

Mosaic: Leveraging Diverse Reflector Geometries for Omnidirectional Around-Corner Automotive Radar

Timothy Woodford
University of California San Diego
San Diego, California, USA
twoodfor@ucsd.edu

Eugene Chai
NEC Laboratories America
Princeton, New Jersey, USA
eugene@nec-labs.com

Xinyu Zhang
University of California San Diego
San Diego, California, USA
xyzhang@ucsd.edu

Karthikeyan Sundaresan
Georgia Tech
Atlanta, Georgia, USA
karthik@ece.gatech.edu

ABSTRACT

A large number of traffic collisions occur as a result of obstructed sight lines, such that even an advanced driver assistance system would be unable to prevent the crash. Recent work has proposed the use of around-the-corner radar systems to detect vehicles, pedestrians, and other road users in these occluded regions. Through comprehensive measurement, we show that these existing techniques cannot sense occluded moving objects in many important real-world scenarios. To solve this problem of limited coverage, we leverage multiple, curved reflectors to provide comprehensive coverage over the most important locations near an intersection. In scenarios where curved reflectors are insufficient, we evaluate the relative benefits of using additional flat planar surfaces. Using these techniques, we more than double the probability of detecting a vehicle near the intersection in three real urban locations, and enable NLoS radar sensing using an entirely new class of reflectors.

CCS CONCEPTS

• **Hardware** → *Wireless devices*; **Sensor devices and platforms**; *Digital signal processing*; **Sensor applications and deployments**.

KEYWORDS

Automotive sensing, NLoS radar, around-the-corner radar, lidar, ADAS

ACM Reference Format:

Timothy Woodford, Xinyu Zhang, Eugene Chai, and Karthikeyan Sundaresan. 2022. Mosaic: Leveraging Diverse Reflector Geometries for Omnidirectional Around-Corner Automotive Radar. In *The 20th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '22)*, June 25–July 1, 2022, Portland, OR, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3498361.3538944>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
MobiSys '22, June 25–July 1, 2022, Portland, OR, USA
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9185-6/22/06.
<https://doi.org/10.1145/3498361.3538944>

1 INTRODUCTION

Non-line-of-sight (NLoS) radars, which can “look around” obstacles and corners to detect oncoming traffic that will otherwise be obscured to vision-based sensors, is an important component in the integrated suite of sensors found on vehicles in the future. With the capability to detect objects around blind corners, such NLoS radars are crucial in reducing accidents and injuries due to collisions at such challenging intersections. In the US, collisions at intersections account for more than 23% of all motor vehicle crashes, and are responsible for 26% of loss of personal functional due to serious injury [27]. Prior work in the field of automotive safety shows that sight line obstructions are a key factor in a large fraction of these cases. For example, 18% of vehicles that crashed into pedestrians while turning had an obscured view of the scene. In addition, detailed reconstructions of real-world crashes have demonstrated that due to sightline obstructions, even an idealized automatic braking system would be unable to prevent around 30% crashes between left vehicles and opposite-direction traffic [4].

Current approaches to NLoS radar sensing [38, 42, 44] have focused on leveraging reflections from *planar* reflectors in the environment to “look around the corner”, to sense the presence of an occluded object. However, such approaches offer insufficient coverage of the obscured region to provide adequately accurate and early detection of vehicles and pedestrians at intersections. Given that a radar must detect vehicles anywhere in the region 0–30 m from the intersection (§ 2.1) to provide adequate early-warning for current safety-systems [4, 37], Figure 1 illustrates how techniques that rely only on planar reflectors fall short, and we further describe these shortcomings in § 2.2. In particular, these conventional NLoS radar systems only have a narrow window of coverage, and fail to identify vehicles that are within 5 m of entering the intersection, and often fail to provide any meaningful coverage when the radar is more than 15 m from the intersection. Worse, these systems rely on the availability of large, planar surfaces to act as relay reflectors, which may not be available at all in many non-urban scenarios. With such poor sensing coverage, it is impossible to design reliable safety systems based on such NLoS radars.

Unfortunately, the challenge of NLoS automotive sensing cannot be adequately solved with other radar/non-radar approaches. Camera-based NLoS sensing [13] works in limited scenarios where there are optical reflectors or indirect lighting [28, 49]; Lidar-based

NLoS sensing requires very high power, costly, specialized equipment that exceeds eye safety limits [24], and thus are impractical for current automotive deployments; Through-wall radars [9, 20] take advantage of the fact that radar can pass through certain materials that block visible light, but obstructions on a road are made from dense materials such as metal or concrete, which are opaque to radar at most frequencies; Cooperative sensing using V2X or V2V infrastructure [10] requires significant infrastructure development in both onboard and roadside sensors and thus faces economical and technical barriers to widespread deployment in the near/medium term.

In this paper, we re-examine the principles of NLoS radar sensing, and answer two key questions to accelerate their deployment in practical automotive scenarios. First, *how can radars "see" around corners by using non-planar, non-ideal reflection surfaces?* Most candidate reflectors (e.g., vehicles, street lighting, trees) that can redirect radar signals around obstacles have complex, non-planar surfaces. Our measurements (§ 2.2) show that if we can harness cylindrical reflectors in addition to planar ones, then we can expand the blind-spot coverage by 41% in an urban scenario, thus making it feasible to build safety systems based on NLoS radars. Second, analogous to how mirrors and traffic signals are installed in key locations to aid drivers in making blind turns safely, *how can we augment the environment to make it "radar-friendly" to enhance NLoS radar performance?*

In addressing these questions, we introduce MOSAIC, an integrated sensing suite that employs LiDARs to map the positions of the LoS reflectors, and automotive radars to process secondary reflections from those surfaces so as to detect people and vehicles around obstacles and blind corners. MOSAIC contains three key contributions:

Diversity-aware radar detection. We develop an algorithm that handles sensing paths where the AoA and AoD of the radar signals are not aligned, *i.e.*, the outgoing signal and returning signal follow different paths, which is common in practical NLoS environment. Our algorithm handles both noise and mobility within the sensing frame. We show that this technique improves the chance of detecting an occluded object within 20 m of the intersection by 20%.

Radar ambiguity resolution. We develop a method for resolving the location of objects when the path of the radar signal is ambiguous. Inspired by trilateration-based localization techniques, we build an algorithm that resolves ambiguous reflections to a single point, giving priority to paths with the highest signal strength. Using this technique, we can reliably localize objects to within 2 m using reflections from multiple non-planar reflectors (e.g., metal cylinders), without any advance knowledge of the reflection angle.

Radar-friendly environment design. In practice, not all the environments contain proper reflectors to enable NLoS radar. Our measurement profiling and analytical model lead to a deeper understanding of the performance limits of a joint LiDAR+Radar system for NLoS sensing. Furthermore, we derive guidelines on placing different types of artificial reflectors that can enable accurate and reliable NLoS sensing.

2 BEYOND-LOS RADAR MODEL

Conventionally, NLoS radar works as follows. Consider a monostatic radar and a moving target. The radar provides a heat map of RF power reflected back to it in terms of AoD, AoA, range, and doppler. The AoD and AoA can be independently resolved by using separate arrays of transmit and receive antennas. Further, suppose the moving target is occluded by an object. To sense the target location, we introduce a *relay reflector* which creates a NLoS path between the radar and the target. The radar signal traverses this path from the radar to the target, and then returns along the same path. The resulting radar heat map then includes an image of the target at a location on the far side of the reflector. By reflecting this part of the image across the reflector plane, we can determine the target location, as long as the pose and position of the reflector is known a priori.

Since millimeter-wave (mmWave) automotive radar only detects a sparse set of points, the radar alone may not capture the entire surface of a potential relay reflector. To build a more complete model of potential reflectors, prior research proposed either using a known map of nearby buildings' walls [44] or using LiDAR data [38]. Since an existing map of the scene is unlikely to capture mobile reflectors such as parked or moving vehicles, we opt to use onboard LiDAR data in MOSAIC. Lidar generates a much denser point cloud than a mmWave radar because it has a much shorter wavelength and finer angular resolution, leading to increased levels of backscatter. This means that we can use a single Lidar snapshot to build a model of the available *first-order* LoS reflectors in real-time.

2.1 Desired Coverage

Now we develop a model of the possible locations that fall in the blind spots of the radar due to NLoS occlusion. The NHTSA has gathered extensive data on car crashes in the US, and identified a number of crash types as disproportionately caused by sightline obstructions [27], as summarized in Table 1. In total, there are about twice as many crashes into *lateral* traffic as there are into *oncoming* traffic at intersections. This implies there are similar numbers of crashes from the left, right, and opposite sides of the intersection.

Our next task is to determine the exact regions of the road that a NLoS radar needs to cover to prevent collision. First, we determine how far in advance the radar must be able to find the occluded object that may cause a crash. Based on prior work on left turn into oncoming traffic crashes, we know that an automatic emergency braking (AEB) system must be able to sense a potential collision at least 1.5 sec to prevent the majority of crashes. A system that warns the driver without automatically applying the brakes benefits from sensing the object 2.5 sec or more before the potential collision [4]. For cross-traffic crashes, real-world crash reconstructions show that around half of all crashes can be avoided or mitigated if the cross-traffic vehicle is detected at least 1.5 sec before collision when AEB is used [37].

Urban speed limits in the US generally range from 15 mph (24 km/h, 7 m/s) to 35 mph (56 km/h, 16 m/s) [14]. Crashes involving turns may involve stationary or nearly-stationary vehicles near the edge of the intersection, so we also must be able to detect vehicles up to the edge of the intersection. To detect vehicles travelling at up to 16 m/s at least 1.5 s before a potential collision,

Crash Type	% Crashes	% FYL
Car from oncoming direction	6.9%	8.5%
Car from lateral direction	16.2%	17.2%
Crash into cyclist while turning	0.3%	0.4%
Crash into pedestrian while turning	0.3%	0.9%

Table 1: Prevalence of crash types in US disproportionately caused by sightline obstructions. FYL (Functional Years Lost) indicates severity of resulting injuries [27].




Figure 1: Blind spot problem for a NLoS radar that relies on planar reflectors. Due to specular reflections, both the pedestrian and magenta vehicle are beyond the radar coverage.

an AEB system should be able to detect vehicles in all directions 0–24 m from the edge of the intersection, from a vehicle located up to 24 m from the edge of the intersection. Likewise, a driver-warning system should ideally detect vehicles from 0–40 m from the intersection to allow for 3 sec of advance warning.

2.2 NLoS Radar Coverage Model

2.2.1 Ray tracing model of NLoS coverage. Next, we develop a ray tracing model to examine *the relationship between the locations of relay reflectors and the regions of the scene where the moving target can be sensed*. To demonstrate the problems inherent in current NLoS radar modalities, we focus on a relatively common geometry where two roads intersect at a right angle. The most common large, planar reflectors are either the walls of the buildings at the corners of the intersection, or the long sides of vehicles parked or driving on the road. Thus, the planar reflectors will generally be aligned to the axis of one of the two perpendicular roads. In addition, the width of road lanes, the minimum gaps between the road, and the potential locations of parked cars are governed by road standards, accessibility and zoning laws, and road-related laws. *Therefore, although there are many possible road configurations, we can generate representative scenarios without exhaustively measuring large numbers of real intersections.*

The width of the planar reflector blind spot for cross traffic on either the left or right of the intersection is

$$w_{blind} = x_{refl} - \left(\frac{x_{block}}{y_{block}} + \frac{x_{refl}}{y_{refl}} \right) y_{block},$$

where (x_{block}, y_{block}) are the coordinates of the corner of the blockage nearest to the center of the intersection, and (x_{refl}, y_{refl}) are the coordinates of the corner of the planar reflector closest to the center of the intersection. The width of the blind spot w_{blind} is a function of y_{blind} , the y coordinate at which we measure the blind spot.

To find the possible coordinate values, we note that the standard lane width for new roads in the US is 3.7 m, and that buildings must be set back at least 1.5 m from the roadway in order to leave enough space for a sidewalk that complies with the Americans with Disabilities Act [29]. We use these standard dimensions to construct intersections that correspond with common urban street grids.

We consider 3 representative intersection sizes: (i) two 2-lane roads, (ii) a 2-lane road and 4-lane road, and (iii) two 2-lane roads with an additional lane of on-street parking in each direction. At each intersection, we place buildings with smooth concrete walls on all four corners, at the minimum possible setback from the road. A vehicle turning left blocks the view of oncoming traffic from the perspective of the vehicle with the radar, as shown in Figure 1. For both the vehicle carrying the radar and the target vehicle, we consider a set of locations from the edge of the intersection to 32 m from the intersection, at intervals of 1 m, as discussed in § 2.1.

We then determine percentage of cases where the radar vehicle can sense the target vehicle in three setups: (i) Using only large, planar reflectors, as in prior studies. (ii) Adding traffic light poles as relay reflectors, each placed 0.5 m from the edge of the adjacent traffic lane. (iii) Considering all paths including those where the NLoS signal returns along a different relay reflector than the one along which it departed. We use our measurements of reflection loss from § 2.4 for the relay reflectors, and use the measurements from [26] to represent the vehicle’s reflection strength, *i.e.*, its radar cross section (RCS) values.

We find benefits from both the additional of cylindrical reflectors and of paths with differing AOA and AOD values, with all but one scenario showing over 20% of coverage improvement. Cylindrical reflectors are especially useful in sensing oncoming traffic, where the ordinary placement of planar reflectors leaves a large coverage gap near the intersection. *Therefore, the key to comprehensively sensing occluded objects in automotive scenarios is to use as many reflection paths as possible, and in particular nonplanar relay reflectors and paths where AOA and AOD do not match.* Using large, planar reflectors and assuming that the AoA and AoD are equal, as in prior work, significantly limits the cases in which NLoS radar is useful.

	Planes	+Cylinders	All Paths
2/2 Lane Oncoming	57%	73%	78%
2/2 Lane Cross	58%	62%	73%
2/4 Lane Oncoming	56%	81%	85%
2/4 Lane Cross	60%	71%	84%
4/4 Lane Oncoming	75%	71%	89%
4/4 Lane Cross	55%	86%	95%

Table 2: Ray tracing-based coverage estimate for three different intersections. X/Y denotes intersection between an X-lane and Y-lane road.

2.2.2 Demonstration: Urban Blind Spots. To demonstrate that the coverage limitations we discovered with the analytical ray-tracing




Figure 2: The T Junction beyond-LoS sensing scenario. The pedestrian is not detected by NLoS radar beginning 4 m from the corner of the blockage.

model correspond with real-world coverage limitations, we conduct a measurement study at a real intersection. In the first scenario, shown in Fig. 2, we place the radar near a 3-way intersection with buildings on two corners. We have a pedestrian move at different distances from the intersection, and identify the locations where the pedestrian can be detected. We find that the blind spot in this T-junction has a width of 4.1 m, which is within a 2% error of the theoretical blind spot size of 4 m based on the dimensions of the intersection.

In the second scenario, shown in Fig. 10, we consider a more extreme case where buildings are set back significantly from the road. The vehicle with a radar is attempting to turn right, but its view of cross traffic is blocked by a left vehicle, and thus may collide with another vehicle crossing the intersection in the perpendicular direction. We simulate the vehicle blockage using an RF absorber, and verify that the LoS path is blocked by ensuring that all the radar returns associated with the moving vehicle are at a range at least 5 m greater than the direct path. We find that the occluded moving vehicle may be detected at distances of over 40 m using only reflections from the traffic light pole. However, due to the large gap between the road and the buildings, we do not find reflections from any of the building walls.

2.3 Generalized NLoS Radar Sensing Model

Recall that a NLoS radar operates by observing reflections from a *target* scatterer, such as a pedestrian or motor vehicle, reflected from a *relay reflector*. Based on the intuitions gained from our coverage analysis, we expand our MOSAIC radar model beyond existing around-the-corner model [38, 42, 44] by: (1) allowing any possible reflective object, not just large planes; (2) considering that some objects may allow both transmission and reflection, *i.e.*, the radar may receive both reflections from in front of the relay and returns from target scatterers behind the relay; (3) decoupling the transmission and return paths.

Based on this expanded radar model, the received signal strength from a NLoS radar path will be

$$P_{rx} = \frac{P_{tx} G_{tx} G_{rx} \lambda^2 \sigma_{tg} \prod_i \sigma_{refl,i}}{(4\pi)^{N+1} \prod_i R_i^2}, \quad (1)$$

where R_i is the length of the i^{th} segment of the path, N is the total number of path segments, G_{rx} and G_{tx} are the receive and transmit antenna gains, σ_{tx} is the radar cross section (RCS) of the target, and σ_{refl} are the effective RCS values of the relay reflectors. The RCS of an object determines the proportion of incoming RF

energy reflected in a given direction, and is defined as $\sigma = A_{refl} \Gamma$, where A_{refl} is the effective cross-sectional area of the reflector and Γ is the efficiency of the reflector material. Note that for large objects such as walls, the area A_{refl} depends on the size of the surface illuminated by the radar. An object can be detected as long as $P_{rx} > S_{min}$, *i.e.*, the received power is greater than the radar's sensitivity level.

Since many objects in the roadside environment may be at least partially transparent to mmWave radar signals, we also consider the case where the target object is *behind* an object within LoS, rather than reflected from it. The received signal strength after passing through such an object will be

$$P_{rx} = \frac{P_{tx} G_{tx} G_{rx} \lambda^2 \sigma_{tg} T_{occl}^2}{(4\pi)^3 R_1^4 R_2^4},$$

where T_{occl} represents the total transmission coefficient of the occluding object.

This expanded model introduces *localization ambiguities*. Specifically, *interaction type ambiguities* arise where an object (*e.g.*, windows, wood walls, plastic boards, *etc.*) may either reflect incidental signal or allow it to pass through. *Reflection angle ambiguities* originate where we cannot determine the angle of reflection at the relay reflector. This occurs when the relay reflector's normal direction varies rapidly enough that the radar's angular resolution is insufficient to determine the normal direction at the point of reflection. We will provide additional details on the nature of these ambiguities and our solution in § 2.5.

2.4 Measurement Profiling of Diverse Reflectors

We have shown that the large, stationary relay reflectors favored by previous studies frequently have large blind spots, even in near-ideal urban scenarios. Covering these blind spots will require additional non-planar reflectors such as moving vehicles and lamp poles, which are not considered by prior work [38, 42, 44]. However, compared with large planar reflectors, these smaller and more geometrically complicated relay reflectors may have significantly different reflection characteristics, which we will profile in this section.

Our test proceeds in three phases. First, we measure the maximum NLoS sensing range of a moving pedestrian using reflections from the given relay reflector. The maximum range depends on the RCS of the relay reflector, and the distances between the relay and the target and radar. Since the RCS may vary with the angle between the incident and reflected rays, we conduct this experiment where the angle varies between 0° (the monostatic RCS), 20° , and 45° . To ensure that we only detect NLoS paths in this phase, we point the radar away from the target and place RF absorbers adjacent to the radar to completely block LoS paths. The monostatic radar we use cannot directly measure the bistatic RCS. Instead, we measure the NLoS detection range of a scattering target when varying the angle between the radar-reflector ray and the reflector-target ray. Using Eq. (1), and assuming for now the transmitted path is the same as the return path, we have

$$\sigma_{refl}^2 > R_1^4 R_2^4 \frac{S_{min} (4\pi)^5 \sigma_{tg}}{P_{tx} G_{tx} G_{rx} \lambda^2}, \quad (2)$$




Figure 3: The relay reflectors evaluated: (a) glass wall, (b) rough concrete, (c) smooth concrete, (d) smooth brick, (e) guardrail, (f) lamp post, (g) tree, (h) traffic light pole, (i) sedan, (j) pickup truck, and (k) small SUV.

such that we can estimate the relay reflector’s RCS by finding the maximum detection range of a target assisted by that reflector. To make this measurement more intuitive, Table 3 shows the RCS in terms of the *worst-case detection range* of a pedestrian, *i.e.*, the values of R_1 and R_2 such that the total range $R_1 + R_2$ is minimized. In addition, we calculate the detection range of a car by using the RCS measurements of a sedan from [26].

In the final phase, we evaluate the viability of the specular reflection model used in prior work [38] for localizing a moving target. We use the cell-averaging CFAR approach to find the range and Doppler values associated with the target, and then estimate the corresponding azimuth angle of the radar path using maximum likelihood estimation. We estimate the target location by using its surface normal determined using the 20 nearest neighboring points from a lidar point cloud (using the lidar in § 6.1). We compare the target location estimated by a NLoS radar to the ground-truth obtained using lidar. For comparison, we also include the estimation error for a pedestrian moving within the radar’s LoS.

Table 4 shows the mean square error (MSE) corresponding to the various relay reflectors in Figure 3. The first set of reflectors are the outdoor, street-level walls of various buildings, typically used in prior work [38, 42, 44]. Note that very similar materials, such as concrete walls with two different finishing types, rough and smooth, may have substantially different RCS values. The rough surfaces result in more scattering rather than specular reflections, which reduces the SNR for target sensing. The second set of measurements are stationary curved objects typically placed on the side of the road. The building corner is a 1mx1m concave right-angle corner in an exterior wall made of smooth concrete. The lamp post is a 14 cm diameter smooth metal pole, of a size typically used to light pedestrian areas, with a similar size as the poles used to mount pedestrian signals at intersections [5, 11]. The traffic light pole is a 32 cm diameter smooth metal pole, of a similar diameter and structure to other poles that are placed at the corners of intersections to support signal lights on mast arms [5, 11, 12, 33]. The concrete pillar is a 65 cm diameter cylindrical structural support for a building. The tree trunk is a 90 cm diameter rough cylinder, larger than typical trees planted near roads. This suggests that even very

	Pedestrian (m)	Car (m)
Glass wall	50 – 55	63 – 99
Smooth concrete wall	31 – 33	39 – 59
Smooth brick wall	21 – 25	26 – 45
Rough concrete wall	12 – 14	15 – 25
Guard rail	12–18	15 – 32
Square lamp post (side)	26 – 30	33 – 60
Metal fence	13 – 16	16 – 29
Small metal sign	16 – 18	20 – 36
Building corner (concave)	17–20	21–40
Metal bollard	7 – 9	9 – 16
Circular sign post	5 – 6	6 – 11
Square lamp post (corner)	3 – 5	4 – 9
Circular lamp post	8 – 10	10 – 18
Traffic light pole	12–13	15 – 24
Concrete pillar	5 – 6	6 – 11
Tree trunk	4 – 5	5 – 9
Sedan (side)	16 – 17	20 – 31
Sedan (front)	8 – 10	10 – 18
Sedan (rear)	10 – 12	12 – 22
SUV (side)	28 – 30	35 – 54
SUV (front)	5 – 6	6 – 11
SUV (rear)	11 – 12	14 – 22
Pickup truck (side)	18	23 – 32
Pickup truck (front)	10–11	13 – 20
Pickup truck (rear)	12	15 – 22

Table 3: Reflector object range measurements. We include planar/directional reflectors, quasi-omnidirectional reflectors, and vehicles in separate sections.

large trees will have lower range than metal cylinders available near roads.

2.5 A Typology of NLoS Reflectors

For the purposes of NLoS sensing, we categorize all objects in the environment into three classes: quasi-planar, quasi-omnidirectional, and non-reflective. Quasi-planar objects were the focus of previous NLoS radar work. In contrast, MOSAIC can employ quasi-omnidirectional reflectors as well. Non-reflective objects such as foliage, ground clutter, and outdoor furniture may be detected on the radar, but do not have a sufficient RCS to create a NLoS path.

Based on the reflector profiling in Tables 3 and 4, we develop a typology of typical relay reflectors available near roads, and describe how MOSAIC estimates the NLoS paths created by each reflector type. Note that MOSAIC only relies on reflectors within the radar’s LoS, and the reflector type can be detected by a co-located lidar sensor (details in § 6.1).

	MSE (m)
Baseline (LoS)	0.57
Glass wall	0.75
Sedan (side)	0.91
Sedan (front)	1.02
Traffic light pole	2.89

Table 4: Target localization errors with diverse reflectors.

Directional reflectors. Directional reflectors along the sides of roads generally take the form of planar surfaces such as walls, fences, and other architectural features. Such reflectors bear a one-to-one relationship between incident and reflection angles. Therefore, we can determine the original location unambiguously. The directional surfaces we measured exhibit a variety of reflection strength values, depending on material, roughness, texture, *etc.* This observation aligns with prior measurements of the mmWave band channel, which found that most materials for exterior walls are rough and cause high transmission loss [21].

Quasi-omni reflectors. Quasi-omni objects reflect radar signals to a wide range of angles which helps cover blind spots. However, they suffer from higher reflection loss compared to flat surfaces. Amongst the cylindrical surfaces we measured, only the metal ones are usable as relay reflectors. The other weakness is that we usually cannot conclusively determine the angle of a signal reflected from the curved reflector. This is a result of limitations in both radar and lidar sensors. State-of-the-art automotive radar has a best-case angular resolution of a few degrees (*e.g.*, 1.4° in our experimental radar, § 6.1). At a range of just 3 m, this corresponds to a 7 cm uncertainty in the actual reflection location. For a 32 cm diameter traffic light pole, it means that there is a 39° uncertainty in reflection angle, even assuming perfect knowledge of the locations and geometries of reflectors! For this reason, Table 4 shows that *curved reflectors such as the traffic light pole and front of the sedan have significantly higher localization error than flat surfaces.*

Vehicles. Our measurements show that, when using vehicles as relay reflectors, the detection ranges vary widely depending on the vehicle form factor and orientation. For example, the fronts of most consumer vehicles have low RCS values due to the sloped and/or irregular surfaces. As a result, these surfaces had ranges comparable to a small metal lamp post at best, or a concrete pillar at worst. The rear section of most vehicles had measured relay ranges in between a lamp post and a traffic light pole. The sides of vehicles showed the greatest variability, with sedans providing 16 m of range, and SUVs 28 m, comparable to a smooth wall.

Vehicles may be modelled as a combination of curved and planar reflectors. We thus begin by using plane fitting to model the sides of the vehicle. The remaining surfaces are neither circular nor planar, so we instead determine the normal values by fitting a line to the k -nearest neighbors in the point cloud. As with cylindrical reflectors, this curved surface will have an ambiguous reflection angle, so we model this ambiguity as an ellipse, and narrow this ellipse to an arc of points corresponding with the possible normal angles within the uncertainty of the radar's azimuth angle estimate.

Random surfaces. Whereas the continuous variation of curved and planar surfaces creates reflections that tend to add constructively, the discontinuous surfaces on irregular shaped objects usually do not add coherently. Therefore, we could not find any significant NLoS reflection paths from objects such as chairs, foliage, benches, and do not consider them as usable relay reflectors.

3 AOA/AOD DECOUPLING

We now introduce the practical steps in MOSAIC to leverage the generalized NLoS sensing model. The first aspect is to use *asymmetric sensing paths* with distinct AoA and AoD.




Figure 4: Benefits of AOA/AOD decoupling: (a) range improvement from capturing forward-scattering paths and (b) improved localization using path diversity.

3.1 AoA and AoD Decoupling Algorithm

Traditional monostatic radar systems assume the AoA approximately equals the AoD, *i.e.*, the signal propagates directly from the radar to the target, and then back following the reverse LoS path. By making this assumption, a radar can achieve high angular resolution by leveraging transmit and receive arrays with differing antenna spacing. This raises a problem when $\text{AoA} \neq \text{AoD}$: *how can we detect separate AoA and AoD values without a significant loss of angular resolution and/or increase in aliasing?*

Detecting asymmetric paths. Our core insight in solving this problem is that the *radar paths are reciprocal*. In other words, a path with AoA θ_a and AoD θ_b will have a reverse path with AoA θ_b and AoD θ_a . Therefore, if we can match these two paths, we can achieve similar angular resolution to cases where $\theta_a = \theta_b$. To accomplish this, we leverage the MIMO beamforming capability of modern automotive radars to sweep through a series of AoAs, and simultaneously receive the returned signals at multiple antennas to estimate the AoA.

First, we isolate the signal path by selecting a single range value and single doppler value. Then, we create a 2D AOA/AOD matrix using a standard matched-filtering approach to beamform to each azimuth angle of arrival. We define $A_{i,j}$ to be the resulting beamforming power at AoD index i and AoA index j . To detect AOA/AOD azimuth angle pairs using this matrix, we extend the standard cell-averaging CFAR algorithm [36]. In the second stage, we estimate the noise power $P_n = \frac{1}{N} \sum_{m=1}^N x_m$ using N adjacent bins AOA/AOD bins. Next, we detect asymmetric paths using two criteria: (i) The beamforming power at both points exceeds a threshold, *i.e.*, $A_{n,m} > T_p$ and $A_{m,n} > T_p$. (ii) The power of the reciprocal paths is sufficiently close, *i.e.*, $|A_{n,m} - A_{m,n}| < T_m$. This proceeds from the assumption that the reciprocal paths have the same path loss. Finally, we localize the target using simple geometrical models. If at least one of the reflectors is a plane, we can resolve the target location by finding the intersection of paths reflected from the planes. For cylindrical reflectors, we generate a set of ambiguous points falling along an ellipse, as described in § 2.5.

As with standard CFAR techniques, proper selection of the detection threshold is crucial. We set the beamforming power threshold $T_b = \alpha_b P_N$, where α_b is a coefficient selected for the desired false-alarm probability P_{fa} , and $\alpha_b = N(P_{fa}^{-1/(2N)} - 1)$ following classical CFAR [36]. Note that we add an exponent of $1/2$ to reflect the fact

that the additive noise in $A_{n,m}$ and $A_{m,n}$ are independent when $m \neq n$, which reduces the false positive probability by a power of 2 compared to a single value. The matching threshold T_m is also set based on the measured noise of the angular spectrum, *i.e.* $T_m \geq T_b$. However, since different AoDs are active at different points in time, we must also consider target and environment dynamics, as follows.

Handling target/radar mobility. To ensure an acceptable false negative rate, we must also consider the mobility of the target. Our starting assumption was that the two reciprocal paths have approximately equal beamforming power. However, if the NLoS sensing paths between the radar and target change substantially over the course of the radar sensing frame, then this assumption may not hold true. For example, if the total range of the sensing path increases between the time when the radar transmitter beamforms to azimuth θ_a and when it beamforms to azimuth θ_b , the path with AoD θ_b will have a lower beamforming power than that with azimuth θ_a . Similar cases may occur where a target with nonuniform RCS rotates with respect to the relay reflectors, or where a new blockade begins to occlude the NLoS path.

For our purposes, we assume that changes in range during the radar frame will be the most common and significant problem. Define t_f as the length of a radar frame, the time between the beginning of the first chirp of the first beamformed angle and the end of the final chirp of the final beamformed angle. When a target moves a distance Δx , the maximum change in total NLoS path length will be $2\Delta x$ [38]. Therefore, a radar with frame duration t_f will have a maximum change in beamformed power of $40(\log_{10} R_1 - \log_{10}(R_1 - t_f v))$, where R_1 is the initial radar range and v is the maximum possible velocity of the target. To ensure a false negative probability of P_{fn} or lower for the angle matching stage, the radar frame time should satisfy:

$$t_f < \frac{R_1}{v} \left(1 - 10^{-T_m/40} \right) \quad (3)$$

Therefore, we must limit the frame time to 225 ms to ensure a 1% false negative rate where the maximum velocity is 16 m/s (the maximum typical urban speed limit found in § 2.1), the minimum range is 5 m (the minimum range in § 2.1), and the matching threshold limit is up to 40 dB. State-of-the-art automotive radar can easily achieve a shorter frame time (*e.g.*, 90 ms in the TI radar we use) to meet the requirement.

3.2 Microbenchmark Validation

Angular resolution. To determine the angular resolution of the above method, we place two cylindrical reflectors side by side and move them apart until they are distinguishable in the radar's range FFT spectrum. We use the standard transmit beamforming method used by the TI Cascade radar [16] as a baseline. The baseline method achieves an angular resolution of 3° and only works when AoA=AoD. Whereas our method performs the same when AoA=AoD, the resolution is only slightly coarser (5°) when AoA≠AoD.

Detection performance. The AoA/AoD decoupling is especially effective in an oncoming traffic scenario with relay reflectors on both sides of the road. The signal therefore is transmitted toward a reflector on one of the road, forward scattered toward the reflector on the other side of the road, and returns along a different





Figure 5: AOA/AOD decoupling experiment setup. **Figure 6: SNR improvement with AOA/AOD decoupling.**

angle than it departed. Prior measurement profiling of vehicles in anechoic chambers showed that the RCS within 10° of the exact front of a vehicle is around 4-8 dBsm higher than that at an offset angle of $10-45^\circ$ [26, 39]. Therefore, if there are two reflecting paths from opposite sides of the road, such that the paths correspond to a specular reflection from the high-RCS region of the front of the vehicle, the resulting SNR is much higher than an AoA=AoD path, which reflects from a lower-RCS corner of the vehicle.

First, we verify that our algorithm works in an artificial environment with an arrangement of reflectors similar to an oncoming traffic scenario. We place two 6 cm radius metal poles 1 m each from the radar. Holding a 0.4 m^2 metal plane pointed directly toward the midpoint between the two cylinders, we walk toward and away from the cylinders (Fig. 5). We find that considering paths with separate AoA and AoD and using our detection algorithm yields an average SNR improvement of 3.8 dB and improves the maximum range by an average of 27%.

In addition, we conduct an experiment at a public intersection using a set of four bollards as relay reflectors. We point the radar towards the bollards, to ensure that the entire intersection is outside of the LoS region. We collect data from 5 approaching vehicles. Again, we find that considering paths with separate AoA and AoD to detect oncoming traffic yields an average SNR improvement of 3.4 dB and improves the maximum range by an average of 16%.

Second, due to the difficulty of collecting large amounts of sensing data from the center region of public roads, we verify that this improvement holds across many other possible scenarios by conducting a simulation across many oncoming traffic sensing scenarios. We consider the intersection scenarios from § 2.2, and find all possible paths from the radar vehicle to the oncoming target vehicle, using both traffic light poles and the walls of adjacent buildings. We use the measured vehicle RCS values in [26, 39] and derive bistatic RCS values using [18]. Figure 6 shows the CDF of the SNR difference obtained by the including NLoS paths vs. LoS alone. Table 5 shows the coverage improvement. We find SNR improvement in about 40% of cases, showing our AoA/AoD decoupling method significantly expands the region of coverage of left-side and oncoming traffic.

	Avg. SNR Diff (dB)	Coverage Diff (%)
Oncoming	2.0 dB	20%
Left Cross	1.7 dB	30%
Right Cross	0.5 dB	9%

Table 5: Simulated SNR and coverage area improvements for vehicles within 20 m of intersection.

4 RESOLVING PATH AMBIGUITIES

A disadvantage of using omnidirectional or nonplanar reflectors is that they greatly reduce the angular resolution, which greatly

reduces the accuracy with which the radar can localize the object of interest, as we previously demonstrated in Table 4. To compensate for this reduction in accuracy, we proceed in two phases: *first*, we develop a model of the reflection ambiguities that result from the use of nonplanar reflectors, and *second* we describe an algorithm that can find the most-likely true location of the target.

First, we define the set of *ambiguous locations* as an isorange ellipse defined by

$$||\mathbf{l} - \mathbf{r}_1|| + ||\mathbf{l} - \mathbf{r}_2|| = 2R - ||\mathbf{s} - \mathbf{r}_1|| - ||\mathbf{s} - \mathbf{r}_2||,$$

where \mathbf{l} is the ambiguous target location, \mathbf{s} is the location of the radar sensor, R is the pseudorange measured by the radar, and \mathbf{r}_1 and \mathbf{r}_2 are the locations of the AoA and AoD relay reflections, respectively. We then narrow this set of ambiguous points by only including reflection angles corresponding to points within the angular uncertainty of the radar. We describe the method to fully resolve the target location in § 4.1.

We begin with the set of reflectors described in § 2.5. Each reflection angle ambiguity or interaction type ambiguity corresponds to a set of points on an ellipse defined by the focal points located at the relay reflectors. Uniquely resolving such an ambiguity to a single point of target location may require up to three reflection paths. If the AOA and AOD estimation has been decoupled as described in § 3, we may obtain the necessary number of paths using only two reflectors, by including a path that departs using one reflector and returns using the second reflector. Three paths are not necessary in all scenarios, as the possible set of points may be limited by either constraints on possible reflection angles, or because some points on the ellipse fall within the direct LoS of the radar.

4.1 Ambiguity Resolution Algorithm

To resolve the location of a target vehicle from a set of ambiguous reflections, we begin by mapping each detected range/angle point to a reflector. We then generate a heatmap grid with cell size equal to the expected maximum width of the target to ensure that multiple scattering points on a single target will resolve to a single location. For example, for pedestrian localization, we use a grid size of 0.5 m. For each reflector, we draw the set of ambiguous points to the heatmap as either a set of points or an antialiased [48] arc or ellipse. The magnitude of each of these sets of points is equal to the CFAR-estimated SNR [36] of the detected point. Where multiple point sets overlap, we add the magnitudes together. Next, we resolve each NLoS path. We generate a coarse estimate the target location for each NLoS path by selecting the grid cell with the maximum magnitude, beginning with the highest-SNR path, then proceeding to lower-SNR paths. Finally, for each grid cell containing more than one resolved path, we refine the target location estimate by using a least-squares method [8].

However, even if the radar paths may algebraically resolve to a single point, error and uncertainty in range measurements and reflector locations may still introduce localization error in our least-squares estimate. The exact nature of this error is dependent on the geometric locations of the relay reflectors. For example, as shown in Figure 7, if the relay reflectors are closely spaced, range errors will have a significant effect on target localization accuracy compared to a near-optimal placement. To quantify these accuracy limitations, we adapt the concept of geometric dilution of precision




Figure 7: Influence of reflector locations on localization error. Reflectors on the left have a small angle with respect to the target, with larger GDOP than reflectors on right.




Figure 8: Diagram of the GDOP microbenchmark experiment. The radar points toward the three metal cylinders, which reflect the radar signal toward the pedestrian.

(GDOP), which is used to determine geometry-related accuracy characteristics in localization systems. As in the standard GDOP formulation, we linearize the equations of the ellipses about the intersection point to derive a least-squares problem so as to find the intersection point [8]:

$$\delta \hat{\mathbf{t}} = (H^T H)^{-1} H \delta \mathbf{t}, \text{ where}$$

$$H = \begin{bmatrix} e_{x1} & e_{y1} \\ \vdots & \vdots \\ e_{xn} & e_{yn} \end{bmatrix}, \quad e_{xi} = \frac{\hat{x} - X_i}{\hat{r}_i}, \quad e_{yi} = \frac{\hat{y} - Y_i}{\hat{r}_i}, \quad (4)$$

$\hat{\mathbf{t}} = (\hat{x}, \hat{y})$ is the estimated target location, (X_i, Y_i) are the relay reflection points, and $\hat{\mathbf{y}} = (r_1, r_2, \dots, r_n)$ are the estimated ranges between the relay reflectors and the target. Again following the standard GDOP formulation, we determine the error covariance matrix R_{tt} as a function of the range error covariance matrix R_{rr} .

$$R_{tt} = (H^T R_{rr} H)^{-1}$$

This formula may be used to estimate the uncertainty of a given target location estimate, as shown in § 4.2, providing a downstream ADAS with a confidence level that may be used to provide appropriate false-positive rates. It also provides insights into the ideal placement of curved reflectors in an environment. Namely, we can use this model to show that the variances of the target coordinates will be minimized when the relay reflectors are at approximately right angles with respect to the target. Finally, this model shows that additional relay reflectors will only improve the localization result, not degrade it [43].

4.2 Microbenchmark: Accuracy and GDOP

To evaluate the effect of reflector geometry on the accuracy of the location estimates, we use a set of three 6 cm radius metal lamp poles

GDOP Values	Localization MSE (m)
10 ± 10	0.5 ± 0.2
30 ± 10	1.5 ± 0.6
50 ± 10	2.0 ± 0.9

Table 6: Estimated localization error for different ranges of GDOP values. The MSE uncertainty is given in *std*.

spaced 1.5 m apart to localize a pedestrian, moving linearly from 3-10 m from the poles at random angles. The pedestrian stands behind the radar to stand outside of its LoS field of view, as shown in Figure 8. We use the target localization technique described above to localize the pedestrian, and the procedures described in § 6.1 to compare the location of the moving pedestrian to a ground-truth position obtained using Lidar. In Table 6, we show the results of this evaluation, with the GDOP defined as $GDOP = \text{tr}((H^T H)^{-1})$. This shows an approximately linear increase in the localization MSE as the GDOP value increases, as expected. It signifies the importance of the locations of the set of reflectors we use to resolve reflection angle ambiguities. In addition, it shows that our MOSAIC radar system can detect scenarios where the available reflector geometry may not provide sufficient accuracy.

5 RADAR-AWARE ENVIRONMENT DESIGN

Even with the expanded set of NLoS reflection paths from AoA/AoD decoupling and ambiguity resolution, many intersections may not have a sufficient set of reflectors to enable reliable NLoS sensing. For example, a suburban or rural intersection may only have a few trees and a small number of vehicles near the intersection to act as reflectors. Given that trees have very short NLoS sensing ranges (§ 2.4), and that vehicles provide only very limited geometric coverage of the environment (§ 2.2), these environments are unlikely to offer comprehensive NLoS coverage.

To remedy this problem, we propose to deploy *passive reflectors* to inexpensively expand the NLoS radar coverage at intersections lacking sufficient ambient reflectors. Since these reflectors must be deployed in a way that does not unduly disrupt the environment, we consider reflectors that already exist in many roadside locations, and may be either added or re-positioned to improve radar coverage. For instance, we consider (i) metal signs, which act as directional reflectors, and (ii) metal poles (*e.g.*, street lights, traffic light supports, etc.), which as quasi-omnidirectional reflectors.

To determine the size and type of reflector that should be placed at a given intersection, we must consider both the *range* and the *field of view* each reflector can provide. Quasi-omnidirectional reflectors have a wide field of view (at least 45° according to our measurements), but a low RCS, which we measured in § 2.4.

Planar reflectors, by contrast, offer a much higher RCS, but more limited field of view. Assuming that the radar is in the far field of the reflector, the RCS of a planar reflector is $4\pi A_e^2/\lambda^2$, where A_e is the effective area of the metal reflector. For our radar to sense an approaching vehicle at ranges of at least 24 m from the intersection (see § 2.1), and using the signal strength model from Eq. (1), the metal reflector must have an area of at least 347 cm^2 . To sense a pedestrian in a crosswalk adjacent to the intersection, a metal reflector must have an area of at least 167 cm^2 . This is the size of a small road sign. In practice, field of view requirements

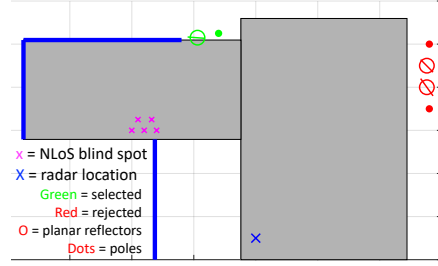


Figure 9: Demonstration of the reflector placement algorithm. Red objects are pruned reflectors. Dots are cylinders. Selected reflector orientations are inside each circle.

often result in much larger planar reflector size than RCS requirements alone. Planar reflectors have a field of view of approximately $2 \arctan\left(\frac{w}{2d}\right)$, where d is the distance from the radar to the artificial reflector, and w is the width of the reflector. Note that because the arc of the width of the region covered by a relay reflector is approximately $\frac{R}{d}w$, where R is the total range, the required width of the planar reflector is minimized as the distance from the radar to the reflector and increases when region of interest is further from the reflector. For example, to enable sensing pedestrians inside a 3m wide crosswalk area on a two-lane road using a radar at a fixed location at the edge of the intersection, a planar reflector must be at least 1.8m wide. A planar reflector designed to detect cross traffic at this intersection would be at least 1.1m wide, and a reflector used to detect occluded oncoming traffic would only need to be 0.6m wide. By contrast, a 32 cm diameter traffic light pole such as the one we measured previously would be sufficient to enable NLoS sensing in the crosswalk area, which requires much less horizontal space than a 1.8m planar reflector.

To determine the orientation θ_i of a planar reflector, we consider the a polygon enclosing the occluded region with points $\{O_i\}$. For each occluded point O_i there are a range of possible orientations $\theta_{i,j,min} < \theta_i < \theta_{i,j,max}$ that reflect the radar signal into the occluded region. To find these values, we calculate the angle of reflection on each corner of the reflector to find the minimum and maximum reflection angles. Let (x_c, y_c) be the location of the center of the reflector, and w be the reflector width. Then the reflection location associated with the maximum orientation angle θ_{max} will be located at $(x_c + \frac{1}{2}w \cos \theta_{max}, y_{B,c} + \frac{1}{2}w \sin \theta_{max})$, whereas the reflection location associated with the minimum angle θ_{min} will be located at $(x_c - \frac{1}{2}w \cos \theta_{max}, y_{B,c} - \frac{1}{2}w \sin \theta_{max})$, where w is the width of the reflector. We can then numerically determine the minimum and maximum orientation angles for each reflector.

Due to the tradeoff between field of view and RCS, an ideal NLoS automotive radar-friendly environment consists of a combination of cylinders and planar surfaces. By placing a combination of carefully-oriented road signs and traffic light or lamp posts on all four corners of an intersection, we can enable NLoS sensing at an intersection *without using any purpose-built reflectors*.

5.1 Demonstration: T Junction Coverage

We demonstrate the usefulness of this algorithm at the T junction where we previously showed a coverage gap in § 2.2.2. We previously found a set of points that were not covered by our NLoS



Figure 10: Intersection used to conduct the highway coverage evaluation, consisting of a two-lane driveway crossing a four-lane divided road.

radar, and now we use the method described above to fill in the missing coverage in the intersection by leveraging the placement of ordinary traffic infrastructure.

First, we demonstrate the placement of a cylindrical reflector. To maximize signal strength, we select the available location with the shortest total range from the expected sensor location to the reflector and then to the occluded target region. In this case, the opposite corner of the intersection is the closest possible location. To ensure the reflector has a sufficiently large RCS, we select a metal cylinder that matches the size of the traffic light pole. We place this cylinder at the far corner of the junction from the sensor, and conduct a new set of measurements of the pedestrian moving in the occluded region, as expected.

Second, we place a metal sign to demonstrate the use of a directional reflector. We again select the location on the opposite side of the driveway from the location of the vehicle. Since the vehicle moves toward this location nearly directly, this reflector placement minimizes the change in field of view as the vehicle approaches the intersection. Using the plane size calculations shown above, we find that the relay reflector should have a width of at least 33cm, given a fixed sensor location. To emulate a road sign, we use a 42cm wide metal sheet, which allows up to $\pm 3^\circ$ error in the orientation of the surface.

6 FIELD TEST OF MOSAIC NLOS RADAR

We have verified the core MOSAIC algorithms through microbenchmark experiments in previous sections. Now we proceed with a field test in representative environment.

6.1 Experimental Configuration

For our experiments throughout this paper, we use a combination of a Texas Instruments 76 GHz cascaded mmWave imaging radar [16] and Velodyne Puck Lite 16-line surround Lidar. We use three radar chips, with a total of 9 transmit antennas with a two-wavelength horizontal spacing and 12 receive antennas with non-uniform spacing consisting of three groups of 4 antennas with half-wavelength horizontal spacing. Our radar configuration has a maximum range of 75 m and a range resolution of 30 cm. It has a theoretical maximum angular resolution of 1.4° , although we find that our radar configuration has a 3° angular resolution in practice. (see § 3.2). Practical automotive radars are generally placed in the grill or bumper of a vehicle, rather than the hood, dash, or roof locations generally used for experimental radar mounts. Therefore, we mount our radar on a portable cart at a height of 50 cm, which is within the ordinary range of automotive bumpers and grills [3].




Figure 11: Garage entrance intersection scenario. The car locations in the diagram represent the sensing locations. Plotted locations show where a passing vehicle was detected. Inset photo shows perspective of departing vehicle.

We use a Velodyne 16-line Puck LITE surround lidar [45] for reflector detection and ground truth measurements. This lidar has an angular resolution of about 0.2° for the azimuth and 2° for elevation, with a 30° vertical field of view. Its range measurements have a typical accuracy of ± 3 cm. We conduct two separate measurements for reflector detection and for ground truth target locations. First, we place the lidar next to the radar to find the reflectors available in the scene. This measurement gives no knowledge of the occluded parts of the scene. Second, we place the lidar in a location with a view of both the occluded target to find the ground truth location of the occluded target. We conduct this measurement synchronously with the radar data to determine the ground truth location of the occluded target at each point in time. In both cases, we determine the location offset between the radar and lidar by finding the location of the radar device in the lidar data.

6.2 Evaluation in the Wild

We verify our estimates of the NLoS radar coverage (as defined in § 2.2) at a set of intersections and corners on active roads. First, we consider a large intersection between a public four-lane divided road and a two-lane road, shown in Fig. 10. This intersection may experience occlusion both from the perspective of the two-lane road, where adjacent landscaping blocks our ability to sense cross traffic, and from the perspective of a vehicle turning left from the four-lane road, where oncoming vehicles may occlude the view of adjacent lanes. Although there are buildings on two corners of this intersection, they are blocked by ground-level objects and are not usable as relay reflectors. The only static reflectors usable at this intersection are various metal poles used for traffic signals. Second, we consider a parking garage entrance which intersects a two-lane road at a right angle, shown in Figure 11. It has walls blocking the sight lines on both the left and right sides, meaning that stopped vehicles cannot see cross traffic or pedestrians crossing the street. In this location, we consider both the scenario of a vehicle turning into the entrance, and the case of a vehicle exiting the location, which must sense approaching vehicles on the cross-road. Third, we consider a sharp corner inside a below-ground parking garage, shown in Figure 12. In this scenario, both the walls of the structure and vehicles moving in the opposite direction may block the view of pedestrians at a nearby crosswalk, along with vehicles maneuvering within the parking area around the corner. This location contains a combination of metal bollards and concrete walls that are usable as




Figure 12: Interior corner scenario. The primary reflectors in this scenario are color-coded. Inset photo shows perspective of vehicle.

	Large Planar	Small Planar	Quasi-Omni
Highway	0%	0%	100%
Entrance	31%	46%	23%
Corner	22%	40%	38%

Table 7: The percentage of occluded targets discovered using each reflector type during evaluation.

reflectors, although neither provides complete coverage of the area alone.

We consider two types of coverage: *detection coverage*, where we can detect a moving object and localize it to the direction of approach (e.g. oncoming traffic, left cross traffic, or right cross traffic); and *localization coverage*, where we can additionally localize the object to within the width of a lane, e.g. an error of less than half the standard lane width of 1.9 m. The former is the minimum requirement for automatic emergency braking, whereas the latter enables a more detailed image of the occluded scene for applications such as autonomous driving.

We collect radar returns from ambient traffic in the locations described above. At the highway intersection, we collect radar data for 5 busses, 7 cars (sedans and hatchbacks), and 3 pedestrians as they move through the intersection. At the garage entrance, we collect radar data from 9 cars, 2 trucks, and 3 pedestrians. At the corner, we collect data involving 10 cars and 5 pedestrians. We measure the coverage in terms of the distance from entering the intersection at which each coverage threshold (detection and localization) begins. The results are shown in Table 8.

We observe that Mosaic can detect a bus with the optimal 1.5 s advance notice for automatic emergency braking [4]. It can also detect the cars with this optimal advance notice if the car is travelling at up to 25 mph (40 kmh). Finally, pedestrians are detected by the time they enter the intersection. This is sufficient because it is usually not possible to plan evasive action around a pedestrian still on the sidewalk. Note, however, that cars are not localized to a particular lane until they are much closer to the intersection. To achieve our localization threshold, we generally need NLoS paths from two quasi-omni relay reflectors, or a single planar reflector. The localization range is much shorter at the highway intersection because there is only a single large cylinder within the radar’s field of view, so we must wait for the vehicle to come into range of a smaller metal pole. Placing metal signs to the intersection using the method in § 5 would extend the localization and detection ranges.

		Detect (m)	Localize (m)
Highway	Car	15–26	4–8
	Bus	32–47	15 – 16
	Ped	0 – 1	-2
Entrance	Car	10 – 18	10–15
	Ped	2–3	2–3
Cross-road	Car	14–27	14 – 23
	Truck	39–42	30
Corner	Car	18–24	1–19
	Ped	6–7	0–7

Table 8: Coverage of Mosaic in practical environments. Ranges are measured from first possible collision location.

At the garage entrance, which consists primarily of planar reflectors, the localization range is primarily limited by the angular resolution of the radar and the accuracy of planar reflector reconstructions. This means that we can generally achieve accurate localization at similar ranges as the initial detection.

Finally, we consider how frequently different types of objects are usable as relay reflectors in each scenario. We categorize each object as a large planar object ($> 1\text{m}$ wide), small planar object ($< 1\text{m}$ wide), or quasi-omni. We find that whereas prior studies considered only large planar reflectors, such as walls, the sides of vehicles, and guard rails, *the vast majority of useful reflections in our tests come from smaller objects*. This holds even in the urban locations with many nearby walls. In fact, our system more than doubles the probability of detecting an occluded vehicle compared to methods that only use large planes. Further, we find that even in scenarios with many small and large planar reflectors, leveraging quasi-omni reflectors increases the probability of detecting an occluded object by 30% or more.

7 DISCUSSION

Sensing Limitations. Mosaic’s ambiguity reduction algorithm assumes only a single vehicle is within range of any given reflector. In practice, since the total range of these reflectors is limited to about 24m, about 4-5 car lengths, we expect this condition to be easily satisfied in light traffic. Even in heavy traffic, simply detecting a moving occluded object may be sufficient to issue a driver warning.

Separating LoS and NLoS paths. Mosaic does not include a method for differentiating LoS paths from NLoS paths. We suggest that since Mosaic uses lidar to detect reflectors, we can prior work in radar-lidar fusion [35] to find paths that are detected by radar but not lidar, since lidar rarely detects NLoS paths.

Sensing from a Moving Vehicle. Our primary experiments were conducted with a stationary radar. Since we use doppler shift to separate targets of interest from static clutter, if Mosaic moves, we must compensate for the nonzero doppler shift of static objects. To do this, we estimate the velocity of the radar from lidar data (or from the speedometer and steering data of a vehicle), then determine the doppler shift of static objects at each azimuth angle. Then, for each detected point, we compare the doppler shift to the static doppler estimate for the detected angle of arrival to remove static paths. We have conducted an experiment to verify this method using a wall as a relay reflector. We move a cart to maximum a range of 10m from the wall at 2.5 m/s, while a pedestrian moves within an

occluded region. We find that the occluded pedestrian can still be detected in 97 % of all radar snapshots while using this method.

Processing Latency. To demonstrate the feasibility of using these techniques in a real-time environment, we evaluate the additional radar processing latency required by our technique. To avoid fully re-implementing of a real-time radar stack, we only profile the steps unique to our system. We find that the time to process the lidar data to find possible reflectors is 220 ± 20 ms, and that the additional time to process to resolve the target location was 2 ms or less. Combined, this may reduce the maximum speed at which we can avoid crashes by up to 13%. However, since both steps are currently unoptimized and written in Matlab, we expect both processes may be rewritten in C/C++ to achieve real-time performance.

Deployment Considerations. How a vehicle responds to the sensing results from MOSAIC (e.g., emergency braking, warnings) is beyond the scope of this paper. However, we suggest using the uncertainty estimate developed in § 4 to avoid engaging collision avoidance in cases with a high probability of a false positive. Our system used a Puck LITE lidar sensor to detect nearby reflectors, but our system will work similarly with any sensor that can build a point cloud with an azimuth resolution similar to the lidar sensor.

8 RELATED WORK

Prior research has attempted to use around-corner radar in automotive applications. One such work [38] searched for planar surfaces of with a width of 1m or more to use as relay reflectors, and then used a machine learning approach to accurately detect and track pedestrians and cyclists behind occlusions. However, we show in § 2.2 that such a planar reflector model gives only very limited coverage of many real intersections. Other work [44] used a matched-filtering approach to localize a pedestrian using up to six relay reflections. Using this larger number of reflections improves coverage, especially on cross streets, but requires a priori knowledge of reflectors in the occluded region, which is unlikely to be available in dynamic scenes. Whereas these prior papers primarily focus on accurately localizing and tracking an occluded object, our work instead focuses on ensuring that we can detect an occluded vehicle in as many scenarios as possible. It expands the NLoS reflector model to make around-corner radar usable at real intersections, rather than improving raw accuracy metrics. Our work may therefore be seen as complementary to existing work that focused on NLoS tracking. Around-corner radars have also been demonstrated in an indoor environment [15]. This work used a very similar model to [38], assuming that there is a single large planar reflector that is used as a relay reflector. Finally, one prior work [41] demonstrated the use of a small, passive planar reflector to improve around-corner radar coverage. However, it did not consider alternate artificial reflector types, such as cylindrical reflectors, and did not provide any algorithm to determine its placement.

A different approach to NLoS radar sensing is *through-wall radar*. Radars using up to cm-wave frequencies can pass through many types of walls, allowing the radar to sense a moving or stationary person behind the wall [22, 31]. Our experiments reveal that the mmWave radar can effectively pass through the windows of some vehicles, allowing for a similar type of sensing. However, many objects in a road environment are made of materials such

as metal and concrete, which severely attenuate radio signals at most frequencies, as we show in § 2.4. For this reason, we find that relay reflections are a more reliable approach to NLoS sensing in automotive scenarios. Prior work used radar multipath signals to improve the localization accuracy in through-wall radar sensing, in a similar manner to our technique in § 4.1. However, these techniques generally require a LoS path to work. In addition, they do not estimate the uncertainty resulting from the reflector geometry, which is crucial in automotive applications, where the system may make safety-relevant decisions.

Other work has used different types of sensors for around-corner sensing, including lidar [6, 7, 24, 30, 46], acoustic sensors [1, 23], thermal IR [25], visible light cameras with active illumination [17, 40], and visible light with passive illumination [28, 49]. Compared with radar, methodologies involving active illumination, such as lidar, generally require impractical illumination levels above eye safety limits. Existing passive illumination methods generally only work in controlled environments, and are therefore not suitable for dynamic road environments. Specifically, visible-light solutions do not work as well in low-light scenarios, and do not provide velocity information as reliably as radar can. Similarly, prior acoustic imaging solutions have primarily been demonstrated in controlled environments.

Finally, there is limited prior work in using passive reflectors to enable NLoS radar sensing [41] and to improve NLoS mmWave communications [2, 19, 32, 47]. The work in [19, 32] considers multiple reflector shapes for communications, modelling the relative advantages of planar and nonplanar reflector types for communications coverage. The work in [2] proposes a convex optimization problem to place metal planar reflectors for communications purposes. Our reflector placement algorithm is uniquely designed for radar path tracking, and embraces non-planar reflectors. Emerging intelligent reflecting surfaces [34] can create anomalous reflection patterns to diversify mmWave signal paths and expand coverage. Extending our algorithm to leverage such unique capabilities is an interesting area for future research.

9 CONCLUSION

Existing NLoS radar systems have significant coverage gaps when applied to sensing scenarios at many real-world intersections. We show that these coverage gaps can be ameliorated by considering a much broader set of possible sensing paths between the radar and target, encompassing paths from nonplanar reflectors and paths with asymmetric departure and arrival paths. This results in coverage improvements of 30% or more in urban areas, as well as NLoS sensing capability in sparse highway scenarios where NLoS sensing was impossible using existing techniques. We then show that adding or moving a limited number of roadside objects, such as poles and signs, can further reduce coverage gaps in environments that do not have sufficient reflectors available.

ACKNOWLEDGMENTS

We appreciate the insightful comments and feedback from the anonymous reviewers and shepherd. The work reported in this paper is supported in part by the NSF under Grants CNS-1901048, CNS-1925767, and CNS-2128588.

REFERENCES

- [1] Inkyu An, Doheon Lee, Jung-woo Choi, Dinesh Manocha, and Sung-eui Yoon. 2019. Diffraction-aware sound localization for a non-line-of-sight source. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 4061–4067.
- [2] Chethan Kumar Anjinappa, Fatih Erden, and Ismail Güvenc. 2021. Base Station and Passive Reflectors Placement for Urban mmWave Networks. *IEEE Transactions on Vehicular Technology* 70, 4 (2021), 3525–3539. <https://doi.org/10.1109/TVT.2021.3065221>
- [3] David Aylor, Danny L. Ramirez, Matthew Brumelow, and Joseph M. Nolan. 2005. Limitations of Current Bumper Designs and Potential Improvements. *SAE Transactions* 114 (2005), 1591–1598. <http://www.jstor.org/stable/44725186>
- [4] Max Bareiss, John Scanlon, Rini Sherony, and Hampton C. Gabler. 2019. Crash and injury prevention estimates for intersection driver assistance systems in left turn across path/opposite direction crashes in the United States. *Traffic Injury Prevention* 20 (2019). <https://doi.org/10.1080/15389588.2019.1610945>
- [5] Caltrans. 2018. 2018 Standard Planes.
- [6] Piergiorgio Caramazza, Alessandro Boccolini, Daniel Buschek, Matthias Hullin, Catherine F Higham, Robert Henderson, Roderick Murray-Smith, and Daniele Faccio. 2018. Neural network identification of people hidden from view with a single-pixel, single-photon detector. *Scientific reports* 8, 1 (2018), 1–6.
- [7] Susan Chan, Ryan E Warburton, Genevieve Gariepy, Jonathan Leach, and Daniele Faccio. 2017. Non-line-of-sight tracking of people at long range. *Optics express* 25, 9 (2017), 10109–10117.
- [8] Chien-Sheng Chen. 2015. Weighted geometric dilution of precision calculations with matrix multiplication. *Sensors* 15, 1 (2015), 803–817.
- [9] Pin-Heng Chen. 2010. *Design and performance analysis of a portable real-time through-wall noise radar*. The Pennsylvania State University.
- [10] Junil Choi, Vutha Va, Nuria Gonzalez-Prelcic, Robert Daniels, Chandra R. Bhat, and Robert W. Heath. 2016. Millimeter-Wave Vehicular Communication to Support Massive Automotive Sensing. *IEEE Communications Magazine* 54, 12 (2016), 160–167. <https://doi.org/10.1109/MCOM.2016.1600071CM>
- [11] Maryland DOT. 2007. Book of Standards for Highway and Incidental Structures.
- [12] New York City DOT. 2020. Traffic Signal Standard Drawings.
- [13] Daniele Faccio, Andreas Velten, and Gordon Wetzstein. 2020. Non-line-of-sight imaging. *Nature Reviews Physics* 2, 6 (2020), 318–327.
- [14] Gerry Forbes, Teresa Gardner, Hugh W McGee, Raghavan Srinivasan, et al. 2012. *Methods and practices for setting speed limits: An informational report*. Technical Report. United States. Federal Highway Administration. Office of Safety.
- [15] Shisheng Guo, Qingsong Zhao, Guolong Cui, Songlin Li, Lingjiang Kong, and Xiaobo Yang. 2020. Behind Corner Targets Location Using Small Aperture Millimeter Wave Radar in NLOS Urban Environment. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 13 (2020), 460–470. <https://doi.org/10.1109/JSTARS.2020.2963924>
- [16] Texas Instruments. [n.d.]. mmWave cascade imaging radar evaluation module.
- [17] Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature photonics* 8, 10 (2014), 784–790.
- [18] R.E. Kell. 1965. On the derivation of bistatic RCS from monostatic measurements. *Proc. IEEE* 53, 8 (1965), 983–988. <https://doi.org/10.1109/PROC.1965.4077>
- [19] Wahab Khawaja, Ozgur Ozdemir, Yavuz Yapici, Fatih Erden, and Ismail Güvenc. 2020. Coverage Enhancement for NLOS mmWave Links Using Passive Reflectors. *IEEE Open Journal of the Communications Society* 1 (2020), 263–281. <https://doi.org/10.1109/OJCOMS.2020.2969751>
- [20] Chieh-Ping Lai and Ram M Narayanan. 2010. Ultrawideband random noise radar design for through-wall surveillance. *IEEE Trans. Aerospace Electron. Systems* 46, 4 (2010), 1716–1730.
- [21] B. Langen, G. Lober, and W. Herzig. 1994. Reflection and transmission behaviour of building materials at 60 GHz. In *5th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, Wireless Networks - Catching the Mobile Future*, Vol. 2. 505–509 vol.2. <https://doi.org/10.1109/WNCMF.1994.529141>
- [22] Jing Li, Zhaofa Zeng, Jiguang Sun, and Fengshan Liu. 2012. Through-wall detection of human being's movement by UWB radar. *IEEE Geoscience and Remote Sensing Letters* 9, 6 (2012), 1079–1083.
- [23] David B. Lindell, Gordon Wetzstein, and Vladlen Koltun. 2019. Acoustic Non-Line-Of-Sight Imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [24] David B. Lindell, Gordon Wetzstein, and Matthew O'Toole. 2019. Wave-Based Non-Line-of-Sight Imaging Using Fast f-k Migration. *ACM Trans. Graph.* 38, 4, Article 116 (July 2019), 13 pages. <https://doi.org/10.1145/3306346.3322937>
- [25] Tomohiro Maeda, Yiqin Wang, Ramesh Raskar, and Achuta Kadambi. 2019. Thermal Non-Line-of-Sight Imaging. In *2019 IEEE International Conference on Computational Photography (ICCP)*. 1–11. <https://doi.org/10.1109/ICCPHOT.2019.8747343>
- [26] Isamu Matsunami, Ryohei Nakamura, and Akihiro Kajiura. 2012. RCS measurements for vehicles and pedestrian at 26 and 79GHz. In *2012 6th International Conference on Signal Processing and Communication Systems*. 1–4. <https://doi.org/10.1109/ICSPCS.2012.6508004>
- [27] W. Najm, J. Smith, and M. Yanagisawa. 2007. Pre-Crash Scenario Typology for Crash Avoidance Research.
- [28] Felix Naser, Igor Gilitschenski, Alexander Amini, Christina Liao, Guy Rosman, Sertac Karaman, and Daniela Rus. 2019. Infrastructure-free NLoS Obstacle Detection for Autonomous Cars. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 250–257. <https://doi.org/10.1109/IROS40897.2019.8967554>
- [29] James Nobles and Roger Brooks. 1997. *Lane and Shoulder Width Standards*. State of Minnesota.
- [30] Matthew O'Toole, David B Lindell, and Gordon Wetzstein. 2018. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature* 555, 7696 (2018), 338–341.
- [31] John Jr E Peabody, Gregory L Charvat, Justin Goodwin, and Martin Tobias. 2012. *Through-wall imaging radar*. Technical Report. Massachusetts Institute of Technology-Lincoln Laboratory Lexington United States.
- [32] Zhangyou Peng, Linxiao Li, Miao Wang, Zhonghao Zhang, Qi Liu, Yang Liu, and Ruoran Liu. 2016. An Effective Coverage Scheme With Passive-Reflectors for Urban Millimeter-Wave Communication. *IEEE Antennas and Wireless Propagation Letters* 15 (2016), 398–401. <https://doi.org/10.1109/LAWP.2015.2447734>
- [33] Jay A. Puckett, Robert G. Erikson, and John P. Peiffer. 2010. Fatigue Testing of Stiffened Traffic Signal Structures. *Journal of Structural Engineering* 136, 10 (2010), 1205–1214. [https://doi.org/10.1061/\(ASCE\)ST.1943-541X.0000229](https://doi.org/10.1061/(ASCE)ST.1943-541X.0000229)
- [34] Kun Qian and Xinyu Zhang. 2022. MilliMirror: 3D Printed Reflecting Surface for Millimeter-Wave Coverage Expansion. In *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*.
- [35] Kun Qian, Shilin Zhu, Xinyu Zhang, and Li Erran Li. 2021. Robust Multimodal Vehicle Detection in Foggy Weather Using Complementary Lidar and Radar Signals. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [36] M. A. (Mark A.) Richards. 2005. *Fundamentals of radar signal processing*. McGraw-Hill, New York.
- [37] John M. Scanlon, Rini Sherony, and Hampton C. Gabler. 2016. Preliminary potential crash prevention estimates for an Intersection Advanced Driver Assistance System in straight crossing path crashes. In *2016 IEEE Intelligent Vehicles Symposium (IV)*. 1135–1140. <https://doi.org/10.1109/IVS.2016.7535532>
- [38] Nicolas Scheiner, Florian Kraus, Fangyin Wei, Buu Phan, Fahim Mannan, Nils Appenrodt, Werner Ritter, Jürgen Dickmann, Klaus Dietmayer, Bernhard Sick, et al. 2020. Seeing Around Street Corners: Non-Line-of-Sight Detection and Tracking In-the-Wild Using Doppler Radar. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [39] Tom Schipper, Joaquim Fortuny-Guasch, Dario Tarchi, Lars Reichardt, and Thomas Zwick. 2011. RCS measurement results for automotive related objects at 23–27 GHz. In *Proceedings of the 5th European Conference on Antennas and Propagation (EUCAP)*. IEEE, 683–686.
- [40] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. 2005. Dual Photography. In *ACM SIGGRAPH 2005 Papers* (Los Angeles, California). Association for Computing Machinery, New York, NY, USA, 745–755. <https://doi.org/10.1145/1186822.1073257>
- [41] Dmitrii Solomitskii, Carlos Baquero Barneto, Matias Turunen, Markus Allén, Yevgeni Koucheryavy, and Mikko Valkama. 2020. Millimeter-Wave Automotive Radar Scheme With Passive Reflector for Blind Corner Conditions. In *2020 14th European Conference on Antennas and Propagation (EuCAP)*. 1–5. <https://doi.org/10.23919/EuCAP48036.2020.9135926>
- [42] A Sume, M Gustafsson, A Jänis, S Nilsson, J Rahm, and Å Örbom. 2009. Radar detection of moving objects around corners. In *Radar Sensor Technology XIII*, Vol. 7308. International Society for Optics and Photonics, 73080V.
- [43] ER Swanson. 1978. Geometric dilution of precision. *NAVIGATION, Journal of the Institute of Navigation* 25, 4 (1978), 425–429.
- [44] Khac-Phuc-Hung Thai, Olivier Rabaste, Jonathan Bosse, Dominique Poullin, Israel D. Hinojosa Sáenz, Thierry Letertre, and Thierry Chonavel. 2019. Detection-Localization Algorithms in the Around-the-Corner Radar Problem. *IEEE Trans. Aerospace Electron. Systems* 55, 6 (2019), 2658–2673. <https://doi.org/10.1109/TAES.2019.2897031>
- [45] Velodyne. 2016. Puck LITE Surround LiDAR.
- [46] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mouni G Bawendi, and Ramesh Raskar. 2012. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature communications* 3, 1 (2012), 1–8.
- [47] Teng Wei, Anfu Zhou, and Xinyu Zhang. 2017. Facilitating Robust 60 GHz Network Deployment by Sensing Ambient Reflectors. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation (NSDI)*.
- [48] Xiaolin Wu. 1991. An Efficient Antialiasing Technique. *SIGGRAPH Comput. Graph.* 25, 4 (July 1991), 143–152.
- [49] Adam B Yedidia, Manel Baradad, Christos Thrampoulidis, William T Freeman, and Gregory W Wornell. 2019. Using unknown occluders to recover hidden scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12231–12239.