Practical Reflection Manipulation on mmWave-Based Physical Intrusion Detection

Aliu Akinwale *, Wentao Gao *, Jiawei Li†, Ang Li ‡, Xiaojun Shang*, Yanchao Zhang†, Dianqi Han*
* University of Texas at Arlington, † Arizona State University, ‡ University of Michigan-Dearborn
aliu.akinwale@uta.edu, wentao.gao@uta.edu, jwli@asu.edu, angli@umich.edu, xiaojun.shang@uta.edu, yczhang@asu.edu, dianqi.han@uta.edu

***Abstract*—mmWave sensing has emerged as a promising technique for physical intrusion detection, offering reliable performance across diverse environments and benefiting from the growing availability of low-cost commodity mmWave devices. However, its integration also raises critical security concerns, particularly on vulnerabilities to signal manipulation attacks. While previous studies have demonstrated the feasibility of such attacks using advanced software-defined radios, practical and low-cost signal manipulation attacks remain largely unexplored, highlighting a significant research gap in understanding the security of mmWave-based physical intrusion detection systems.**

In this study, we identify and experimentally evaluate a novel reflection manipulation attack that costs less than 100 USD to implement. Specifically, an array of low-cost passive corner reflectors is mounted on the intruder. The reflectors are uniformly arranged in a circular formation and rotated by a motor. This simple yet effective design manipulates the CFAR detection threshold, thereby concealing both the intruder and the reflectors. Experimental results confirm that the attack is highly effective under certain system configurations but fails under others. Motivated by this observation, we further propose several defenses to effectively mitigate the threat posed by the attack.

***Index Terms*—mmWave sensing, wireless security.**

I. Introduction

Physical intrusion detection systems designed to detect unauthorized entry into constrained physical spaces have been extensively explored in mission-critical contexts. For example, detection systems deployed in border regions could effectively detect and intercept illegal crossings [1]. Critical infrastructures rely on these systems to prevent unauthorized access [2]. Additionally, battlefield deployments enhance situational awareness by identifying potential threats [3]. Reliable and robust detection is critical in those applications.

Millimeter-wave (mmWave) sensing has emerged as a promising technology for physical intrusion detection [4], [5]. Unlike conventional camera-based approaches that often suffer degraded performance in low-visibility conditions (e.g., fog, heavy rain, or nighttime) [6], [7], mmWave radars utilize high-frequency radio signals to achieve reliable detection across diverse environmental conditions [8], [9]. The increased availability of cost-effective commercial mmWave radar devices has further facilitated practical deployment of mmWave-based systems [10]. Recent demonstrations confirm the effectiveness and potential of mmWave-based physical intrusion detection for real-world applications [11]–[13].

Despite these advantages, the integration of mmWave sensing introduces critical security concerns, particularly regarding the vulnerability to signal manipulation attacks. Prior research has demonstrated that attackers may forge reflective signals to mmWave radar in order to create phantom objects or conceal real ones in detection results [14], [15]. However, those attacks typically require expensive software-defined radios (SDRs) costing thousands of U.S. dollars for achieving nanosecond-level synchronization with the targeted radar and precise signal control [16], rendering them less practical. Conversely, low-cost signal manipulation attacks remain largely unexplored, leaving a significant research gap in fully understanding the security of mmWave-based physical intrusion detection.

In this work, we address this gap by investigating mmVeil, a novel reflection manipulation attack executable with hardware costing less than 100 U.S. dollars. mmVeil exploits the Constant False Alarm Rate (CFAR) algorithm widely employed in mmWave sensing to compromise the detection system. In particular, the CFAR algorithm computes an adaptive detection threshold of reflection signal based on the background noise to reliably differentiate real object reflections from ambient noises, maintaining a reliable false alarm rate regardless of noise conditions [17]. The mmVeil attack utilizes low-cost reflectors to generate signals interpreted by CFAR as environmental noise, thereby artificially elevating detection thresholds and effectively concealing intruders from radar detection.

Although the idea is straightforward, the successful implementation faces critical challenges. Specifically, reflectors naturally produce strong reflection signals, likely exposing themselves to detection. To address this challenge, mmVeil exploits a set of reflectors with a carefully designed layout so that the CFAR detection threshold on each reflector is increased due to interference from other reflectors. This design could effectively hide all the reflectors from detection. For practical implementation, we propose deploying those reflectors in a rotating circular formation driven by a motor. With properly configured parameters,

979-8-3315-3856-9/25/$31.00 ©2025 IEEE

this design ensures distinct ranges and velocities across different reflectors that are measured by the radar, which is essential for successful attacks. We further develop a lightweight algorithm to search for effective settings of the reflector array.

We implemented the mmVeil attack using a reflector array and experimentally evaluated its concealment performance against a commodity mmWave radar. The results show that the attack's effectiveness is highly dependent on the radar configuration. Under vulnerable configurations, the mmVeil attack can effectively conceal both human and robot intruders from detection in indoor and outdoor environments. Moreover, the attack exhibits robustness to minor variations in the intruder's speed under those conditions. These findings highlight the need for mmWave-based physical intrusion detection systems to avoid the identified vulnerable configurations. To further mitigate the threat posed by mmVeil, we also propose several effective defense strategies.

Our contributions are threefold. First, we identify a low-cost yet effective concealment attack against mmWave-based physical intrusion detection systems that has been overlooked in existing studies. Second, we implement the attack and conduct a comprehensive evaluation of its performance through real-world experiments. Third, based on our experimental findings, we propose effective defenses against this attack.

The rest of this paper is organized as follows. Section II introduces mmWave-based physical intrusion detection. Section III presents the system and adversary models. Section IV explains the rationale behind the attack and demonstrates the feasibility study. Section V details the attack design. Section VI evaluates the attack through comprehensive experiments and proposes effective defenses against the attack based on the experimental findings. Section VII reviews related work.

II. mmWave Sensing for Physical Intrusion Detection

mmWave-based physical intrusion detection typically explores Frequency Modulated Continuous Wave (FMCW) radars, which offer strong reliability and are increasingly available as commercial off-the-shelf devices. The detection process generally proceeds through the following key steps.

a) IF signal acquisition. The radar periodically emits a linear chirp signal with a sweeping bandwidth of B and a duration of T, which is formulated as $s_{tx}(t) = A\exp[j(2\pi f_c t + \pi B t^2/T)]$. f_c denotes the carrier frequency. The reflection signal from an object at range R to the radar is received by the radar as $s_{rx}(t) = A\exp[j(2\pi f_c(t-\tau) + \pi B(t-\tau)^2)/T]$, which includes a delay $\tau = 2R/c$. Here, c denote the speed of radio signals. s_{tx} and s_{rx} are mixed to acquire the Intermediate Frequency (IF) signal as $s(t) = s_{tx}(t)s^*_{rx}(t)$.

b) Range measurement. Object ranges are derived from the frequency components of the IF signal. In particular, the reflection from an object at range R results in a component with frequency $(2B/cT)R$. The Fast Fourier transform (FFT), known as the range-FFT, is applied to the IF signal to resolve frequency components associated with different objects and their corresponding ranges. The frequency resolution of the range-FFT determines the range granularity, which is given by $\Delta R = c/2B$. B denotes the sweeping bandwidth.

c) Velocity measurement. Velocities are measured by analyzing the phase changes of IF components across chirps. In particular, the varying propagation delay of reflection signals from a moving object induces a phase change in the corresponding IF components. This phase change is proportional to the object's radial velocity relative to the radar. To capture this, the radar transmits a series of chirps within a sensing frame and inserts idle periods between frames to mitigate inter-frame interference. A second FFT, known as the Doppler-FFT, is applied to IF components across different chirps to extract Doppler frequencies for velocity estimation. The granularity of velocity measurement determined by the frequency resolution of the Doppler-FFT is given by $\Delta v = \lambda/2T$, where λ denotes the wavelength of the mmWave sensing signal.

d) CFAR-based intrusion detection. Constant False Alarm Rate (CFAR) algorithms are widely employed in mmWave-based object detection to ensure robustness against noise. The range and velocity measurements from mmWave sensing are typically represented as a range-Doppler heatmap, which visualizes the distribution of reflected signal strength across different range and velocity values. Fig. 1b shows an example of the heatmap. Specifically, the range and velocity axes are divided into intervals according to their respective resolutions, forming a grid of cells. The intensity of each cell indicates the strength of reflected signals within the corresponding range and velocity intervals.

To prevent noisy cells from being mistakenly identified as objects, CFAR algorithms compute an adaptive detection threshold for each cell based on the intensity of its surrounding cells. Cells located immediately adjacent to the target cell are excluded from this calculation to mitigate the influence of signal spillover from the same object. These cells are referred to as guard cells. The cells used for threshold estimation are called reference cells. The selection of reference cells is defined by two parameters, N_{gd} and N_{ref}, representing the number of guard and reference cells along each axis, respectively. Figure 1a illustrates an example with $N_{gd} = 2$ and $N_{ref} = 2$. Black and red lines are used to indicate the boundaries of guard and reference cells, respectively. The detection threshold is calculated by aggregating the intensities of the reference cells, for which various approaches have been proposed [18], [19]. CFAR algorithms enable the mmWave-based detection to adapt to varying noise levels and maintain stable detection performance.

Further steps may be taken before making a final decision in physical intrusion detection. For example, the detection system may utilize the antenna array commonly integrated into mmWave radars to extract the direction of arrival and 3D point cloud data of objects. These additional spatial features can then be incorporated into the final decision-making process.

III. System and Adversary Model

A. System Model

We consider a mmWave-based physical intrusion detection system that follows the procedure outlined in Sec. II to identify potential intruders. The system comprises a front-end mmWave radar for data acquisition and a back-end server for processing. Systems that integrate mmWave radars with other sensing modalities are beyond the scope of this study. The system employs the widely used Cell-Averaging (CA)-CFAR algorithm for object detection. The effectiveness of the mmVeil attack on alternative CFAR variants is discussed in Sec. VI. The CA-CFAR detection threshold $\mathcal{T}_{CA}$ is computed as

$$\mathcal{T}_{CA} = N_c(P_{fa}^{-1/N_c} - 1)I_{avg}, \tag{1}$$

where N_c and I_{avg} denote the number of reference cells and the average intensity of reference cells, respectively. P_{fa} denotes the desired false alarm rate of detection, i.e., the possibility that any empty cells are detected as including objects by mistake.

B. Adversary Model

The mmVeil attack is designed to conceal an intruder from the detection system. The intruder may be a malicious human, a trained animal, or a robot remotely controlled by the attacker. Following adversary models established in prior literature [14], [15], we make the following assumptions about the attacker. The attacker cannot compromise the mmWave radar or the backend server to manipulate the detection process. However, the attacker is able to obtain the radar's range and velocity resolutions by, for example, deploying a sniffing device to extract the chirp bandwidth and duration. The attacker also knows the radar's physical location and the CFAR algorithm parameters, including N_{gd}, N_{ref}, and P_{fa}. Unlike existing studies that assume strong attackers capable of generating precisely synchronized spoofing signals, we consider a more practical adversary incapable of precise signal crafting.

IV. Design Rationale and Feasibility Study

A. Design Rationale

The intruder concealment is considered effective if the detection system reports no object in the range-Doppler cells associated with the intruder. In practice, mmWave-based detection systems often adopt lower range resolutions to maximize coverage area while complying with spectrum regulations. For example, the U.S. Federal Communications Commission (FCC) permits radar operation over a wide band (57–71 GHz) with a maximum transmit power of +10 dBm, or over a narrower band (61–61.5 GHz) with a higher transmit power of up to +40 dBm [20]. To benefit from the increased power and thereby extended sensing range, physical intrusion detection systems typically adopt the narrowband 61–61.5 GHz configuration, which results in a range resolution around 0.3 meters as demonstrated in Sec. II. As a result, an intruder may occupy only a few range-Doppler cells. For the simplicity of demonstration, we assume that the intruder occupies a single range-Doppler cell. The performance on human intruders that occupy multiple cells is evaluated in Sec. VI.

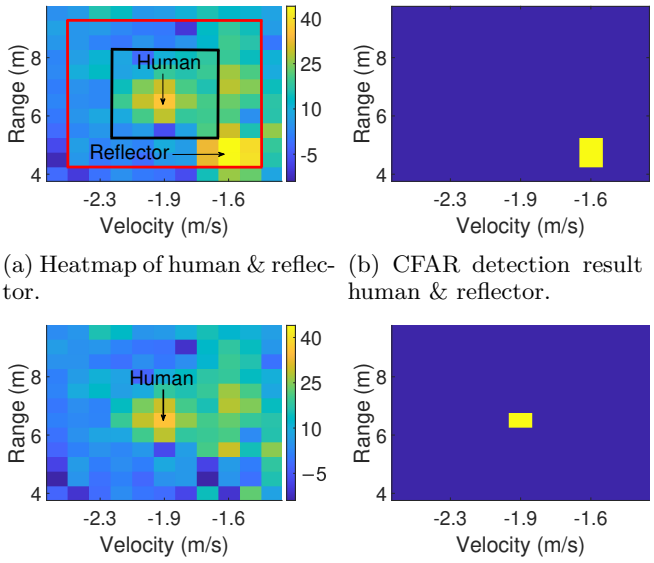

(a) Heatmap of human & reflector.

(b) CFAR detection result on human & reflector.

(c) Heatmap of human only.

(d) CFAR detection on human.

Fig. 1: Reflector's impact on the Range-Doppler heatmap and CFAR detection results.

To conceal the intruder, the attacker manipulates the CFAR detection threshold in the corresponding range-Doppler cell. Specifically, the CFAR algorithm sets the detection threshold based on the aggregated intensities of surrounding reference cells. By inducing strong reflections in these reference cells, the attacker raises the threshold above the signal intensity of the intruder's cell, thereby preventing its detection. This effect is achieved using highly reflective materials, such as smooth metal surfaces. Since reference cells cover a local region in the range-velocity domain, the reflector must exhibit only slight differences in range and velocity relative to the intruder. To satisfy this constraint, we propose physically attaching the reflector to the intruder.

B. Feasibility study.

We demonstrate the feasibility of the mmVeil attack through an experiment. Specifically, a TI IWR6843 mmWave radar was deployed for object detection. A human

subject walked in front of the radar while holding a rod with a 20 cm×20 cm metal surface attached to the other end. The metal surface was positioned between the human and the radar but offset from their line-of-sight path to avoid blocking direct reflections from the human. To ensure effective retroreflection, the metal surface was oriented to directly face the radar. The radar was configured with a range resolution of 0.5 meters and a velocity resolution of 0.12 m/s, and the reflector's position was carefully selected to influence the reference cells of the human.

Figure 1a shows the range-Doppler heatmap from a sensing frame, highlighting the cells corresponding to the human and the reflector. The boundaries of the guard and reference cells are marked based on $N_{gd} = 2$ and $N_{ref} = 2$. We further process the heatmap using the CA-CFAR algorithm with a false alarm rate of 0.00001, a commonly adopted setting in mmWave-based object detection. The resulting detection output is shown in Fig. 1b. Due to the reflector's influence on the reference cells, the human is effectively concealed from the radar.

To verify the cause of this concealment, we manually eliminate the reflector's influence by replacing the intensity values of the reflector cells with those of adjacent, object-free cells. The modified heatmap and corresponding detection results are shown in Figs. 1c and 1d. The successful detection of the human in the absence of the reflector confirms the effectiveness of the mmVeil attack.

C. Challenges and Solutions.

Although experimental results confirm the feasibility of concealing an intruder from detection, realizing the mmVeil attack in practice introduces several technical challenges.

1) Reflector concealment: A key challenge in executing the mmVeil attack lies in hiding the reflectors themselves. Specifically, the moving reflectors introduced to manipulate the reference cells of the intruder could also be detected as moving objects, thereby exposing the attack. To address this issue, mmVeil employs an array of reflectors that not only cooperatively conceal the intruder but also hide one another by leveraging the same CFAR manipulation strategy. In particular, if a sufficient number of reflectors contribute to the reference cells of a given reflector, the resulting CFAR threshold exceeds the reflection strength of that reflector, thereby preventing its detection.

Achieving mutual concealment requires that the reflectors be well separated in the range-Doppler domain. This translates to that reflectors must exhibit sufficiently distinct ranges and velocities. To satisfy this requirement, we design a practical and effective mechanism: reflectors are uniformly arranged in a circular formation, which is rotated by a motor. In addition, we develop a lightweight algorithm to determine the optimal rotation speed, array radius, and number of reflectors needed to achieve effective concealment. Implementation details are provided in Sec. V.

2) Consistent reflection from rotating reflectors: The second challenge lies in maintaining consistent reflection strength despite the continuous rotation of the reflector array. Due to the narrow aperture of the radar, a smooth metal surface can reflect sensing signals back to the radar only when the incident angle is close to zero, i.e., when the surface is nearly perpendicular to the incoming wavefront. However, the surface cannot maintain this orientation throughout the rotation, resulting in fluctuating reflection strength and reduced concealment performance.

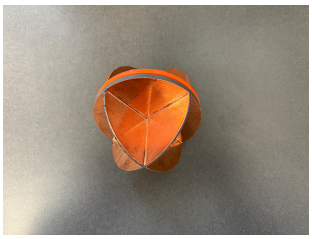

Fig. 2: Reflection module.

To address this challenge, the mmVeil attack employs corner reflectors. A corner reflector typically consists of three mutually perpendicular metal surfaces forming a cubic corner. This structure reflects incoming signals back toward the source regardless of the incident angle so long as the signal enters the interior of the corner. To ensure omnidirectional coverage during rotation, we construct a reflection module shown in Fig. 2. The module composed of eight corner reflectors with complementary orientations, enabling consistent retroreflection across all angles. Additionally, we select rounded-edge designs to stabilize the Radar-Cross Section (RCS) under rotation, ensuring minimal variation in reflected signal strength. For simplicity, we refer to this module as a reflector throughout the rest of the paper.

V. Attack Design

A. Overview

The attacker attaches an array of reflectors to the intruder to conceal it from detection. Fig. 3 illustrates the implementation on an intruder robot. Specifically, all the reflectors are connected to a motor and arranged uniformly in a circle. The motor is mounted on the robot and drives the rotation of the reflector array. The reflector array is characterized by three key parameters: the number of reflectors N_R, the circle radius r, and the rotation speed ω. Additionally, we denote the central angle between two adjacent reflectors as θ, which is given by $\theta = 2\pi/N_R$ due to the uniform placement.

The success of the mmVeil attack relies on configuring these parameters to satisfy two conditions: (1) a sufficient number of reflectors induce reflections in the reference cells of the intruder to conceal it, and (2) each individual reflector must also be concealed by having enough other

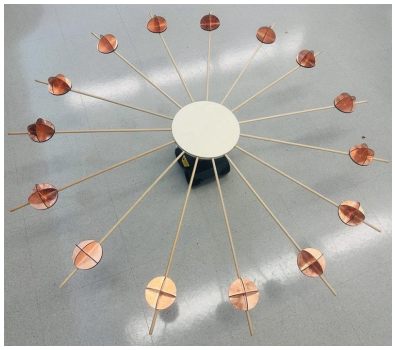

Fig. 3: Attack demonstration on a robot intruder.

reflectors contribute reflections to its own reference cells. To identify such configurations, we model the radar measurements and perform a grid search over parameter settings to guide the selection of effective configurations.

As discussed later, the effectiveness of a particular configuration of N_R, r, and ω also depends on the intruder's location, making it challenging to maintain concealment over an extended trajectory. To address this, we propose dynamically adjusting the motor's rotation speed, allowing the system to adapt to varying intruder positions while keeping N_R and r fixed.

B. Radar Measurement Modeling

To facilitate accurate spatial representation, we construct a 3D Cartesian coordinate system with the origin located at the radar's position. The x-axis is defined to align with the radar's boresight direction. The z-axis is oriented vertically upward, perpendicular to the ground plane. The y-axis is determined accordingly to complete a right-handed coordinate system, pointing to the left when facing along the x-axis. For ease of presenting the attack design, we make the following assumptions, which do not affect the general effectiveness of the proposed attack. We assume the intruder is initially located at $(x_0, 0, z_0)$ and moves toward the radar along the x-axis with a constant speed v. The array of reflectors rotates within the x–y plane.

Accordingly, we formulate the radar measurements on the intruder and reflectors. At time t, the radar measures the range of the intruder, denoted as $R_I(t)$, and the range of the i-th reflector, denoted as $R_i(t)$, as

$$R_I(t) = \sqrt{(x_0 - vt)^2 + z_0^2} \tag{2}$$

and

$$R_i(t) = \sqrt{(x_0 - vt - r\sin\theta_i(t))^2 + (r\cos\theta_i(t))^2 + z_0^2} \tag{3}$$

respectively. Here, $\theta_i(t)$ denotes the angular displacement of the ith reflector at time t, which is defined as the angle between the vector pointing from the center of the reflector array to the i-th reflector and the y-axis direction. $\theta_i(t) = \theta' + \omega t$, where θ' denotes the initial angular displacement of the reflector at time $t = 0$.

By taking the derivative, we derive the radial velocity of the intruder, denoted as $V_I(t)$, and that of the i-th reflector, denoted as $V_i(t)$, as

$$V_I(t) = \frac{v(x_0 - vt)}{\sqrt{(x_0 - vt)^2 + z_0^2}} \tag{4}$$

and

$$V_i(t) = \frac{(x_0 - vt - r\sin\theta_i(t)) \cdot (v + r\omega\cos\theta_i(t))}{R_i} + \frac{r^2\omega\cos\theta_i(t)\sin\theta_i(t)}{R_i} \tag{5}$$

C. Grid Search for Effective Parameter Settings

An effective configuration of N_R, r, and ω for successful concealment must fulfill two requirements. First, there are at least N_{ic} reflectors inducing reflections within the reference cells of the intruder to elevate the detection threshold and thereby conceal the intruder. The value of N_{ic} depends on the relative reflection strengths of the intruder and the reflectors. Specifically, N_{ic} is given by

$$N_{ic} = \arg\min_{N\in\mathbb{Z}^+} \left\{ N \;\middle|\; NS_R(P_{fa}^{-1/N_c} - 1) \geq S_I \right\}, \tag{6}$$

where S_R and S_I denote the reflection strength of the reflector and intruder, respectively. Second, there are at least N_{rc} reflectors inducing reflection within the reference cells of each individual reflector, ensuring that all the reflectors are themselves concealed from detection. Since all reflectors are of identical size and generate similar reflection strength, N_{rc} depends only on the CFAR algorithm parameters, and can be computed as

$$N_{rc} = \arg\min_{N\in\mathbb{Z}^+} \left\{ N \;\middle|\; N(P_{fa}^{-1/N_c} - 1) \geq 1 \right\}. \tag{7}$$

For example, assuming the CA-CFAR settings of $N_{gd} = 2$, $N_{ref} = 2$, and $P_{fa} = 0.00001$, a reflector can be effectively concealed with at least five other reflectors affecting its reference cells.

Guided by the radar measurement model, we perform a grid search to identify effective configurations of N_R, r, and ω. Specifically, we evaluate a discrete set of candidate values for each parameter, selected within a specified range and spaced at equal intervals. For each combination of the assessed parameter values, we estimate the range and velocity measurements of the intruder and reflectors based on the models presented in Sec. V-B, and then evaluate the corresponding concealment performance. Owing to the symmetric layout of the reflector array, we assess reflector concealment by focusing on a single target reflector. If this target reflector can be successfully concealed, the remaining reflectors will exhibit equivalent concealment performance.

Search Space. The selection of assessed values is guided by mmWave sensing principles and constrained by hardware limitations. The range and step size for the rotation speed ω are determined by the capabilities of the motor. For example, the electric gear reducer motor used in our

experiment supports rotation speeds ranging from 1 RPM to 90 RPM with a configurable step size of 1 RPM. This limitation defines the search space on ω. The search for the number of reflectors N_R begins at $N_{rc}+1$ and extends up to 30, which we found sufficient based on experimental observations. For effective reflector concealment, the array radius r must be large enough to separate reflectors in the range domain by more than $2 \cdot \Delta R$, ensuring that reflected signals from other reflectors fall outside the guard cells. Accordingly, the search range for r starts from ΔR and extends up to 2 meters, beyond which the array becomes too large to be stably mounted on the intruder. The step size for the radius search is set to $0.05 \cdot \Delta R$, as smaller variations typically have negligible impact on concealment performance. This step size provides a balance between resolution and computational efficiency, ensuring that the search does not miss any effective radius values.

Configuration Assessment. We further evaluate each configuration within the search space. According to the range and velocity measurement models introduced in Sec. V-B, the angular displacement of a reflector has a critical impact on its range and velocity measured by the radar, therefore affecting its concealment effectiveness. To reliably assess concealment performance, we evaluate 360 angular displacements of the target reflector, spanning from 1 to 360 degrees. For each angular displacement, we compute the ranges and velocities of the intruder and all the reflectors using eqs. (2) to (5) and determine their corresponding range-Doppler cells based on the system's range and velocity resolutions. We then estimate the resulting range-Doppler heatmap using the experimentally measured reflection strengths of the intruder and the reflectors. The CA-CFAR algorithm is subsequently applied to determine the detection result. A configuration is considered effective if both the intruder and the target reflector are consistently concealed across all 360 evaluated angular displacements.

D. Concealment over Extended Trajectory

As indicated by eqs. (2) to (5), the range of the intruder robot also affects the radar measurements and the effectiveness of concealment. A configuration that successfully conceals the intruder and target reflector at one range may fail at another. To address this limitation, we leverage the configurable rotation speed of the motor. Given a planned movement trajectory of the robot, we divide it into short segments, each of length ΔR. For the starting point of each segment, we acquire all the effective configurations through the grid search. We then identify values of N_R and r that could achieve effective concealment across all the segments when combined with different ω values. The reflector array is constructed using one such N_R and r configuration. During operation, the attacker adapts ω in real time based on the current location of the intruder robot, ensuring consistent concealment throughout its movement.

VI. Evaluation

A. Implementation

We implemented a mmWave-based physical intrusion detection system using the TI IWR6843 radar [21]. The radar board was connected to a TI DCA1000EVM data capture board [22], which was further connected to a Dell Latitude desktop via an Ethernet cable to stream the captured IF data for processing. On the desktop side, we implemented the standard data processing procedure described in Sec. II using MATLAB.

We further implemented the reflector array for mmVeil. Specifically, the reflector frames were fabricated using an Ultimaker 3D printer and coated with metal foil to create smooth high-reflectivity surfaces. Alternatively, the frames can be constructed from rigid paper or plastic boards as a low-cost substitute. The reflectors were mounted on wooden rods, which were connected to an electric gear-reducer motor. The motor was powered by an external power supply and controlled the rotation of the reflector array.

B. Evaluation Metrics

Physical intrusion detection systems may adopt different algorithms to identify potential intrusion, which may result in varying outcomes on the same mmWave sensing data. To support a more general and algorithm-independent evaluation, we propose a generic metric for the assessment of concealment performance. Specifically, we process the IF data from all mmWave sensing frames and apply the CFAR detection algorithm. The attack is considered successful on a frame if no moving object is detected in any range-Doppler cell. Otherwise, the attack is considered to have failed for that frame. We define the Ratio of Frames with Successful Concealment (RFSC) as the proportion of successful frames over the total number of frames and employ it as the metric for performance evaluation.

C. Performance Evaluation

For a comprehensive evaluation, we experimentally assessed the effectiveness of the mmVeil attack under a variety of settings. To avoid repetitive descriptions, this section first presents the evaluation under a specific configuration, which serves as the default experimental setting. When analyzing the impact of an individual factor, we vary that factor while keeping all other parameters fixed at their default values.

Basic experiment. We conducted the basic experiment in an outdoor environment, which is an empty parking lot. The radar was configured to reflect the narrowband settings commonly used in physical intrusion detection systems. In particular, the radar parameters were set as follows: a sweep bandwidth of 500 MHz ranging from 61.0 GHz to 61.5 GHz, a chirp duration of 25.60 μs, and 128 chirps per frame. Additionally, the radar transmitted 40 frames per second by adjusting the inter-frame interval accordingly.

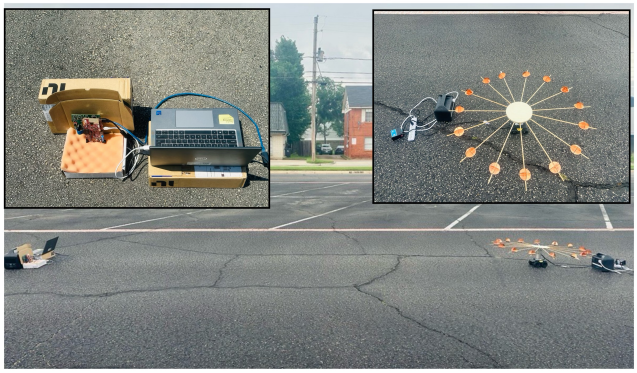

Fig. 4: The default setup of experiments.

Under this configuration, the range and velocity resolutions were 0.3 meters and 0.12 m/s, respectively. We employed the CA-CFAR algorithm for object detection, with $N_{ref} = 2$, $N_{gd} = 2$, and $P_{fa} = 0.00001$.

We further employed a mobile robot to emulate the intruder and mounted reflectors on it to implement the attack. The robot was programmed to move toward the radar at approximately 0.3 m/s, starting from a distance of 15 meters and stopping at 10 meters. Based on the grid search results, we constructed an array of 15 reflectors arranged in a circular formation with a radius of 0.73 meters. During the robot's movement, the rotation speed of the motor was reconfigured every 0.3 meters according to the optimal settings identified through the grid search.

We processed all the sensing frames captured during the robot movement using the standard data processing pipeline described in Sec. II, and identified successful frames in which no moving object was detected. This experiment was repeated 10 times, resulting in an average RFSC of 0.442. Fig. 5 illustrates the range-Doppler heatmap of a successful frame and the corresponding CFAR result. The rotating reflector array introduced significant noise in the reference cells around the robot, effectively concealing it from detection. For comparison, we removed the reflector array from the robot and repeated the experiment. In this case, the robot was detected in all frames, further validating the effectiveness of the mmVeil attack.

Performance on Human Intruders. We conducted experiments to evaluate the concealment performance on human intruders. As illustrated in Fig.6, a human subject held the reflector array to emulate an intruder. The subject walked toward the radar at an approximate speed of 0.3 m/s, starting from 15 meters away and stopping at 13 meters. The average RFSC across 10 experiment runs was 0.148.

While the attack remained effective, the concealment performance on human intruders showed noticeable degradation compared to that on robot intruders. This degradation can be attributed to two primary reasons. First, the human body is larger than the robot, resulting in stronger reflections. Second, human movement is more complex, with different body parts exhibiting varying velocities.

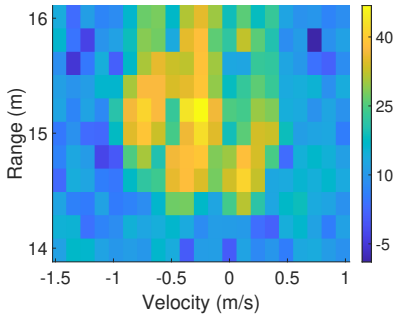

(a) Range-Doppler heatmap.

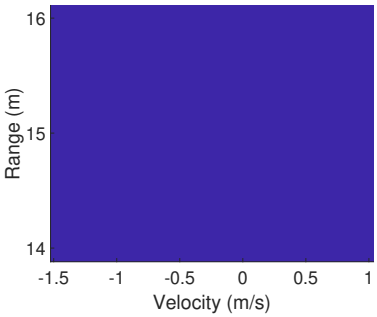

(b) CFAR detection result.

Fig. 5: The range-Doppler heatmap and CFAR detection result of a successful frame.

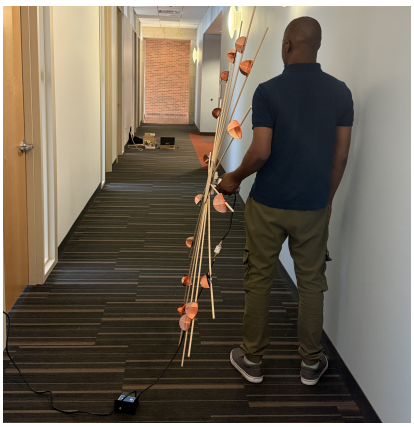

Fig. 6: The implementation of the human intruder.

To enhance concealment performance, we propose suppressing reflections from the human body using absorbing materials. Specifically, we attached a 24 cm×24 cm C-RAM MMW-1 microwave-absorbing foam [23] to the subject's chest and repeated the experiment. The average RFSC across ten runs increased significantly to 0.393, which is comparable with the performance on robot intruders.

Impact of Environmental Factors. We further evaluated the performance of the mmVeil attack in an indoor environment to examine the impact of environmental factors.

Specifically, the system was deployed in a corridor approximately 1.5 meters wide, while all other configurations remained identical to those in the default setting. The average RFSC measured across ten experiment runs was 0.34, indicating a degradation in concealment performance. We speculate that this decline is attributed to the rich multipath environment inherent to indoor settings. In particular, beyond the line-of-sight reflection path, mmWave signals also propagate via reflections from walls, the ceiling, and the floor, causing energy to spill over into neighboring cells. This spillover disturbs the precise relative range and velocity differences between the intruder and reflectors, which are critical for effective concealment.

Impact of Velocity Variations. Accurate speed control of the intruder is often challenging in practice, particularly for human intruders. Therefore, it is critical to assess the robustness of the attack against intruders with varying speeds. We used the robot for this experiment. The reflector array was configured for an intruder speed of 0.3 m/s as in the basic setting. We then configured the robot to move toward the radar at different speeds, including 0.26 m/s, 0.28 m/s, 0.32 m/s, and 0.34 m/s. For each speed, the experiment was repeated 10 times, and the average RFSC was recorded. The results are summarized in Tab. I. We also list the RFSC of the basic experiment for reference. Experimental results indicate that minor variations in velocity have negligible impact on the effectiveness of the mmVeil attack, thereby confirming its robustness in real-world scenarios.

Velocity (m/s)	0.26	0.28	0.3	0.32	0.34
RFSC	0.436	0.425	0.442	0.434	0.42

TABLE I: The impacts of speed variances.

Impact of Range and Velocity Resolutions. We further evaluated the impact of the range resolution ΔR and velocity resolution Δv of the radar system. Specifically, we tested various resolution settings and experimentally measured the corresponding RFSC values. The results are presented in Tab. II, where the units for range and velocity resolutions are meters and m/s, respectively. For certain configurations, we were unable to identify any reflector configuration to achieve effective concealment, and the RFSC is reported as 0 for those cases. Additionally, we measured the minimum distance at which the intruder could approach the radar while maintaining effective concealment, defined as an RFSC above 0.3. This distance is reported in the parenthesis next to the corresponding RFSC value.

The results suggest that the attack is more effective against radar systems with finer range resolution. Configurations combining coarse range resolution with fine velocity resolution—or vice versa—tend to be more robust against the attack. However, we were unable to derive a precise principle for assessing specific configurations. We recommend experimentally evaluating a configuration before its deployment in a real-world detection system.

Effectiveness on CFAR Variants. Finally, we experimentally evaluated the performance of the mmVeil attack against various CFAR algorithms. In addition to the commonly used CA-CFAR, we tested several popular alternatives, including Greatest-Of (GO)-CFAR, Smallest-Of (SO)-CFAR, and Ordered Statistics (OS)-CFAR. Our results indicate that mmVeil can achieve slightly better concealment performance against GO-CFAR and SO-CFAR, as these algorithms compute detection thresholds using fewer reference cells, making them more susceptible to manipulation. In contrast, OS-CFAR exhibits significantly greater resilience to the attack. Specifically, the Median-CFAR which determines the detection threshold based on the median intensity of the reference cells can nearly eliminate the effectiveness of mmVeil. This is because the attack can only manipulate a limited number of reference cells, resulting in minimal impact on the median intensity value used by the OS-CFAR algorithm.

D. Defenses Against mmVeil

Inspired by our experimental findings, we propose several defense strategies to effectively mitigate the threat of the mmVeil attack. First, the detection system can adopt radar configurations resulting in range and velocity resolutions that have been experimentally validated to be more robust against the attack. Second, the system may adopt OS-CFAR algorithms to improve resilience, as their reliance on the median intensity of reference cells makes them less susceptible to manipulation. However, a more sophisticated reflector deployment may still compromise the system by influencing a larger portion of the reference cells associated with the intruder. Finally, the system may employ alternative detection techniques, such as deep learning-based analysis, instead of CFAR algorithms. As demonstrated in our experiments, physical intrusions can be easily detected in the heatmap even with interference from reflectors. Nevertheless, this approach incurs significantly higher computational and data collection overhead.

VII. Related Work

Motivated by the rapid advancement of mmWave sensing technology, there has been extensive research on the security of mmWave sensing systems. Many of these studies focus on physical-layer security. For instance, transceivers have been exploited to forge reflection signals based on the received signals, enabling the creation of fake objects detectable by mmWave radars [24], [25]. However, such attacks are generally ineffective at concealing real objects from detection. Software-defined radios have also been successfully employed to launch signal spoofing attacks on mmWave radars, enabling both the injection of fake objects and the concealment of real ones [14], [15], [26]. Nevertheless, these attacks require specialized hardware that costs thousands of dollars, limiting their practicality in real-world scenarios.

ΔR \ Δv	0.062	0.092	0.12	0.24	0.37
0.2	0.362 (5 meters)	0.46 (3.4 meters)	0.125 (NA)	0.103 (NA)	0 (NA)
0.3	0.143 (NA)	0.264 (NA)	0.443 (4.7 meters)	0.26 (NA)	0.1 (NA)
0.4	0 (NA)	0.06 (NA)	0.175 (NA)	0.396 (5.6 meters)	0.297 (NA)

TABLE II: The impacts of speed variances.

The most relevant study to our work is TileMask [27], which employs 3D-printed reflective materials to mislead deep learning-based radar object detection. This attack can effectively cause autonomous vehicles to either ignore real objects or react to non-existent ones. However, TileMask is specifically designed for deep learning-based detection and is not effective on physical intrusion detection systems that may not rely on such methods. For similar reasons, other studies focusing on the security of deep learning-based mmWave sensing systems are considered orthogonal to our work [28], [29].

Acknowledgement

This work was supported in part by the U.S. National Science Foundation under grant CNS-2325563 and by STARs program of University of Texas System under grant AR911847. Research was also sponsored in part by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0225. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the University of Texas System, the Army Research Laboratory, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

[1] DoD support to U.S. border security. [Online]. Available: https://www.northcom.mil/BorderSecurity/
[2] Physical security monitoring for critical infrastructure in apac. [Online]. Available: https://shorturl.at/ZHVS1
[3] Military iot - internet of military things. [Online]. Available: https://shorturl.at/8Nbun
[4] T. Gu, Z. Fang, Z. Yang, P. Hu, and P. Mohapatra, "Mmsense: Multi-person detection and identification via mmwave sensing," in *ACM Workshop on Millimeter-wave Networks and Sensing Systems*, Los Cabos, Mexico, October 2019.
[5] H. Cui and N. Dahnoun, "High precision human detection and tracking using millimeter-wave radars," *IEEE Aerospace and Electronic Systems Magazine*, vol. 36, no. 1, pp. 22–32, 2021.
[6] K. Qian, S. Zhu, X. Zhang, and L. E. Li, "Robust multimodal vehicle detection in foggy weather using complementary lidar and radar signals," in *IEEE/CVF CVPR*, June 2021.
[7] Y. Wan, H. Yao, J. Liu, C. Sun, A. Ma, and Y. Zhong, "Low-light and infrared multimodal remote sensing in nighttime rescue mission: A review of anomaly detection methods," *IEEE Transactions on Geoscience and Remote Sensing*, 2025.
[8] A. Chen, X. Wang, K. Shi, Y. Huo, J. Chen, and Q. Ye, "Towards weather-robust 3d human body reconstruction: Millimeter-wave radar-based dataset, benchmark, and multi-modal fusion," *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.
[9] X. Peng, M. Tang, H. Sun, K. Bierzynski, L. Servadei, and R. Wille, "4d mmwave radar for sensing enhancement in adverse environments: Advances and challenges," *arXiv preprint arXiv:2503.24091*, 2025.
[10] Texas instruments mmwave radar sensors. [Online]. Available: https://shorturl.at/DyXoF
[11] T. Shi, P. Guo, R. Wang, Z. Ma, W. Zhang, W. Li, H. Fu, and H. Hu, "A survey on multi-sensor fusion perimeter intrusion detection in high-speed railways," *Sensors*.
[12] H. Abdelnasser, M. Heggo, O. Pang, M. Kovac, and J. A. McCann, "Radro: Indoor drone tracking using millimeter wave radar," *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2024.
[13] J. Zhang, X. Na, R. Xi, Y. Sun, and Y. He, "mmhawkeye: Passive uav detection with a cots mmwave radar," in *IEEE SECON*, Madrid, Spain, September 2023.
[14] X. Chen, Z. Li, B. Chen, Y. Zhu, C. Lu, Z. Peng, F. Lin, W. Xu, K. Ren, and C. Qiao, "Metawave: Attacking mmwave sensing with meta-material-enhanced tags," in *NDSS*, 2023.
[15] A. K. Q. Z. C. T. Hunt, D. and M. Pajic, "Madradar: A black-box physical layer attack framework on mmwave automotive fmcw radars," in *NDSS*, 2023.
[16] NI USRP products. [Online]. Available: https://www.ettus.com/products/
[17] M. Nguyen, R. Feger, T. Wagner, and A. Stelzer, "Analysis of 2d ca-cfar for ddma fmcw mimo radar," in *IEEE EuRAD*, Berlin, Germany, September 2023.
[18] H. Rohling, "Ordered statistic cfar technique-an overview," in *IEEE IRS*, Kansas City, Missouri, USA, May 2011.
[19] A. Jalil, H. Yousaf, and M. I. Baig, "Analysis of cfar techniques," in *IEEE IBCAST*, Islamabad, Pakistan, January 2016.
[20] Fcc empowers short-range radars in the 60 ghz band. [Online]. Available: https://shorturl.at/l3qPh
[21] Iwr6843 single-chip 60-ghz to 64-ghz intelligent mmwave sensor integrating processing capability. [Online]. Available: https://www.ti.com/product/IWR6843
[22] Dca1000 evaluation module for real-time data capture and streaming. [Online]. Available: https://www.ti.com/tool/DCA1000EVM
[23] Cuming microwave c-ram mmw - high frequency radar absorber. [Online]. Available: https://shorturl.at/DeMvh
[24] J. Panyavaraporn, M. N. Bhutta, and Q. Mirza, "A low-cost replica-based distance-spoofing attack on mmwave fmcw radar," in *Proceedings of the 25th Annual International Conference on Mobile Computing and Networking*, Florence, Italy, October 2019.
[25] R. Vennam, I. K. Jain, K. Bansal, J. Orozco, P. Shukla, A. Ranganathan, and D. Bharadia, "mmspoof: Resilient spoofing of automotive millimeter-wave radars using reflect array," in *2023 IEEE Symposium on Security and Privacy (SP)*, San Francisco, CA, USA, May 2023.
[26] Z. Sun, S. Balakrishnan, L. Su, A. Bhuyan, P. Wang, and C. Qiao, "Who is in control? practical physical layer attack and defense for mmwave-based sensing in autonomous vehicles," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 3199–3214, 2021.
[27] Y. Zhu, C. Miao, H. Xue, Z. Li, Y. Yu, W. Xu, L. Su, and C. Qiao, "Tilemask: A passive-reflection-based attack against mmwave radar object detection in autonomous driving," in *ACM CCS*, Copenhagen, Denmark, November 2023.
[28] Y. Xie, R. Jiang, X. Guo, Y. Wang, J. Cheng, and Y. Chen, "Universal targeted adversarial attacks against mmwave-based human activity recognition," in *IEEE INFOCOM*, Hoboken, New Jersey, May 2023.
[29] A. Singha, Z. Bi, T. Li, Y. Chen, and Y. Zhang, "Securing contrastive mmwave-based human activity recognition against adversarial label flipping," in *ACM WiSec*, Seoul, Republic of Korea, May 2024.