

On the Realism of LiDAR Spoofing Attacks against Autonomous Driving Vehicle at High Speed and Long Distance

Takami Sato^{*†}, Ryo Suzuki^{*‡}, Yuki Hayakawa^{*‡}, Kazuma Ikeda[‡], Ozora Sako[‡], Rokuto Nagata[‡],
Ryo Yoshida[‡], Qi Alfred Chen[†], Kentaro Yoshioka[‡]

[†]University of California, Irvine, Department of Computer Science

[‡]Keio University, Department of Electronics and Electrical Engineering

Abstract—The rapid deployment of Autonomous Driving (AD) technologies on public roads presents significant social challenges. The security of LiDAR (Light Detection and Ranging) is one of the emerging challenges in AD deployment, given its crucial role in enabling Level 4 autonomy through accurate 3D environmental sensing. Recent lines of research have demonstrated that LiDAR can be compromised by LiDAR spoofing attacks that overwrite legitimate sensing by emitting malicious lasers to the LiDAR. However, previous studies have successfully demonstrated their attacks only in controlled environments, yet gaps exist in the feasibility of their attacks in realistic high-speed, long-distance AD scenarios. To bridge these gaps, we design a novel Moving Vehicle Spoofing (MVS) system consisting of 3 subsystems: the LiDAR detection and tracking system, the auto-aiming system, and the LiDAR spoofing system. Furthermore, we design a new object removal attack, an adaptive high-frequency removal (A-HFR) attack, that can be effective even against recent LiDARs with pulse fingerprinting features, by leveraging gray-box knowledge of the scan timing of target LiDARs. With our MVS system, we are not only the first to demonstrate LiDAR spoofing attacks against practical AD scenarios where the victim vehicle is driving at high speeds (60 km/h) and the attack is launched from long distances (110 meters), but also we are the first to perform LiDAR spoofing attacks against a vehicle actually operated by a popular AD stack. Our object removal attack achieves $\geq 96\%$ attack success rates against the vehicle driving at 60 km/h to the braking distances (20 meters). Finally, we discuss possible countermeasures against attacks with our MVS system. This study not only bridges critical gaps between LiDAR security and AD security research but also sets a foundation for developing robust countermeasures against emerging threats.

I. INTRODUCTION

The rapid implementation of Autonomous Driving (AD) on public roads poses emerging challenges for our society. AD vehicles with roof-mounted sensor stacks are also a common sight [1, 2] in many cities around the world. Waymo’s robotaxi services are available in several cities, such as Phoenix, San Francisco, and Los Angeles [3]. LiDAR (Light Detection and

Ranging) plays a significant role in enabling such a rapid implementation of AD, especially for driverless Level 4 AD [4], which heavily relies on accurate 3D environmental information obtained via LiDAR to achieve production-level object detection and localization. Meanwhile, the essential roles of LiDAR in AD have motivated extensive research efforts to ensure its security due to the potentially fatal safety implications. One of the major security concerns is the robustness against LiDAR spoofing attacks [5]–[13], which project malicious lasers against target LiDARs to compromise their distance measurements by overwriting legitimate measurements.

In the context of object detection, LiDAR spoofing attacks have been demonstrated to have two attack effects: object injection attacks [6]–[11] and object removal attacks [6, 10]–[12]. These attacks have demonstrated successful results at the object detection model level and in low-speed, short-distance lab environments. However, none of the prior works have successfully proven their effectiveness in practical high-speed, long-range autonomous driving scenarios, despite suggesting potential safety and security impacts on AD vehicles. Table I provides an overview of existing LiDAR spoofing attacks demonstrated in the physical world. As listed, prior work predominantly focuses on stationary lab-level setups or dynamic but impractical low-speed setups (e.g., at most 5 km/h), and none of prior work has evaluated LiDAR spoofing attacks against a vehicle actually controlled by an AD software stack. Based on our survey, we identify the following three research limitations that potentially prevent demonstrating LiDAR spoofing attacks in practical AD scenarios:

Lack of practical detection and tracking system capable of high speeds and long distances: Precise and long-range detection and tracking system is essential to keep LiDAR spoof attacks effective on moving AD vehicles, especially for removal attacks, which need to be effective for at least several consecutive frames. The only previous attempt is camera-based detection and tracking with a pan-tilt system PTX-ATX18 [12, 14]. However, this system has not been evaluated in practical scenarios at long distances, in which we found that camera-based detection cannot be effective as in §V-A1.

Lack of practical spoofing devices capable on public roads: None of the prior works have demonstrated their attacks in realistic environments. Most of them have only been evaluated in indoor lab-level or outdoor setups where AD vehicles are driving slowly (< 5 km/h), such as parking lots. To deploy Li-

^{*}co-first authors

DAR spoofing attacks on real roads, the hardware requirements will be significantly higher to enable high-speed and long-range attacks. For example, when spoofing at long distances such as 100m, even small distortions in optics, errors in detection results, and delays in motor control could lead to an error of a few meters in aiming.

Lack of practical spoofing attacks against recent LiDARs:

As also pointed out in [11], recent LiDARs, so-called New-Gen LiDARs, equip security-related features, such as timing randomization and pulse fingerprinting, which can directly foil essential assumptions in prior attacks [7, 8, 12]. The HFR attack [11] is the only attack that demonstrates the removal attack capability even against New-Gen LiDARs. However, the HFR attack has not shown high attack effectiveness against LiDARs with pulse fingerprinting.

Motivated by these limitations in prior work, we pursue the following research question:

Can LiDAR spoofing attacks actually have end-to-end safety impacts in practical AD scenarios?

To answer this question, we design a novel Moving Vehicle Spoofing (MVS) system that enables us to deploy LiDAR spoofing attacks in practical AD scenarios. Our MVS system consists of 3 subsystems: infrared (IR) camera-based detection and tracking system, auto-aiming system, and LiDAR spoofing system with significant improvements over prior work. Furthermore, we design a new attack named adaptive high-frequency removal (A-HFR) attack that can be effective against LiDARs with pulse fingerprinting. The A-HFR attack leverages gray-box knowledge of the target LiDARs to avoid overheating the laser diode in the spoofer; it allows the emission of laser pulses at a frequency more than 25 times higher than that of the HFR attack while using the same laser.

This paper is structured as follows: In §II, we summarize previous efforts in LiDAR spoofing attacks and formulate our threat model to attack AD vehicles driving at high speeds from long distances. In §III, we describe the design details of our MVS system consisting of 3 subsystems. In §IV, we introduce the methodology of our A-HFR attack and describe how the A-HFR attack can be effective against LiDARs with pulse fingerprinting. In §V, we evaluate the attack capability of our MVS system and the effectiveness of our A-HFR attack. We find that the prior vision-based detection can only detect the LiDAR at up to 5 meters. On the other hand, our IR camera-based detection can detect the target LiDAR even at distances ≥ 100 meters regardless of the LiDAR types. We also explore multiple tracking designs including the Kalman filter. We evaluate the attack effectiveness on three commercial LiDARs and find that for all LiDARs, the A-HFR attack can successfully remove over 96% of the point cloud within a 20° horizontal and a 16° vertical angle.

In §VI, we demonstrate LiDAR spoofing attacks against the victim vehicle driving at high speed, up to 60 km/h. We find that the HFR attack achieves $\geq 96\%$ attack success rates until the braking distance (20 meters) at 60 km/h. The A-HFR attack also achieves 100% attack success rates until 40 meters away. The injection attacks are also successful with $\geq 1k$ injected points at 60 km/h. In §VII, we perform an end-to-end closed-loop attack evaluation on an AD vehicle operated by Autoware.ai [15]. We demonstrate that both object injection

and removal attacks with our MVS system can cause serious safety consequences. Particularly for the object removal attack, the victim fails to detect an SUV car-sized object in front of it. In §VIII, we finally discuss the findings and limitations of this study, including potential countermeasures.

In summary, our study has the following contributions:

- We design a novel MVS system that can conduct LiDAR spoofing attacks in practical AD scenarios. The MVS system consists of 3 subsystems: IR camera-based detection and tracking system, auto-aiming system, and LiDAR spoofing systems. Our MVS systems can aim at a target vehicle driving at 60 km/h from 110 meters away.
- We design a new object removal attack, A-HFR, which can be effective against LiDARs with pulse fingerprinting by utilizing a gray-box knowledge of the target LiDAR to avoid overheating the laser diode. The A-HFR attack can remove $\geq 96\%$ of points within attack angles even under pulse fingerprinting.
- We are the first to demonstrate LiDAR spoofing attacks against practical AD scenarios where the victim vehicle is driving at high speeds and the attack is launched from long distances. Our object removal attack achieves $\geq 96\%$ attack success rate against vehicles driving at 60 km/h up to their braking distance (20 meters).
- We are the first to deploy LiDAR spoofing attacks against a vehicle actually controlled by a popular AD stack (Autoware.ai [15]) and demonstrate the end-to-end serious safety consequences (e.g., hard crash into a parking mock car).

Project Website: Demo videos and more detailed hardware configurations are released on project websites from each collaborating institutions <https://sites.google.com/keio.jp/keio-csg/projects/AttackonDrivingVehicle> (Keio University side) and <https://sites.google.com/view/av-ioat-sec/real-av-lidar-attack> (UCI side)¹.

II. BACKGROUND

A. LiDAR Spoofing Attacks

LiDAR spoofing attacks [5]–[13] compromise the distance measurements of LiDAR sensors by overwriting legitimate laser signals with higher-power malicious lasers. Table I lists an overview of the existing LiDAR spoofing attacks demonstrated in the physical world. We can taxonomize LiDAR spoofing attacks into two types based on the attack goals.

1) *Object Injection Attacks:* This type of attack is designed to inject ghost objects that do not actually exist. The relay attack [5] sends back recorded laser signals to the victim LiDAR. However, the impact of this attack in AD contexts is limited because it cannot inject objects closer than the attacker [11]. To address this limitation, synchronized injection attacks [6, 8, 9] have been designed to more effectively compromise AD vehicles in practical scenarios. This attack requires a “white-box” [11] knowledge of the target LiDAR’s deterministic scan pattern and its current state, i.e., where the LiDAR is currently scanning. With this *white-box* knowledge, the attacker can emit malicious lasers to overwrite legitimate

¹We have two websites for publicizing purposes from all collaborating institutions; their contents are jointly developed and identical.

Table I. Literature survey on prior works that evaluate LiDAR spoofing attacks in the physical world. We are the first to demonstrate the LiDAR spoofing attack in practical high-speed and long-range AD scenarios and the first to perform the attacks against a vehicle controlled by an AD software stack. ✓ : Covered, - : Not Covered

	Attack on Moving Target	Attack on AD Vehicle	Maximum Speed	Attack from Roadside	Attack Goals		Maximum Attack Range
					Injection	Removal	
Ours	✓	✓	60 km/h	✓	✓	✓	110 m
Cao et al. [12]	✓	-	5 km/h	-	-	✓	10 m
Cao et al. [14]	✓	-	0.4 km/h	-	✓	-	4 m
Jin et al. [10]	✓	-	0 km/h (running parallel)	-	✓	✓	15 m
Petit et al. [5]	-	-	-	-	✓	-	1 m
Shin et al. [6]	-	-	-	-	-	✓	5 m
Sun. et al. [8]	-	-	-	-	✓	-	5 m
Hallyburton et al. [9]	-	-	-	-	✓	-	5 m
Sato et al. [11]	-	-	-	-	✓	✓	10 m

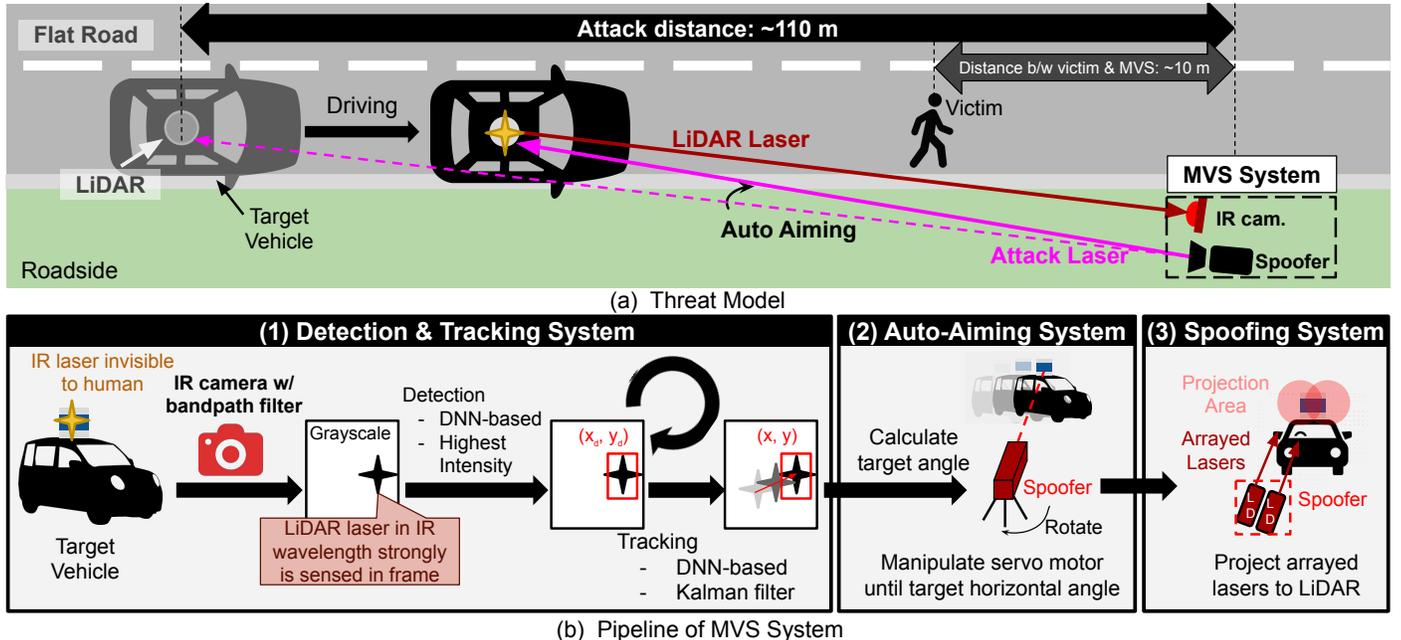


Figure 1: Overview of our threat model and pipeline of our MVS system consisted of 3 subsystems: (1) detection and tracking system with IR camera, (2) auto-aiming system with high-precision servo motor, and (3) spoofing system with arrayed lasers to achieve a wider attack projection area. Our MVS system can attack a driving AD from ≥ 110 meters away.

scans based on the obtained pattern. However, synchronized injection attacks heavily rely on a deterministic LiDAR scanning pattern; therefore, they can be directly foiled by laser scan timing randomization, which is a common feature in recent New-Gen LiDARs [11]. Considering the worst-case scenario in which some AD vehicles may still use old-generation LiDARs, we also evaluate the effectiveness of synchronized injection attacks in driving scenarios in §VI and §VII.

2) *Object Removal Attacks*: These attacks aim to prevent object detectors from identifying actual objects. Synchronized removal attacks [10, 12] remove objects by moving all points of the object to a distance far away or within the area below the minimum distance threshold. However, as discussed in §II-A1, these *white-box* attacks, which rely on synchronization, are directly ineffective against New-Gen LiDARs. Asynchronized (or “black-box”) removal attacks [5, 6, 11] do not depend on these assumptions. Particularly, the high-frequency removal (HFR) attack [11] has shown even higher effectiveness than synchronized attacks. The HFR attack made 5 sedan cars undetected and showed potential effectiveness against AD in a driving simulation. However, the HFR attack has very limited

effectiveness against LiDARs with pulse fingerprinting. This thus motivates us to design a new attack, A-HFR attack, in §IV.

B. Pulse Fingerprinting for LiDAR

Pulse fingerprinting [16] is a technique that authenticates whether the reflected laser pulses are emitted by the LiDAR itself or by other LiDARs. While pulse fingerprinting is installed originally for anti-interference purposes to allow multiple LiDARs to operate at close distances, it also demonstrates a high defense capability against LiDAR spoofing attacks [11]. Pulse fingerprinting can be found in several New-Gen LiDARs such as Livox Mid-360 [17], Hesai XT32 [18], and AT128 [19]. While the detailed implementation of pulse fingerprinting is not publicly released, it is considered to be encoded in the interval between two consecutive pulses; therefore, high-frequency pulses can accidentally match the interval and then bypass the authentication [11]. Naturally, higher frequency leads to higher chances of hitting the interval. However, it is not trivial to increase the frequency because a higher pulse frequency causes overheating of the laser diode and degrades the peak power of the laser, which must be higher than the

peak power of the legitimate laser. To overcome this trade-off, our A-HFR attack utilizes *gray-box* knowledge of the LiDAR scan pattern to know when to cool down the diode (§IV). We would like to note that such a naive authentication in New-Gen LiDARs should be an inevitable design rather than a random choice because more complex authentication requires higher laser power per time, which may harm human eyes (also discussed in [11]). For example, 2 times more pulses for complex authentication will double the laser power per time. Moreover, the detection range of LiDAR will also be degraded if LiDAR uses more power for authentication.

C. Prior Attempts to Attack Moving Vehicles

So far, there has been no successful demonstration of LiDAR spoofing attacks on AD vehicles driving in realistic AD scenarios. As shown in Table I, the majority of prior works do not consider attacks on moving targets. Several prior attempts [10, 12, 14] have targeted moving vehicles, but these works have the following three critical limitations that prevent them from effectively attacking driving AD vehicles: First, the LiDAR detection systems in prior work are not capable of operating at long ranges. Jin et al. [10] just manually aimed at the target LiDAR, and thus their approach cannot be applied to long-range scenarios since humans cannot even see a LiDAR from a far point, such as from 100 meters away. Cao et al. [12, 14] detect the target LiDAR with a vision-based approach with a YOLOv3 [20] model trained to directly detect the target LiDAR. While this is a systematic approach, we find that the vision-based approach cannot handle long-range LiDAR detection even at 10 meters away since the LiDAR appears too small to be detected in the captured image frame, as detailed later in §V-A1.

Second, we also find that prior works did not have mature-enough spoofing attack devices that can automatically and precisely aim at a far-away target with attack lasers with sufficiently high power. The only prior attempts [12, 14] used a generic pan-tilt system [21] to control the laser emitter direction. However, these attack devices are unlikely to be effective against fast-moving vehicles at long distances since such generic pan-tilt systems are not originally designed for precise targeting of far and small objects (e.g., LiDAR), but only for coarse vision tracking of photo subjects. Finally, the majority of prior attacks that require synchronization with the target LiDAR cannot be effective against the recent New-Gen LiDARs since timing randomization makes synchronization virtually impossible [11]. HFR attack [11] remains effective even under timing randomization but is highly mitigated by pulse fingerprinting, and thus its end-to-end attack impact on AD vehicles has not been fully validated. To address these limitations, we design our MVS system that can practically deploy effective LiDAR spoofing attacks against driving AD vehicles, as detailed later in §III.

D. Threat Model

We generally follow the same threat model adopted in prior works [7]–[9]. We particularly target the threat model called “Spoofer placed in environment” threat model [9], in which the attacker deploys LiDAR spoofing attacks against a driving AD vehicle from a roadside. We consider that this threat model is the most practical but also the most challenging among the

ones discussed in [9]. The other threat models assume that the attack is launched from a front or side vehicle traveling alongside the victim vehicle to avoid tracking. Fig. 1 illustrates the bird’s-eye-view of our threat model. The victim vehicle driving at high speeds (e.g., 60 km/h) is under attack from ≥ 110 meters away. We assume that the attack site is a straight and flat road, which is the most common, and assume that the height of the target LiDAR is fixed and known for each attack scenario. The attack device is placed at the roadside and starts attacking the victim from as far away as possible (e.g., ≥ 110 meters away).

E. Major Updates from Our Preprint Papers

From our preprint works as WIP (Work-in-Progress) papers at the VehicleSec symposium [22, 23], this full paper has major updates in not only new design contributions enabling the longer-range attacks with the parabolic mirror antenna designs (§III-D) to receive and track the lasers from the LiDAR even at 110 meters away, but also has more comprehensive evaluation including a quantitative comparison between the vision and IR-based detection methods (§V-A1), driving evaluation at higher speeds such as 60 km/h (§VI), injection attack evaluation against the high-speed vehicle (§VI-D), A-HFR attack evaluation against the high-speed vehicle (§VI-B2), and end-to-end evaluation with the actual AD system (§VII). Our preprint works focused on the initial testing of our attack devices and attack design validity at lower speeds.

III. MVS SYSTEM DESIGN

A. Overview

To practically deploy long-range LiDAR spoofing attacks against AD vehicles driving at high speeds, we design a Moving Vehicle Spoofing system (MVS system). Fig. 1 illustrates the overview of our MVS system consisting of three subsystems: (1) IR camera-based detection and tracking system, (2) auto-aiming system, and (3) LiDAR spoofing system. For the IR camera-based detection and tracking system, we design a novel methodology with an IR camera to accurately localize the target LiDAR even at longer ranges (e.g., 100 m) than the prior vision-based method, which can only detect a LiDAR up to 5 m (§V-A1). For the auto-aiming system, we design a responsive and accurate auto-aiming system that can precisely aim at any horizontal angle with a high-precision servo motor. For the LiDAR spoofing system, we build a novel spoofer upon the existing LiDAR spoofers [7, 8, 10]–[12, 14] to handle long-range attack scenarios. We further improved the electronics and optics as detailed in Fig. 2. Particularly, we install a parabolic mirror antenna before the PD (photodiode) to receive the laser from LiDAR even at long distances.

B. Infrared Camera-based Detection and Tracking System

Prior work [12, 14] utilizes a vision-based approach based on YOLOv3 [20] to detect the target LiDAR. However, as shown in Fig. 3 (left), the LiDAR appears too small in the camera frame for reliable detection at long distances (e.g., 15 m). A naive solution would be to use an expensive camera and telephoto lens to get a higher resolution image, but this approach not only incurs high costs but will introduce challenges in real-time processing due to the high-resolution image and manipulation of a heavy-weight camera.

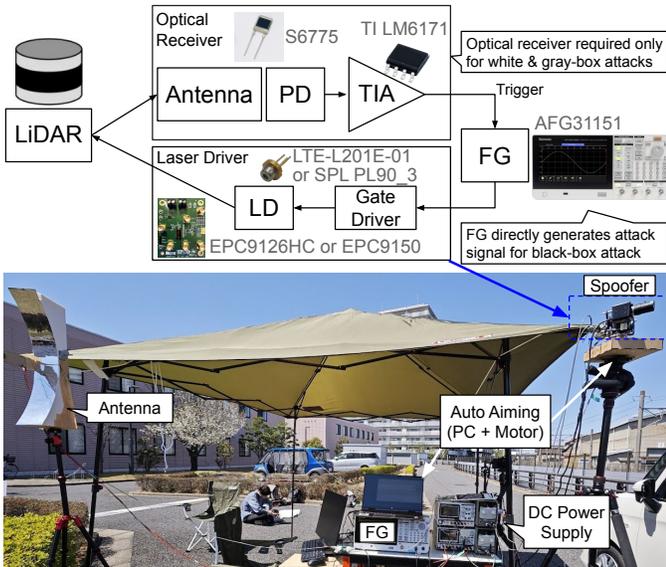


Figure 2: Overview of the hardware setup of our MVS system. The antenna is only used for white-box and gray-box attacks to obtain the current state of the target LiDAR.

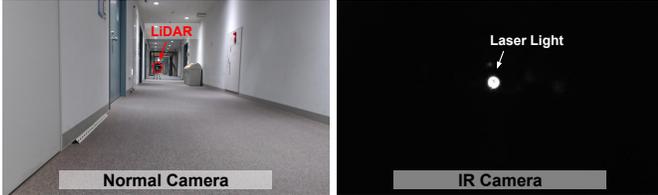


Figure 3: Comparison of vision and IR-camera frames at the same sight 15 m away from the target LiDAR. The IR camera only senses the light in the IR wavelength and thus enables accurate LiDAR-agnostic detection of the LiDAR location.

1) *Target LiDAR Detection with IR Camera*: To address this challenge, we design a simple yet novel detection system inspired by InfraRed Search and Track Systems (IRST Systems) designed for military applications, enabling long-distance detection and tracking of objects like enemy fighter jets [24]. By design, LiDAR must keep emitting IR lasers in all directions to obtain 3D point cloud data. Intuitively, we should be capable of achieving stable detecting and tracking of the IR laser source, similar to how an anti-aircraft missile operates. Fig. 3 (right) shows the IR camera frame at the same location as the left figure. As the shutter time of the IR camera is significantly longer than that of the LiDAR scans, the IR camera can very likely capture the laser trace during the shutter time like the distinct white circle in the figure.

This approach can achieve three key advantages over prior vision-based systems: First, we can *directly* determine the location of the laser trace. As the laser and its reflection travel in a straight line, the origin of the laser emission must coincide with the location where the reflection was sensed. Therefore, the attacker can directly target the source of the laser, as indicated by the white circle in Fig. 3 (right). On the contrary, the prior vision-based approach requires an additional step to determine the laser emission location even after detecting the bounding box of the target LiDAR. Furthermore, the detected bounding box itself has localization errors and jittering, especially for small object detection.

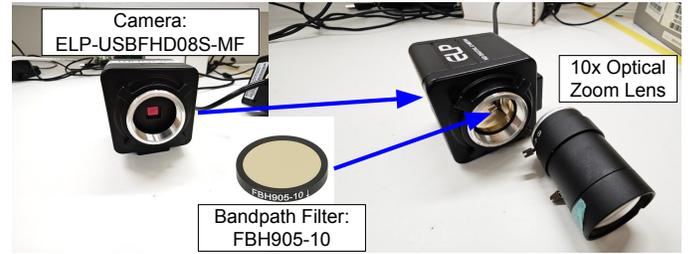


Figure 4: IR Camera Configurations. We install the bandpass filter in front of the image sensor and place the 10 times optical zoom lens on it.

Secondly, our approach is *LiDAR model-agnostic* because this approach relies only on the fundamental property inherent to all LiDARs, which inevitably emits lasers into the scanning area. In §III-B1, we demonstrate that our method can be universally applied to three LiDAR models without requiring additional training or adaptation. On the other hand, the prior vision-based approach requires a large number of various images of the target LiDARs to train a dedicated object detector. If the attacker wants to handle multiple scenarios (e.g., multiple LiDAR models), the number of required images could increase exponentially. Finally, our approach can be robust against different environmental conditions. For example, the prior vision-based approach is not generally capable of nighttime scenarios due to the dependence on passive visible lights from other sources. Our approach can handle such challenging environmental conditions because this approach merely relies on the active IR light emitted from the target LiDAR. The other wavelengths from other sources will be filtered out by the bandpass filter.

To eventually localize the LiDAR in the IR camera frame, we design two approaches: DNN (Deep Neural Network)-based object detection and highest-intensity methods. For the DNN-based object detection, we train an object detector (e.g., YOLOv5 [25]) to detect the LiDAR laser-induced white circle. For the highest-intensity method, we simply select the pixel with the highest intensity in the frame. In §V-A2, we find that the DNN-based detection shows superior robustness compared to the highest-intensity method. While the highest-intensity method works well in indoor environments, it does not consistently perform well in outdoor environments where the sunlight and its reflection often have the highest intensity. In §V-A1, we confirm that our DNN-based method achieves significantly higher detection accuracy than the prior vision-based approach while the vision-based detection can only detect a LiDAR at most 5 meters away. Our IR camera-based detection can LiDAR-agnostically detect different LiDARs from distances exceeding 20 meters with $\geq 90\%$ success rates and can effectively handle the 110-meter outdoor attack scenarios in §VI.

IR Camera Configurations: Fig. 4 shows the configuration of our IR camera (ELP-USBFHD08S-MFV [26]) with 10 times optical zoom. We removed the pre-installed low-pass filter in front of the CMOS sensor to accept the IR wavelength. We then placed a bandpass filter (FBH905-10 [27]) that can selectively receive the IR wavelength (905 ± 5 nm), which is used in the majority of LiDARs [28]. If the attack wants to attack other LiDARs using different wavelengths, the adversary can use another bandpass filter corresponding to the wavelength range covering the wavelength.

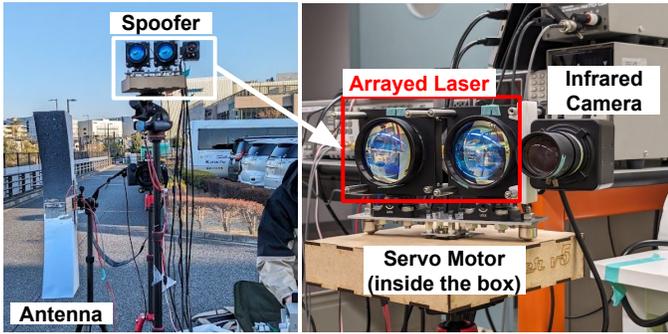


Figure 5: Overview of our spoofing system. The antenna is only used for white-box and gray-box attacks to know the current state of the target LiDAR.

2) *Robust Tracking with Misdetected Frames*: After the LiDAR is localized in each frame, the next step is to track the LiDAR with prior detection results and maintain the LiDAR location for the frame with failed or outlier detection. We note that tracking is particularly beneficial for our IR camera-based detection since we cannot always sense the laser from the target LiDAR. For example, VLP-16 [29] scans each point at around 10 Hz, meaning we can at most detect the LiDAR location every 0.1 seconds. To handle this, we implement two methods: the Kalman filter [30] and DNN-based tracking. Kalman filter is one of the most widely used methods for tracking purposes. Kalman filter estimates the state of a linear dynamic system from a series of noisy measurements. We implement the Kalman filter with outlier elimination based on the Mahalanobis distance [31]. For the DNN-based tracking, we simply feed the current and a few prior frames as the channel of the input image of the DNN-based detection. We denote the number of frames to feed as N_p . In §V-A2, we find that both methods generally achieve high tracking performance and that the combination of the two methods shows the highest performance. Particularly, the DNN-based tracking with $N_p=3$ shows high robustness against the ghosting effect [32] caused by strong lasers from the LiDAR.

C. Auto-Aiming System

After the current position of the target LiDAR is estimated with tracking, the auto-aiming system calculates the required rotation angle to aim at the LiDAR. We assume the attack site is on a flat road and that the height of the LiDAR is known and constant, as discussed in our threat model (§II-D). Therefore, we first set our spoofer at the same height as the target LiDAR and the auto-aiming system only adjusts the horizontal angle on the fly. To accurately aim at the target horizontal angle, we find that the precision of the servo motor has a major impact on attack success because a 1° error leads to around 1.7 m error at 100 m away. We eventually selected a high-precision servo motor, Dynamixel MX-28, which has 0.088° angle resolution with PID control. To smoothly rotate the entire MVS system, we enclose the motor on a box and the system rotates on the box with two bearings as shown in Fig. 5.

D. LiDAR Spoofing System

Table II lists the comparison between our MVS system and prior setups to attack moving targets. Our MVS system has two major improvements over prior works: (1) We adopt

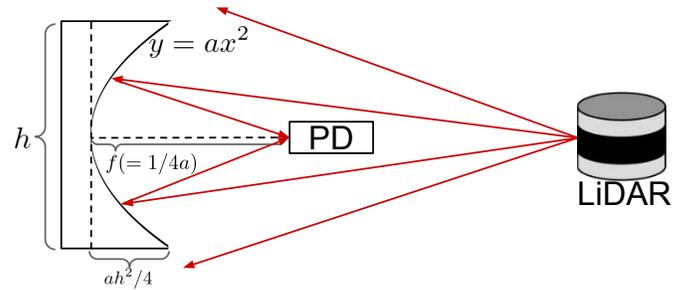


Figure 6: Illustration of how the vertical parabolic mirror antenna helps the PD receive more lasers from LiDAR.

two arrayed lasers with 2-inch lenses to cover a significantly larger (100 times larger) area than prior work. The two lenses are horizontally arrayed. While we could place more arrayed lasers to cover a larger beam area, doing so could harm the responsiveness of the auto-aiming system due to the increased weight, and it also introduces additional costs. In §VI-C, we find that the setup with the two 2-inch lenses can achieve sufficient performance to attack a vehicle driving at 60 km/h. (2) We introduce a vertical parabolic mirror antenna before the photodiode (PD) to reliably receive the laser from the LiDAR. Prior work has used a bare PD, but it cannot work in long-range attack scenarios because the lasers from LiDAR spread radially and sparsely at long distances, and thus the bare PD is less likely to receive them. The parabolic shape can collect signals arriving from any direction at a single point (focal point). Fig. 6 shows how the vertical parabolic mirror antenna helps the PD receive lasers from LiDAR at long distances. The parabolic mirror surface reflects incoming lasers and concentrates them at the focal point where the PD is located. While we can also gather such sparse lasers by using a convex lens, the parabolic antenna has clear advantages in the lightweight design and cost efficiency as evidenced by their widespread use in space exploration applications [33].

For our threat model, we design a one-dimensional vertical parabolic mirror as shown in Fig. 5. We use a focal length $f = 200$ mm and an antenna height $h = 600$ mm. For the horizontal dimension, we adopt a thick design, consisting of a stack of parabolic shapes, because the strong directivity of the parabola is not suitable for the horizontal dimension where the target vehicle moves quickly. This design allows us to eliminate the need to move the antenna during the attack. Once the MVS system receives the laser from LiDAR with the antenna, we can launch not only white-box attacks [6, 8, 9] by synchronizing the LiDAR scan pattern, but also our gray-box A-HFR attack, which will be introduced in §IV. For the LiDAR spoofing methodologies themselves, we follow prior works (except for the A-HFR attack part) since our major contributions of the MVS system are in how to deploy LiDAR spoofing attacks in high-speed and long-range scenarios.

E. Cost of the MVS system

Fig. 2 illustrates the entire components of our MVS system, which is built completely using off-the-shelf components. The system has a total cost of \$2.3k, including the \$1.1k laser driver, \$36 optical receiver, \$700 detection and auto-aiming systems, and \$500 in miscellaneous parts including jigs. Specifically, the laser driver component includes two laser boards ($\$450 \times 2$), two 2-inch lenses ($\54×2), and two

laser diodes (LD) (\$36 x 2). The optical receiver consists of a PD (\$2), a TIA (\$4), and an antenna (\$30). The detection and auto-aiming systems include the IR camera (\$70), the bandpass filter (\$165), the servo motor (\$270), and the box with a turning table (\$200). Additionally, to run the MVS system, we use a function generator (FG) (\$6.7k), 4-channel DC power supplies (\$3k), and a laptop (\$1.7k). We note that the majority of the cost is due to the expensive FG, which can be replaced with Analog Discovery (\$400) [34] or FPGAs (~\$200) with a certain level of engineering efforts. As this study is motivated to explore the feasibility of LiDAR spoofing against moving vehicles, we used the FG to flexibly evaluate various parameters and setups with less effort.

IV. NEW ATTACK: A-HFR ATTACK

To bypass pulse fingerprinting in recent New-Gen LiDARs, we design a new removal attack, named Adaptive HFR (A-HFR) attack. We discover that the pulse fingerprinting can be bypassed if we can achieve much higher frequency pulses than the ones of the prior HFR attack [11]. We achieve 25 times higher frequency than the HFR attack at the target attack range by utilizing the *gray-box* knowledge of the scan pattern of target LiDAR, to avoid the overheating issue of the laser diode.

A. Basic Concept to Bypass Pulse Fingerprinting

Prior work [11] points out that the pulse fingerprinting in recent New-Gen LiDARs presumably uses a pair of pulses to measure a point and embeds the fingerprint into the interval of a pair of pulses. The pulse intervals of each pair randomly vary and the LiDAR can authenticate the reflected laser based on whether the interval matches the one the LiDAR emitted or not. The HFR [11] attack demonstrated that it can still bypass pulse fingerprinting LiDARs with 2.1% attack success rate and hypothesizes that their high-frequency pulses can occasionally match the correct interval and bypass the authentication as described in Fig. 7. Another possibility is that the overheating may differ the laser wavelength to out of LiDAR’s acceptable wavelength. Based on the datasheet [35], the effect is very small (3 nm per 10°C) and thus should be negligible.

The hypothesis assumes the existence of T_α , which is a tolerance error time used to authenticate the fingerprinting. With T_α , the attack success rate of the HFR attack in an ideal case can be written by $\min\left(1, \frac{T_\alpha}{T_A}\right)$, where T_A is the interval of attack pulses. In other words, the attack always succeeds if the interval of attack pulses is less than the tolerance error, and stochastically succeeds otherwise.

If the hypothesis is correct, a higher pulse frequency should result in higher attack effectiveness. As shown in Fig. 11, we confirmed that the hypothesis is correct through the experiment, i.e., a sufficiently high pulse frequency can bypass the fingerprinting in recent New-Gen LiDARs. The remaining challenge lies in achieving such a high pulse frequency. However, we find that it is not trivial to increase the frequency due to a trade-off between pulse frequency and peak power. Specifically, emitting a pulsed laser at a higher frequency causes overheating of the laser driver and the laser itself, significantly degrading the laser’s peak power. To overpower the legitimate laser power emitted by the target LiDAR, the peak power of the attack laser must be greater

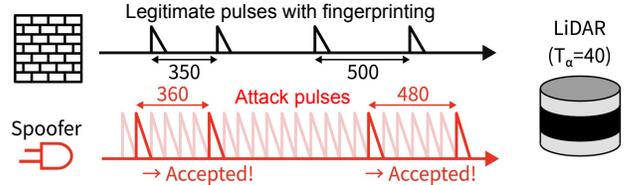


Figure 7: Requirements to bypass fingerprinting: (1) The interval of the attack pulses (T_A) should be short enough and close to the tolerance error (T_α); (2) the peak power of the attack laser must be higher than the legitimate laser of LiDAR.

than that of the legitimate laser. The HFR attack achieves the highest attack success rate at around 1 MHz when attacking LiDARs without pulse fingerprinting. Ideally, attacking with higher frequencies should increase the effectiveness of HFR, but overheating causes a reduction in laser power, significantly lowering the attack success rate.

B. Attack Design

To overcome the trade-off between pulse frequency and peak power, we design the adaptive-HFR (A-HFR) attack. As the name implies, the A-HFR attack adaptively changes its attack laser frequency to avoid overheating. As shown in Fig. 8, the A-HFR attack strategically boosts the frequency only when the target LiDAR scans the specific object the attacker wants to hide. This design allows the diode to rest most of the time and effectively cool down during the rest. To know when the LiDAR scans the target objects, we utilize a photodiode (PD) similar to the existing *white-box* (§II-A1) LiDAR spoofing attacks that require synchronization with the target LiDAR. However, unlike *white-box* attacks that require precise knowledge of when the LiDAR scans each point and a predictable scan pattern, A-HFR only needs a coarse-grained understanding of the victim LiDAR’s state. Due to this reduced information requirement compared to *white-box* synchronization attacks, we classify the A-HFR attack as a *gray-box* attack and term the coarse-grained requirement as *weak synchronization*.

Obtaining Gray-box Knowledge: We can roughly estimate when LiDAR scans each horizontal angle based on the laser from the LiDAR with PD. As LiDAR is rotating and scans each angle evenly, the interval between two consecutive lasers from the LiDAR corresponds to the time required to scan the full 360°. We can calculate the scan timing of each relative angle between the PD and the target object. Although the relative angle will change as the victim AD vehicle moves, we can account for this by using a wider margin in the attack angle. In §V-B3, we evaluate the maximum attack angle that can maintain attack effectiveness. Similarly, we can also rest the attack when the LiDAR is scanning less important vertical angles. For example, pedestrians and other vehicles generally do not exist at a height ≥ 3 meters. If we can focus only on the area within 20° of horizontal and vertical angles, the laser diode can rest 97% of the total LiDAR scan time. In Appendix A, we evaluate the impact of limiting vertical and horizontal attack angles and find that limiting both vertical and horizontal attack angles is necessary to achieve more than 25 MHz frequency.

Table II. Comparison between our MVS system and the prior setups to attack moving targets.

	Number of Lasers	Beam Diameter	Total Beam Area	Maximum Tracking Distance	Tracking	Detection Strategy (target, using device)	LiDAR Agnostic	Vertical Antenna
Ours	2	60 cm	5654.7 cm ²	110 m	Auto	Detect LiDAR laser w/ IR Camera	Yes	Yes
Cao et al. [12, 14]	1	2.54 cm	5.1 cm ²	5 m	Auto	Detect LiDAR w/ RGB Camera	No	No
Jin et al. [10]	1	8cm	50.3 cm ²	15 m	Manual	-	-	-

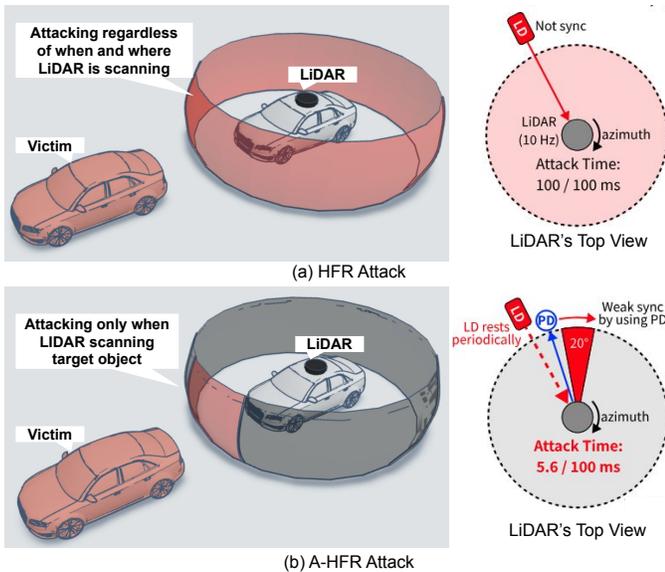


Figure 8: Comparison between (a) HFR attack and (b) A-HFR attack. A-HFR can keep high peak power at high effective frequencies by limiting the attack angle with weak synchronization to know where to boost the frequencies.

V. ATTACK CAPABILITY EVALUATION

We first evaluate the attack capability of our MVS system and then evaluate A-HFR attack through static indoor and outdoor experiments to validate our proposed designs.

A. MVS System Evaluation

Of the three subsystems of the MVS system, we begin with an evaluation of the IR camera-based detection and tracking system because the others require actual high-speed driving scenarios to adequately evaluate their capabilities, which will be covered in §VI and §VII.

1) *Detection Capability Evaluation – Vision v.s. IR Camera:* We evaluate the detection capability of our IR camera-based target LiDAR detection (§III-B1) through the comparison with the prior state-of-the-art vision-based detection [12, 14]. We calculate the detection success rate of each method at 4 different distances between the target LiDARs and the vision or IR camera. We note this is an indoor static experiment, i.e., both the LiDAR and camera remain stationary. For the vision-based detection, we train YOLOv5 [25] with 200 images of the target VLP-32c [36] LiDAR. For our IR camera-based detection, we train YOLOv5 with 532 images of the attack traces like those shown in Fig. 3 (right). The resolution of all vision and IR images is 640x640. Since the IR camera cannot always sense the emitted laser from the target LiDAR, we define the attack success rate as whether a detector can correctly output a point within the area of the LiDAR with the previous 10 camera frames. The camera operates at 60 Hz, meaning 10 frames span 0.17 seconds. For both vision-

Table III. Detection success rates of vision- and IR camera-based methods at different distances between the target LiDAR and the cameras.

	LiDAR	5 m	10 m	15 m	20 m
Vision-based Detection	VLP-32c	100%	0%	0%	0%
IR Camera-based Detection (Ours)	VLP-32c	100%	100%	100%	100%
	Horizon AT128	90%	90%	100%	90%

and IR camera-based detection, we finally output the averaged LiDAR position of the frames excluding the frames that failed to detect the target LiDAR.

Results: Table III lists the detection success rates of vision- and IR camera-based methods at different distances between the target LiDAR and the cameras. As listed, the vision-based detection fails to detect the target LiDARs even at 10 m away from the camera. On the other hand, our IR camera-based detection achieves $\geq 90\%$ detection rates for 3 different LiDARs. Among the 3 LiDARs, AT128 [19] is shown to be slightly more difficult to stably receive its laser by the IR camera. We suspect that AT128 has some irregular scan patterns for the vertical angles. Nevertheless, $\geq 90\%$ detection rates are already sufficiently high, and we find that the remaining 10% of frames should be complemmentable by the tracking system as evaluated in the next section.

2) *Tracking Capability Evaluation:* We then evaluate the tracking performance for each combination of the detection and tracking methods. Since a dynamic experimental setup is required to evaluate the performance of tracking, we record a video as shown in Fig. 9 where we put a target LiDAR (AT128) on a push cart and manually push it toward an IR camera from 40 meters away at around 8 km/h speed. We apply each method for this video and calculate the tracking success rate defined by the ratio of whether the tracking point is within the LiDAR area or not for each frame.

Results: Table IV lists the tracking success rates of different combinations of detection and tracking methods for each 2-meter interval bin for AT128 LiDAR. As listed, YOLOv5 with $N_p = 3$ achieves the highest tracking rates regardless of the existence of the Kalman filter. This result indicates that the DNN-based tracking is highly effective in our threat model. As shown in Fig. 9, the highest-intensity methods and YOLOv5 with $N_p = 1$ cannot adequately handle several cases. For example, at a distance of 15 meters away, the pixel with the highest intensity is not the LiDAR but rather a street light reflecting sunlight into the IR camera. Similarly, at 7 meters away, YOLOv5 with $N_p = 1$ misdirects the ghost reflection [32] caused by strong direct light from the LiDAR entering the lens. Since such strong direct light does not occur in consecutive frames, YOLOv5 with $N_p = 3$ appears unaffected by ghost reflections. For these methods, the Kalman filter sometimes helps to correct the LiDAR localization but not always since the Kalman filter also sometimes causes the long-last effect of misdetections at the current frame for

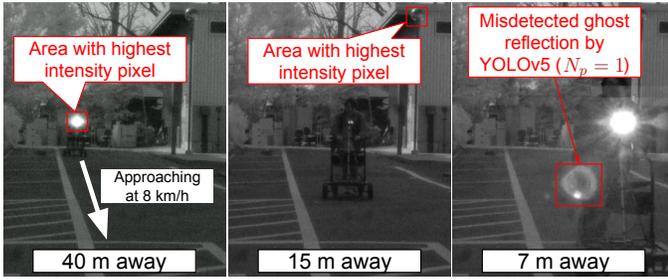


Figure 9: IR camera frames at 40, 15, and 7 meters away from the target LiDAR. The highest intensity pixel in the 15-meter frame is the reflection of the sunlight on a street light. The 7-meter frame has ghosting misdected by YOLOv5 with $N_p=1$.

Table IV. Tracking success rates of different combinations of detection and tracking methods for AT128 LiDAR.

	Distance (m)								Avg.
	2-4	4-6	6-8	8-10	10-12	12-14	14-16	16-40	
Highest Intensity	100%	100%	91%	91%	75%	63%	73%	58%	68%
+ Kalman Filter	0%	0%	91%	100%	100%	100%	100%	100%	89%
YOLOv5 ($N_p = 1$)	77%	0%	14%	100%	100%	100%	100%	100%	89%
+ Kalman Filter	0%	0%	95%	100%	100%	100%	100%	100%	89%
YOLOv5 ($N_p = 3$)	55%	70%	100%	100%	100%	100%	100%	98%	95%
+ Kalman Filter	55%	70%	100%	100%	100%	100%	100%	98%	95%
# of frames	22	20	21	23	20	16	15	145	

Table V. Three production LiDARs with pulse fingerprinting functionalities based on our measurements or official documentation. n/a means that we could not measure it.

	Livox Mid-360 [17]	Hesai XT32 [18]	Hesai AT128 [19]
Scanning Type	Prism Rotating	Rotating	Mirror Rotating
T_{min}	1250 ns	250 ns	n/a
T_{max}	1550 ns	450 ns	n/a

the following frames. Although other state-of-the-art tracking methods (e.g., extended Kalman filter [37], unscented Kalman filter [38], particle filter [39]) may overcome the tracking issues including the lower performance at close distances, we adopt the tracking method with YOLOv5 ($N_p = 3$) and the Kalman filter in the later experiments since it already has almost 100% accuracy until 6 meters away, which is much shorter than the braking distances (20 meters) at 60 km/h.

B. A-HFR Attack Evaluation

We evaluate the attack effectiveness and robustness of the A-HFR attack against LiDARs with pulse fingerprinting by comparing the attack results of the HFR attack, which is the current state-of-the-art removal attack [11].

1) *Evaluating LiDARs with Pulse Fingerprinting*: We identify three mass-produced LiDARs with pulse fingerprinting as listed in Table V. For XT32 [18] and Mid-360 [17], we find that the two LiDARs emit a pair of lasers for a single distance measurement. The pulse interval of each pair varies between the minimum (T_{min}) and maximum (T_{max}). For AT128 [19], For AT128 [19], we confirm the fingerprinting features through their official document, although we could not measure accurate T_{min} and T_{max} with our measurement

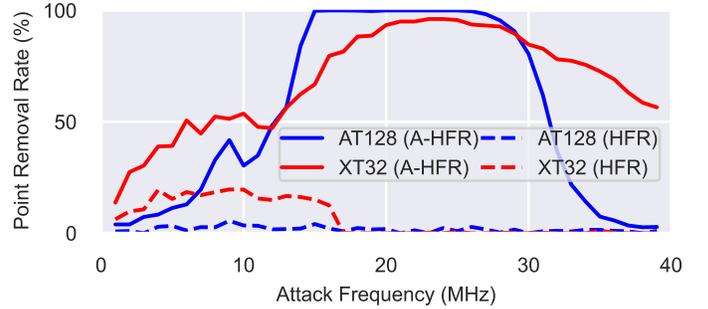


Figure 10: Point removal rates of the HFR and A-HFR attacks at different frequencies for LiDARs with pulse fingerprinting.

environments. We mainly evaluate the A-HFR attack against XT32 and AT128 because the fingerprinting in Mid-360 [17] does not have sufficient complexity and can be bypassed by the HFR attack alone, i.e., the pulse frequency of the normal HFR attack is already high enough to bypass Mid-360's fingerprinting. We further discuss it in Appendix B.

2) *Attack Effectiveness*: Fig. 10 shows the point removal rates for HFR and A-HFR attacks at different frequencies. We evaluate the ratio of removed points within the angle of 20° horizontally and 16° vertically based on our preliminary analysis that shows that $\geq 95\%$ of objects located more than 6 meters away, the minimum LiDAR detection range, in the KITTI dataset [40] fit within this area. We placed the spoofer 2 meters away from the LiDAR and 3 meters away from the target room wall. For XT32, we restrict the vertical attack angle to half (16°) to have more rest. As shown, our A-HFR attack can achieve significantly higher attack success rates compared to the current state-of-the-art HFR attack. For AT128 and XT32, due to the overheating issue, the HFR attack can achieve a maximum success rate of only 20% at around 10 MHz, and the attack success rate even starts dropping around 17 MHz. In contrast, the A-HFR attack can remove 100% of the points for AT128 at 15 MHz for AT128 ($T_\alpha = 67$ ns) and 96% of the points at 24 MHz for XT32 ($T_\alpha = 42$ ns).

Fig. 11 demonstrates the effectiveness of the HFR and A-HFR attacks against AT128. For the HFR attack, the majority of the points still remain, although we can see some points have been removed and the shape appears blurred. With the A-HFR attack, almost all points of the person have been completely removed. We note that this attack can be sustained for a long time, such as >100 seconds.

3) *Robustness to Wide Attack Horizontal Angles*: We evaluate the robustness of the A-HFR attack when the attack angle is increased. Although horizontal 20° can cover most road objects, wider attack angles can help compensate for spoofing errors and sudden movements of the target vehicle. Table VI lists the point removal rates of the HFR attacks at different attack horizontal ranges for AT128 and XT32. As shown, the point removal rate decreases as the attack angle increases, especially for XT32, which requires a higher frequency to attack as shown in §V-B2. As a higher frequency leads to greater overheating of the diode, XT32 is more robust against wider attack angles than AT128 since the A-HFR attack is also affected by the overheating issue, at wider attack ranges. Nevertheless, for both LiDARs, the A-HFR attack can cover at least 20° , which is wide enough to hide at least one pedestrian.

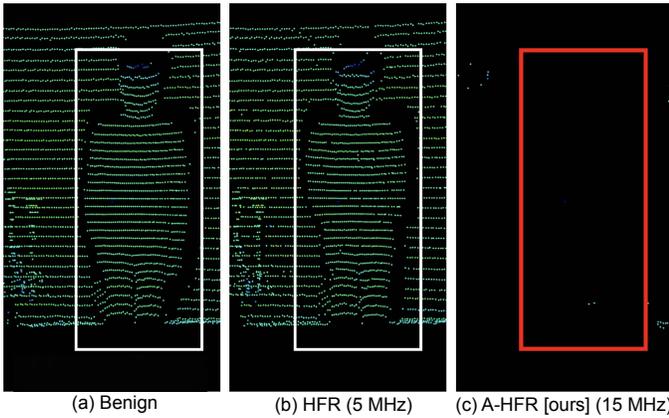


Figure 11: Comparison of the results of HFR and A-HFR attacks on AT128. A-HFR attack can completely remove almost all points (right).

Table VI. Point removal rates of the A-HFR attack with different attack horizontal ranges LiDARs with pulse fingerprinting.

	Attack Freq.	10°	20°	30°	40°	50°	60°
AT128	15 MHz	97%	100%	99%	98%	90%	90%
XT32	24 MHz	97%	96%	78%	58%	47%	38%

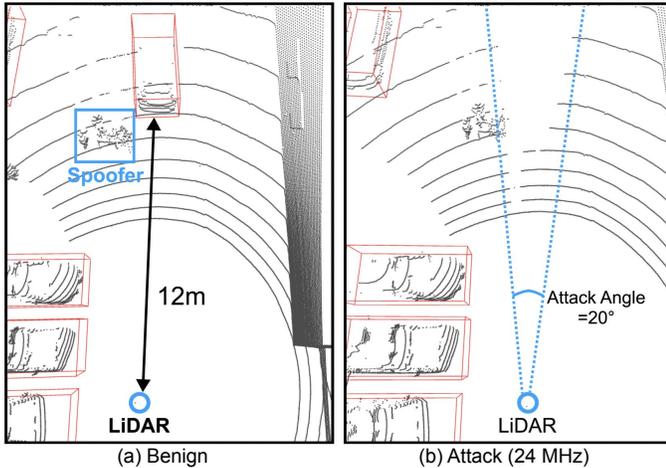


Figure 12: A-HFR attack on Hesai XT32 to hide a real vehicle. We limit the attack angle to 20° horizontally and 16° vertically. The red bounding box shows the detection results generated by Pointpillars in Apollo. When under attack, the vehicle becomes undetected with a 98% success rate for over 15 seconds.

4) *Outdoor Evaluation:* To further evaluate the effectiveness of the A-HFR attack, we conducted an outdoor experiment aiming to remove an entire car point cloud on XT32 LiDAR. We place the target car at a distance of 12 meters away. The car occupies around the area within 10° horizontally and 8° vertically in the LiDAR point cloud. As shown in Fig. 12, the target vehicle is entirely removed from the point cloud with a 98% success rate by the PointPillars of Apollo 6.0 [41] in 15 seconds. The several failures are due to the failure in weak synchronization to receive the laser from the LiDAR. Currently, we trigger the attack only when the laser correctly is received. We consider that this issue could be addressed by a speculative execution. Detailed discussions are in §VIII-5.

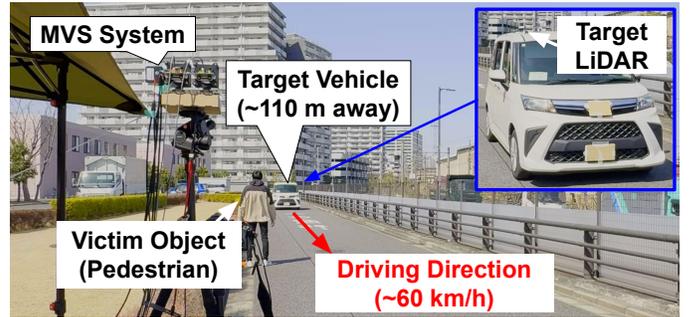


Figure 13: Experimental setup of high-speed driving evaluation with real car driving at ≤ 60 km/h. Our MVS system starts deploying LiDAR spoofing attacks around 110 meters away.

VI. HIGH-SPEED DRIVING EVALUATION WITH REAL CAR

We evaluate the attack feasibility and performance of LiDAR spoofing attacks with our MVS system against a real car driving at high speeds (e.g., 60 km/h). For removal attacks, we assess the effectiveness of HFR and A-HFR attacks. For injection attacks, we evaluate the synchronized injection attack.

A. Experimental Setup

Fig. 13 illustrates the overview of our experimental setup. We follow our threat model described in §II-D: We rent a private testing road with exclusive-use permission to ensure a controlled environment. We placed the MVS system device 2 m away from the driving lane and positioned the victim pedestrian 10 m away from the MVS system device. To prioritize safety, the victim pedestrian is also 2 m away from the driving lane, not directly in the lane. The target vehicle with a mounted LiDAR on its roof is approaching from 100 m away from the victim pedestrian, i.e., the attack distance between the MVS system and the target vehicle is up to 110 m. We measure the attack effectiveness starting from 70 m away from the victim, where the target vehicle reaches the desired speed. We evaluate the system at 6 different speeds from 10 km/h to 60 km/h. For each speed, we collected data from 4 different traces. We divided the measurements from all trials into distance bins and calculated evaluation metrics for each bin by combining all the trials.

B. Removal Attack Evaluation

1) *HFR Attack at High Speed:* We evaluate the HFR attack against a high-speed driving vehicle with Livox Horizon [42], which is one of the New-Gen LiDARs whose manufacturer releases an official object detector, Livox Detection v2.0 [43] working with Livox Horizon. We selected 0.25 as the confidence threshold for Livox Detection v2.0, ensuring that around 90% detection rate is achieved in benign scenarios.

Table VII lists the attack success rates of the HFR attack for different speeds and the detection rates in the benign traces. As shown, the HFR attack with our MVS system can hide the victim pedestrian with almost 100% attack success rates until the vehicle is 20 meters away from the pedestrian. We do not observe evident performance degradation due to the higher vehicle speeds, and higher speeds actually yield even greater attack success rates. This result indicates that our MVS system has sufficiently high targeting precision to effectively aim at a fast-driving vehicle.

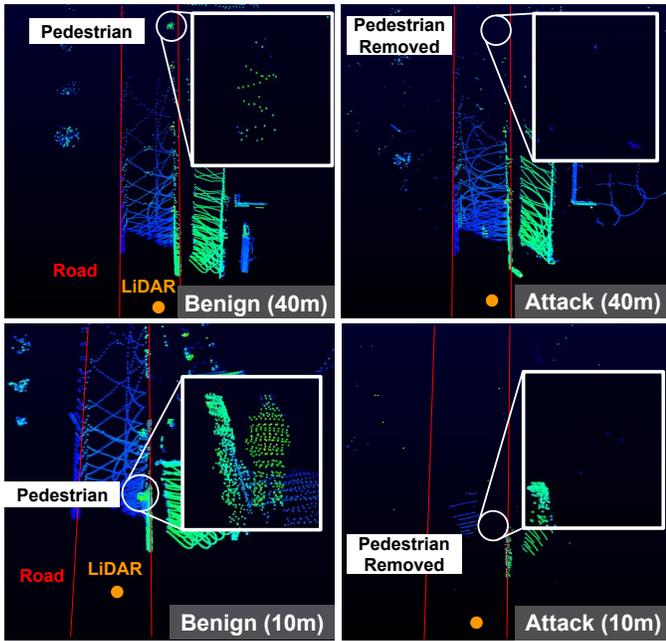


Figure 14: An example of LiDAR point clouds visible from a vehicle moving at 60 km/h. The HFR attack with our MVS system successfully eliminated a pedestrian 40 m away, and it is possible to continue to remove it until it approaches 10 m.

Table VII. Attack success rates of the HFR attack against Livox Horizon for different speeds. The “benign” row lists the detection success rates of the target pedestrian without attacks.

	Distance between victim and target (m)						
	0-10	10-20	20*-30	30-40	40-50	50-60	60-70
10 km/h	59%	83%	97%	100%	100%	100%	100%
20 km/h	65%	76%	99%	100%	100%	100%	100%
30 km/h	67%	78%	100%	100%	100%	100%	100%
40 km/h	71%	76%	100%	100%	100%	100%	100%
50 km/h	93%	97%	100%	100%	100%	100%	100%
60 km/h	85%	85%	96%	100%	100%	100%	100%
Benign	50%	100%	100%	100%	89%	94%	76%

*: the braking distance at 60 km/h

Table VIII lists the point removal rates of the HFR attack. We calculate the rate by dividing the number of points belonging to the victim pedestrian by the number of points in the corresponding benign frames at the same location. We manually annotate the area of the victim for all frames. As listed, the HFR attack succeeds in removing the majority of the victim’s points, but it appears that $\geq 95\%$ of the points need to be removed to completely fool the object detector. Fig. 14 shows examples of the LiDAR point clouds for both benign and attack scenarios at 10 and 40 meters away from the victim pedestrian. As shown, the majority of the victim’s points have been removed by the HFR attack.

Notably, the attack success until 20 meters away has a significant impact on the safety implications considering that the braking distance at 60 km/h is 20 meters without the reaction distance [44]. To more rigorously estimate the impact on safety, we will conduct a closed-loop end-to-end evaluation with a popular AD stack in §VII. We also discuss possible countermeasures in §VIII-2.

Table VIII. Point removal rates of the HFR attack against Livox Horizon for different speeds by comparing the points in the benign scenarios at the same distance.

	Distance between victim and target (m)						
	0-10	10-20	20*-30	30-40	40-50	50-60	60-70
10 km/h	90%	89%	99%	100%	100%	100%	100%
20 km/h	79%	90%	98%	100%	100%	100%	100%
30 km/h	75%	92%	100%	100%	100%	99%	95%
40 km/h	75%	90%	97%	100%	100%	100%	100%
50 km/h	87%	97%	100%	100%	100%	100%	93%
60 km/h	87%	93%	97%	100%	100%	100%	97%

*: the braking distance at 60 km/h

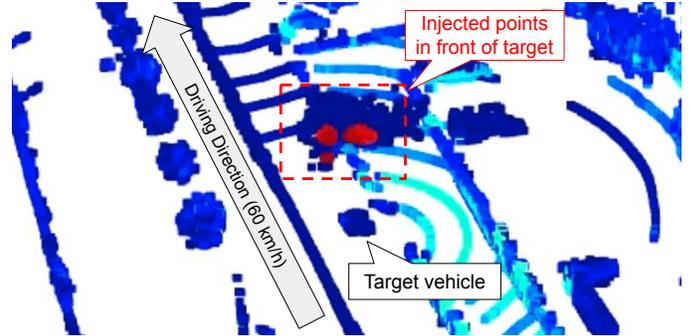


Figure 15: Point cloud under the synchronized injection attack in the high-speed driving scenario at 60 km/h.

2) *A-HFR Attack at High Speed*: We evaluate the A-HFR attack against a high-speed driving vehicle with AT128 [19], which is one of the New-Gen LiDARs with pulse fingerprinting. The attack frequency is set at 15 MHz. For the object detector, we use the PointPillars used in Apollo 6.0 [41]. Table IX lists the attack success rates of the A-HFR attack for 6 different speeds. As listed, the A-HFR attack is quite successful at distances of up to 40 meters away from the victim, with attack success rates of 100% at 60 km/h. Similar to the observations in the HFR experiment, there was no strong correlation between the vehicle’s speed and the attack success rate, indicating that accurate tracking was achieved. However, the attack success rate drops to 50% when the vehicle is 20 meters away. This degradation is mainly due to the occasional failure to receive the laser from the LiDAR in the optical receiver during certain LiDAR scans. Currently, for each LiDAR scan, the A-HFR attack is triggered only if the laser is received by the PD; consequently, the attack inevitably fails for scans in which the LiDAR laser is undetected by the optical receiver. Even with a parabolic antenna, the laser light from a LiDAR mounted on a distant vehicle follows a sparse pattern, making it prone to occasional misses by the receiver. This limitation can be handled by two approaches: improving the optical receiver sensitivity or implementing speculative execution. Further details are provided in §VIII-5.

C. Ablation Study on Hardware Configuration of MVS Systems

We evaluate the validity of our MVS system design introduced in §III in practical high-speed and long-range attack scenarios. Table X lists the point removal rates of the HFR attack for different numbers of beams and sizes of lenses in the scenario with a 60 km/h driving speed. As shown, our setup with 2 beams and 2-inch lenses achieves $\geq 97\%$ attack success rates until 20 meters, which corresponds to the braking

Table IX. Attack success rates of the A-HFR attack against AT128 for different speeds. The “benign” row lists the detection success rates of the target pedestrian without attacks.

	Distance between victim and target (m)						
	0-10	10-20	20*-30	30-40	40-50	50-60	60-70
10 km/h	42%	61%	46%	84%	95%	100%	100%
20 km/h	37%	70%	42%	80%	94%	100%	100%
30 km/h	44%	67%	45%	86%	97%	100%	100%
40 km/h	44%	66%	40%	80%	91%	100%	100%
50 km/h	34%	61%	50%	73%	81%	100%	100%
60 km/h	27%	72%	50%	87%	100%	100%	100%
Benign	32%	50%	49%	53%	70%	100%	100%

*: the braking distance at 60 km/h

Table X. Point removal rates of the HFR attack against the target vehicle driving at 60 km/h for the different numbers of beams and sizes of lenses in the MVS system.

# of beams	Lens size	Distance between Pedestrian & LiDAR (m)						
		0-10	10-20	20-30	30-40	40-50	50-60	60-70
1	1 inch	49%	63%	61%	76%	88%	94%	99%
1	2 inches	49%	43%	63%	69%	89%	100%	100%
2	2 inches	87%	93%	97%	100%	100%	100%	97%

Table XI. Averaged number of points and their angles injected by the synchronized attack against VLP-32c for each 10-meter bin in the 60 km/h scenario.

# of injected points	Distance between injected wall and target (m)					
	0-10	10-20	20-30	30-40	40-50	50-60
Point Range	4.7	6.6	10.4	10.2	8.1	3.1

distance at 60 km/h [44]. We consider that the beam size of roughly 4 inches x 4 inches proves a sufficiently large total beam size to cover inaccuracies in the current MVS system, as the improvement observed when doubling the lens size is not as significant as the improvement achieved by doubling both the lens size and the number of beams.

D. Injection Attack Evaluation

We evaluate the synchronized injection attack in the high-speed driving scenarios against VLP-32c [36] with a deterministic scan pattern, which is a requirement for the white-box attack. We inject a ghost wall in front of the target vehicle at a fixed position in the world coordinate by following the methodology in [10] as shown in Fig. 15. Table XI lists the average number of points and their angles injected by the synchronized attack against VLP-32c for each 10-meter bin from the injected wall to the target vehicle in the 60 km/h scenario. As shown, the attack can inject $\geq 1k$ points at around 20 meters away, which is already the braking distance at 60 km/h [44], meaning that hard braking is inevitable to avoid the injected wall. Unlike the removal attack, the injection attack may cause serious safety consequences (e.g., hard braking) even if the attack is only successful for a few frames. We will evaluate the system-level impact of the injection attack in the next section (§VII).

VII. END-TO-END EVALUATION ON AD VEHICLE

We finally evaluate the attack effectiveness of LiDAR spoofing attacks with our MVS system against an actual AD vehicle. We deploy the attacks against Autoware.ai [15], a

popular open-sourced AD stack, installed on PIXKIT [45], an autonomous driving development platform.

1) *Experimental Setup*: Fig. 16 (b) shows the experimental setup for the end-to-end attack evaluation on an AD vehicle, PIXKIT [45] with a VLP-32c [36] LiDAR controlled by Autoware.ai version 1.14.1 [15]. For the LiDAR object detection, we use the LiDAR Euclidean cluster detection, which is officially supported by Autoware.ai. We placed our MVS system on the roadside of our private road with exclusive use permission. The AD vehicle started accelerating 40 meters away from our MVS system. For the removal attack, we placed an SUV-sized inflatable mock car 5 meters away from the MVS system and tried to make the mock car undetected by the HFR attack. For the injection attack, we injected a ghost wall as shown in Fig. 17 to see if the AD vehicle would stop before the ghost wall. We accelerate the AD vehicle up to 5 km/h and 15 km/h for removal and injection attacks, respectively. The 15 km/h for injection attacks is the maximum speed we could test on the testing road. The 5 km/h for removal attacks is an acceptable speed to avoid damage to the facilities. To ensure safety, the experiments were conducted only during nighttime. Demo videos and perception results are available on <https://sites.google.com/keio.jp/keio-csg/projects/AttackOnDrivingVehicle> (Keio University side) and <https://sites.google.com/view/av-loat-sec/real-av-lidar-attack> (UCI side)².

2) *Removal Attack Results*: As shown in Fig. 16 (a) the AD vehicle adequately stopped 7 meters before the front mock vehicle in the benign scenario. The “Follow” word on the AD vehicle indicates the current driving mode, meaning that the AD vehicle is following a front object and must stop before it. On the other hand, the AD vehicle failed to detect the mock car and crashed into it in the attack scenario as shown in Fig. 16 (c), as evidenced by the driving mode “Foward”, meaning the AD vehicle is decided to drive forward. These results directly support the end-to-end safety impacts of object removal effects of LiDAR spoofing attacks; the HFR attack with our MVS system has a high potential to cause a hard crash into objects on roads, such as cars and pedestrians, in the real world.

3) *Injection Attack Results*: Fig. 17 shows the results of the synchronized injection attack. In the benign scenario (the left figure), the AD vehicle keeps driving forward at 15 km/h. In the attack scenario (the right figure), the AD vehicle permanently stops before the injected ghost wall during the attack. These results indicate that the injection attack can make an AD vehicle stop on a road and thus has a high potential to cause severe safety and transportation concerns such as hard braking and traffic jams.

VIII. DISCUSSIONS AND LIMITATIONS

1) *Safety Implications*: The safety implications (e.g., hard crash into a parked car and hard braking) demonstrated in §VII are very likely to be feasible for AD vehicles driving at fast speeds, with our MVS system. Although we could not evaluate the attack on a high-speed vehicle actually controlled by an AD stack due to the limitations in our testing facility, the attack effectiveness in the high-speed driving scenarios has been addressed in §VI, which shows that the HFR attack with

²We have two websites for publicizing purposes from all collaborating institutions; their contents are jointly developed and identical.

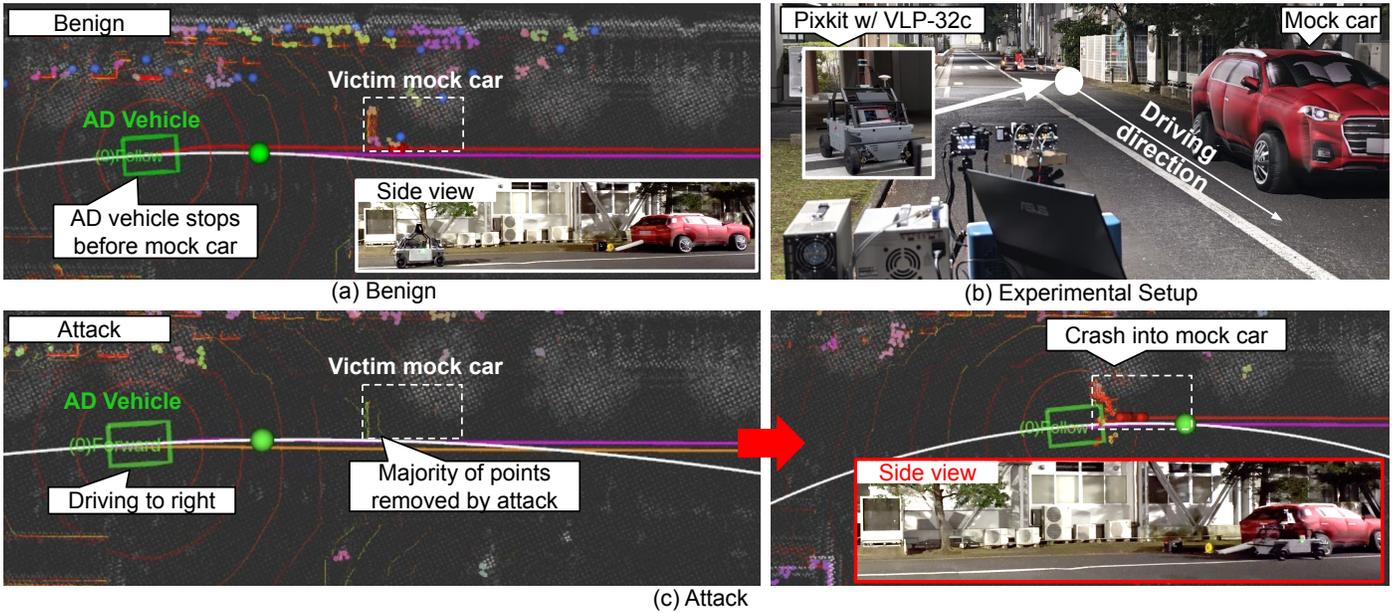


Figure 16: Experimental setup and removal attack results on AD vehicle. We deploy the HFR attack with our MVS system against PIXKIT with VLP-32c controlled by Autoware.ai. PIXKIT drives from 40 meters away from the MVS system.

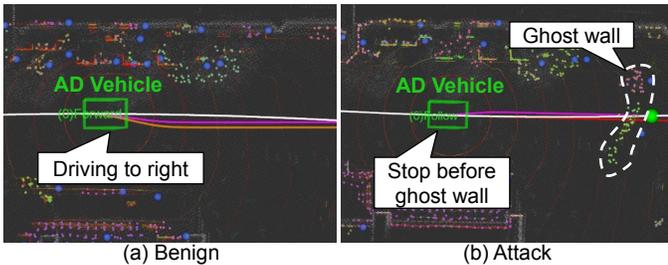


Figure 17: Injection attack results on AD vehicle. We deploy the synchronized injection attack with our MVS system against the PIXKIT with VLP-32c controlled by Autoware.ai.

our MVS system can achieve $\geq 96\%$ attack success rates at 60 km/h until the braking distance, 20 meters away from the victim. We believe that compared to prior works, this work is able to much more clearly show the safety implications of LiDAR spoofing attacks against AD vehicles in real-world scenarios. Although our MVS system is costly, we so far do not see major technical challenges to launch LiDAR spoofing attacks against real AD vehicles on public roads with our MVS system. Different from rock throwing or gun shooting, as LiDAR spoofing attacks can directly attack the AD technology, adversaries may implement the MVS system to damage the reputation and reliability of AD services. However, private AD vehicles such as Waymo robotaxi may have their original countermeasures. Thus, before the publication of this paper, we conducted a responsible vulnerability disclosure for popular Level-4 and Level-3 AD companies using LiDARs due to the potential high security and safety impacts of this work on the AD industry.

2) *Possible Countermeasures*: We consider that the use of New-Gen LiDARs is mandatory to be more robust against LiDAR spoofing attacks, as also mentioned in [11]. For example, New-Gen LiDARs with timing randomization can effectively defend against synchronized injection attacks. For removal attacks, pulse fingerprinting shows high attack mitiga-

tion capability. While the A-HFR attack can potentially bypass it, this attack increases the hardware requirements of the MVS system. However, as discussed in §II-B, timing randomization and pulse fingerprinting are employed originally for anti-interference, not for security purposes. We encourage LiDAR manufacturers to enhance the magnitude of randomization and fingerprinting as high as possible to make them work as security features. We also strongly recommend that AD companies develop availability-check and fail-safe methods. For example, Tesla shows an alert when all three front cameras do not correctly work [46]. Particularly, the HFR and A-HFR attacks work like “jamming”. It should not be hard to detect their distinctive randomized point pattern as in Fig. 11 and also shown in [11], e.g., by an entropy-based method [47]. For fail-safe, adequate recovery is highly scenario-dependent and an open problem, but even simple measures may mitigate the fatality. For example, slowing down is likely to be effective as it is advised even for human drivers under uncertainty.

3) *Multi-Sensor Fusion*: Multi-sensor fusion could be an effective mitigation strategy since the current MVS systems are designed to track only one LiDAR sensor on a vehicle. However, it is not technically difficult to attack multiple LiDARs by driving multiple MVS systems simultaneously for the HFR or A-HFR attacks. Furthermore, LiDARs struggle with scanning duplicated areas due to interference, meaning each area is likely scanned by only a single LiDAR. The MVS system thus may not need to attack multiple LiDARs even if the target AD has multiple LiDARs. Fusion with different types of sensors such as cameras, radars, and ultrasonic sensors could be a mitigation strategy. However, the current AD vehicles, especially Level-4 AD, heavily rely on LiDARs as our attack works on the actual AD stack as demonstrated in §VII. Sensor fusion generally improves robustness, but it is currently far from being sufficient as a defense. Prior research has also identified a wide variety of vulnerabilities in these sensors, cameras [48, 49], radars [50, 51], ultrasonic sensors [51]. Furthermore, more sensors result in increased costs and may

even open new attack channels.

4) *Attack Deployment on Uneven Roads:* Our current threat model only assumes a flat and straight road as an attack site, i.e., the target LiDAR height is constant and preknown, and thus the auto-aiming is only designed for the horizontal direction. Thus, the current MVS system has a limitation in vertical tracking. Meanwhile, we note that the majority of high-speed scenarios should be only on flat and straight roads because the vehicle cannot accelerate on uneven and winding roads. In low-speed scenarios, tracking both for horizontal and vertical directions should not be challenging as demonstrated in prior work [12, 14]. Furthermore, our arrayed laser design (§III-D) can even eliminate the need for vertical tracking by vertically arraying more lasers. However, large number of arrayed lasers may harm the portability of the MVS system, and thus the attacker needs to decide the number of lasers based on the tracking performance in the target attack scenarios.

5) *Speculative Execution for A-HFR Attack:* As discussed in §VI-B2, we find that the current optical receiver often fails to receive lasers at many frames and thus fails to launch the attack. There are two main approaches to address this issue. The first is to improve the optical receiver by either enlarging the antenna to receive more signals or integrating it into the auto-aiming system to better align with the LiDAR. Another approach is to implement speculative execution of the attack at each frame based on the attack timing from the previous frame. Currently, we trigger the A-HFR attack only when the laser is received at each frame, meaning the attack fails if the laser is not received. To address this, we could start the attack even if the laser is not received, based on the interval patterns of previous frames. Since the interval between LiDAR frames is periodic, such speculative execution should improve the A-HFR attack. This might result in shorter rest periods and more overheating, but it should be still more preferable than failing to launch the attack. This type of speculative execution should be performed in real-time on an FPGA.

6) *Ethical and Safety Considerations:* The experiments were safely carried out in controlled conditions on a private road with exclusive-use permission. A human with a driving license drove the experimental vehicle, and the area was surveilled to keep people off the road. During the experiments, participants potentially exposed to the attack laser wore protective goggles for eye safety. While the IRB process is not mandatory for our study since we do not analyze data derived from human subjects, we have strictly followed the advice of the authorities at our testing site to ensure the safety of all individuals involved in the experiments, including human drivers and the target pedestrians standing at the roadside.

IX. CONCLUSION

In this study, we investigate the safety and security impact of LiDAR spoofing attacks against AD vehicles driving at high speeds in practical long-range attack scenarios. We first identify 3 research limitations in prior work to prevent them from deploying their LiDAR spoofing attacks in practical AD scenarios. To address the limitations, we design a novel MVS system that can detect, track, and aim at the target LiDAR moving at high speeds, and also design a new practical removal attack, A-HFR attack, which can be effective against New-Gen LiDARs even with pulse fingerprinting. We demonstrate

that our designs in the MVS system can achieve significantly longer attack distances (e.g., 110 meters) in a LiDAR-agnostic way. We demonstrated that both injection and removal attacks can be deployed against high-speed vehicles driving at 60 km/h, and the attacks can have end-to-end attack impacts on a popular AD stack. This study bridges critical gaps between LiDAR security and AD security research. We hope that our study can give a new perspective in this research space to more adequately understand the safety and security implications of LiDAR in AD perception.

ACKNOWLEDGEMENTS

This research was supported in part by the NSF CNS-2145493, CNS-1929771, USDOT under Grant 69A3552348327 for the CARMEN+ University Transportation Center, JST CREST JPMJCR23M4, JST PRESTO JPMJPR22PA, JSPS KAKENHI 24K02940, and Amano Institute of Technology.

REFERENCES

- [1] “RoboSense Releases The Latest Version Of Its Sensor Evaluation System For Autonomous Driving,” <https://www.robosense.ai/en/news-show-1473/>, 2021.
- [2] “LiDAR: New “Eye” for Vehicle Autonomy,” <https://www.mobilityengineeringtech.com/component/content/article/43672-sae-ma-02983>, 2018.
- [3] “Waymo Has Launched its Commercial Self-Driving Service in Phoenix,” <https://www.businessinsider.com/waymo-one-driverless-car-service-launches-in-phoenix-arizona-2018-12>, 2018.
- [4] “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles,” https://www.sae.org/standards/content/j3016_202104/, 2021.
- [5] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, “Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and Lidar,” *Black Hat Europe*, vol. 11, p. 2015, 2015.
- [6] H. Shin, D. Kim, Y. Kwon, and Y. Kim, “Illusion and Dazzle: Adversarial Optical Channel Exploits Against Lidars for Automotive Applications,” in *International Conference on Cryptographic Hardware and Embedded Systems*. Springer, 2017, pp. 445–467.
- [7] Y. Cao, C. Xiao, B. Cyr, Y. Zhou, W. Park, S. Rampazzi, Q. A. Chen, K. Fu, and Z. M. Mao, “Adversarial Sensor Attack on Lidar-Based Perception in Autonomous Driving,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019, pp. 2267–2281.
- [8] J. Sun, Y. Cao, Q. A. Chen, and Z. M. Mao, “Towards Robust LiDAR-based Perception in Autonomous Driving: General Black-box Adversarial Sensor Attack and Countermeasures,” in *USENIX Security Symposium*, 2020.
- [9] R. S. Hallyburton, Y. Liu, Y. Cao, Z. M. Mao, and M. Pajic, “Security Analysis of Camera-LiDAR Fusion Against Black-Box Attacks on Autonomous Vehicles,” in *USENIX Security Symposium*, 2022.
- [10] Z. Jin, J. Xiaoyu, Y. Cheng, B. Yang, C. Yan, and W. Xu, “PLA-LiDAR: Physical Laser Attacks against LiDAR-based 3D Object Detection in Autonomous Vehicle,” in *IEEE Symposium on Security and Privacy (SP)*, 2023, pp. 710–727.
- [11] T. Sato, Y. Hayakawa, R. Suzuki, Y. Shiiki, K. Yoshioka, and Q. A. Chen, “LiDAR Spoofing Meets the New-Gen: Capability Improvements, Broken Assumptions, and New Attack Strategies,” in *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2024.
- [12] Y. Cao, S. H. Bhupathiraju, P. Naghavi, T. Sugawara, Z. M. Mao, and S. Rampazzi, “You Can’t See Me: Physical Removal Attacks on LiDAR-based Autonomous Vehicles Driving Frameworks,” in *USENIX Security Symposium*, 2023.
- [13] Y. Cao, N. Wang, C. Xiao, D. Yang, J. Fang, R. Yang, Q. A. Chen, M. Liu, and B. Li, “Invisible for Both Camera and LiDAR: Security of Multi-Sensor Fusion based Perception in Autonomous Driving under Physical-World Attacks,” in *IEEE Symposium on Security and Privacy (SP)*, 2021, pp. 176–194.

- [14] Y. Cao, J. Ma, K. Fu, R. Sara, and M. Mao, "Automated Tracking System for LiDAR Spoofing Attacks on Moving Targets," in *ISOC NDSS Workshop on Automotive and Autonomous Vehicle Security (AutoSec)*, 2021, p. 1.
- [15] S. Kato, S. Tokunaga, Y. Maruyama, S. Maeda, M. Hirabayashi, Y. Kitsukawa, A. Monroy, T. Ando, Y. Fujii, and T. Azumi, "Autoware On Board: Enabling Autonomous Vehicles with Embedded Systems," in *ICCPs'18*. IEEE Press, 2018, pp. 287–296.
- [16] M. Yu, M. Shi, W. Hu, and L. Yi, "FPGA-based Dual-pulse Anti-interference Lidar System Using Digital Chaotic Pulse Position Modulation," *IEEE Photonics Technology Letters*, 2021.
- [17] "Livox Mid-360," <https://www.livoxtech.com/mid-360>.
- [18] "XT32 — Mid-Range Mechanical Lidar — HESAI Technology," <https://www.hesaitech.com/product/xt32/>.
- [19] "AT128 Auto-Grade Ultra-High Resolution Long Range Lidar — HESAI Technology," <https://www.hesaitech.com/product/at128/>.
- [20] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [21] "PhantomX Robot Turret Kit," <https://www.trossenrobotics.com/>.
- [22] Y. Hayakawa, T. Sato, R. Suzuki, K. Ikeda, O. Sako, R. Nagata, Q. A. Chen, and K. Yoshioka, "WIP: An Adaptive High Frequency Removal Attack to Bypass Pulse Fingerprinting in New-Gen LiDARs," in *ISOC Symposium on Vehicle Security and Privacy (VehicleSec)*, 2024.
- [23] R. Suzuki, T. Sato, Y. Hayakawa, K. Ikeda, O. Sako, R. Nagata, Q. A. Chen, and K. Yoshioka, "WIP: Towards Practical LiDAR Spoofing Attack against Vehicles Driving at Cruising Speeds," in *ISOC Symposium on Vehicle Security and Privacy (VehicleSec)*, 2024.
- [24] S. Hari, Babu, L. Y.B., S. Ram, and K. Ashok, "Airborne Infrared Search and Track Systems," *Defense Science Journal*, vol. 57, no. 5, pp. 739–753, 2007.
- [25] G. Jocher, "YOLOv5," <https://github.com/ultralytics/yolov5>, 2020.
- [26] "ELP-USBFHD08S-MFV with 10x Optical Zoom," <https://www.elpctv.com/20-megapixel-high-speed-10x-zoom-550mm-webcam-usb-cam-era-for-live-driving-teaching-system-p-281.html>, 2022.
- [27] "FBH905-10," <https://www.thorlabs.com.jp/thorproduct.cfm?partnumber=FBH905-10>, 2022.
- [28] "LiDAR for Automotive and Industrial Applications 2021," <https://www.i-micronews.com/products/lidar-for-automotive-and-industrial-applications-2021/>.
- [29] "VLP-16 User Manual," <https://velodynelidar.com/wp-content/uploads/2019/12/63-9243-Rev-E-VLP-16-User-Manual.pdf>.
- [30] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, 1960.
- [31] G. Chang, "Robust Kalman Filtering based on Mahalanobis Distance as Outlier Judging Criterion," *Journal of Geodesy*, 2014.
- [32] "Lens Characteristics: Flare, Ghosting and Aberration," <https://av.jpn.support.panasonic.com/support/global/cs/dsc/knowhow/knowhow15.html>, 2024.
- [33] "Antenna Design for Space," <https://www.rantecantennas.com/blog/antenna-design-for-space/>, 2024.
- [34] "Analog Discovery 3: 125 MS/s USB Oscilloscope, Waveform Generator, Logic Analyzer, and Variable Power Supply," <https://digilent.com/shop/analog-discovery-3/>, 2016.
- [35] "SPL PL90 3," https://look.ams-osram.com/m/249c8a382f5faab2/original/SPL-PL90_3.pdf, 2018.
- [36] "Ultra Puck Surround View Lidar Sensor — Velodyne Lidar," <https://velodynelidar.com/products/ultra-puck/>.
- [37] M. S. Grewal and A. P. Andrews, "Applications of Kalman Filtering in Aerospace 1960 to the Present [Historical Perspectives]," *IEEE Control Systems Magazine*, vol. 30, no. 3, pp. 69–78, 2010.
- [38] E. A. Wan and R. Van Der Merwe, "The Unscented Kalman Filter for Nonlinear Estimation," in *Proceedings of the IEEE 2000 adaptive systems for signal processing, communications, and control symposium (Cat. No. 00EX373)*. Ieee, 2000, pp. 153–158.
- [39] P. Del Moral, "Nonlinear Filtering: Interacting Particle Resolution," *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics*, vol. 325, no. 6, pp. 653–658, 1997.
- [40] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision Meets Robotics: The KITTI Dataset," *International Journal of Robotics Research*, 2013.
- [41] "Baidu Apollo," <https://github.com/ApolloAuto/apollo>.
- [42] "Livox Horizon User Manual," <https://www.livoxtech.com/3296f540ecf5458a8829e01cf429798e/assets/horizon/Livox%20Horizon%20user%20manual%20v1.0.pdf>.
- [43] "Livox Detection V2," https://github.com/Livox-SDK/livox_detection.
- [44] "Stopping Distances on Wet and Dry roads, Queensland Government," <https://www.qld.gov.au/transport/safety/road-safety/driving-safely/stop-ping-distances/graph>.
- [45] "PIXXIT: An Autonomous Driving Development and Education Kit," <https://www.pixmoving.com/pixkit>, 2021.
- [46] W. Wang, Y. Yao, X. Liu, X. Li, P. Hao, and T. Zhu, "I Can See the Light: Attacks on Autonomous Vehicles Using Invisible Lights," in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2021, pp. 1930–1944.
- [47] Y. Xingwei, L. Dawei, Z. Jun, and W. Jianwei, "Radar Jamming Detection based on Approximate Entropy and Moving-cut Approximate Entropy," 2012.
- [48] T. Sato, S. H. Bhupathiraju, M. Clifford, T. Sugawara, Q. A. Chen, and S. Rampazzi, "Invisible Reflections: Leveraging Infrared Laser Reflections to Target Traffic Sign Perception," in *Network and Distributed System Security Symposium (NDSS)*, 2024.
- [49] Y. Zhao, H. Zhu, R. Liang, Q. Shen, S. Zhang, and K. Chen, "Seeing isn't Believing: Practical Adversarial Attack Against Object Detectors," in *ACM SIGSAC Conference on Computer and Communications Security (ACM CCS)*, 2019, p. 1989–2004.
- [50] N. Miura, T. Machida, K. Matsuda, M. Nagata, S. Nashimoto, and D. Suzuki, "A Low-Cost Replica-based Distance-Spoofing Attack on Mmwave Fmcw Radar," in *Proceedings of the 3rd ACM Workshop on Attacks and Solutions in Hardware Security Workshop*, 2019.
- [51] C. Yan, W. Xu, and J. Liu, "Can You Trust Autonomous Vehicles: Contactless Attacks Against Sensors of Self-Driving Vehicle," *Def Con*, vol. 24, no. 8, p. 109, 2016.

APPENDIX

A. Impact of Vertical and Horizontal Attack Angles

We evaluate the laser pulse peak power at different attack frequencies to evaluate the improvements of the A-HFR attack over the HFR attack. For the A-HFR attack, we evaluate two different vertical attack angles: 32° ($v=100\%$) and 16° ($v=50\%$). Fig. 18 shows the laser pulse peak power at different attack frequencies. We calculate the average power of the attack laser with a laser power meter and then divide this value by the laser frequency to obtain the power per pulse. As designed, the A-HFR attack can achieve higher peak powers at the same attack frequencies compared to the HFR attack. If we limit more vertical attack angles, the A-HFR can achieve higher attack capability by reducing the vertical attack angle as well. The A-HFR attack between 10 and 20 MHz can achieve substantially higher power, with a 10 times increase. This attack angle reduction approach of the A-HFR is necessary to achieve high frequencies such as 25 MHz.

Fig. 19 shows the efficacy of A-HFR's vertical attack angle reduction. Since the angle reduction boosts the peak power under high attack frequencies (Fig. 18), this expands the point cloud removal angle by about 20%.

B. HFR Attack on Mid-360

Fig. 20 shows the efficacy of the conventional HFR attack on Mid-360, equipped with pulse fingerprinting. We found that even the HFR attack can fully remove points in 20° horizontally and vertically. The Mid-360's vulnerability can

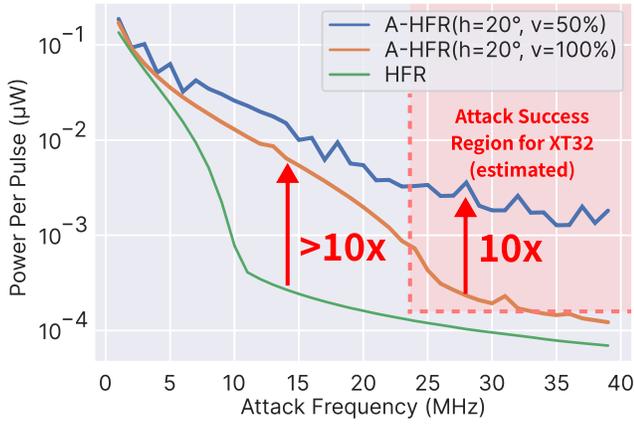


Figure 18: Laser pulse peak power of the HFR and A-HFR attacks at different attack frequencies. The A-HFR attack can achieve higher laser power at the same frequency.

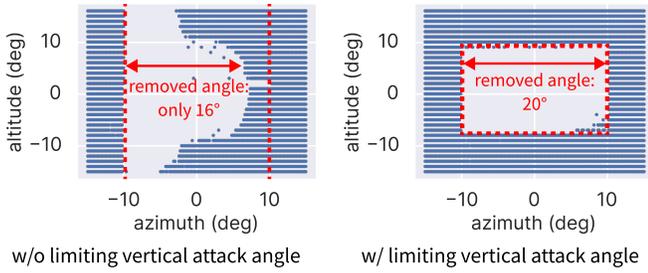


Figure 19: 2D projection of point cloud of the A-HFR attack on XT32 with horizontal 20°. Blue points mean the remaining points, and the red dotted line indicates the attack angle.

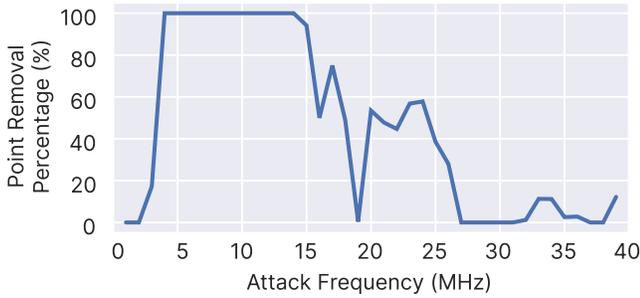


Figure 20: Point removal percentage by HFR attack against Livox Mid-360, under the same condition as Fig. 10.

be attributed to its relatively less secure authentication system compared to other LiDARs. This suggests that a larger T_α results in potentially lower power in legitimate pulses, indicating that certain LiDAR models might have more lenient T_α settings and hence weaker security features.