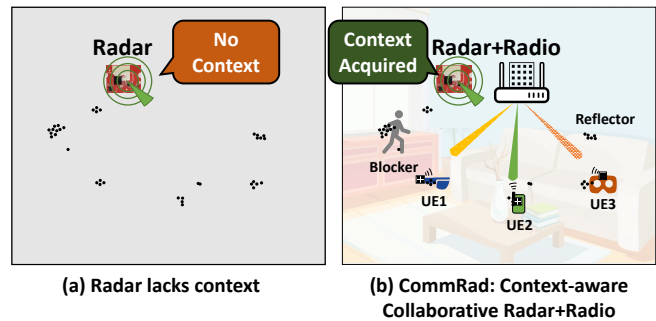


Figure 1: CommRad is a collaborative bi-directional learning framework for sensing-driven mmWave networks.

1 Introduction

Millimeter-wave access is vital for the next-generation wireless network to connect the last mile devices with extremely high multi-Gbps data rates with ultra-low latency [1, 2]. With such high data rates, future autonomous vehicles, loaded with sensors such as Lidar, radars, or cameras, can share their data with the infrastructure or other vehicles for safer driving; the wireless AR/VR devices would enjoy high-quality XR streaming for an immersive experience [3]. However, a major challenge in supporting these use cases is user mobility – anything mobile is an issue for mmWave links. As millimeter-wave links deploy hundreds of antennas to create narrow beams to increase coverage range, they inadvertently fail to maintain these directional links in mobile scenarios. When the target users are mobile, the link requires tracking the users to maintain directional connectivity. Or when other objects in the environment are mobile, they can obstruct or block communication when they enter across an established link. While there are procedures for beam maintenance in 5G NR, such as beam training, where radio scans multiple beams to find the best direct or reflected-path beam for the users, these techniques suffer from high overhead manifested in their long duration and frequent repetitions. The more time a radio spends in beam training, the less time is available for communication, resulting in a high overhead of up to 25% (5 ms duration of beam training repeated every 20 ms, according to 5G NR protocol [4]). To make it worse, the overhead increases with the number of users, as each user demands separate time and frequency resources for feedback in the form of CSI reports, thus lacking scalability in multi-user scenarios.

To overcome the overhead, sensing-driven communication is envisioned as a pivotal technology where a dedicated sensor is deployed to aid in beam maintenance. The sensors can track users, identify good reflectors, proactively avoid environmental blockers, and instruct the radio about the best direct or reflected-path

beam at a given time. Among many choices of sensors, a millimeter-wave radar sensor stands out for low-cost, privacy-preserving, and weather-independent sensing modality in both indoor and outdoor scenarios [5–9]. However, radar-sensor-driven communication suffers from a fundamental cold start problem as depicted in Figure 1. The problem is that the radar yields extensive reflections from many objects in the environment without having any context for those objects – which objects are mmWave users that need to be tracked, or which objects are reflectors/blockers that can assist or obstruct communication links? Due to this contextual deficiency, prior implementation of radar-aided communication is limited to a straightforward clutter-free environment with only one user and does not support reflector-assisted links [10–14]. These works rely on data-driven machine learning techniques to perform end-to-end training from the radar data to predict communication beam patterns directly, skipping the important step of acquiring context. The hope is that the model will somehow learn the context by feeding extensive data while training. To the best of our knowledge, no prior work on radar-driven communication experimentally demonstrated context-aware multi-user tracking while supporting a reflected-path link dynamically and blockage adaptation.

CommRad: This paper presents the design and implementation of CommRad, a sensing-driven communication framework driven by a fundamental insight: the communication radio can serve as a complementary sensing modality, providing the contextual awareness that radar sensing currently lacks. CommRad represents a *collaborative bi-directional learning framework* that seamlessly integrates wireless sensing (radar) and the communication radio as sensing modalities to collectively comprehend the environment and user mobility within a relevant context. In this framework, we first deploy the radio as a sensor to provide the required context to the radar through the mandatory beam training procedure. We then deploy the context-aware radar sensor to track any movement of objects and aid in maintaining the beams while the radio is in the communication phase, as shown in Figure 2. This bi-directional radar-radio collaborative learning reduces the overhead due to frequent beam training and ensures that radar has appropriate context to maintain the beams, ensuring reliable and high throughput mobile mmWave links.

A key observation within CommRad is the recognition that the radio and radar units provide ‘*complementary information*’ essential for comprehending the environment and user mobility. **Radar as a Sensor:** The radar unit measures reflections from the environment and users through *mono-static sensing*, where the radar Rx is co-located and synchronized with the radar Tx. **Radio as a Sensor:** The radio unit conducts *bi-static sensing*, capturing signal transmissions toward users. The radio signal is reflected toward the user transceiver via reflectors instead of returning directly to the radio unit. Leveraging this observation, we develop a collaborative learning framework that combines wireless sensing (radar) and the communication link (radio) to gain a holistic understanding of the environment and user mobility, particularly in the context of beam management. We then develop beam tracking and blockage prediction mechanisms to ensure a reliable link.

Acquiring Contextual Information: A key question is: what is the context that Radar lacks? We define context as the information about the users and environment that radar needs to aid

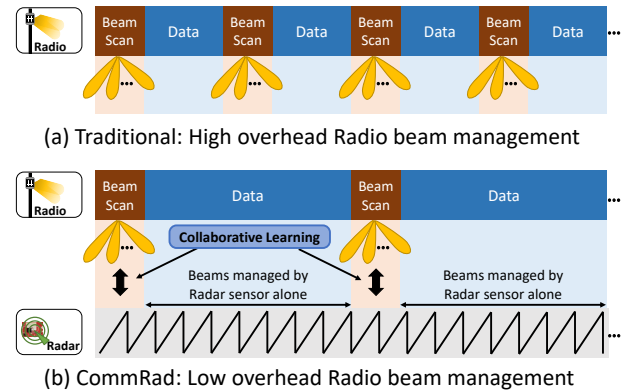


Figure 2: CommRad’s bi-directional radar-radio collaborative learning framework to improve data communication efficiency by reducing radio beam scan overhead.

in communication beam management. For active users, this context information includes their angular and distance parameters relative to the base station, while for reflectors, it is their precise location, orientation, and size (Sec 3.1). Having this context will ensure that radar can identify the user/reflectors in its view and distinguish them from other irrelevant objects in the environment. The next question is how to acquire this context from radio beam training. The base station radio emits a known preamble across multiple beams, received and utilized by active users to estimate Channel State Information (CSI) across angle and bandwidth. We developed algorithms that leverage this two-dimensional CSI to estimate user and reflector attributes and map them to the radar’s view as contextual information as described in Sec 3.2.

Context-aware Beam Management: After acquiring context, the radar can distinguish active users and reflectors from other environmental objects. We then developed context-aware beam management by actively tracking the direct and reflected path angle towards all active users for the duration when the radio is busy with data communication. Our tracking solution is inspired by the spatial-temporal consistency of user motion, i.e., a certain maximum speed of motion typically bounds the user location, and further uses Kalman filtering to reduce the impact of noise and provide accurate tracking (Sec 3.3). A key challenge, however, is when the direct path is blocked, the radio needs to find a reflected path beam to maintain a reliable connection. To ensure reliability, we track the reflected path angle via all reflectors identified in the previous radio beam training. We estimate the reflector parameters (location, orientation, and size) and develop a geometry transformation to derive the reflected path angle as a function of the user’s direct path angle and known reflector parameters. This way, we track the direct and reflected path angles for all time. The base station uses these angle estimates to maintain a reliable beam towards the user via the direct or reflected paths, even in the presence of intermittent blockages. We developed proactive blockage management schemes that detect blockages in advance and instruct the radio to switch the communication beam to one of the unblocked paths to maintain a reliable connection. This is only possible with CommRad since

we actively maintain both direct and reflected path angles using context-aware radar tracking.

Evaluation overview: We evaluate CommRad with an experimental setup that consists of a 28 GHz mmWave radio and 24 GHz COTS ISM band radar with four mobile users. We synchronously collect radio beam scan measurements and radar frames for the duration of the experiment and evaluate CommRad and other baselines on this dataset. The results show that CommRad maintains a reliable and high throughput link while tracking mobile users and proactively handles blockages using radar measurements while periodically collaborating with the radio sensor. The median throughput is improved by 2.5x, while the worse-case 20th percentile throughput is improved by 8x while keeping the communication overhead as low as 0.5%.

Contributions: CommRad makes the following contributions:

- Proposed, for the first time, the requirements and framework for collaborative bi-directional learning between mono-static radar and bi-static radio sensor.
- Developed user and reflector identification technique using collaborative radio beam scan and radar range-angle sensing.
- Designed novel algorithms for both direct and reflected path tracking integrated with proactive blockage mitigation.
- Built a synchronized radar-radio platform with real-world experiments demonstrating 1.5x median and 8x worst-case throughput improvements over a non-collaborative baseline. The artifacts for CommRad are available at <https://wcsng.ucsd.edu/commrad>.

2 Related Work

Radar-aided mmWave systems: Radar provides high-resolution sensing of various objects' location, distance, and velocity in the environment [15–17]. They are low-cost and can operate in ISM bands such as 24 GHz, 60 GHz, or 77 GHz, which does not require spectrum auctioning [18]. FMCW radars have been used for tracking mmWave radio beams using theoretical stochastic geometry models [19], and machine learning [10–14]. These works exploit an end-end machine learning approach where raw range-angle and range-Doppler images are fed into the ML model to get the beam index. Here, only the direct path angle is tracked using the radar, and a single-user system is assumed, which does not address the complexity of multi-user scenarios with reflectors. GPS [12] or mmWave tags [9, 20, 21] can be used to identify object's identity and location, but GPS is unreliable and tags are unavailable on all user devices, reflectors and blockers. All these works are still primitive in handling multiple devices, lack beam tracking over time and integration with mmWave beam scan procedure, and do not focus on system integration. CommRad is the first end-to-end radar-driven mmWave communication system that leverages collaborative learning, user tracking, reflector modeling, and blockage mitigation to maintain a high throughput and reliable link.

Other sensors-aided mmWave systems: LiDARs [22–24] and cameras [25–30] are demonstrated to track directional beams using signal processing and machine learning models. However, these models demand re-training for diverse environments, often involving labor-intensive data collection efforts with ground truth information, yet may still fall short of obtaining the requisite contextual

awareness. In fact, the foundation of contextual awareness laid in CommRad can be applied to these sensing modalities to further improve their performance and generalize them to any indoor or outdoor scenarios. Other sensor-aided communication using device positioning systems [31, 32], light sensors [33], and out-of-band WiFi [34, 35] are less accurate and are mainly developed for indoor WiFi setting which is not common for NextG cellular deployment scenarios.

High-overhead beam training: There is extensive work in beam training methods to reduce the codebook size (scan fewer beams), but since they have to cover an entire angular space to detect a new user or reflector, —which can be anywhere in the angular space— they make a compromise of higher misdetection rates to reduce scanning overhead, especially in multi-user mobility scenarios [36–42]. To reduce this overhead, [43, 44] relies on a pre-determined set of beams that are obtained for a site-specific location (e.g., road or highway) for V2X scenarios, but it does not generalize to arbitrary 6G cellular deployments.

Reflector Modeling: The authors of [45–53] leverage reflective surfaces to improve indoor mmWave connectivity & coverage using only the radio data at multiple static locations. These works require extensive offline training and site modeling using crowdsourcing data collection with ground truth information of user locations to model the reflectors. In contrast, CommRad is designed to operate in real-time without requiring any offline data collection or ground truth information. CommRad achieves this by leveraging complementary information from bi-static radio and mono-static radar sensing.

Blockage Prediction: The authors of [54] developed radar-aided blockage prediction using an end-end machine learning model. BeamSpy [55], Terra [56] handle blockage by stationary objects without modeling mobile links and cannot adapt dynamically to blockage arrival patterns. Multi-connectivity [57] and handovers [58] are alternative ways of recovering from permanent blockages. In contrast to these works, we explore dynamic blockage events where a mobile blocker crosses the link or the mobile user crosses a static/mobile blocking object. Unlike [54], our end-end framework not only predicts blockages but also integrates with direct/reflect-path beam tracking with contextual awareness.

3 CommRad Design

In this section, we discuss the design and implementation of CommRad, a sensing-driven mmWave communication system to aid in mmWave beam management. Managing directional beams in mobile multi-user scenarios requires maintaining the beam direction towards the user either via a direct path or via a reflected path when the direct path is not available due to blockage events. The key information needed for this beam management at the radio is a time series of the best beam angles for all users at all times. The radio can obtain this information through the periodic beam training phase, but this information gets outdated when the radio is not doing beam training, i.e., it is communicating with active users. During this time, the radio relies on the radar sensor to manage its beams. The radar sensor is great for detecting objects in the environment with high resolution in the distance and angle estimates thanks to the high bandwidth available (over GHz) in free ISM bands such

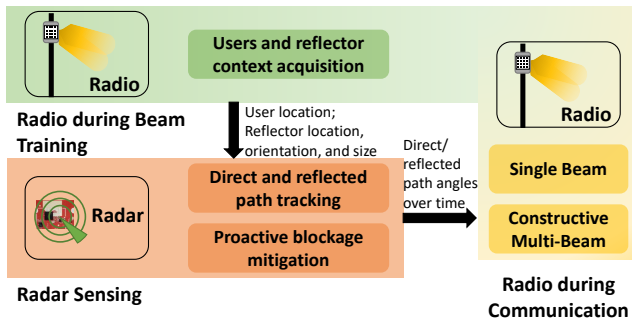


Figure 3: Overview of CommRad’s end-end system implementation of collaborative bi-directional learning with radar+radio integration.

as 24 GHz, 60 GHz, or 77 GHz center frequencies. Therefore, in CommRad, we explore integrating a radar sensor with the radio to aid beam management. To aid in radio beam management, the radar sensor must provide the time series of the best beam angles (direct or reflected path angles) for each user at all times. However, the radar sensor receives an extensive amount of reflections from many objects in the surroundings, and it cannot distinguish which objects are relevant for beam management, such as users, reflectors, and blockers.

In CommRad, we developed a three-stage process to acquire this crucial context and build a bi-directional collaborative framework for beam management, an overview of which is shown in Figure 3. **(Stage 1) What is context?** The first step is to define the context for the radar so that it can aid in beam management. The context should be relevant for identifying users, reflectors, and blockers in the environment and should be understandable by the radar, i.e., mapped into the radar’s range-angle view. We discuss this context and how it is hard to obtain by the radar in Section 3.1. **(Stage 2) How to acquire context?** The second stage is to acquire this context for the radar. More importantly, the context of the users and reflectors should be updated periodically since the radar may lose this context in a dynamic environment where new users may appear or existing users may disappear. We discuss our algorithms for acquiring this context through the periodic radio beam training process in Section 3.2. **(Stage 3) How to do context-aware beam management?** Finally, as the radar is equipped with context, it has to continuously aid mmWave radio in maintaining the beams when the beam training is not immediately available. We discuss how radar aids beam management by tracking the user’s direct and reflected path and proactively detecting blockers in Section 3.3. Actively tracking the reflected path along with the direct path enables us to implement advanced constructive multi-beam techniques [59, 60] to improve link reliability. The overall system is implemented through a synchronized COTS radar and radio platform and evaluated in real-world settings, as discussed in upcoming sections.

3.1 What is the Context that Radar Lacks?

The radar sensor detects different objects in the environment in the form of their distance and angle estimates. It creates a distance-angle view that shows multiple points of reflection from objects with varying strengths. We define the context for the radar as the set of points that belong to the relevant objects for the communication radio, such as users, reflectors, and blockers. Other objects in the environment that do not fall into this category are not relevant to the scope of this paper. **(1) Context for users:** For the users, the context lies in the points in the distance-angle view that belong to an active user for which the radio has to maintain a directional beam for communication. Here, we define context as the approximate {distance, angle} of points that belong to an active user. **(2) Context for reflectors:** For the reflectors, the location is crucial too, but not enough. A reflector is modeled as a line segment, which is characterized by the location (any point on the reflector), orientation (slope of the line segment), and size (the start and end points on the reflectors). Learning these reflector parameters is essential to compute the reflected-path beam angle. **(3) Context for blockers:** We define blockers as any object (e.g., other users in the network, mobile vehicles, or a reflector etc) that is relatively mobile with respect to the user to cause blockage for a certain duration. We define the context for the blockers by their relative location, size, and speed, which is necessary to estimate the arrival time and duration of blockage events. Without this context information, the radar sensor cannot aid mmWave radio in direct/reflected path beam maintenance.

3.1.1 How Radar Lacks Context for Users?

The radar sensor detects multiple objects in the environment, but it does not know which objects belong to an active mmWave user. Without this crucial context, Radar cannot track user mobility and aid in beam maintenance. To demonstrate this limitation, consider a scenario with two mobile users in the environment crossing each other, i.e., one user is moving left to right, and the other is moving right to left in front of the radar. The radar can detect two users and track their angle over time. Without context information, the radar cannot distinguish between the two users and incorrectly assumes that both users continue along the same trajectory after the crossing, as shown in Figure 4(a). In reality, they move in opposite directions. This confusion can be resolved by periodically incorporating contextual information from the radio sensor, resulting in accurate mapping of user trajectories to user ID, as demonstrated in Figure 4(b).

3.1.2 How Radar Lacks Context for Reflectors?

The next challenge for radar is determining which objects in the radar’s view serve as useful reflectors. Simply gauging signal strength on the radar to locate strong reflectors is not feasible, as objects that appear as a strong reflector for radar could be weak for communication. To illustrate this, we conducted an over-the-air experiment with our testbed involving varied reflector orientations. We plotted the Received Signal Strength (RSS) at both the radar and radio, as depicted in Figure 5(a). When the reflector is positioned at 0 degrees orientation, it optimally serves the user with the highest RSS, whereas for other reflector orientations, RSS diminishes as shown in Figure 5(b). In contrast, the radar obtains its highest RSS when

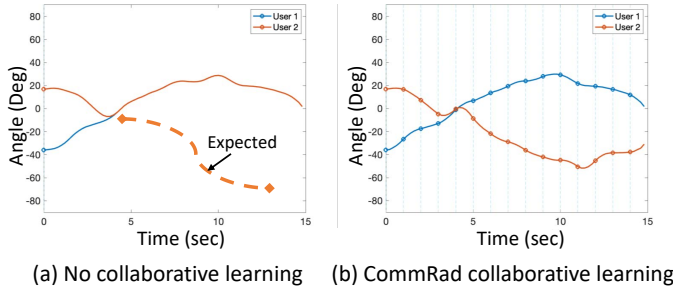



Figure 4: Challenges with lack of context for user identification. Radar could not distinguish between two users when they cross each other.



(a) Varying reflector orientation experiment setup (b) RSS with varying reflector orientation

Figure 5: Challenges with lack of context for reflector identification. Radar cannot detect useful environmental reflectors without collaboration with Radio.

the reflector is oriented at 60 degrees, a configuration that does not effectively serve the user through the reflected path, adhering to the law of reflection. We show in later sub-sections how radar-radio collaboration helps in obtaining the context for the reflectors and aiding in reflected beam maintenance even under user mobility.

3.2 Acquiring Context through Radar+Radio Collaboration

So far, we have seen that the radar sensor alone cannot obtain contextual information about users and reflectors independently. This section describes how contextual information can be obtained using our radar+radio framework. This contextual information is then used to provide context to radar to map different objects as users or reflectors.

3.2.1 Acquiring Context for Users

Acquiring user context is a process of identifying which points in the radar’s distance-angle view are active users. We use the radio beam scan procedure to estimate the user’s angle and distance from the base station and then map it to the radar’s view. First, consider a simple multi-user scenario without any reflectors and blockages and later analyze their impact. During the beam scan process, the base station transmits the preamble through multiple beams in different directions. Each user captures this signal and estimates the RSS per beam. The beam index that gives the highest RSS is identified as the best beam that serves the individual user.

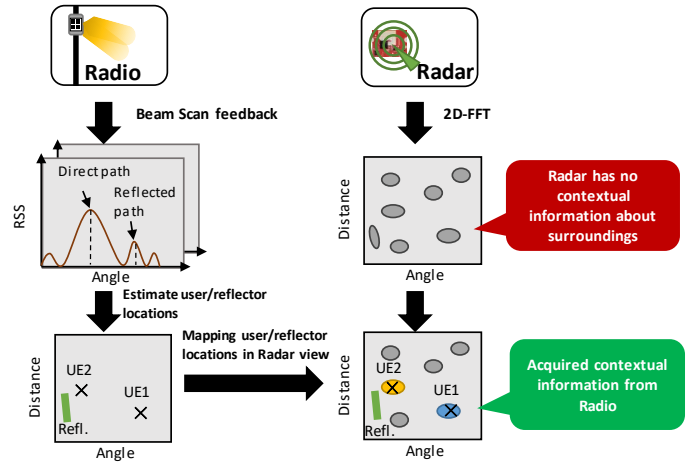


Figure 6: Identifying users and reflectors in radar view using contextual information from radio.

The direction at which this beam points is identified as the user’s angle. We further estimate the user distance using the observation that the distance is inversely proportional to the RSS using Friss equation [61]. These estimates are used to map the user’s location in radar view as shown in Figure 6.

3.2.2 Acquiring Context for Reflectors

Reflector identification is a crucial process in CommRad, as reflectors are key in maintaining reliable communication [59, 60]. With its monostatic sensing capabilities, radar faces challenges in independently identifying reflectors, which contrasts with the bistatic nature of communication systems. In CommRad, we leverage distinctive attributes from both radar and radio to acquire reflector identification.

Reflectors can be categorized as point reflectors (e.g., poles and lamp posts) or surface reflectors (e.g., buildings, glass walls, and monitor screens). Surface reflectors are particularly relevant for mobile links since their reflections are well-defined and predictable, unlike point reflectors, which are better suited for static users as they struggle to sustain mobile links. Therefore, our focus in this paper is primarily focused on surface reflectors.

We represent surface reflectors as line segments in Cartesian coordinates, characterized by their location (any point on the reflector), orientation (slope of the line), and size (start and end points on the reflector). We estimate these parameters leveraging two-dimensional channel measurements across angle and bandwidth, following a three-step process as follows:

(i) Estimating reflected path angle and distance: As a first step, we must separate the direct path and reflected paths in the radio RSS measurements. We apply a variation of the MUSIC algorithm, “Subspace MUSIC,” to separate the two paths across the angle domain, as we do not have per-antenna measurements in mmWave systems with analog beamforming arrays. We then apply MUSIC across bandwidth to obtain relative Time of Flight (ToF) information between the two paths, following the classic procedure in localization [62]. The direct path is identified as the one with the lower ToF, corresponding to the shorter signal travel distance. By




Figure 7: Reflector location estimation using radio beam scan measurement.




Figure 8: Reflector orientation and size estimation.

utilizing the relative ToF and actual user location obtained through RSS measurements, we determine the angles and distances of both the direct and reflected paths. The procedure is explained in Figure 7.

(ii) **Estimating reflector orientation:** After obtaining the angles and distances of each path, we construct an ellipse with the base station and user location as its foci. The reflector’s position locus lies on this ellipse, as shown in Figure 8(a). To pinpoint the exact reflection point on the ellipse, we find the intersection between the ellipse and the line extending from the base station towards the reflector angle. This intersection point serves to identify the orientation of the reflector as the slope of the tangent line to the ellipse at that location.

(iii) **Estimating reflector size:** To complete the reflector characterization, we need to estimate its size, which is crucial for tracking the reflected path angle. We address this challenge through a learning process over time, as demonstrated in Figure 8(b). We collect multiple beam scan measurements from various user locations as they move around and repeat steps (i) and (ii) to identify all the points on the reflectors. In this way, we obtain the reflector size as a set of two endpoints on the line segment obtained in this process. Notably, this is a one-time process for static reflectors but can be repeated if changes occur in the reflector configuration.

Problem of unreliable RSS measurements: It is essential to note that the reflector identification process may accumulate errors if underlying parameters are inaccurate. A key source of error is distance estimation from RSS measurements, especially in multipath environments. We have observed that radar distance measurements are more accurate, relying on time-of-flight measurements from a synchronized Tx and Rx architecture. Thus, our approach initially

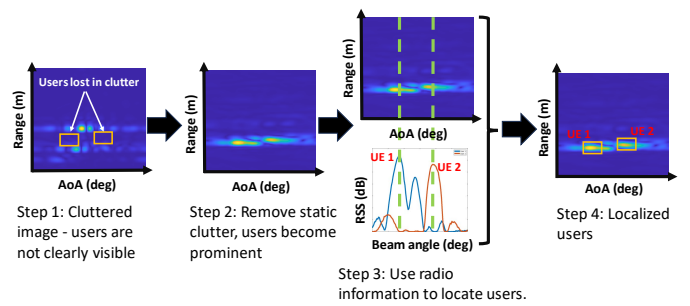


Figure 9: Clutter removal and user localization.

obtains course distance estimates from radio measurements and then refines them by identifying the nearest object in the radar view. This dual-step *radar-radio collaborative learning* process ensures accurate distance and angle measurements, enhancing reflector identification accuracy.

3.3 Context-Aware Beam Management using Radar Sensor

After identifying users and reflectors in the radar view through radar-radio collaborative learning, we next use the radar sensor to track the user motion through changes in both the direct path and reflected path over time. It is important for radar sensors to maintain and aid direction beams while the radio is occupied with data communication and not performing beam scans. Our approach to the direct path and reflected path tracking is described in this subsection.

Before discussing our tracking algorithm, we present a primer on the choices of range and resolution, followed by the techniques used to counter the challenges described in the previous sections.

Maximum Range and Range-Resolution:

The range resolution for FMCW radar is given by $\Delta R = \frac{c_0}{2B}$. Comm-Rad uses a bandwidth of 200MHz, thus resulting in a range resolution of 0.75m. With a resolution under 1m, we are able to capture the movement of the user across time accurately. The maximum range of the FMCW radar is given by $R_{max} = \frac{c_0}{2B} N$, where N is the number of samples corresponding to the duration observed. For an N value of 1000, we get a maximum range of 750 m.

Angle resolution: The angle resolution of a antenna array with N_{ant} is given by, $\theta_{res} = \frac{\lambda}{N_{ant}}$. For a 16 virtual antenna system, angular resolution is 7.16° , whereas for a 8 antenna system, angular resolution is 14.32° . By using a virtual antenna array, we achieve better angular resolution, which helps improve the localization and tracking accuracy of the system.

3.3.1 Static Clutter removal with frame subtraction

The Range-AoA representation from radar contains high-energy points corresponding to all environmental reflections, including users, reflectors, blockers, etc. While tracking the user, all static reflections coming from the environment can be considered as clutter. The presence of clutter causes ambiguity in distinguishing the users over time. To remove clutter, we use the static nature of those objects and cancel them across time. Each frame, corresponding to one chirp, is 1ms apart from the consecutive frame. Depending on

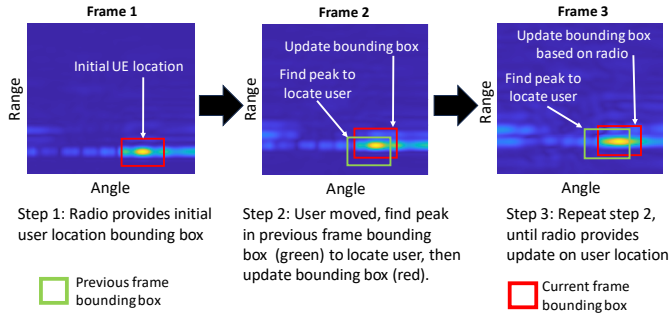


Figure 10: Proposed user tracking algorithm.

the type of mobility (users in indoor scenarios have lesser mobility than in outdoor vehicular scenarios), we perform an average over multiple chirps and subtract the Range-AoA profile of the current frame from the averaged Range-AoA profile. This removes the static objects, and user tracking can be done accurately and reliably. We show the above-mentioned workflow in Figure 9. We first obtain the clutter-removed image using frame subtraction and then use the radio data to localize the users within this image.

3.3.2 Direct path beam tracking

For direct path tracking using the radar sensor, we have devised a three-step process illustrated in Figure 10. The first step involves identifying the user's initial location within the radar view and creating a bounding box around each user. Creating a local bounding box around a user allows us to remove the dynamic clutter in the environment, like other users, blockers, etc., that may not be present near the user. In the second step, we examine the subsequent radar frame to locate the user with the highest signal strength while retaining the same bounding box as in the previous step. Finally, in step 3, we adjust the bounding box to align with the new user location identified in step 2. These last two steps are repeated for all consecutive frames to ensure continuous user tracking. By tracking the user with radar, we can still maintain directional beams toward the user, even when the radio resources are unavailable for beam training. Periodic updates of user location from the radio can then help maintain the tracking accuracy of the radar, especially when the radar loses the context of the users who might be too close.

3.3.3 Reflected path beam tracking

Reflector tracking is more complicated than user tracking, as mobile users change the angle of the reflected path. We need to know the reflector's location, size, and orientation to track the reflector. Once this information is available, we can create a virtual representation of the base station relative to the reflector as an object behind a mirror. The line connecting this virtual base station to the user defines the angle of the reflected path. Let the user's location at time t_1 be denoted as (x_1, y_1) , and at a subsequent time t_2 , the user's location changes to (x_2, y_2) (See Figure 11). Assuming the radar and radio system is located at the origin $(0, 0)$ and has a virtual location of (x_v, y_v) , we can estimate changes in the reflected path angle, denoted as $\Delta\phi_r$, as follows:

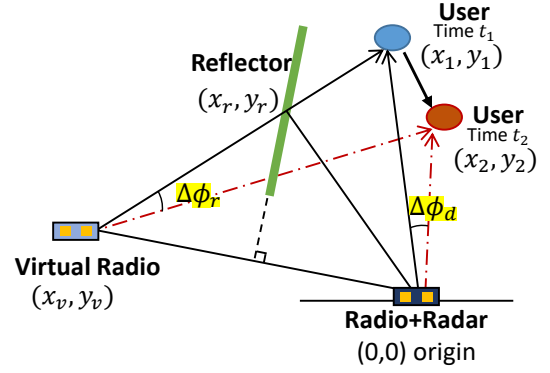


Figure 11: Tracking both direct path and reflected path.




Figure 12: Blockage tracking and proactive prevention.

$$\Delta\phi_r = \tan^{-1} \frac{y_2 - y_v}{x_2 - x_v} - \tan^{-1} \frac{y_1 - y_v}{x_1 - x_v} \quad (1)$$

Here, the virtual location of the radar/radio system x_v, y_v is unknown, which is dependent on the reflector's location and orientation, defined as:

$$(x_v, y_v) = (\tan(\phi), -1) \times \frac{2(y_r - \tan(\phi)x_r)}{1 + \tan(\phi)^2} \quad (2)$$

Where (x_r, y_r) represents an arbitrary point on the reflector, and ϕ denotes the reflector's orientation angle.

With the estimate of the reflector orientation, we can track the change in reflected path angle $\Delta\phi_r$, which can be used as an alternate beam path to aid communication when the direct path is blocked or unavailable for communication.

3.3.4 Context-Aware Blockage Mitigation using Radar Blockage causes trouble to the mmWave link by reducing the signal strength and bringing the link to an outage. Therefore, to revive the link under blockage, alternate reflected paths that are not blocked are utilized. When all the paths are blocked, the link is unrecoverable. To recover the link for partial blockage, it is important to learn the blockage arrival and departure process promptly. Prior work on blockage mitigation is either reactive in nature, which acts when the link is already down, or they rely on radio alone, compromising accuracy in a large multi-user environment. Our proposed collaborative learning framework takes care of the blockage events. By leveraging unique features of radio and radar sensing, we create

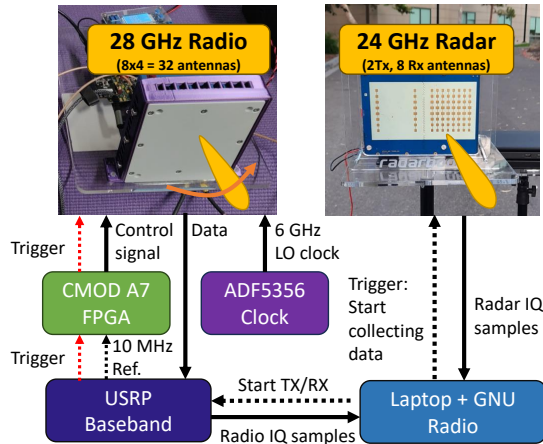


Figure 13: Base station with synchronized Radio and Radar setup. We deploy 28 GHz (n257 band) Radio and 24 GHz (ISM band) Radar.

a signal processing pipeline to promptly detect blockage and take precautionary actions such as switching beams to a reflected path.

Defining a blockage-prone region: The first step in detecting blockage with the help of radar is to determine a blockage-prone region. We define a blockage-prone region as a rectangle patch between the user and the base station, as shown in Figure 12. Any blocker in this region can cause blockage. The radar tracks an object’s velocity over frames to estimate the time-of-arrival of the blocker into the blockage-prone region. The blockage duration is estimated from blocker velocity and blocker size. A blocker of length l_B and velocity v_B would block the link for the duration of l_B/v_B .

Proactive blockage mitigation: With the blocker velocity and duration known, the base station alerts the radio of the upcoming blockage event. The response to the blockage depends on the duration of the blockage event. If the blockage event is too slow and lasts longer, then the base station switches to one of the reflected beams. The reflected beam is served until the blockage event lasts. After the blockage is gone, the base station switches back to the original link.

4 Implementation

The implementation of the entire system is made in a way that enables experiments that can be conducted both indoors and outdoors. We describe the hardware and software used, the implementation of each subsystem, and the integration of these subsystems to create the joint synchronized platform as shown in Figure 13.

4.1 Radio Hardware and Software

The mmWave radio system is powered by a combination of commercial software-defined radios (SDRs) and custom-designed phased arrays. The SDRs are connected to the PC using SFP optical cables or USB cables. UHD and GNU-Radio are used to control the SDRs.

Phase Array specifications and control hardware: The phased arrays have 32 antennas in a 4x8 configuration [63]. Each antenna is dual-polarized, so we have 64 antennas. The antenna RF

input and outputs are connected through SMA connectors to the SDRs and operate at 28GHz, which is up/down-converted from/to the 4GHz baseband signal with a 24GHz UDC. We use an Artix-7-based FPGA [64] for programming the phased array registers and activating beam training commands using the SPI protocol. The registers are modified to set the phase and gain values (supports 6-bit phase and 6-bit gain values) obtained from MATLAB corresponding to a particular beam. Look Up Tables (LUTs) present in the phased arrays allow for storing the phase and gain values in a continuous register location that can be loaded rapidly without microcontroller intervention every time. The phased array is capable of performing a beam switch under $6\mu s$. The FPGA, SDRs, and phased arrays are fed with a common 10MHz reference signal generated from the Ettus OctoClock clock distribution module that ensures sample level synchronization within the radio. A power regulator board takes 12V input [65] and distributes power to phased arrays, clock modules [66] and fans.

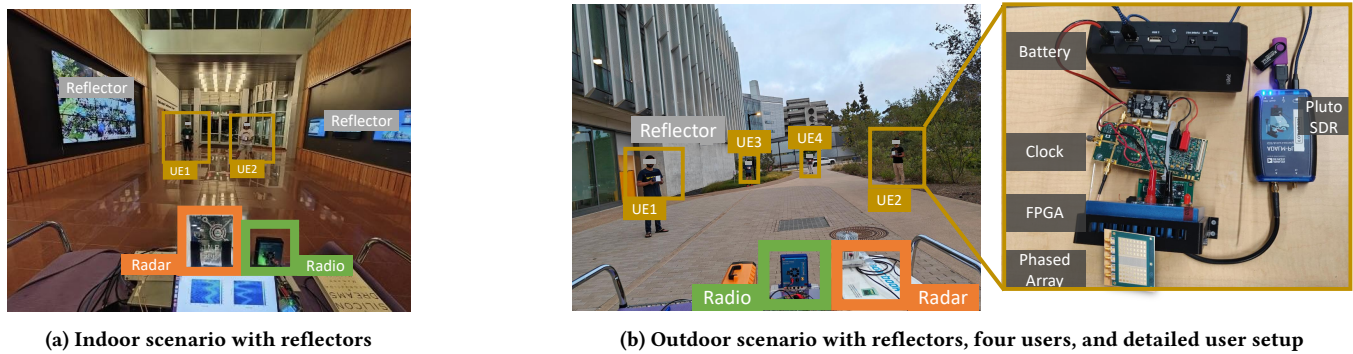
Base station setup: For the base station, we use Ettus N310 USRP [67] that is connected to a PC through SFP optical cables. We operate the USRP at 4.05GHz center frequency and a sampling rate of 30.72MHz using GNU-Radio software [68]. We program the pre-calculated phase and gain values for the required angles for an experiment into the LUTs present in the phased arrays. When the USRP starts transmitting, a trigger is sent to the FPGA, which then activates the beam training, and the phased array loads the register values from the LUTs sequentially.

User setup: We developed a compact mobile user equipment (UE) setup, consisting of a mmWave phased array, clock, Arduino micro-controller, a battery, and PlutoSDR as demonstrated in [69–71]. We choose plutoSDR [72] for the user because it is compact and allows streaming data directly from a USB stick. The UE stick contains raw IQ samples with OFDM waveform that can be transmitted using PlutoSDR without needing a laptop, thus making it easier to scale to mobile multi-UE experiments. We use frequency-division multiplexing to avoid interference among multiple UEs. The phased array is controlled through an Arduino microcontroller to set a fixed quasi-Omni-directional beam at each user.

Waveform design: We use a modified 5G NR beam management protocol with a $12.5\mu s$ OFDM symbol with the preamble and data blocks. We choose this duration of OFDM symbol to get an integer number of samples at the sampling rate of 30.72MHz and allocate enough samples for the preamble. The preamble contains a typical 64-sample pseudorandom sequence, which is repeated twice and has a 32-sample cyclic prefix (CP). The preamble is followed by data that contains the repeated beam IDs that are also coded with a pseudorandom sequence to avoid the high Peak-to-Average-Power-Ratio (PAPR) prevalent in OFDM systems.

4.2 Radar Hardware and Software

We used an INRAS RadarBook FMCW radar, which operates at 24GHz [18]. The radar is connected to the same PC as the base station with a 1 Gbit Ethernet cable. The radar has 2 TX and 8 RX antennas that are combined in post-processing to give a 16-antenna virtual array. The radar uses a bandwidth of 200MHz, with one chirp containing 1000 samples, with a slope of $2\text{ MHz}/\mu s$ having a ramp-up time of $100\mu s$ and a total duration of $200\mu s$. Each data frame



(a) Indoor scenario with reflectors

(b) Outdoor scenario with reflectors, four users, and detailed user setup

Figure 14: Experiment scenarios with the static radar-radio platform with four mobile users. We consider diverse scenarios indoors and outdoors with various reflectors such as concrete walls, monitor screens, glass walls, etc.

contains 200 chirps, and one frame lasts 200 ms. The chirps are split equally between the two TX antennas, i.e., 100 chirps for each antenna. For the above parameters, we have a range resolution of 0.75m, velocity resolution of 0.75 m/s (14 Hz frequency resolution), and angular resolution of 7.16 degrees.

■ **Synchronization and calibration:** Radio and radar each save timestamps using the NTP protocol during data collection for synchronization purposes. The radio achieves sample-level synchronization via the USRP, which sends triggers to the FPGA precisely when transmission or reception starts. This, in turn, triggers the phased array to begin beam-sweeping. By chaining triggers from the USRP to the FPGA to the phased array, we align the beam training with the initial sample transmission or reception. Figure 13 illustrates the clock and trigger signals used for synchronization and control. Additionally, calibration is necessary because the radar and radio, positioned about 15 cm apart, have different local coordinate systems. We establish a global coordinate system at the radio’s location and apply a linear transform to convert radar points into this global coordinate system, similar to the well-known multi-camera calibration method [73].

5 Evaluation

We evaluated CommRad with up to 4 mobile users at a range from 5m to 80m in various indoor and outdoor scenarios as shown in Figure 14. The users are human volunteers who move at a walking or running speed of around 3 to 6 mph. Each volunteer carry our mobile UE setup which consists of a PlutoSDR with USB stick that continuously stream OFDM waveform in uplink with a fixed Quasi-Omni-directional beam pattern. We evaluate uplink scenarios while scanning beam patterns at the base station in the range of -60 to 60 degrees with an increment of 1 degree to collect accurate ground truth data.

■ **Oracle:** We consider Oracle as an entity that knows the best beam at all times for all users. The Oracle is obtained by performing a fast beam-scan using the base station phased array. The phased array scans 121 beams, with each beam lasting 12.5 μ s. Thus, the total time required for a beam scan is 1.5 ms. Note that this beam scan is faster than 5G NR protocol which allow only 64 beams over an interval of 5 ms. This procedure is executed for the entire duration of the experiment to capture the robust ground truth /

Oracle data from the perspective of the radio. The Oracle perfectly captures the reflected beam when the direct path is blocked.

■ **Baseline:** We implement a non-collaborative baseline that does not utilize the periodic contextual information from the radio for aiding radar tracking. To ensure fairness, we provide a one-time context for radar measurement at the beginning of tracking and let the radar track the user without context for the remaining duration of the experiment. Here, the radar follows a similar approach for tracking by creating bounding boxes around the users after the radio gives the initial locations. Then, the radar tracks the user by updating the bounding box, but it does not recalibrate with the radio periodically. It essentially does the tracking on its own. We also compare the results with a reactive baseline, where the radio holds the beam until the beam is updated in the next beam training and has no input from the radar.

We first present end-end results with CommRad and then benchmark individual components of CommRad separately.

5.1 Overall Throughput Improvement

The goal of this evaluation is to show the throughput improvement of CommRad in diverse multi-user multipath environments. We obtain radio context information periodically, ensuring the radar can track and maintain beams in this duration. We obtain a CDF of throughput measurements of CommRad and compare it against the non-collaborative baseline and Oracle in Figure 15(a). Here, we combined all the measurements across different environments and mobility patterns. The median throughput of CommRad is 1000 Mbps, while the baseline offers only 400 Mbps median throughput, thus providing a 2.5x throughput gain over the baseline while performing almost similar to the Oracle system. The throughput improvement is more significant for the worst case, the 20th percentile of the cases, improving from 50 Mbps to 400 Mbps, providing a staggering 8x throughput gain. These results show that CommRad can help maintain a reliable link in difficult situations like device mobility and blockages. These improvements are due to robust radar tracking and blockage mitigation enabled by context awareness in CommRad’s radar-radio collaborative framework.

Further, we evaluate the throughput improvement using traditional single-beam and constructive multi-beam techniques [41, 60, 74–77]. Constructive multi-beam uses two beams simultaneously,




Figure 15: End-end results of CommRad: Improving median throughput by 2.5x compared to non-collaborative baseline. Supporting both single-beam (SB) and Constructive multi-beam (CMB) schemes with blockage mitigation. Reducing beam management overhead in 5G NR.

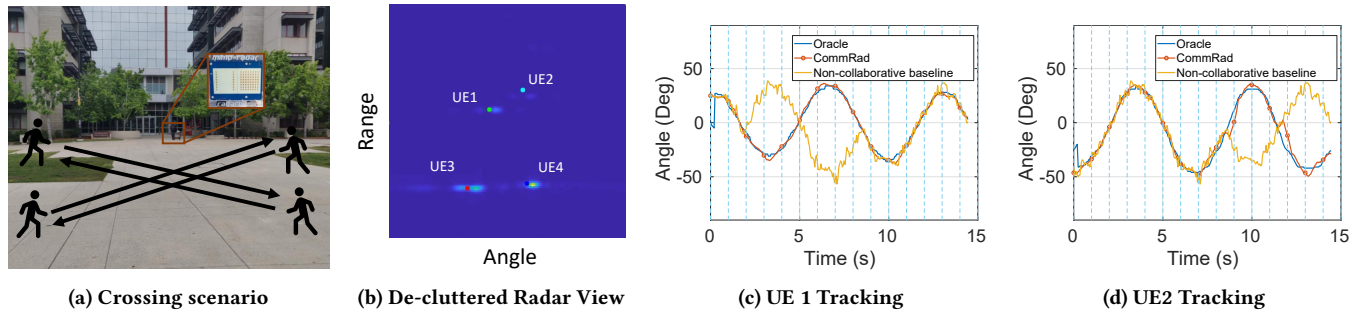


Figure 16: Scenario: Four users crossing each other. We remove the clutter and then show angle-tracking accuracy with collaborative learning.

one in the direct path and the other in the reflected path direction, to provide reliable communication during blockages [60]. The box plots in Figure 15(b) show that the mean throughput increases by about 100 Mbps without blockers. With blockage, the traditional single-beam experiences a significant drop in throughput due to the direct LOS path being blocked. In contrast, the constructive multi-beam, which uses the reflected path to support the communication even during blockage to direct LOS path, experiences a drop in throughput by half, which is less significant compared to almost a 10x drop in the single beam case.

5.2 Overhead Reduction

Next, we emphasize the advantage of using radar sensing to aid mmWave communication to reduce beam scan overhead. In cellular protocols like 5G NR, the base station (gNB) needs periodic beam scans and constant feedback for establishing and maintaining the beams towards the user, which creates a significant amount of overhead in terms of radio resources, affecting throughput and latency. Based on the periodicity of beam training in 5G NR, this overhead could be between 2.5% to 25%. CommRad can avoid this high overhead in many scenarios encountered in 5G NR. During initial access, the radar can indicate the arrival of a new object in the environment and its angle, thereby allowing gNB to scan only a small subset of angles. This can serve as a substitute for the

beam refinement procedure suggested in 5G NR. During mobility, CommRad provides the optimal angle to the gNB, thereby reducing user feedback overhead. Finally, during blockages, while CommRad can provide alternate reflected paths, during the absence of strong reflectors, radar can provide the user location immediately without needing to re-establish the link. We show that CommRad significantly reduces this overhead down to 0.5% when it relies on radar-based beam management with a periodicity of 0.5 seconds. At a higher periodicity of 2 seconds, this overhead is further reduced by 0.25%. This shows CommRad as a low-overhead mechanism for the radio to spend more time on communication than spending time for beam scanning.

5.3 Angle Tracking Accuracy

Accurate user tracking is the most important aspect of the system on which many features like reflector estimation, tracking, and blockage prediction impinge on. One of the main reasons for throughput gain for CommRad is because of its capability to track user angle accurately. We evaluate this by creating a link between the user and the base station radio and performing experiments where the user moves in various directions and speeds. We use radar to track the mobile user along with periodic updates from the radio.

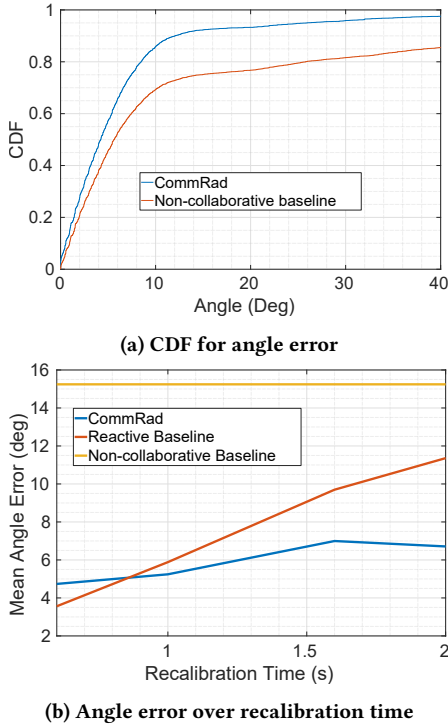


Figure 17: CommRad’s end-end angle tracking accuracy.

An example is shown in Figure 16 with 4 users crossing each other diagonally. Here, we compare CommRad with a non-collaborative baseline comprising only radar tracking without collaboration with the radio. We observe that in complex scenarios where two users overlap, the non-collaborative baseline loses track of the user’s identity and reports the same path for both users. With the periodic collaboration with radio, this effect is mitigated in CommRad.

To show the accuracy in user angle tracking, we plot the CDF of the error in angle prediction for CommRad compared to the non-collaborative baseline in Figure 17(a). The angle error is the absolute value of the difference in Oracle angle and CommRad/baseline angle. We show that the median angular error of CommRad is 4 degrees while the baseline suffers from 8 degrees of median error. The 90th percentile error is more significant, where the CommRad error of 10 degrees is significantly better than the baseline error of more than 40 degrees. Figure 17(b) shows the trade-off between angle error and recalibration time. The higher the recalibration time, the higher the angular error, consistent with the results expected from radar tracking objects without context for a long time. We also observe that around 2s, the angle error is around 7°, within the beamwidth of 10° of the radio with 8 antennas. Therefore, an ideal trade-off between overhead reduction and reliable communication can be achieved when we update the user location at 2s intervals, causing the overhead reduction up to 0.25% as estimated previously.

5.4 User/Reflector Tracking Accuracy

We benchmark the accuracy of our radar tracking algorithm for both direct and reflected paths. For comparison, we obtain the ground




Figure 18: CommRad acquires ground truth by performing continuous beam training at all times. We show that the angle acquired by this beam scan (Oracle) matches the radar prediction.




Figure 19: CommRad adaptation to dynamic blockages.

truth angle through a continuous beam training process described earlier. Figure 18(a) shows the RSS power across time and angle. This represents the ground truth (Oracle) information, where the high energy corresponds to the angle path traveled by the user over time. The second higher energy path corresponds to the reflected path seen by the radio. In Figure 18(b), we compare CommRad’s angle tracking with the Oracle angle and show that they match with high accuracy. The reflected path tracking is remarkable as the radar could accurately predict when the reflector is available to the user at the 2-sec mark and when it vanishes at the 7-sec mark due to user mobility. This high accuracy is achieved thanks to accurate context acquired by radar in reflector location, orientation, and size.

5.5 Blockage Prediction and Mitigation

Blockage for mmwave links reduces SNR and throughput and causes link outages. Utilizing alternate paths to maintain the link while the direct path is blocked is one of the proven solutions to prevent link outages. Since the potential reflected path can be anywhere in the entire angular space, it becomes tedious to do a full beam scan. Instead, we rely on radar to detect the blockage event and identify potential reflected paths to serve the user during the blockage. To demonstrate this, we show a 10-sec trace of adapted angles and corresponding RSS in Figures 19(a) and 19(b), respectively. We compare CommRad’s angle adaptation with a direct path baseline without any adaptation. As we can see, there is a sharp drop in

angle twice at the 3-sec and 5-sec mark, which indicates a blockage event and beam switching to a reflected path. This adaptation in CommRad helped maintain high RSS, while the non-adaptive baseline suffered an RSS loss of more than 10 dB due to blockage. This shows the feasibility of radar-aided beam management and blockage adaptation in mmWave systems.

5.6 Cost and Computation Complexity

CommRad deploys an additional radar sensor to aid in mmWave base station beam management, which leads to additional cost and computation complexity. In the radar processing pipeline, the computation is dominated by 2D-FFT to get the Range and Angle-of-Arrival profile. The computation complexity for 2D-FFT is $O(MN \log_2(MN))$, where M is the number of distance bins and N is the number of angle bins. Considering $M = 1000$ and $N = 120$, we get 0.02 ms computation latency on a typical 100 GFLOPS CPU. Other computational tasks such as estimating reflector size and orientation are based on closed-form expression, thus having $O(1)$ complexity. We discuss the additional cost, power, and resolution with integrated radar sensors at the base station and compare them with alternative sensors such as Lidar and cameras in Table 1.

	Cost	Power	Resolution
Radar [18, 78]	\$200 – \$500	1 – 2 W	Low ($\sim 6^\circ$)
Lidar [79]	\$300 – \$1000	5 – 15 W	High ($\sim 1^\circ$)
Camera [80]	\sim \$300	\sim 2 W	Medium ($\sim 14px/deg$)

Table 1: Cost, power, and performance comparison.

6 Limitation and Future work

We discuss potential limitations and future directions of CommRad:

5G NR Compatibility: 5G NR requires an SSBurst between 20 ms and 160 ms to enable fast initial access, but at the cost of high beam scan overhead. To reduce this overhead while maintaining 5G NR compatibility, CommRad can be adapted to perform full beam scan over 1-2 seconds and perform a low overhead partial beam scan every 160ms. The partial beam scan could scan fewer beams per SSBurst guided by the radar sensor. CommRad also requires complex CSI from the UE for each beam, which may utilize additional signals such as CSI-RS in 5G NR [81]. We leave this exploration to make CommRad 5G NR compatible for future work.

LoS Assumption: We assume that the LoS path between the BS and the UE is available for our path estimation algorithm to work. If the LoS path is not available, then CommRad could misclassify an NLoS path as a direct path and result in wrong estimation. After this initial reflector estimation process, CommRad works in NLoS scenarios with proactive blockage mitigation.

Mobile Reflectors and Scatterers: CommRad models static reflectors as a plain surface of finite size to estimate their location, orientation, and size. Mobile reflectors or scatterers (e.g. corner reflectors) could also help with NLoS link. However, mobile reflectors are not persistent and cannot provide long-term link reliability and throughput since they could move out of the communication range

in a short period of time, and may lead to an increased number of false beam switchings.

Static and Mobile Blockers: CommRad proactively detects dynamic blockage events and estimates the relative speed and size of the blockers in relation to the user. It adapts to various scenarios, including a static blocker with a mobile user (e.g., the user moving behind buildings), a mobile blocker with a static user (e.g., an obstacle crossing an established link), or a combination of both. By leveraging relative velocity, CommRad can estimate the arrival and departure times, as well as the duration, of blockage events. This insight makes CommRad scalable to real-world multi-UE and multi-blocker environments.

3D Beamforming: CommRad is tested with only azimuth direction beamforming at the base station in scenarios where the elevation angle is broad enough to cover the coverage range. The UE is Quasi-Omni directional due to power and size constraints for hand-held devices. Expansion of CommRad to 3D beamforming at the base station and UE is an important future direction.

Sensor-Fusion with Contextual Awareness: CommRad deploys a 24 GHz Radar sensor in the ISM band to aid in beam management with contextual awareness. While Radars stand out as low-cost, privacy-preserving, and weather-resistant sensing modality, they can be fused with other sensors, such as radars at 60 GHz or 77 GHz ISM bands or light-based sensors such as cameras or lidars, to further improve beam management. Notably, CommRad’s framework for contextual awareness can be generalized to other sensing modalities. Deploying these sensing-driven solutions in dense mmWave network requires additional multi-site collaboration and interference management.

7 Conclusion

In this paper, we investigate a collaborative bi-directional learning framework through the collaboration of radar and radio to improve next-generation mobile mmwave links. We identify complementary sensing information from mono-static radar and bi-static radio and combine them under our framework to get the best of both worlds. As a result, we improve mmwave link quality for mobile users even under environmental vulnerabilities such as blockages and beam misalignment due to mobility. The future work would explore advanced algorithms such as machine learning together with collaborative learning to further improve the tracking accuracy, develop multi-modal sensing with additional sensors such as cameras, and explore more challenging scenarios such as extremely high vehicular mobility.

Ethical concerns: Does not raise any ethical issues.

8 Acknowledgements

We are thankful to WCSNG lab members at UC San Diego for insightful feedback and to Raghav Subbaraman for their contribution to the initial phase of this project. This work was generously funded by Qualcomm Innovation Fellowship to Ish Kumar Jain and Raghav Subbaraman.

References

- [1] 5G; Study on channel model for frequencies from 0.5 to 100 GHz. TR 138 901 V14.0.0, 3rd Generation Partnership Project (3GPP), 2017-05.
- [2] 5G NR Frequency Bands. <https://www.5gradio.com/5g-technology/5g-nr-frequency-bands/>, September 2023.
- [3] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. Enabling high-quality untethered virtual reality. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 531–544, 2017.
- [4] 3GPP. 5G NR (New Radio) Release 16. 3GPP. <https://www.3gpp.org/release-16>, Oct 2019.
- [5] Kai Zheng, Kun Qian, Timothy Woodford, and Xinyu Zhang. Neuroradar: A neuromorphic radar sensor for low-power iot systems. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, pages 223–236, 2023.
- [6] Kaiyan Cui, Qiang Yang, Leming Shen, Yuanqing Zheng, and Jinsong Han. Integrated sensing and communication between daily devices and mmwave radars. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, pages 831–832, 2022.
- [7] Kshitiz Bansal, Keshav Rungta, Siyuan Zhu, and Dinesh Bharadia. Pointillism: accurate 3d bounding box estimation with multi-radars. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 340–353, 2020.
- [8] Kai Zheng, Wuqiong Zhao, Timothy Woodford, Renjie Zhao, Xinyu Zhang, and Yingbo Hua. Enhancing mmwave radar sensing using a phased-mimo architecture. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*, pages 56–69, 2024.
- [9] Manideep Dunna, Kshitiz Bansal, Sanjeev Anitha Ganesh, Eamon Patamasing, and Dinesh Bharadia. R-fiducial: Millimeter wave radar fiducials for sensing traffic infrastructure. In *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, pages 1–7. IEEE, 2023.
- [10] Umut Demirhan and Ahmed Alkhateeb. Radar aided 6g beam prediction: Deep learning algorithms and real-world demonstration. In *2022 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 2655–2660. IEEE, 2022.
- [11] Hao Luo, Umut Demirhan, and Ahmed Alkhateeb. Millimeter wave v2v beam tracking using radar: Algorithms and real-world demonstration. *arXiv preprint arXiv:2308.01558*, 2023.
- [12] Canan Aydogdu, Fan Liu, Christos Masouros, Henk Wymeersch, and Mats Rydström. Distributed radar-aided vehicle-to-vehicle communication. In *2020 IEEE Radar Conference (RadarConf20)*, pages 1–6. IEEE, 2020.
- [13] Leyan Chen, Kai Liu, Zhibo Zhang, and Baoqi Li. Beam selection and power allocation: Using deep learning for sensing-assisted communication. *IEEE Wireless Communications Letters*, 2023.
- [14] Zhanghua Chen and Jianxiao Xie. Radar aided active prediction based on deep learning. In *2023 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, pages 1–5. IEEE, 2023.
- [15] Andre Pearce, J Andrew Zhang, Richard Xu, and Kai Wu. Multi-object tracking with mmwave radar: A review. *Electronics*, 12(2):308, 2023.
- [16] Hansol Ku, Jinyue Song, Ding Zhang, Prasant Mohapatra, and Parth Pathak. Characterizing real-time radar-assisted beamforming in mmwave v2v links. In *2023 20th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pages 168–176. IEEE, 2023.
- [17] Rohith Reddy Vennam, Ish Kumar Jain, Kshitiz Bansal, Joshua Orozco, Puja Shukla, Aanjan Ranganathan, and Dinesh Bharadia. mmspoof: Resilient spoofing of automotive millimeter-wave radars using reflect array. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 1807–1821. IEEE, 2023.
- [18] Inras radarbook at 24 ghz and 77 ghz. <https://inras.at/en/radarbook2/>, June 2024.
- [19] Yasser Nabil, Hesham ElSawy, Suhail Al-Dharrab, Hassan Mostafa, and Hussein Attia. Beamwidth design tradeoffs in radar-aided millimeter-wave cellular networks: A stochastic geometry approach. *IEEE Access*, 2024.
- [20] Ryu Okubo, Luke Jacobs, Jinhua Wang, Steven Bowers, and Elahe Soltanaghaj. Integrated two-way radar backscatter communication and sensing with low-power iot tags. In *Proceedings of the ACM SIGCOMM 2024 Conference*, pages 327–339, 2024.
- [21] Mohammad Hossein Mazaheri, Alex Chen, and Omid Abari. Mmtag: A millimeter wave backscatter network. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 463–474, 2021.
- [22] Timothy Woodford, Xinyu Zhang, Eugene Chai, Karthikeyan Sundaresan, and Amir Khojastepour. Spacebeam: Lidar-driven one-shot mmwave beam management. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys '21*, page 389–401, New York, NY, USA, 2021. Association for Computing Machinery.
- [23] Aldebaro Klautau, Nuria González-Prelcic, and Robert W Heath. Lidar data for deep learning-based mmwave beam-selection. *IEEE Wireless Communications Letters*, 8(3):909–912, 2019.
- [24] Shuaifeng Jiang, Gouranga Charan, and Ahmed Alkhateeb. Lidar aided future beam prediction in real-world millimeter wave v2i communications. *IEEE Wireless Communications Letters*, 12(2):212–216, 2022.
- [25] Muhammad Alrabeiah, Andrew Hredzak, and Ahmed Alkhateeb. Millimeter wave base stations with cameras: Vision-aided beam and blockage prediction. In *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, pages 1–5. IEEE, 2020.
- [26] Gouranga Charan, Muhammad Alrabeiah, Tawfik Osman, and Ahmed Alkhateeb. Camera based mmwave beam prediction: Towards multi-candidate real-world scenarios. *arXiv preprint arXiv:2308.06868*, 2023.
- [27] Shoab Imran, Gouranga Charan, and Ahmed Alkhateeb. Environment semantic aided communication: A real world demonstration for beam prediction. *arXiv preprint arXiv:2302.06736*, 2023.
- [28] Iftikhar Ahmad, Ahsan Raza Khan, Rao Naveed Bin Rais, Ahmed Zoha, Muhammad Ali Imran, and Sajjad Hussain. Vision-assisted beam prediction for real world 6g drone communication. In *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 1–7. IEEE, 2023.
- [29] Gouranga Charan and Ahmed Alkhateeb. User identification: A key enabler for multi-user vision-aided communications. *IEEE Open Journal of the Communications Society*, 2023.
- [30] Jiali Nie, Quan Zhou, Junsheng Mu, and Xiaojun Jing. Vision and radar multimodal aided beam prediction: Facilitating metaverse development. In *Proceedings of the 2nd Workshop on Integrated Sensing and Communications for Metaverse*, pages 13–18, 2023.
- [31] Teng Wei and Xinyu Zhang. Pose information assisted 60 GHz networks: Towards seamless coverage and mobility support. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 42–55. ACM, 2017.
- [32] Vutha Va, Xinchun Zhang, and Robert W Heath. Beam switching for millimeter wave communication to support high speed trains. In *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*, pages 1–5. IEEE, 2015.
- [33] Muhammad Kumail Haider, Yasaman Ghaseempour, Dimitrios Koutsonikolas, and Edward W Knightly. Lister: Mmwave beam acquisition and steering by tracking indicator leds on wireless aps. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 273–288. ACM, 2018.
- [34] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. WiFi-assisted 60 GHz wireless networks. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 28–41. ACM, 2017.
- [35] Thomas Nitsche, Adriana B Flores, Edward W Knightly, and Joerg Widmer. Steering with eyes closed: mm-wave beam steering without in-band measurement. In *2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 2416–2424. IEEE, 2015.
- [36] Cheol Jeong, Jeongho Park, and Hyunkyuu Yu. Random access in millimeter-wave beamforming cellular networks: issues and approaches. *IEEE Communications Magazine*, 53(1):180–185, 2015.
- [37] Liang Zhou and Yoji Ohashi. Efficient codebook-based MIMO beamforming for millimeter-wave WLANs. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on*, pages 1885–1889. IEEE, 2012.
- [38] C Nicolas Barati, S Amir Hosseini, Marco Mezzavilla, Thanasis Korakis, Shivendra S Panwar, Sundeep Rangan, and Michele Zorzi. Initial access in millimeter wave cellular systems. *IEEE Transactions on Wireless Communications*, 15(12):7926–7940, 2016.
- [39] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 60 GHz indoor networking through flexible beams: A link-level profiling. In *ACM SIGMETRICS Performance Evaluation Review*, volume 43, pages 71–84. ACM, 2015.
- [40] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. Fast millimeter wave beam alignment. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 432–445. ACM, 2018.
- [41] Irmak Aykin, Berk Akgun, and Marwan Krunz. Smartlink: Exploiting channel clustering effects for reliable millimeter wave communications. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 1117–1125. IEEE, 2019.
- [42] Yuqiang Heng and Jeffrey G Andrews. Grid-free mimo beam alignment through site-specific deep learning. *IEEE Transactions on Wireless Communications*, 23(2):908–921, 2023.
- [43] Song Wang, Jingqi Huang, and Xinyu Zhang. Demystifying millimeter-wave V2X: Towards robust and efficient directional connectivity under high mobility. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–14, 2020.
- [44] Yuqiang Heng, Yu Zhang, Ahmed Alkhateeb, and Jeffrey G Andrews. Site-specific beam alignment in 6g via deep learning. *arXiv preprint arXiv:2403.16186*, 2024.
- [45] Xia Zhou, Zengbin Zhang, Yibo Zhu, Yubo Li, Saipriya Kumar, Amin Vahdat, Ben Y Zhao, and Haitao Zheng. Mirror mirror on the ceiling: Flexible wireless links for data centers. *ACM SIGCOMM Computer Communication Review*, 42(4):443–454, 2012.
- [46] Zulkuf Genc, Umar H Rizvi, Ertan Onur, and Ignas Niemegeers. Robust 60 GHz indoor connectivity: Is it possible with reflections? In *2010 IEEE 71st vehicular technology conference*, pages 1–5. IEEE, 2010.
- [47] Man Zhao, Danpu Liu, and Feng Yu. Improving the robustness of 60 GHz indoor connectivity by deployment of mirrors. In *2018 IEEE 29th Annual International*

- Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 188–193. IEEE, 2018.
- [48] Wahab Khawaja, Ozgur Ozdemir, Yavuz Yapici, Ismail Guvenc, and Yuichi Kakishima. Coverage enhancement for mm wave communications using passive reflectors. In *2018 11th Global Symposium on Millimeter Waves (GSMW)*, pages 1–6. IEEE, 2018.
- [49] Yibo Zhu, Zengbin Zhang, Zhinus Marzi, Chris Nelson, Upamanyu Madhoo, Ben Y Zhao, and Haitao Zheng. Demystifying 60GHz outdoor picocells. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 5–16. ACM, 2014.
- [50] Teng Wei, Anfu Zhou, and Xinyu Zhang. Facilitating robust 60 GHz network deployment by sensing ambient reflectors. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 213–226, 2017.
- [51] Anfu Zhou, Shaoyuan Yang, Yi Yang, Yuhang Fan, and Huadong Ma. Autonomous environment mapping using commodity millimeter-wave network device. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 1126–1134. IEEE, 2019.
- [52] Anfu Zhou, Shaoqing Xu, Song Wang, Jingqi Huang, Shaoyuan Yang, Teng Wei, Xinyu Zhang, and Huadong Ma. Robotic millimeter-wave wireless networks. *IEEE/ACM Transactions on Networking*, 28(4):1534–1549, 2020.
- [53] Anfu Zhou, Xinyu Zhang, and Huadong Ma. Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*, pages 1–9. IEEE, 2017.
- [54] Umot Demirhan and Ahmed Alkhateeb. Radar aided proactive blockage prediction in real-world millimeter wave systems. In *ICC 2022-IEEE International Conference on Communications*, pages 4547–4552. IEEE, 2022.
- [55] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. Beamspy: enabling robust 60 GHz links under blockage. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*, pages 193–206, 2016.
- [56] Santosh Ganji, Jaewon Kim, and PR Kumar. Terra: blockage resilience in outdoor mm-wave networks. In *Proceedings of the SIGCOMM'22 Poster and Demo Sessions*, pages 55–57. 2022.
- [57] Ding Zhang, Mihir Garude, and Parth H Pathak. mmchoir: Exploiting joint transmissions for reliable 60GHz mmwave w lans. In *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 251–260. ACM, 2018.
- [58] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. Towards scalable and ubiquitous millimeter-wave wireless networks. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 257–271. ACM, 2018.
- [59] Theodore S Rappaport, Shu Sun, Rimma Mayzus, Hang Zhao, Yaniv Azar, Kevin Wang, George N Wong, Jocelyn K Schulz, Mathew Samimi, and Felix Gutierrez. Millimeter wave mobile communications for 5g cellular: It will work! *IEEE access*, 1:335–349, 2013.
- [60] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. Two beams are better than one: Towards reliable and high throughput mmwave links. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 488–502, 2021.
- [61] Joseph A Shaw. Radiometry and the friis transmission equation. *American journal of physics*, 81(1):33–37, 2013.
- [62] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. Spotfi: Decimeter level localization using wifi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 269–282, 2015.
- [63] Extreme waves 5g phased arrays. <https://www.extreme-waves.com/5g>, Jul 2023.
- [64] Diligent cmod artix a7-35t fpga. <https://digilent.com/reference/programmable-logic/cmod-a7/start>, Jul 2023.
- [65] 12v battery. <https://www.amazon.com/65Watts-Portable-Charger-Battery-Notebooks/dp/B07P87JTVT>, Jul 2023.
- [66] Analog devices inc., adf5356 module. <https://www.analog.com/en/products/adf5356.html>, Jul 2023.
- [67] Ettus usrp b210. <https://www.ettus.com/all-products/ub210-kit/>, Jul 2023.
- [68] Gnu radio. <https://www.gnuradio.org/>, Jul 2023.
- [69] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. A compact and real-time millimeter-wave experiment framework with true mobility capabilities. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, pages 1–3, 2023.
- [70] Ish Kumar Jain, Raghav Subbaraman, Tejas Harekrishna Sadarahalli, Xiangwei Shao, Hou-Wei Lin, and Dinesh Bharadia. mmobile: Building a mmwave testbed to evaluate and address mobility effects. In *Proceedings of the 4th ACM Workshop on Millimeter-Wave Networks and Sensing Systems*, pages 1–6, 2020.
- [71] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. Demo and dataset for mmwave multi-beam tracking using mmobile 28 ghz testbed. In *Proceedings of the 12th ACM Wireless of the Students, by the Students, and for the Students (S3) Workshop*, pages 8–8, 2021.
- [72] Analog devices inc., adalm pluto sdr. <https://www.analog.com/en/design-center/evaluation-hardware-and-software/evaluation-boards-kits/adalm-pluto.html>, Jul 2023.
- [73] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11):1330–1334, 2000.
- [74] Ish Kumar Jain, Rohith Reddy Vennam, Raghav Subbaraman, and Dinesh Bharadia. mmflexible: Flexible directional frequency multiplexing for multi-user mmwave networks. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*, pages 1–10. IEEE, 2023.
- [75] Irmak Aykin, Berk Akgun, and Marwan Krunz. Multi-beam transmissions for blockage resilience and reliability in millimeter-wave systems. *IEEE Journal on Selected Areas in Communications*, 37(12):2772–2785, 2019.
- [76] Joan Palacios et al. Adaptive codebook optimization for beam training on off-the-shelf IEEE 802.11 ad devices. In *MOBICOM*, pages 241–255, 2018.
- [77] Shu Sun and Theodore S Rappaport. Multi-beam antenna combining for 28 GHz cellular link improvement in urban environments. In *2013 IEEE Global Communications Conference (GLOBECOM)*, pages 3754–3759. IEEE, 2013.
- [78] Texas Instruments. Ti iwr6843isk, radar development kit. <https://www.ti.com/tool/IWR6843ISK>.
- [79] Sabbir Rangwala. Automotive lidar has arrived. <https://www.forbes.com/sites/sabbirrangwala/2022/05/23/automotive-lidar-has-arrived/>, May 2022.
- [80] Intel Corporation. Intel realsense depth camera. <https://www.intelrealsense.com/compare-depth-cameras/>.
- [81] sharetechnote. 5g csi report in detail. https://www.sharetechnote.com/html/5G/5G_CSI_Report.html, May 2024.