MulDar: Unleashing the Potential of Distributed COTS mmWave Radar by Exploiting Cross-Device Channels

Xinghua Sun
University of Washington
Seattle, USA
xinghua@uw.edu

Qiancheng Li
University of Washington
Seattle, USA
qiancl2@uw.edu

Akshay Gadre
University of Washington
Seattle, USA
gadre@uw.edu

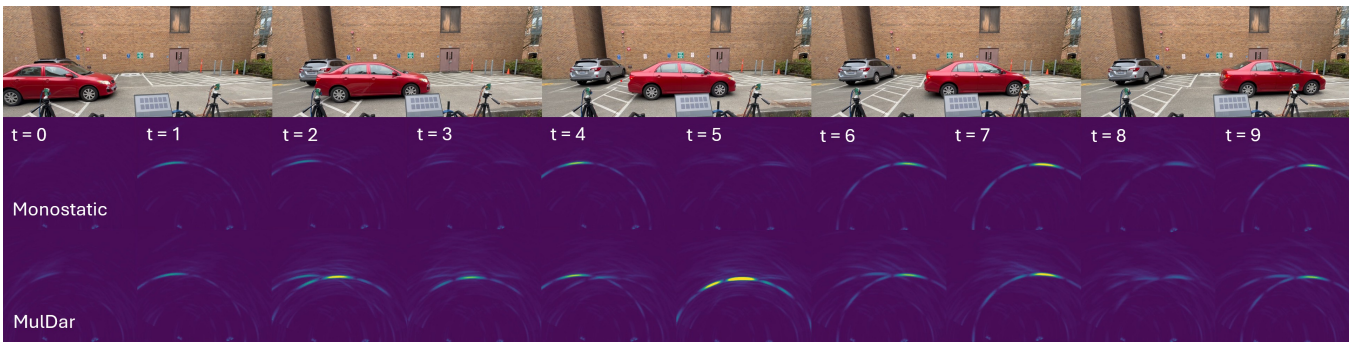

Figure 1: MulDar develops a synchronous multi-static radar system that addresses the issue of specular reflection and ghost artifacts in the scene. Check out the video for a recorded demo: https://youtu.be/Dgcf0vv0X7E.

Abstract

mmWave radar based sensors are increasingly being deployed for robotics and automobile applications to sense the environment. Unfortunately, specular reflections of smooth and planar surfaces makes mmWave sensing unreliable at detecting objects and vehicles in the real world. In fact, due to the small wavelength of mmWave radars, most surfaces are smooth reflectors for the radar signals. This paper addresses this fundamental limitation of detecting specular reflectors by capturing the energy reflected away from the radar.

This paper presents MulDar, a multi-static distributed mmWave radar imaging solution that leverages cross-radar channels to detect specular reflectors and remove ghost multi-bounce artifacts. At the core of MulDar's design lies the engineering of a distributed phase-synchronization approach using the direct path between the radars to perform imaging. The multi-static radar is then leveraged to combine the mono-static and bi-static measurements across pairs of radars to create a richer scene free of unseen and ghost objects. Our evaluation across materials and real world scenarios demonstrate a 66.79% reduction in Chamfer distance between the ground truth point clouds and the generated mmWave point clouds.

Code and datasets are available at: https://github.com/xsun2445/MulDar.

MobiSys '26, Cambridge, United Kingdom

ACM ISBN 979-8-4007-2027-7/26/06
https://doi.org/10.1145/3745756.3809206

CCS Concepts

• **Computer systems organization** → **Sensor networks**; • **Computing methodologies** → **Perception**; • **Human-centered computing** → **Mobile computing**.

Keywords

Wireless Sensing, mmWave, Perception

ACM Reference Format:
Xinghua Sun, Qiancheng Li, and Akshay Gadre. 2026. MulDar: Unleashing the Potential of Distributed COTS mmWave Radar by Exploiting Cross-Device Channels. In *The 24th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '26), June 21–25, 2026, Cambridge, United Kingdom.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3745756.3809206

1 Introduction

Millimeter-wave (mmWave) radar is an indispensable sensor for autonomous systems, particularly in self-driving vehicles [14, 43] and smart home applications [24, 25]. Its wide bandwidth provides fine-grained resolution, while its resilience to adverse lighting and weather conditions makes it suitable for robust perception [19, 30, 35, 36]. Despite its widespread adoption, a fundamental limitation of conventional mmWave radar threatens its reliability: the specular reflection [33, 34]. Most of current mmWave radar systems operate in a mono-static configuration, where the transmitter and receiver are co-located. This design implicitly assumes that the target is rough enough to scatter energy back toward the source. However, many common materials found in urban environments behave as specular reflectors at 77 GHz, such as the flat metal side of a truck, the glass wall of a building, or even concrete walls at home [21].

This physical property creates a safety-critical blind spot as shown in Fig. 1. If a large, flat surface is oriented at an oblique incident angle, vast majority of the energy will be reflected away from the transmitting direction. Therefore, this obstacle is completely invisible to a mono-static radar system. Additionally, specular surfaces can create multi-bounce shadows in the range profile which creates 'ghost' targets [20], confusing the downstream perception algorithm. These effects caused by reflective surface can cause significant safety issues in automobile and industrial applications.

Previous works [7, 41] explored to address the problem of specular reflection and limited field-of-view (FoV) by deploying and combining multiple radars in the environment, but primarily limited to the data received from mono-static responses and constrained by the radar placement. Further, other research [12, 19, 30, 32, 35, 37, 42, 45, 48, 49, 52–54] have explored mmWave scene reconstruction using single-bounce and multi-bounce assumptions, but typically assume the reflector itself is visible and/or its geometry is known. Hence, scene reconstruction and robust sensing using mmWave radars remains a challenging problem due to the sparse and noisy nature of responses received by mono-static radars.

This paper presents MulDar, a distributed radar imaging system that aims to address both the visibility and ghost problem in mmWave sensing by leveraging spatial diversity. MulDar's key idea observes that while the signals reflected from these objects do not return back to the radar, they do reflect to other directions in the environment. If we can observe those reflected signals, we can image these invisible objects and disambiguate the ghost reflections. MulDar designs a multi-static distributed mmWave imaging system that synchronizes distributed radars to capture these reflected signals observing cross-device channels. MulDar then combines the information aggregated across the mono-static and bi-static measurements to improve scene reconstruction. MulDar can be directly applied to already-deployed distributed radars without requiring additional RF devices or hardware modification. MulDar also enables enhanced improvements from spatial diversity as its multi-static nature increases the number of channels quadratically when adding new radars in the environment.

The above system design requires MulDar to address several unique challenges to enable multi-static imaging on commodity off-the-shelf mmWave radars:

1. Synchronizing commodity mmWave radars (Sec. 5): Given the distributed design of MulDar, the transmitters (TX) and receivers (RX) across radars need to be synchronized to receive other radars' transmissions. However, modern radars are specifically designed to reject interference from other radars by severely undersampling the IF signal received. What's more, limited hardware access to synchronize and calibrate the modern mmWave radars further complicates the problem. Further, since chirps from separate radars are not synchronized, every mixed bi-static IF signal is subject to a random frequency and amplitude offset. MulDar solves this by designing a 2 stage bi-static synchronization technique. First, we ensure signal reception from different radars by a one-time chirp calibration using the ambient environment. Second, MulDar leverages the direct path between TX and RX radars as the reference distance for finer synchronization.

2. Multi-static imaging and ghost removal algorithm (Sec. 6): Different from beamforming algorithm that works for uniform linear mono-static antenna arrays, antennas in our system are arbitrarily distributed in the space. Hence, we developed a new multi-static sparse imaging algorithm using back projection matched filtering to calculate for the amplitude of a pixel in the 2D space with previously synchronized bi-static radars. We further analyze the coverage and resolution limit of MulDar along with its intrinsic ability to remove multi-bounce ghost reflections.

We implement MulDar using 3 commodity TI AWR2243 mmWave radars and extensively evaluated MulDar's ability to image in indoor and outdoor environments. We compare MulDar with the mono-static combination of the individual radars as the baseline. We evaluate MulDar's ability to image surfaces and objects with different materials and curvatures in indoor environments as well as a case study on detecting corners in corridors, imaging pillars and outdoor automobile imaging. Our results demonstrates:

- 66.79% reduction in Chamfer Distance between the observed and true point clouds.
- ability to detect real world objects invisible in mono-static radars in practical indoor settings.
- ability to continuously monitor and observe vehicle robustly in practical outdoor settings.

Our contributions are as follows:

- We propose a novel distributed and bi-static imaging system capable of imaging specular objects using multiple commodity mmWave radars in the environment.
- A novel bi-static radar synchronization, imaging algorithm and signal processing pipeline to generate more reliable 2D image.
- Extensive real-world evaluation demonstrating MulDar's ability to image objects indoors and outdoors with case studies in corridor and outdoor environments.

2 Related Work

There has been much work on mmWave radar sensing for scene reconstruction and detecting multi-bounce reflections.

High-Resolution mmWave Radar Sensing: Given the low angular resolution and mono-static design of the mmWave radars, the problem of removing multi-bounce reflections to achieve high resolution spatial sensing has been explored in many prior works. Hydra [32] explores leveraging multi-bounce reflections for NLOS imaging while RadarHD [35] combines the radar with LIDAR-based learning to extract spatial resolution. Further work on leveraging multi-modal sensing by combining mmWave radar with cameras [23, 29, 38, 39, 44] to detect objects and sense vital signs has shown tremendous promise. Further, new approaches [9, 11, 12, 17, 26, 31, 46, 50, 51] that combine mmWave radar signals with motion to improve imaging of objects in NLOS and occluded environments have been presented. However, all of the above works leverage a single mono-static-radar (combined with other modalities) and leverage sparsity of object response to achieve super-resolution.

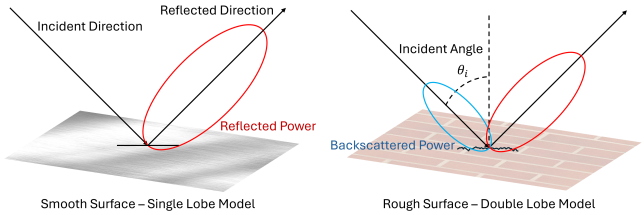

Figure 2: The single and double lobe scattering model for smooth and rough surface

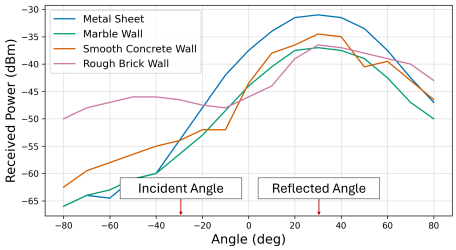

Figure 3: Simulated received power with $\theta_i = -30°$ [21]

In contrast, our work designs a multi-static mmWave radar imaging system that combines cross-channel measurements at multiple mmWave radars to improve imaging resolution.

Multi-View mmWave Radar Sensing: Limited to the Field-of-View (FoV) of radar, recent work has also explored combining multiple radars' signal for extended sensing ability. Radars are placed at different locations. Prior work such as [7, 41] has leveraged multiple mmWave radars to capture the scene from different vantage points to detect missing objects. [22] combines multi-view of distributed mmWave radars for improved human gait recognition and [13] has shown higher accuracy on human mesh reconstruction by placing 2 mmWave radars on both side of user's head and facing down to the body as wearable sensors. Further, some work has also explored utilizing specular reflections as additional views to extend the FoV for improved sensing capability[8, 17, 26, 27]. However, all of these works primarily combine multiple mono-static measurements captured across radars at different orientations for expanding the FoV and improve the received information. However, such systems would first require a visible LoS path of the specular reflector which is not very common, and would also require additional prior knowledge of the specular reflector for correct mapping of the multi-path to the real location.

Multi-Static Radar Solutions: While multi-view solutions treat separate radars as independent transceivers, multi-static radars separates TXs and RXs that every pair of TX and RX can communicate. The most obvious benefit is that, the number of channels increases quadratically compared to linear increment of multi-view radars. Those additional channels carry new and richer information from the environment and can improve the quality of sensing. Prior work have developed different frequency and scale of such multi-static system [6, 10, 15, 18, 47] for enhanced target detection, primarily for military and climate applications. Yet, those systems are usually build upon expensive and specialized RF devices for synchronizing spatially separated TX and RX. Additionally, [16] takes an approach similar to MulDar that uses frequency duplex and distance referencing for enabling cross-radar channel and range finding. Still, the system is based on advanced devices and only need to deal with a constant frequency offset during synchronization. In the mobile computing community, recent work [28] has shown promise of multi-static synchronization for vital sign sensing using UWB radars. However, the system synchronization is achieved by transmitting the clock signal over-the-air making the synchronization approach infeasible at 77 GHz due to range limits. [40] enables polarimetric imaging with multiple linearly polarized radars. But the system is built in a confined manner for precise SAR imaging. Different from prior works, MulDar enables multi-static channels on COTS devices in a plug-and-play manner for already-deployed distributed radars in the environment.

3 Motivation Study

In this section, we highlight the blind-spots of mmWave imaging using mono-static radars and motivate the effectiveness of our multi-static design with preliminary results.

3.1 Vulnerability of Mono-Static Radar to Specular Reflectors and Multipath Effects

Modern mmWave radar systems are mostly designed to be mono-static, meaning that the transmitters (TX) and receivers (RX) are co-located in space, usually on the same PCB board. The typical assumption in radar sensing is that, obstacles in the environment are considered as collections of points that will scatter the energy omnidirectionally. This enables the mono-static radar to detect these reflections to sense all the objects in the environment. Previous research [21] has shown this assumption is wrong and actually only very few objects in the environment will backscatter enough energy back to the radar to get detected. Hence, the response from the environment can be very sparse for mmWave radars, which usually contains false information in the imaging space, such as line-of-sight blind spots or multipath ghosts.

The primary reason behind this vulnerability is most material reflections are specular at mmWave frequencies. To illustrate this, Fig. 3 shows the scattering pattern of common materials under 77 GHz. As shown in the figure, for common materials like concrete and drywall, only a small amount of the energy will scatter back to the TX and most of the energy is reflected to the reflecting angle. This is because the roughness of the surface does not match the wavelength of the 77 GHz signal causing most of the energy to go undetected. As the incident angle increases, this phenomenon is more significant and eventually the drywall will disappear in the response signal, and this angle is surprisingly small. Meanwhile, the reflected energy will continue to propagate and some random object in the environment can backscatter this large energy back to the radar creating a ghost object.

This phenomenon means that the line-of-sight obstacles is not always visible which creates lots of blind areas especially for low-angular-resolution single-chip mmWave radars. Although we have discussed the opportunities of non-line-of-sight sensing in Sec. 2, those methods require high angular resolution and the first reflection to identify the order and direction of multi-reflection. In contrast, a common 2TX-4RX single-chip radar's range-azimuth map is neither able to provide accurate obstacle information, nor the ability to remap the ghosts original locations.

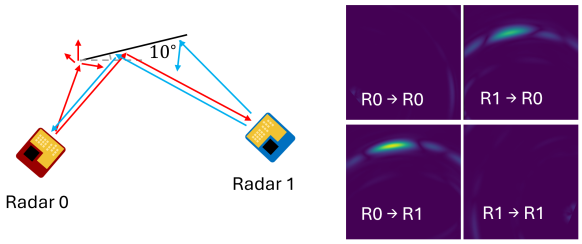

Figure 4: The planar reflector is invisible in mono-static images but visible in bi-static cross-radar images.

3.2 Effectiveness of Distributed Multi-Static Sensing

The easiest solution to overcome the above limitation would be to achieve spatial diversity to image the environment. One potential solution is to simply increase the number of radars in the space. However, Fig. 4 shows that simply adding more mono-static radars does not enable the setup to detect planar reflective surfaces, since the required backscatter incident angle for a reflective surface is very small. A mono-static radar can only see the planar surface across a small range of incident angles around 0° when enough reflected energy reaches the radar. However, any larger angle reflects most of the energy away from the radar removing the ability to detect the surface. Without access to prior knowledge of the environment, this makes it difficult to deploy radars (and sometimes in a robotic or automobile application, it may be impossible) at strategic locations to detect all objects. Moreover, the multipath ghosts of multi-bounce reflections appears in the image and the actual first peak of the reflective surface is invisible, making it impossible to identify the order of reflection. This simple situation in the real world can pose significant safety issue to the mmWave radar sensing, for instance, imaging in the autonomous driving scenario can fail to detect a car or a building causing an accident.

Our approach to address this important blind-spot in mmWave sensing is simple but effective: use the mmWave radars that already exist in the environment to capture the reflected signals. Fig. 4 shows the effectiveness of our method, where the bi-static channels can correctly show the specular-reflecting surface whereas the mono-static channels cannot. Further, the multi-bounce ghosts in the environment can be mitigated by leveraging the overlap across measurements to detect reflectors and remove objects that only show up in mono-static measurements. This results demonstrates the the capacity of multi-static distributed mmWave radars to improve imaging resolution and perform distributed sensing in the real world for detecting smooth objects and detecting ghost reflections.

4 System Overview

The system diagram is shown in Fig. 5. MulDar operates as follows: Upon deployment, the mmWave radars continuously trigger transmissions one-by-one on a pre-scheduled transmission plan to capture the mono-static and across-radar bi-static channel responses. The bi-static channel responses across the radars are then phase calibrated via referencing the direct path across them. These measured impulse responses are collected and streamed to a server

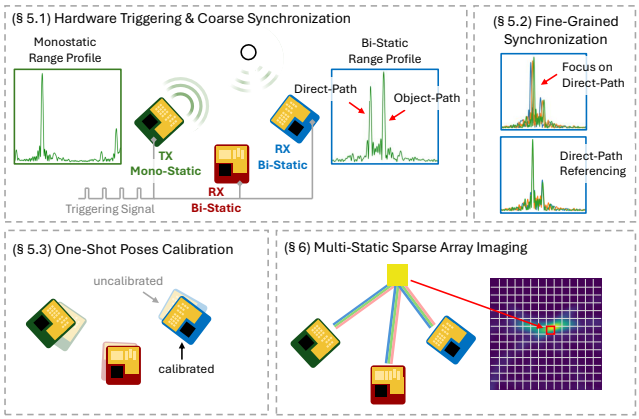

Figure 5: MulDar System Diagram

for real-time operation. The lines between radars to the object represent paths transmitted from the radar with the corresponding color. The scene is then reconstructed by combining these measurements to observe specular reflectors and remove ghost reflections. The above operation happens in real-time continuously as long as the radars are sensing the environment.

MulDar achieves this by addressing two fundamental wireless system design challenges: (1) Synchronizing distributed radars to receive each other's signals and phase-calibrating the received impulse responses using the direct path between the radars **(Sec. 5)**. (2) Combining the distributed mono-static and bi-static measured impulse responses into a coherent scene with maximal expressivity. Further, MulDar also analyzes the coverage and resolution limits of multi-static imaging **(Sec. 6)**.

5 Distributed and Multi-Static Radar

In this section, we first introduce how to build a distributed multi-static radar by combining multiple commodity-off-the-shelf radars. We then calibrate the randomness introduced by such a system. Finally, radar poses are optimized for combine generated images of the scene.

5.1 Coarsely Synchronizing Multiple Radars

Given that most commodity mmWave radars are mono-static[1], the transmit and receive antennas are typically operated by the same clock. In contrast, when we leverage multiple of these mono-static radars to develop a multi-static array and attempt to make them sense each others signal, we fail to receive anything. This originates from the fact that these modern mmWave radars are designed to avoid interference from other radars for safety reasons (particularly in automotive applications).

A frame can contain a diverse combination of chirps with different parameters such as slope, starting frequency, working bandwidth, and etc. Normally radars work on different frame configurations to avoid interference. Further, even if 2 radars has the same chirp parameters, low-pass filters on the IF signal will prevent frequency components higher than 10 MHz to be detected. The order-of-magnitude larger delay offset caused by hardware delays

[1]For simplicity, radar with co-located TX and RX are called mono-static in this paper, for it to be differentiated from bi(multi)-static spatially separated TX and RX.

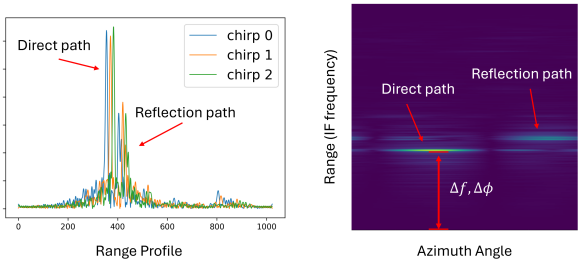

Figure 6: Hardware triggered chirps received from another radar that contain random frequency and phase offsets. Beamforming can identify the direct path and use the direct path peak as a reference distance.

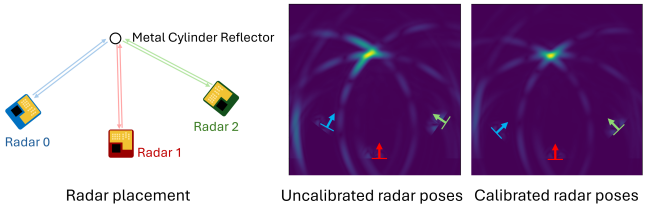

Figure 7: The mono-static image generated by uncalibrated and calibrated radar poses.

and the time-of-flight makes it unlikely to receive any baseband IF signal due to the lack of synchronization. Hence, although the chirp can last for hundreds of microseconds, only the signal that's within 257 nanoseconds of delay can be actually "heard" by another radar.

Luckily, TI mmWave radars provides a hardware trigger (SYNC_IN pin) that can be used for simultaneous triggering of multiple radars by sending pulses to the pin [3]. Further, the TX and RX on the radar operate on separate RF chains, enabling the TX to be turned off. This enables the radars to calmly receive other radars signals without interference from its own transmissions. We send a common pulse train signal of 100 Hz (frame rate) with a 1.5 μs active high duration to the SYNC_IN pin of all the radars with wired connection from an external signal source (Raspberry Pi 4). Fig. 6 shows that it is possible for a radar (including the transmitting radar itself) to receive another radar's chirp.

5.2 Fine-Grained Cross-Radar Phase Synchronization

An immediate problem is that, the range profiles are free-flowing across the whole spectrum since they are not time synchronized. As shown in Fig. 6, three consecutive chirps of exactly the same scene are offseted from each other. In fact, as radars do not share the same PCB board and commodity devices are imperfect, there exist three kinds of synchronization errors that need to be calibrated: frequency, phase, and amplitude. First, the external triggering signal is sent to radars through wires which itself creates a constant time offset of the signal reach separate radars. Since the hardware trigger is used for triggering the 77 GHz FMCW signal, there exists an uncertainty of around 160 nanoseconds, intrinsically in hardware delays. The random time offset is further exacerbated by heat noise and fabrication imperfections. This random time offset between transmitted and received chirps creates a random frequency and phase offset for every triggered chirp. Second, the linear modulation of the frequency is not perfect. The non-linearity of the chirp can create imbalanced noise of the amplitude across all frequencies when down-converting to the IF signal. Finally, the high and low-pass filters for preventing inter-radar and self-interference brings an additional non-linear and IF frequency dependent amplitude distortion. Note that, these offsets and noises always exists. All received chirps are subject to an independent offset that needs to be calibrated. Otherwise, downstream sensing tasks are meaningless since the range profile is dominated by the noise.

Our system solves the above mentioned issues through 3 stages. We first control the coarse time delay by adjusting the starting frequency of the chirp. The configurable starting frequency allow us to coarsely control the distance offset of the mixed received signal. The granularity of starting frequency of the radar in our system is 1 kHz, which can control the range profile roughly to the desired location. This operation is necessary for ensuring that the inter-radar chirp is within the IF signal band so that it can appear in receiving radar's IF signal. Then, we use the direct-path signal in the range-profile to calibrate for the frequency and phase offsets. Since each radar has 4 naturally synchronized receivers, we leverage the beamforming ability of RX radar to focus on the direct-path to the TX radar (Fig 6 (right)). Because we already know the rough location of each radars, the steering vector can be easily calculated. The beamforming can significantly improve the SNR of the direct path and improve the synchronization quality.

Finally, we rotate the signal back to the origin by times $e^{j2\pi(\Delta f t+\Delta\phi)}$ where Δf and $\Delta\phi$ is the direct path's frequency and phase. With the frequency and phase calibrated, consecutive chirps can be coherently averaged to further suppress the amplitude noises.

5.3 One-Shot Pose Calibration of Distributed Radars

Radars work in their own coordinate systems with themselves as origins. The collected information from spatially separated radars need to be combined and projected into the global coordinate system to build a collaborative sensing system. We constraint the spatial locations in this paper to the 2D space for simplicity. The pose of radars can thus be defined as a 3D vector $u = [x, y, \theta]$. The initial location of the radars is manually measured and contains measurement error that range from centimeters to even half meter, considering the scale of the deployed space. Since the range resolution of 76-81 GHz radars is about 3.3 cm, the location calibration of the distributed radars is necessary.

We calibrate the pose vector of the radars by optimizing the location of a common object in the scene. We solve for 3 unknown parameters of radar pose u (with one of the radar as global origin) by capturing the same common object across 3 different locations. Then, MUSIC algorithm is used for a higher resolution and cleaner image of mono-static range-angular map. These estimated locations of the cylinder are used to predict more accurate radar poses for fine-grained distance calibration. The optimization goal is to minimize the distance of cylinder center estimated from different radars. The

loss function can be written as

$$\text{loss} = \underset{u_i \in U}{\text{argmin}} \sum_{i,j} (\hat{p_j^i} - \bar{p_j})^2 \tag{1}$$

Where u_i and U is the i^{th} radar's pose and the set that contains all radar poses. $\hat{p_j^i}$ is the j^{th} cylinder center estimated by i^{th} radar in the global coordinate, and $\bar{p_j}$ is the average location of the estimated cylinder centers. Fig. 7 shows the imaging result generated before and after pose calibration. Note that, the optimized poses still contain millimeter level offsets because of low angular and range resolution of the radar. But that level of error is acceptable for our system and will be discussed in Sec. 6.2 and qualitatively evaluated in Sec. 8.3.

6 Distributed Imaging

In this section, we introduce the bi-static imaging algorithm and analyze the theoretical resolution, and the coverage of the imaging system.

6.1 Bi-Static Imaging

Our imaging algorithm begins with the angular-range profile recovered after performing beamforming using the received signal on 4 channels from the perspective of the receiving radar. However, this distance is a two way distance between the transmitting radar to the target and distance between the target and the receiving radar. Thus, the direct angular-range power profile cannot describe the real 2D scene. Hence, the spatial separation of TX and RX causes a distortion of its own range-azimuth coordinate system from the global 2D coordinate system. While one approach would be to estimate the projection function across all points to detect the potential locations of the reflector, we employ a more efficient backprojection algorithm that performs matched filtering across the global 2D space with the distance from TX to the target grid then to the receiver. Unlike the mono-static setting, since the signals are referenced based on the direct-path, the distance for building the matched filter always needs to be a relative distance that subtracts the length of the direct path. For TX and RX on 2 separate radars, each of the pixel p in the global 2D space can be written as:

$$\mathbf{s}(\mathbf{p}) = \sum_{i=0}^{A} \sum_{j=0}^{B} S_{i,j}(\frac{\rho_{i,j,p}}{\Delta r}) e^{j2\pi\rho_{i,j,p}/\lambda_0} \tag{2}$$

where $s(p)$ is the matched signal at pixel p, A and B are the number of TX and RX antennae, $S_{i,j}$ is the IF signal transmitted from TX_i to RX_j, $\rho_{i,j,p} = r_{i,p} + r_{j,p} - d_0$ is the distance traveled from TX_i to pixel p then to RX_j and removed direct path distance d_0, Δr is the range resolution determined by the bandwidth. λ_0 is the wavelength of the starting frequency of the transmitted chirp.

We demonstrate this geometric projection by plotting the generated global 2D image in Fig. 9 (a-b). Fig. 9(a) considers only 1 TX and 1 RX. Thus, for a fixed object in the space, the image of bi-static setting is shown as an ellipse without taking into account directionality since the sum of the distance to TX and RX is fixed. (b) shows the result retrieved by combining 4 RX antenna via beamforming. The amplitude at the object direction on the same ellipse is highlighted. Note that, the object's direction in the global coordinate and the radar coordinate are the same relative to the radar. We will discuss further in Sec. 6.3.

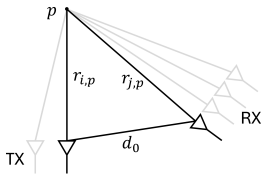

Figure 8: Illustration of Distances in Bi-static setting

6.2 Multi-Static Sparse Array Imaging

The above design demonstrates performing 2D imaging using 1 TX antenna at the transmitter and 4 RX antennas at the receiver by performing beamforming. However, as more radars are added to the system, counterintuitively, the generated images cannot be directly added together for coherent fusion. The primary reason behind this is the trade-off between sampling a sparse radar array: spatial diversity vs. correlation of channels.

A sparsely sampled signal will create aliasing effect in the spatial domain since the the sampling locations are not enough to fully delineate the target object. In fact, this is the reason why the spacing between antenna and synthesized aperture size of typical dense arrays need to be carefully designed. Unfortunately, when the sampling array is too sparse, even worse grating lobes will occur which may cause a incoherent combining effect in the generated image. In MulDar, since we have build a multi-static architecture, every antennas are considered as independent TX and RX elements, decoupling from the concept of independent radars. For n radars with A TX and B RX, there are n^2AB channels in total, and are spatially distributed in the space. Hence, it is very important to correctly combine the images generated from channels overcoming the limitations of sparse sampling.

Correlated Channels: For a conventional uniform linear array (ULA) antenna placement, the signal can be coherently combined with steering vectors. This is possible due to the fact that the received channels across antennas are correlated, meaning that every antenna element will receive almost same responses since they are placed next to each other. The major difference of the responses is their residual phase which is dependent on the traveled distance. Hence, after the matched filter removed the phase difference, the responses has the same phase so the amplitude adds up. This holds true for both far-field or near-field scenarios for ULA, as long as the matched filter are correctly calculated. In MulDar, this can be applied to antennas sit on the same radar since the those TX and RX antennas are naturally placed as ULA. Hence, for chirps that are transmitted from the same TX ULA on radar 1 and received by a RX ULA on another radar 2, the channels are correlated and can be coherently combined to improve the imaging resolution. Additionally, channels from radar 1 to 2 and channels from radar 2 to 1 are also correlated due to the reciprocity. By coherently combining all the correlated channels, the SNR and resolution of the image can be significantly improved. Fig. 9 (a-e) shows how coherent fusion of correlated channels increase resolution and SNR of bi-static channels.

Uncorrelated Channels: However, this correlated channel assumption is not always true for antenna pairs across different pairs

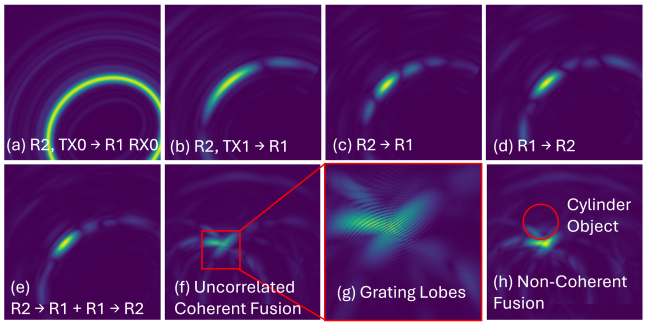

Figure 9: Channel Fusion. R*i* is Radar *i*. (b-e) Channels between R1 and R2 are correlated, can be coherently combined. (f) shows the coherent fusion of all channels which will create grating lobes. (h) is the non-coherent fusion of uncorrelated channels, which shows the shape of cylinder object.

of radars in MulDar. For example, considering 3 radars sparsely distributed in the space, the response from radar 1 to radar 0 and radar 2 to radar 0 can be completely uncorrelated since the virtual antennas are too sparse in the space. Thus, the responses are completely different and the same point reflector most likely will not appear in both kind of channels. Hence, removing the residual phase with matched filters won't work. Although those channels are uncorrelated and cannot improve the resolution or SNR, this uncorrelated behavior also means it contains distinct new information from different perspective. Hence, uncorrelated channels can be non-coherently fused by combining them in the amplitude domain. Fig. 9 (f-h) shows how coherent fusion of uncorrelated channels will hurt the generated image, and how non-coherent fusion can provide more information. Since uncorrelated channels are non-coherently combined, the sub-mm level of localization accuracy is not necessary.

Thus, when calculating a pixel on the MulDar image, channels transmitted and received between a pair of radars will be coherently combined, then all values will be non-coherently combined. We modify the original back-projection method Eq. 2 to:

$$\mathbf{s}(\mathbf{p}) = \sum_{R_a,R_b \in R} \lambda_{R_a,R_b} \left| \sum_{i,j} \left(\frac{\rho_{i,j,p}}{\Delta r}\right) e^{j2\pi\rho_{i,j,p}/\lambda_0} \right| \tag{3}$$

Where R_a and R_b means every 2 radars from all radars R in the system (R_a and R_b could be the same radar and the back-projection method still applies), $i \in \{R_a^{TX}, R_b^{TX}\}$ is one transmitting antenna from R_a and R_b, and $j \in \{R_a^{RX}, R_b^{RX}\}$ is one receiving antenna from R_a and R_b. λ_{R_a,R_b} is a scaling factor depending on whether the channel is mono-static or bi-static. It is necessary for calibrating the amplitude since the mono-static channels are usually suppressed by the low-pass filter (LPF) more significantly.

6.3 Resolution Analysis

Unlike traditional mono-static setup, bi-static radar system separates its transmitter (Tx) and receiver (Rx). This difference also leads to a change in the resolution of imaged scenes. Therefore, we analyze the new resolution limit feasible for MulDar bi-static radar imaging system with both theoretical analysis and experimental results.

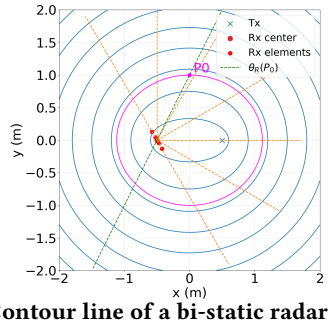

Figure 10: Contour line of a bi-static radar is an ellipse.

Our first difference compared to mono-static setup is the distance traveled by chirps emitted by Tx. Conventionally, Tx and Rx are considered to be at the same place in mono-static setup, so the emitted chirp and the reflected chirp travel the same path, and objects on a circle with a radius R to Tx/Rx will delay the chirp by the same amount of time. However, Tx and Rx are separated in bi-static setup, so the path from Tx to object and the path from object to Rx will be different. If we denote the location of Tx, Rx, and the object as T, R, P, respectively, the total amount of distance ρ traveled by a chirp is $\rho = |P - T| + |P - R|$. We can find that locations that have a constant sum of distance to Tx and Rx form an ellipse, where Tx and Rx are the focuses of this ellipse.

On the other hand, the angular resolution remains unchanged from a mono-static to bi-static setup. It is because angular resolution is defined as the ability of an Rx array to separate two objects in the angular domain, and the Rx array can only distinguish the incident angle from objects, no matter how this angle is generated by the object reflection. As a result, the Tx position should not affect the angular resolution of an Rx array. Together, we have a qualitative contour line of a bi-static radar system in both range and angle, and it is depicted in Fig 10.

Quantitatively, we want to further derive the expression of range resolution at different locations. Since we distinguish two objects based on the time delay resulted by their distances to Tx/Rx, we analyze how the position of object affects the traveled distance in a bi-static setup. To do this, we compute the gradient of total traveled distance ρ with respect to x and y respectively:

$$\frac{d\rho}{dx} = \frac{d\sqrt{(x-x_T^2)+(y-y_T^2)}}{dx} + \frac{d\sqrt{(x-x_R^2)+(y-y_R^2)}}{dx} \tag{4}$$

$$\frac{d\rho}{dy} = \frac{d\sqrt{(x-x_T^2)+(y-y_T^2)}}{dy} + \frac{d\sqrt{(x-x_R^2)+(y-y_R^2)}}{dy} \tag{5}$$

where $\{x, y\}$ is the location of the target P, and $\{x_T, y_T\}$ and $\{x_R, y_R\}$ is the location of Tx and Rx, respectively. This can be finally simplified as

$$|\nabla\rho| = 2 * cos(\frac{\beta(P)}{2}) \tag{6}$$

where β is the angle between PT and PR. Therefore, the range resolution in bi-static radar system is: $\frac{\Delta\rho}{2cos(\frac{\beta(P)}{2})}$. Based on this equation, we notice that a smaller β angle gives a better resolution, while we need a greater distance to differentiate two objects at a larger β angle. Intuitively this makes sense, since a small amount

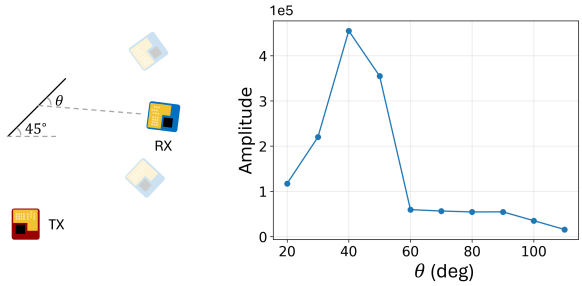

Figure 11: Coverage across angle of radar placement

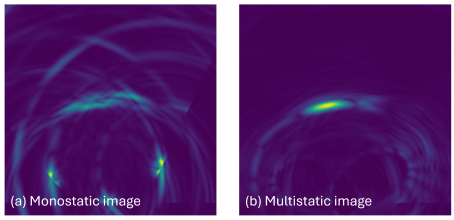

Figure 12: mono-static image contains multipath while multi-static path is sparse.

of change at larger angle will result in a greater change of distance for the signal to travel when the angle is small and the object is far from Tx and Rx. This is shown in Fig 10, as we can notice the gap between the first two ellipses at $x = 0$ is slightly larger than others.

In summary, from mono-static to bi-static, objects will stay at the same angle line but move to a different ellipse because of the change in travel distance of chirps.

6.4 Coverage Analysis

In this section, we explore the limits of the additional coverage provided by MulDar in a bi-static setting. If the reflected signal is too sparse, then our system might only work in the corner cases. Our simulated received energy observed in Sec. 3 estimates that the reflected energy from a planar surface can be received by another radar at a wide range of angles. We validate this by conducting a coverage measurement using bi-static radars shown in Fig. 11(left). For a fixed TX radar location, we measured the received energy reflected from the object (excluding the direct-path) at different reflecting angles. We plot the angles that can receive recognizable amplitude in Fig. 11(right), the radar can receive reflected energy from 20° to 110° which validates the previous theoretical research [21] simulation. Our detailed evaluation across surfaces with varying curvature, materials and shapes is presented in Sec. 8.2 and clearly demonstrates how MulDar outperforms both mono-static-only and bi-static-only baselines.

6.5 Multipath Mitigation in Bi-Static Setting

Fig 12 shows, for mono-static setting, multipath causes "ghost" artifacts, whenever the signal is reflected off planar surfaces. These reflected artifacts from virtual multi-bounce sources can be received at the original true radar location because of the path reciprocity. Although radar responses are usually sparse, each of the virtual source locations will be shared across mono-static responses due to reflection reciprocity. Thus, the "ghost" object will be visible strongly in the mono-static responses across all radars and their

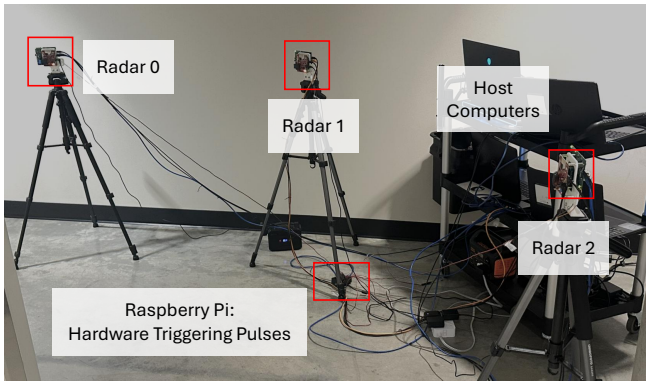

Figure 13: MulDar **system setup**

combination. Further, considering the energy lost, the maximum number of reflections that can be detected is about 3 reflections [32].

However, in the bi-static setting, a multi-reflection path ensures that it requires 2 or more consecutive reflectors that are tangential to the same ellipse with the TX and RX as the focus points. Further, this same virtual reflector has to also follow the prior mono-static constraints. Given this very strict requirement is extremely rare, we can remove the ghost artifacts from combined scene response via simple thresholding. This further highlights the benefits of bi-static setting which is naturally resilient to multipath effects and can provide more information about the LoS obstacles.

7 Implementation and Evaluation

We implement MulDar using multiple TI AWR2243BOOST radars which support three transmit and four receive antennas providing a azimuth angular resolution of 15°. The radar boards are connected with TI DCA1000EVM enabling real-time raw data capture. The radars are hardware-triggered as described in Sec. 5 using a Raspberry Pi 4 to enable loosely synchronized inter-radar transmission and reception. The experimental setup varies across testbeds, yet the chirp configuration is typically the same and configured using TI mmWaveStudio. Our multi-static architecture treat all antennas as independent elements, and it decouples the sensing logic from individual of radar hardware. The chirps are configured to be transmitted in batches that all transmit antennas will be triggered sequentially. Upon each transmission, all receive antennas are also triggered to listen. Each of the radar is connected to a laptop which is only used for configuring the radar with the supporting software. Raw signals are streamed directly from data capture boards to a central server for processing. The server collects all the mono-static and bi-static data, then perform synchronization and multi-static imaging algorithm. Currently, the system runs on a single RTX4080 GPU at up to 100Hz and can be visualized in real-time[2].

Evaluation Baselines: We compare MulDar with two baselines: (1) Combined mono-static Scene Reconstruction (**Mono**) : where all of the mono-static information is combined to reconstruct the scene. (2) Combined bi-static Scene Reconstruction (**Multi**) : where only bi-static cross channel measurements are used to reconstruct the scene. These baselines are compared with MulDar which uses both

[2]All code, data and processing pipeline will be made available upon paper acceptance

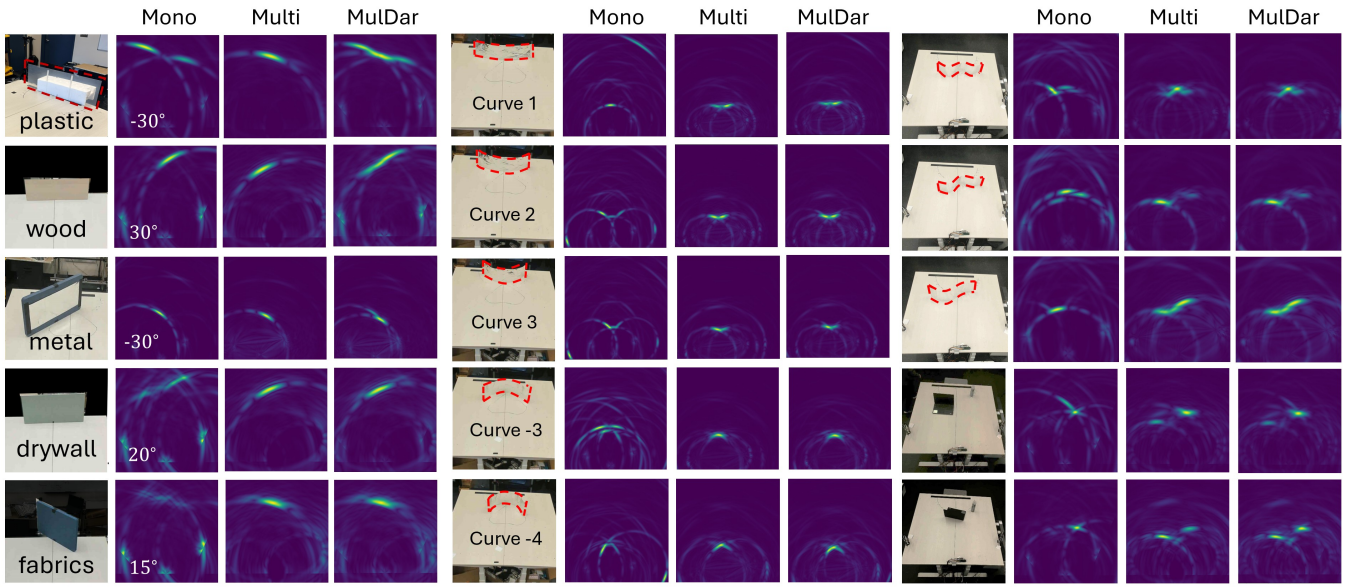

Figure 14: MulDar qualitative results demonstrate the improved sensing capability over Mono and Multi baselines

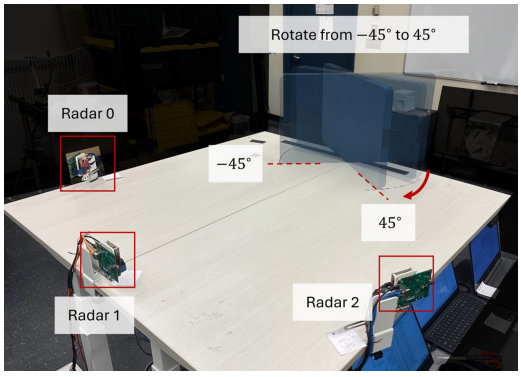

Figure 15: The setup for evaluating the performance on planar surface.

of them to simultaneously image specular reflectors and remove ghost artifacts from the reconstructed scene as described in Sec. 6.

Evaluation Metric: Across all our evaluation, we present the two-way Chamfer distance between the true point clouds and the predicted point clouds. Chamfer distance is typically used for 3D reconstruction predictions and fits perfectly for our application in 2D sensing. The Chamfer distance (CD) between the two 2D point clouds (PC) is calculated as :

$$\text{CD(GT, M)} = \underbrace{\frac{1}{|GT|}\sum_{i\in GT} \min_{j\in M}||i-j||^2}_{\text{Accuracy}} + \underbrace{\frac{1}{|M|}\sum_{j\in M} \min_{i\in GT}||i-j||^2}_{\text{Coverage}}$$

Where the first term is the distance from ground truth (GT) PC to its nearest neighbor in the measured (M) PC that predicts the accuracy of estimated PC. And, the second term is the distance from measured PC to its nearest neighbor in ground truth PC that predicts the coverage of the ground truth.

8 Results

8.1 Qualitative Results

This section evaluates the MulDar in its ability and compares its results with traditional mono-static imaging fusion (monosum) and bi-static imaging fusion (multisum). We present our qualitative result with three examples: images of planar surfaces of different materials, images of an deformable object in various shapes, and random combinations of real world objects.

Setup: We place three radars in the configuration shown in Fig. 15 around the table spanning 1.5 m×1.5 m. (1) *Planar Surface of Different Materials:* We place a single planar surface at the edge of the table made of wood, metal, drywall and fabric to be imaged by MulDar. (2) *Deformed Surfaces:* We place a metal reflector in various curved orientations on the table. (3) *Real Scenes:* We place everyday objects such as box, tripod, laptop and a bottle on the table.

Planar Surface of Different Materials: Fig. 14a shows that the planar object is invisible in the mono-static response of the lateral radars while the front-on radar views it as a point object. However, MulDar's multi-static scene reconstruction is able to perceive the planar object which enables accurate reconstruction of the surface across materials. Further, since MulDar has different path geometry compared to mono-static, the reflecting area on the object is different. Hence, when mono-static and multi-static image can both see the object, the combined MulDar image has significant coverage of the object shown in quantitative evaluations Fig. 16.

Curved Surfaces and Deformed Surfaces: Fig.14b-c shows that the mono-static radars are only able to perceive sections of the deformed surface that is lateral to them creating a very sparse image. Only those angles that are perpendicular to one of the radar can be detected. The complex geometry of those surfaces also create lots of multipath shadows in the mono-static images. In contrast,

Figure 16: Quantitative results of rotated planar surfaces with different materials and curvature.

MulDar's multi-static scene reconstruction is able to perceive the curves of the deformed surface accurately. The multipath effect is also much less in the MulDar image as we have explained in Sec.6.5. MulDar is particularly good at complex geometries since the quadratic more channels not only covers more angle but also improved the imaging resolution with the additional information.

Real Scenes: The mono-static radars are able to see many of the objects placed in front of them, yet fail to see the laptop in all of the test cases (Fig.14c). On the other hand, MulDar reconstructed image is able to clearly see the laptop along with other objects and able to suppress ghost object reflections behind the laptop.

8.2 Quantitative Results

In this section, we present the quantitative results about the accuracy, coverage and overall error using Chamfer distance comparing MulDar with mono-static and multi-static-only baselines. The setups across this evaluation are the same described in Sec. 8.1 except

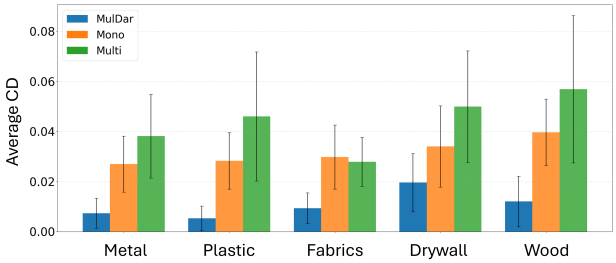

Figure 17: Average Chamfer Distance across planar surfaces made of various materials

for the fact that the rotation test of the planar surface is performed across materials.

Planar surfaces across rotations: Fig.16 plots the accuracy, coverage and combined Chamfer distance across various materials of the planar surface being rotated. Each plot presents the distance across angle of rotation for a particular material. Each row presents a separate material in order (dry wall, fabric, metal, plastic and wood). Our results clearly show the drastic benefit provided by MulDar especially at high angles of the planar surface. MulDar achieves an average 72.62%, 81.03%, 69.47%, 68.47%, and 42.38% reduction in Chamfer distance across angles of rotation for metal, plastic, wooden, fabric, and dry wall surface.

Deformed Metal Curvature: Fig.16 plots the accuracy, coverage and combined Chamfer distance for circular arcs with various curvature coefficients. The mono-static baseline particularly suffers for convex surfaces which it measures as a point reflector. MulDar achieves an average 46.37% reduced Chamfer distance across all curvatures for the object.

8.3 System Resilience to Deployment Inaccuracies

In Sec. 5.3 we have calibrated the manually measured radar positions. We also state that our multi-static imaging method is robust to small centimeter-level errors in manual measurements. We evaluate the robustness qualitatively in Fig. 18, where the first row is the calibrated the second row is uncalibrated. With the same system setup as Fig. 15 and $u = [x(m), y(m), \theta(deg)]$ representing 2D radar orientations, the calibrated orientation of radars using Eq. 1 are: $u_1 = [0, 0, 90]$, $u_2 = [-0.8296, 0.3203, 38.6789]$, $u_3 = [0.7368, 0.4806, 155.5451]$. And, the uncalibrated manually measured radar poses are: $u_1 = [0, 0, 90]$, $u_2 = [-0.84, 0.32, 30]$, $u_3 = [0.79, 0.47, 165]$. The error of manual measurement can be as large as several centimeters and upto ten degrees of orientation. However, since our multi-static algorithm doesn't require precise phase level accuracy (sub-mm), the distortion caused by the manual error is acceptable in the radar image. Also, since the mmWave image is low-resolution (15 degree angular and 3.5 cm range resolution), sometimes the difference is negligible compared to the calibrated image.

8.4 In-the-wild Experiments

We evaluate MulDar in real world environments to demonstrate the feasibility of detecting unseen reflectors in the real world. Our

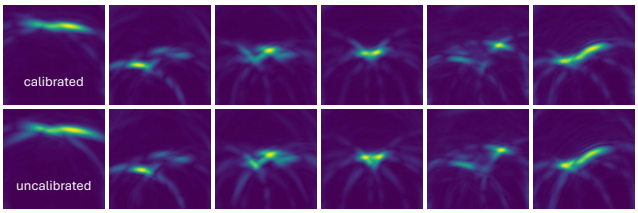

Figure 18: MulDar is robust to manual radar pose measurements. The upper row is generated using calibrated radar poses and the lower row is using uncalibrated manual measurements.

evaluation emulates two real world scenarios of robotics and automobile sensing using mmWave radars.

Corridor Imaging: We implement MulDar using three radars in a corridor environment to emulate indoor SLAM imaging via multiple mmWave radars deployed on robots. While there are several challenges to realize this vision of performing synchronized SLAM, we demonstrate the benefits of leveraging multiple mmWave radars and exploiting cross-device channels for improving scene reconstruction. The three testbeds are shown in Fig. 19a (i) imaging a corner of a corridor, an area with three doors and a circular pillar.

Results: Our results in Fig. 19a (ii–iv) shows the clear benefits provided by MulDar qualitatively. In the corner scenario, the mono-static only image fails to observe some of the corners present in the building while the bi-static combination only sees another surface. The combination however, can see the whole arch of the corridor. Similarly, the MulDar image for the second scenario is clearly able to see the whole corridor with the multiple doors. Finally, the pillar image is retrieved with much lower error in the point cloud generation unlike Mono and Multi baselines. In all of the above cases, the unique combination of the data across the mono-static and multi-static images is key to achieving high-fidelity scene reconstruction.

Automobile Imaging: We evaluate MulDar by combining two bi-static setups to image the automobile using 4 mmWave radars. This example simulates the trajectory of a car turning out of parking adjacent to another car. We evaluate the ability of MulDar with its multiple radar measurements to detect the car adjacent to it.

Results: Our results in Fig.19b shows the two bi-static capture scenarios that a car traverses while moving out of a lateral parking. Our results clearly demonstrate the benefit of capturing multi-static measurements during the unparking motion to capture the outline of the car adjacent to it to avoid collision. In contrast, the mono-static measurements are unable to detect it clearly.

Autonomous Driving: We evaluate a cross-roads scenario of a car passing by perpendicular to the motion of the car waiting at a stop sign. We evaluate the ability of MulDar throughout the car trajectory to detect the car and its motion. Our video demo[3] demonstrates the setup in this scenario quite clearly.

Results: Our results in Fig. 1 while the car is passing across the front of the radars demonstrate that mono-static radars can only perceive the car at select points of the trajectory. In contrast, the combination of both multi-static and mono-static mmWave scene

[3]Demo Video: https://youtu.be/Dgcf0vv0X7E

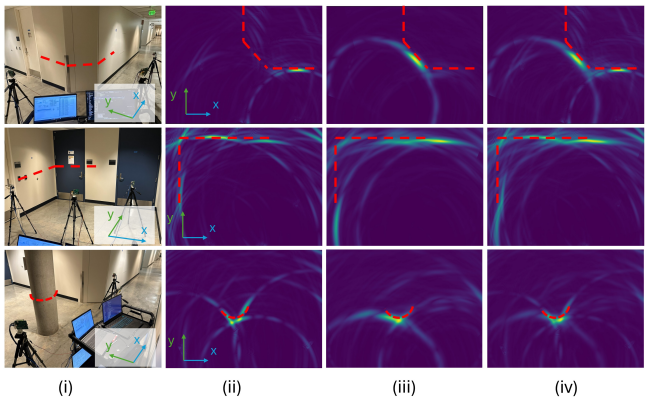

(a) Indoor evaluation: (i) The imaging scene consists of three scenarios: corner of a corridor, an opening space, and a pillar. (i-iv) Comparison of Mono, Multi, and MulDar images.

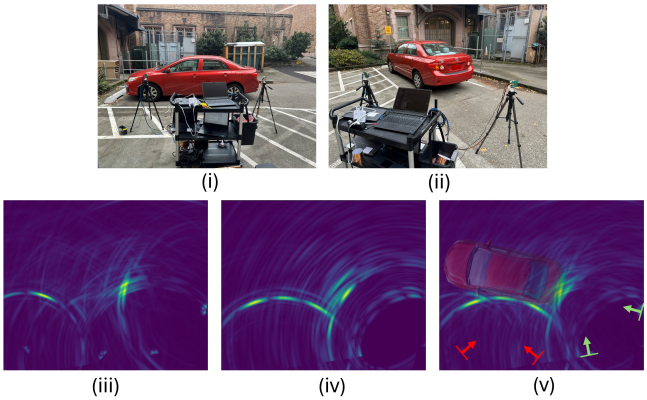

(b) Car evaluation: Red and green arrows are radar locations. (i–ii) are the imaging scene and radar locations. (iii–v) Comparison of Mono, Multi, and MulDar images.

Figure 19: Real-world Evaluations

continuously tracks the trajectory of the car during its complete motion ensuring the car safety of the autonomous vehicle (if it were to only use mmWave radars for sensing).

8.5 Failure Scenarios

Despite the significant improvement in imaging quality provided by MulDar, there still remain several challenging scenarios that limit its applicability.

Firstly, since MulDar relies on spatial diversity provided by multiple radars, its performance degrades when radars are placed too close to each other, thereby lacking good observation angle. As illustrated in Fig 20 (a), the image fails to reconstruct the shape of the corridor when the 3 radars are constrained within a narrow place. They can only detect the door at front but fail to sense the wall on the two sides. Similarly, this degradation in image quality occurs when most of the energy is reflected toward an angle where no receiver is placed. An example is shown in Fig 20 (b). The signal strength is weak, and the reconstructed shape of the plane is blurred.

Secondly, complex multipath effects arising from irregular object shapes will distort the image result, making it hard to identify objects. One example is shown in Fig 20 (d), the extremely curved

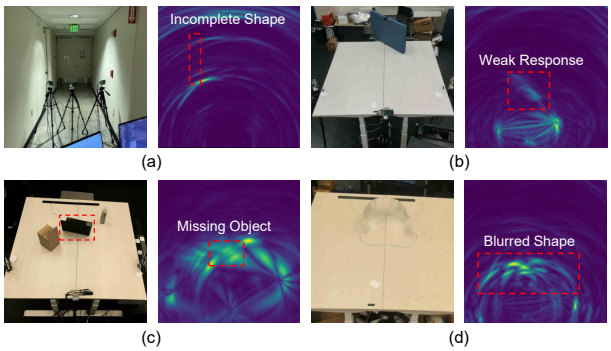

Figure 20: MulDar is constrained by deployment space and amount of multi-path distortions.

metal sheet results in severe propagation distortions, so its U-shape is indistinguishable in the image due to the multipath. Another scenario involving excessive multipath is when multiple objects bounce the signal too many times. As illustrated in Fig 20 (c), the laptop is difficult to visualize from the reconstructed image as the signal passed through the cardboard box and bounced between objects.

9 Discussions and Limitations

Machine Learning for Super-Resolution mmWave Imaging: Recent work [35] has explored leveraging machine-learning models co-trained with LIDAR and camera to explore data-driven super-resolution solutions for static mmWave imaging systems. However, these systems still face the issue of environmental data bias and fail to generalize to our use cases. However, we believe that these type of system designs can be augmented atop MulDar by additionally taking into account multi-static data to scan the environment at high resolution and deal with specular reflections and ghost objects much better.

Number of Radars: We believe that there is an ideal number of radars beyond which MulDar will provide diminishing returns on infrastructure investment due to the sparsity of reflectors in a real world environment. If there are too many radars (or synthetically moved using circular aperture), it is just better to leverage the mono-static responses as the new information in cross-channel measurements will be very limited. In contrast, sparse deployments in automobile and indoor robotic environments which are much more likely, the bi-static cross-device channels provide invaluable information that significantly improves scene reconstruction.

Placement of Radars: Another related concern is the placement of radars. While we used the same and a common radar placement for qualitative and quantitative evaluations shown in Fig. 14-17, we believe that's enough to prove that MulDar can almost always improve the scene reconstruction quality. This means MulDar can be directly applied to lots of already-deployed multi-view radar array without touching the original placement. A key constraint to the MulDar radar placement is the wire connection for hardware triggering. In the future, a dedicated wireless triggering module can be designed for wireless synchronization and the placement strategy considering both mono-static and bi-static FoV can be studied.

Existence of Direct-path: A critical assumption within MulDar design lies the existence of the direct path to perform phase-level synchronization between mmWave radars. This direct path can typically be augmented by deploying the radars at a elevated location in indoor environments or within the car bonnet for automotive applications. This path is critical to perform in-situ synchronization for plug-and-play application which otherwise would require careful calibration which is impractical for a non-expert user. Future work can also explore the use of any structured multipath as long as the geometric relationship can be resolved.

Heterogeneous Radars: We believe our system is broadly applicable as most commodity radar chips featuring a hardware trigger [1–5]. Though using the trigger sometimes require PCB modifications if the pin is not exposed, the cost is negligible compared to adding additional RF synchronization infrastructure. While we haven't evaluated heterogeneous radars in MulDar, we did successfully replaced one of the radar with a different radar model AWR1243 and constrained them to a narrower 4 GHz working bandwidth. The key is to making sure that all radars are working under the same chirp configuration, especially the slope of the ramp and chirp duration. As the number and different model of radar increases in the system, an additional orchestration protocol can be designed for optimizing such a multi-static radar network for improving the sensing quality.

10 Conclusion

Our system MulDar designs a multi-static scene imaging system using distributed COTS mmWave radars overcoming the limitations of specular reflections and multi-bounce ghost effects. Our evaluation in indoor controlled table-top, real-world corridor and automobile environments demonstrate a significant improvement in scene reconstruction and point cloud generation. At an average, MulDar achieves a 66.79% reduction in the average Chamfer distance between the ground truth and the estimated point clouds.

We envision future researchers to build on MulDar's multi-static mmWave radar design to improve wireless sensing and communication. We imagine MIMO beamforming solutions developed on top of MulDar distributed synchronization for powering battery-free wireless tags. Further, the bi-static imaging system can be combined with other modalities such as camera and LiDAR to further improve the accuracy of detecting obstacles in automotive scenarios. Finally, MulDar is a perfect candidate to empower robotics applications involving distributed mmWave radar sensors cooperating to sense the world better.

References

[1] accessed Dec 05, 2025. Andes Series CAL77C844: Next-Generation 77 GHz mmWave Radar SoC. https://www.calterah.com/en/soc-products/fully-integrated-mmwave-radar-so-c/car-soc.
[2] accessed Dec 05, 2025. AWR1243 single-chip 76-GHz to 81-GHz FMCW transceiver evaluation module. https://www.ti.com/tool/AWR1243BOOST.
[3] accessed Dec 05, 2025. AWR2243BOOST AWR2243 second-generation 76-GHz to 81-GHz high-performance automotive MMIC evaluation module. https://www.ti.com/tool/AWR2243BOOST.
[4] accessed Dec 05, 2025. AWR2944 high-performance automotive mmWave sensor evaluation module. https://www.ti.com/tool/AWR2944EVM.

[5] accessed Dec 05, 2025. TEF82xx: 77 GHz RFCMOS Automotive Radar Transceiver. https://www.nxp.com/products/radio-frequency-rf/radar-transceivers-and-socs/fully-integrated-77-ghz-rfcmos-automotive-radar-transceiver:TEF82xx.

[6] Andrew D Byrd, Robert D Palmer, and Caleb J Fulton. 2019. Development of a low-cost multistatic passive weather radar network. *IEEE Transactions on Geoscience and Remote Sensing* 58, 4 (2019), 2796–2808.

[7] Hongyu Chen, Yimin Liu, and Yuwei Cheng. 2024. A robust robot perception framework for complex environments using multiple mmwave radars. *IEEE Journal of Selected Topics in Signal Processing* 18, 3 (2024), 380–395.

[8] Weiyan Chen, Hongliu Yang, Xiaoyang Bi, Rong Zheng, Fusang Zhang, Peng Bao, Zhaoxin Chang, Xujun Ma, and Daqing Zhang. 2023. Environment-aware multi-person tracking in indoor environments with mmWave radars. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 3 (2023), 1–29.

[9] Li-Xuan Chuo, Zhihong Luo, Dennis Sylvester, David Blaauw, and Hun-Seok Kim. 2017. Rf-echo: A non-line-of-sight indoor localization system using a low-power active rf reflector asic tag. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. 222–234.

[10] Cong-Thanh Do and Hoa Van Nguyen. 2019. Tracking multiple targets from multistatic Doppler radar with unknown probability of detection. *Sensors* 19, 7 (2019), 1672.

[11] Laura Dodds, Tara Boroushaki, Kaichen Zhou, and Fadel Adib. 2025. Non-Line-of-Sight 3D Object Reconstruction via mmWave Surface Normal Estimation. In *Proceedings of the 23rd Annual International Conference on Mobile Systems, Applications and Services*. 445–458.

[12] Laura Dodds, Hailan Shanbhag, Junfeng Guan, Saurabh Gupta, and Haitham Hassanieh. 2024. Around the corner mmWave imaging in practical environments. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 953–967.

[13] Di Duan, Shengzhe Lyu, Mu Yuan, Hongfei Xue, Tianxing Li, Weitao Xu, Kaishun Wu, and Guoliang Xing. 2025. Argus: Multi-view egocentric human mesh reconstruction based on stripped-down wearable mmwave add-on. In *Proceedings of the 23rd ACM Conference on Embedded Networked Sensor Systems*. 1–14.

[14] Lili Fan, Junhao Wang, Yuanmeng Chang, Yuke Li, Yutong Wang, and Dongpu Cao. 2024. 4D mmWave radar for autonomous driving perception: A comprehensive survey. *IEEE Transactions on Intelligent Vehicles* 9, 4 (2024), 4606–4620.

[15] Andreas Frischen, Gor Hakobyan, and Christian Waldschmidt. 2020. Coherent measurements with MIMO radar networks of incoherent FMCW sensor nodes. *IEEE Microwave and Wireless Components Letters* 30, 7 (2020), 721–724.

[16] Andreas Frischen, Jürgen Hasch, and Christian Waldschmidt. 2017. A cooperative MIMO radar network using highly integrated FMCW radar sensors. *IEEE Transactions on Microwave Theory and Techniques* 65, 4 (2017), 1355–1366.

[17] Danei Gong, Naiyu Zheng, Binbin Xie, Jie Xiong, Shuai Wang, Yuguang Fang, and Zhimeng Yin. 2025. SeRadar: Embracing Secondary Reflections for Human Sensing with mmWave Radar. In *Proceedings of the 31st Annual International Conference on Mobile Computing and Networking*. 231–246.

[18] Michael Gottinger, Peter Gulden, and Martin Vossiek. 2021. Coherent signal processing for loosely coupled bistatic radar. *IEEE Trans. Aerospace Electron. Systems* 57, 3 (2021), 1855–1871.

[19] Junfeng Guan, Sohrab Madani, Suraj Jog, Saurabh Gupta, and Haitham Hassanieh. 2020. Through fog high-resolution imaging using millimeter wave radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11464–11473.

[20] Junchen Guo, Meng Jin, Yuan He, Weiguo Wang, and Yunhao Liu. 2021. Dancing waltz with ghosts: Measuring sub-mm-level 2d rotor orbit with a single mmwave radar. In *Proceedings of the 20th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2021)*. 77–92.

[21] Yulu Guo, Tongjia Zhang, Shu Sun, Meixia Tao, and Ruifeng Gao. 2025. Measurement and Analysis of Scattering From Building Surfaces at Millimeter-Wave Frequency. In *2025 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6.

[22] Soheil Hor, Shubo Yang, Jaeho Choi, and Amin Arbabian. 2023. MVDoppler: Unleashing the power of multi-view doppler for micromotion-based gait classification. *Advances in Neural Information Processing Systems* 36 (2023), 58064–58074.

[23] Tianshu Huang, Akarsh Prabhakara, Chuhan Chen, Jay Karhade, Deva Ramanan, Matthew O'toole, and Anthony Rowe. 2025. Towards foundational models for single-chip radar. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 24655–24665.

[24] Xu Huang, Hasnain Cheena, Abin Thomas, and Joseph KP Tsoi. 2021. Indoor detection and tracking of people using mmwave sensor. *Journal of Sensors* 2021, 1 (2021), 6657709.

[25] Hao Kong, Cheng Huang, Jiadi Yu, and Xuemin Shen. 2024. A survey of mmwave radar-based sensing in autonomous vehicles, smart homes and industry. *IEEE Communications Surveys & Tutorials* 27, 1 (2024), 463–508.

[26] Haowen Lai, Zitong Lan, and Mingmin Zhao. 2025. Non-Line-of-Sight 3D Reconstruction with Radar. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*.

[27] Haowen Lai, Gaoxiang Luo, Yifei Liu, and Mingmin Zhao. 2024. Enabling Visual Recognition at Radio Frequency. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking (MobiCom)*. 388–403.

[28] Yilong Li, Ramanujan K Sheshadri, Karthik Sundaresan, Eugene Chai, Yijing Zeng, Jayaram Raghuram, and Suman Banerjee. 2025. Medusa: Scalable Multi-View Biometric Sensing in the Wild with Distributed MIMO Radars. In *Proceedings of the 31st Annual International Conference on Mobile Computing and Networking*. 78–92.

[29] Yangyang Liu, Shuo Chang, Zhiqing Wei, Kezhong Zhang, and Zhiyong Feng. 2022. Fusing mmWave radar with camera for 3-D detection in autonomous driving. *IEEE Internet of Things Journal* 9, 20 (2022), 20408–20421.

[30] Chris Xiaoxuan Lu, Stefano Rosa, Peijun Zhao, Bing Wang, Changhao Chen, John A Stankovic, Niki Trigoni, and Andrew Markham. 2020. See through smoke: robust indoor mapping with low-cost mmwave radar. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 14–27.

[31] Jiaxi Lv, Guiyun Fan, Xinyue Fu, Jiahui Sun, Rong Ding, and Haiming Jin. 2025. mmWave-Based Relay Reflector Reconstruction for LiDAR-Free Around-Corner Human Sensing. In *IEEE INFOCOM 2025-IEEE Conference on Computer Communications*. IEEE, 1–10.

[32] Nishant Mehrotra, Divyanshu Pandey, Akarsh Prabhakara, Yawen Liu, Swarun Kumar, and Ashutosh Sabharwal. 2024. Hydra: Exploiting multi-bounce scattering for beyond-field-of-view mmwave radar. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 1545–1559.

[33] Ram M Narayanan, Daniel D Cox, James M Ralston, and Michael R Christian. 1996. Millimeter-wave specular and diffuse multipath components of terrain. *IEEE Transactions on Antennas and Propagation* 44, 5 (1996), 627.

[34] Tasin Nusrat and Stavros Vakalis. 2024. Addressing specularity: Millimeter-wave radar with distributed repeater apertures. *IEEE Transactions on Microwave Theory and Techniques* (2024).

[35] Akarsh Prabhakara, Tao Jin, Arnav Das, Gantavya Bhatt, Lilly Kumari, Elahe Soltanaghai, Jeff Bilmes, Swarun Kumar, and Anthony Rowe. 2023. High Resolution Point Clouds from mmWave Radar. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4135–4142.

[36] Akarsh Prabhakara, Vaibhav Singh, Swarun Kumar, and Anthony Rowe. 2020. Osprey: A mmWave approach to tire wear sensing. In *Proceedings of the 18th international conference on mobile systems, applications, and services*. 28–41.

[37] Kun Qian, Zhaoyuan He, and Xinyu Zhang. 2020. 3D point cloud generation with millimeter-wave radar. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 4 (2020), 1–23.

[38] Xian Shuai, Yulin Shen, Yi Tang, Shuyao Shi, Luping Ji, and Guoliang Xing. 2021. millieye: A lightweight mmwave radar and camera fusion system for robust object detection. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*. 145–157.

[39] Akash Deep Singh, Yunhao Ba, Ankur Sarker, Howard Zhang, Achuta Kadambi, Stefano Soatto, Mani Srivastava, and Alex Wong. 2023. Depth estimation from camera image and mmwave radar point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9275–9285.

[40] Xinghua Sun and Akshay Gadre. 2026. Towards Practical Bi-Static Polarimetric Imaging Using Commodity mmWave Radars for Material Sensing. In *Proceedings of the ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys '26)*. ACM, New York, NY, USA. https://doi.org/10.1145/3774906.3800470

[41] Valentín Barral Vales, Tomás Domínguez-Bolaño, Carlos J Escudero, and José A García-Naya. 2024. An IoT system for smart building combining multiple mmWave FMCW radars applied to people counting. *IEEE Internet of Things Journal* 11, 21 (2024), 35306–35316.

[42] Samuel Viegas, João R Reis, Telmo R Fernandes, and Rafael FS Caldeirinha. 2023. 4D MIMO Radar with 360° field of view: a practical validation. In *2023 International Wireless Communications and Mobile Computing (IWCMC)*. IEEE, 757–762.

[43] Yizhou Wang, Jen-Hao Cheng, Jui-Te Huang, Sheng-Yao Kuan, Qiqian Fu, Chiming Ni, Shengyu Hao, Gaoang Wang, Guanbin Xing, Hui Liu, et al. 2024. Vision meets mmwave radar: 3d object perception benchmark for autonomous driving. In *2024 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2769–2775.

[44] Yingqi Wang, Zhongqin Wang, Jian Andrew Zhang, Haimin Zhang, and Min Xu. 2023. Vital sign monitoring in dynamic environment via mmWave radar and camera fusion. *IEEE Transactions on Mobile Computing* 23, 5 (2023), 4163–4180.

[45] Teng Wei, Anfu Zhou, and Xinyu Zhang. 2017. Facilitating robust 60 {GHz} network deployment by sensing ambient reflectors. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. 213–226.

[46] Yanbo Wen, Shunjun Wei, Xiang Cai, Rong Shen, Mou Wang, Jun Shi, and Guolong Cui. 2024. Cmti: Non-line-of-sight radar imaging for non-cooperative corner motion target. *IEEE Transactions on Vehicular Technology* 74, 1 (2024), 179–190.

[47] Joshua Wurman, Stanley Heckman, and Dennis Boccippio. 1993. A bistatic multiple-Doppler radar network. *Journal of Applied Meteorology and Climatology* 32, 12 (1993), 1802–1814.

[48] Qian Yang, Hengxin Wu, Qianyi Huang, Jin Zhang, Hao Chen, Weichao Li, Xiaofeng Tao, and Qian Zhang. 2023. Side-lobe can know more: Towards simultaneous communication and sensing for mmWave. *Proceedings of the ACM on*

Interactive, Mobile, Wearable and Ubiquitous Technologies 6, 4 (2023), 1–34.
[49] Keisuke Yoneda, Naoya Hashimoto, Ryo Yanase, Mohammad Aldibaja, and Naoki Suganuma. 2018. Vehicle localization using 76GHz omnidirectional millimeter-wave radar for winter automated driving. In *2018 IEEE intelligent vehicles symposium (IV)*. IEEE, 971–977.
[50] Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, and Dina Katabi. 2018. Through-wall human pose estimation using radio signals. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7356–7365.
[51] Mingmin Zhao, Yonglong Tian, Hang Zhao, Mohammad Abu Alsheikh, Tianhong Li, Rumen Hristov, Zachary Kabelac, Dina Katabi, and Antonio Torralba. 2018. RF-based 3D skeletons. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*. 267–281.
[52] Wei Zhao and Rong Zheng. 2023. MARS: a mmWave Rotating Synthetic Aperture Radar System for Indoor Imaging. In *Proceedings of the First ACM Workshop on MmWave Sensing Systems and Applications*. 8–13.
[53] Yanzi Zhu, Yuanshun Yao, Ben Y Zhao, and Haitao Zheng. 2017. Object recognition and navigation using a single networking device. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. 265–277.
[54] Yanzi Zhu, Yibo Zhu, Ben Y Zhao, and Haitao Zheng. 2015. Reusing 60ghz radios for mobile radar imaging. In *Proceedings of the 21st annual international conference on mobile computing and networking*. 103–116.

A Artifact Appendix

A.1 Abstract

The code and instructions of MulDar *is publicly available at https://github.com/xsun2445/MulDar. The system implementation diagram is shown in Fig. 21.*

User can choose the data stream either from live radar data or from a pre-recorded radar data. We have documented detailed instructions on system implementation in the provided repository. We have recorded radar data in different scenarios that have appeared in the paper. We also provide a demo video of cars and every day objects at https://youtu.be/Dgcf0vv0X7E.

A.2 Artifact check-list (meta-information)

- **Dataset:** Dataset and instructions are provided in the github repository.
- **Hardware:** NVIDIA GPU for all evaluations using recorded data. For Live visualization deployed on real radars additional hardware are required: AWR2243BOOST*3, DCA1000EVM*3, Raspberry Pi, PC*4
- **How much disk space required (approximately)?:** 3 GB
- **How much time is needed to prepare workflow (approximately)?:** For software only (using recorded data) evaluations, we have prepared docker pipeline that can automatically configure the environment. We have also prepared python scripts for evaluations. For experiments involving hardware and real-time visualization, the preparation time is about one additional hour for connecting and testing devices.
- **How much time is needed to complete experiments (approximately)?:** Software-only: 10 min. Hardware-involved: 2 hr (mostly on hardware connection and debugging).

A.3 Description

A.3.1 How to access. Our code and dataset is publicly available at https://github.com/xsun2445/MulDar.

A.3.2 Hardware dependencies. To match the configuration in this paper:

- CPU: Intel Core i9-13900KF (13th Gen)
- GPU: NVIDIA GeForce RTX 4090 (24GB VRAM)

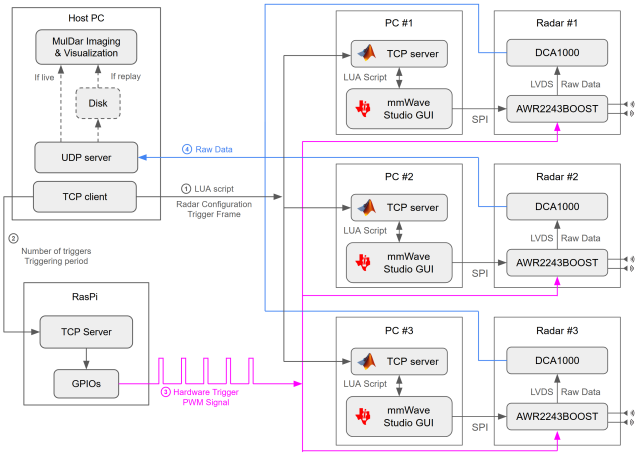

Figure 21: Implementation of MulDar

A.3.3 Software dependencies.

- Python 3.9
- nvidia/cuda >= 12.6.3
- cupy 13.6.0
- additional packages: numpy, matplotlib, scipy, numba, pyserial, pyyaml

A.3.4 Datasets. Datasets are hosted on huggingface: https://huggingface.co/datasets/xinghs/MulDar_Dataset. We also have provided python scripts for downloading the dataset.

A.4 Installation

One line installation (with required NVIDIA driver installed):

```
uv sync
```

Alternatively using docker:

```
docker build -t muldar .
docker run --gpus all -v $(pwd):/app -it muldar bash
```

A.5 Evaluation and expected results

After installation, download the dataset using the provided script:

```
python evaluations/download_dataset.py
```

A.5.1 Chamfer Distance Evaluation (Fig. 16). Evaluates imaging accuracy against ground-truth shapes using the symmetric Chamfer distance metric.

```
# Planar targets
python evaluations/chamfer/eval_chamfer.py \
    --object planar --mode eval
# Curvature targets
python evaluations/chamfer/eval_chamfer.py \
    --object curvature --mode eval
```

Expected results: JSON files with Chamfer distance values and comparison plots are saved in `evaluations/chamfer/chamfer_results/`. The results should show that the combined mono+multi imaging achieves lower Chamfer distance than mono-static-only or multi-static-only configurations across all materials and curvatures.

A.5.2 Car Imaging Evaluation (Fig. 19b). Generates bistatic SAR images of a car from four viewpoints.

```
python evaluations/car/bistatic_car.py
```

Expected results: Per-viewpoint SAR images and combined visualizations are saved in `evaluations/car/out/`. The combined mono+multi images should reveal fuller car contours compared to monostatic-only imaging.

A.5.3 Common Object Imaging (Fig. 14). Generates radar images for deformable, planar, curved, and real-world objects.

```
python evaluations/common_objects/eval_objects.py
```

Expected results: Visualization images are saved in `evaluations/common_objects/images/`. Objects should be clearly distinguishable in the combined imaging output.