WIP: Towards Practical LiDAR Spoofing Attack against Vehicles Driving at Cruising Speeds

Ryo Suzuki[†], Takami Sato[‡], Yuki Hayakawa[†], Kazuma Ikeda[†], Ozora Sako[†], Rokuto Nagata[†],
Qi Alfred Chen[‡], and Kentaro Yoshioka[†]
[†]Keio University, Department of Electronics and Electrical Engineering
[‡]University of California, Irvine, Department of Computer Science

***Abstract*—LiDAR (Light Detection and Ranging) is an essential sensor for autonomous driving (AD), increasingly being integrated not only in prototype vehicles but also in commodity vehicles. Due to its critical safety implications, recent studies have explored its security risks and exposed the potential vulnerability against LiDAR spoofing attacks, which manipulate measurement data by emitting malicious lasers into the LiDAR. Nevertheless, deploying LiDAR spoofing attacks against driving AD vehicles still has significant technical challenges particularly in accurately aiming at the LiDAR of a moving AV from the roadside. The current state-of-the-art attack can be successful only at ≤5 km/h. Motivated by this, we design novel tracking and aiming methodology and conduct a feasibility study to explore the actual practicality of LiDAR spoofing attacks against AD vehicles at cruising speeds. In this work, we report our initial results demonstrating that our object removal attack successfully makes the targeted pedestrian undetectable with ≥90% success rates in a real-world scenario where the adversary at the roadside attacks the victim AD approaching at 35 km/h. Finally, we discuss the current challenges and our future plans.**

I. Introduction

The research and development of Autonomous driving (AD) are rapidly growing year by year. One of the major drivers of the growth is enabled by LiDAR (Light Detection and Ranging) sensors, especially for over Level 4 AD as defined by the SAE [1]. Particularly, object detection and vehicle localization in AD significantly benefit from the 3D sensing capability of LiDARs. However, recent studies have posed the safety and security risks of LiDAR due to its sensitivity to ambient light noises and malicious laser emission, which is known as LiDAR spoofing attack [2]–[10]. The malicious lasers of LiDAR spoofing attacks can overwrite the legitimate measurements and cause 2 attack effects: object injection attacks [3]–[8] and object removal attacks [3], [7]–[9]. Meanwhile, prior work predominantly focuses on stationary lab-level setups or dynamic but impractical low-speed setups (e.g., at most 5 km/h or 3.1 mph [9]) even though they discuss the safety implications of their attacks against AD systems. Motivated by this, we design a novel tracking and aiming methodology to precisely emit attack lasers against the victim vehicle driving at cruising speeds (e.g., 35 km/h).

In this WIP paper, we report our recent progress on the feasibility study attacking a fast-driving vehicle with our novel tracking and aiming system called Moving Vehicle Spoofing system (MVS system). We demonstrate that our attack can successfully remove a pedestrian from object detection results with ≥90% success rate in the real-world scenario that a victim vehicle is approaching at 35 km/h from 45 m away. This result implies potential serious security risks against AD vehicles already operating on public roads such as robotaxi services [11]. The attack demo is available on our website at **https://sites.google.com/view/mvsa-study/**. Finally, we discuss our findings, limitations, and future plans, particularly the potential gaps between our results and actual safety implications in real-world AD vehicles.

II. Background and Related Works

A. LiDAR Spoofing Attacks

LiDAR spoofing attacks [2]–[10] manipulate the distance measurements of LiDAR sensing by overwriting the legitimate lasers with higher-power malicious lasers. Table I lists an overview of the existing LiDAR spoofing attacks demonstrated in the physical world. The attacks can be categorized into two distinct types, each characterized by unique attack methodologies and objectives.

1) Object Injection Attacks: This type of attack injects ghost objects that do not actually exist. To inject malicious lasers effective against LiDARs, relay attack [2] sends back the recorder lasers to the victim LiDAR. By adding a delay on the lasers, the attacker can move an object to a further location. Synchronized injection attacks [3], [5], [6] first learn and synchronize the scanning pattern of the target LiDAR, and then overwrite whatever LiDAR measurements with malicious lasers based on the obtained pattern. However, synchronized injection attacks heavily rely on the predictability of LiDAR scanning patterns, and thus it is known that laser scan timing randomization, a common feature in recent New-Generation LiDARs (NG-LiDARs), can easily prevent these attacks from injection chosen-pattern point clouds [8]. While the attacker can still inject random points under the timing randomization, it is virtually infeasible to cause a designed attack effect. We thus focus on the object removal attacks in this WIP paper.

2) Object Removal Attacks.: These attacks are designed to make actual objects undetected by object detectors. Synchronized removal attacks [7], [9] remove objects by moving all points of the object very far away or within the area below the minimum distance threshold. These attacks demonstrate that they can remove 4,000 points from the detection [7], [9]. Similar to synchronized injection attacks, this type of

Symposium on Vehicles Security and Privacy (VehicleSec) 2024
26 February 2024, San Diego, CA, USA
ISBN 979-8-9894372-7-6
https://dx.doi.org/10.14722/vehiclesec.2024.23023
www.ndss-symposium.org

TABLE I: Comparison of works that conduct spoofing attacks in the physical world. While most previous studies [2], [3], [5], [6], [8] execute experiments against stationary LiDARs, we evaluate the spoofing attack in a real-traffic-like environment.

	Attack on moving target	Relative Speed	Attack from roadside	Attack types		Maximum attack range
				Injection	Removal	
Ours	✓	$\leq$35 km/h	✓	-	✓	$\sim 45\,\mathrm{m}$
Cao et al. [9]	✓	$\leq$5 km/h	-	-	✓	$\sim 10\,\mathrm{m}$
Cao et al. [12]	✓	$\leq$0.4 km/h	-	✓	-	$\sim 4\,\mathrm{m}$
Jin et al. [7]	✓	0 km/h (running parallel)	-	✓	✓	$\sim 15\,\mathrm{m}$
Petit et al. [2]	-	-	-	✓	-	$\sim 1\,\mathrm{m}$
Shin et al. [3]	-	-	-	-	✓	$\sim 5\,\mathrm{m}$
Sun. et al. [5]	-	-	-	✓	-	$\sim 5\,\mathrm{m}$
Hallyburton et al. [6]	-	-	-	✓	-	$\sim 5\,\mathrm{m}$
Sato et al. [8]	-	-	-	✓	✓	$\sim 10\,\mathrm{m}$

✓ : Conducted, - : Not Conducted

TABLE II: Comparison of spoofers that conduct MVS attacks.

	Number of Lasers	Beam Diameter	Total Beam Area	Maximum Tracking Distance	Tracking	Detection Strategy (target, using device)	Multi-LiDAR Compatibility?
Ours	2	60 cm	$5654.7\,\mathrm{cm}^2$	45 m	Auto	Laser Light with IR Camera	Yes
Cao et al. [9], [12]	1	2.54 cm	$5.1\,\mathrm{cm}^2$	5 m	Auto	LiDAR Device with RGB Camera	No
Jin et al. [7]	1	8cm	$50.3\,\mathrm{cm}^2$	15 m	Manual	-	-

attack can be largely mitigated by timing randomization [8]. There is also another type of attack that does not need synchronization with the LiDAR scanning pattern and thus can be effective against recent NG-LiDARs. Saturating attack [3] projects continuous lasers against LiDAR and saturates the capability of receiving lasers. While it is not realistic to keep shooting high-power laser continuously, it demonstrates that can remove a 41×42 cm^2 metal plate invisible in a short time (e.g., 4 sec). High-frequency removal (HFR) attack [8] is similar to the saturating attack but shows significantly higher attack effectiveness by using a high-frequent pulse laser instead of the continuous laser without learning the LiDAR scanning pattern. The HFR attack can remove the majority of points in a $10 \times 10\ \mathrm{m}^2$ and make 5 sedan cars undetected. Considering the high practicality in the physical world, we start with the HFR attacks to evaluate the feasibility of LiDAR spoofing attacks against cruising vehicles as a first step in this WIP paper.

B. Prior Attempt to Attack Moving Vehicles

To date, there has been no successful demonstration of LiDAR spoofing attacks on AD vehicles driving at operational speeds. We call this type of attack the Moving Vehicle Spoofing attack (MVS attack). Prior attempt [7], [9], [12] claims the potential attack effectiveness of the MVS attacks with digital-space simulations and physical-world experiments. However, there are 2 critical research gaps (RGs) to be effective against AD vehicles driving at cruising speeds:

RG1: Lack of Real-Time and Long-Range Detection and Tracking System. To obtain the accurate location of the target AD, high-performance detection and tracking systems to keep identifying the victim AD at long distances are required. In prior work, Cao et al. [9], [12] show that their device with the camera-based tracking system can track the victim at $\leq$5 km/h within a 5 m distance. Jin et al. [7] demonstrate that their attack is successful against the victim driving at 5 km/h with manual aiming, but their relative speed to the victim is 0 km/, as their attack system ran parallel to the victim vehicle. Thus, none of the prior work has shown clear feasibility whether the attacker can actually keep tracking the fast-moving targets coming from a far distance. A vehicle far away appears very small in a camera frame. To zoom in, an additional aiming system for the camera is necessary, as also discussed as a limitation in [9]. Using a very high-resolution image to avoid zooming might compromise real-time performance due to the increased size of the camera frame. To address the limitations, we design a novel detection and tracking system inspired by a military InfraRed Search and Track Systems (IRST Systems) [13]. Our attacking device detects and tracks the target LiDAR based on its emitted lasers captured by an IR camera.

RG2: Real-World Feasibility of Aiming and Laser Emitting Device. Even if the attacker can track the accurate position of the target AD vehicle, it remains unclear whether they can accurately aim and hit it with malicious lasers. Cao et al. [9], [12] use a generic pan-tilt system [14] to control the laser emitter direction. However, its effectiveness against fast-moving vehicles at long range is uncertain, as it was originally developed for coarse vision tracking and not for precise targeting of distant, small objects like LiDAR sensors. To handle this challenge, we designed a new LiDAR spoofing device with an accurate servo motor and arrayed laser diodes, which can cover a wider area and successfully compensate for inaccuracies in aiming.

III. Methodology

To overcome the 2 critical research gaps in prior attempts, we design a Moving Vehicle Spoofing system (MVS system); which is a novel attacking system with the IR-camera-based detection and tracking system, with an aiming system equipped with high-precision servo motor and arrayed laser diodes.

A. Threat Model

We generally follow the same threat model as in previous works [4]–[6], i.e., the attacker fires malicious lasers from their spoofer to the victim LiDAR. As discussed in §II, we add more specifications to be a more realistic threat model for attacking AD vehicles driving at high speeds. Fig. 1 illustrates the bird's-eye-view of the MVS attack scenario. The victim vehicle is driving on a straight road at cruising speeds (e.g., 35 km/h). The attack spoofer is placed at the roadside and starts attacking the victim from as further as possible (e.g., 50 m).

B. Overview of MVS System

Figure 1 illustrates the our MVS system. We generally follow the same optical components and electronics akin to the ones that previous studies [4], [5], [7]–[9], [12] used in static or low-speed setup. Our major improvements focus on the devices to accurately track and aim fast-moving targets.

1) Real-Time and Long-Range Detection and Tracking System : To localize the target LiDAR, RGB cameras are typically employed to capture and feed images directly into an object detector [9], [12]. However, at distances exceeding 20 m, the LiDAR occupies an extremely small portion of the camera's field of view (FOV), typically just a few pixels. This significantly hinders the object detector's ability to extract meaningful features from the camera image (Figure 2). While incorporating a telephoto lens improves distant object visibility, its narrow FOV when capturing a LiDAR 50 m away poses significant challenges for effective tracking. These limitations render RGB cameras unsuitable for our spoofer systems that require long-range operation.

To address this challenge, we design a novel spoofer system with an infrared (IR) camera. We are 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 [13]. This method offers two key advantages over conventional RGB vision-based tracking systems: (1) **Long-range Perception:** Unlike previous methods that rely on RGB-camera object detectors to estimate LiDAR's position, our system directly captures the emitted IR lasers from the target LiDAR. This enables accurate location detection of the LiDAR even at long-range distances, as shown in Figure 2. (2) **LiDAR Model Agnosticism:** Prior work [9], [12] needs to collect the various images of the target LiDARs to train custom object detectors. In contrast, the MVS system can bypass this step as all LiDARs inherently emit IR lasers. We have confirmed that the same detection model can be used to detect both Velodyne VLP-16 [15] and Livox Horizon [16]. With our novel designs, we find that our MVS system can stably detect the position of LiDAR from as far away as 45 m which is 9 times longer than the demonstrated distance in the current state-of-the-art systems [9].

Furthermore, considering the high cost and limited availability of industrial IR cameras, we opted for a more accessible alternative in our experiments. We modified a commercially available web camera, the Logitech c922 [17], by replacing its low-pass filter with a bandpass filter, Thorlabs FBH905-10 [18], tuned to the LiDAR laser wavelength. Though a cost-effective solution, we proved capable of delivering sufficient performance for attack.

2) Handling Image Flickering: The major challenge of our IR-camera-based detection is that the IR camera cannot always detect the laser trace emitted by the LiDAR, which scans each direction at around 10 Hz [15]. This image flickering prevents the stable tracking of target LiDAR without any countermeasures. To address this, we design an object detection architecture fed multiple frames in the channel dimension. We leverage a common practice in video analysis of combining multiple consecutive frames along the channel dimension and feeding them to a CNN [19]. As the IR camera feeds grayscale images at every frame, we use three consecutive latest frames for the single detection, i.e., the channel size is three. We find that at least one IR laser trace can be measured in three frames. This design significantly improves detection stability even under image flickering. We trained YOLOv5 [20] with our original dataset consisting of 710 frames captured at various distances. This dataset was split into 75% training and 25% validation sets, leading to a mAP (mean average precision) @0.5 of 0.97 on the validation set.

3) Stable and High-Precision Spoofer Design: In the context of long-range spoofing attacks, even minimal detection errors can amplify, resulting in substantial inaccuracies. This issue is further compounded when dealing with rough road surfaces, as these conditions induce vibrations in the vehicle body, leading to potential perturbations of the target LiDAR. For an MVS attack to be effective, the spoofer must demonstrate resilience against such variabilities in the target LiDAR's position and orientation. We address this challenge by expanding the laser irradiation area as much as possible.

Firstly, our MVS system adopts an 'arrayed laser' design, as depicted in Figure 3. This setup consists of a parallel arrangement of two pairs of laser diodes and lenses, significantly enhancing the stability of the MVS attack. These lasers emit light in parallel without overlapping their beams, aiming at spatially independent points. By arraying the lasers in the attack device, we effectively double the laser irradiation area without compromising the intensity of the lasers. Furthermore, we have expanded the aperture of each laser as much as possible. Therefore, we have tuned the optical system to the largest feasible beam diameter, which is 60 cm, ensuring that the performance of the HFR attack is not diminished (see §IV-A for details). With these innovations, we achieved a total beam area that is about 110 times larger than that in prior work [7]. This substantial increase in beam area allows us to compensate for vehicle vibrations during road travel and detection errors. Currently, two lasers are parallelized considering the availability of the components used. However, if the motor power and power supply for rotating the optical system are sufficient, there are no significant limitations to increasing the number of lasers in parallel, and the implementation can be readily accomplished.

IV. Evaluation

A. Optical System Optimization

LiDAR spoofing attacks involve the injection of false point clouds into a LiDAR system by emitting laser light at a higher power than the light normally reflected from the standard laser. However, crafting a robust attack requires navigating a delicate balance: ideally, the laser diameter should be as wide as possible to maximize the irradiation area, but there is an inherent trade-off between the laser's diameter and its power. To identify the optimal laser diameter for our experiments, we first assessed the minimum laser power necessary to execute effective removal attacks. In our experiments, we employed the HFR attack method [8] to remove point clouds from a wall

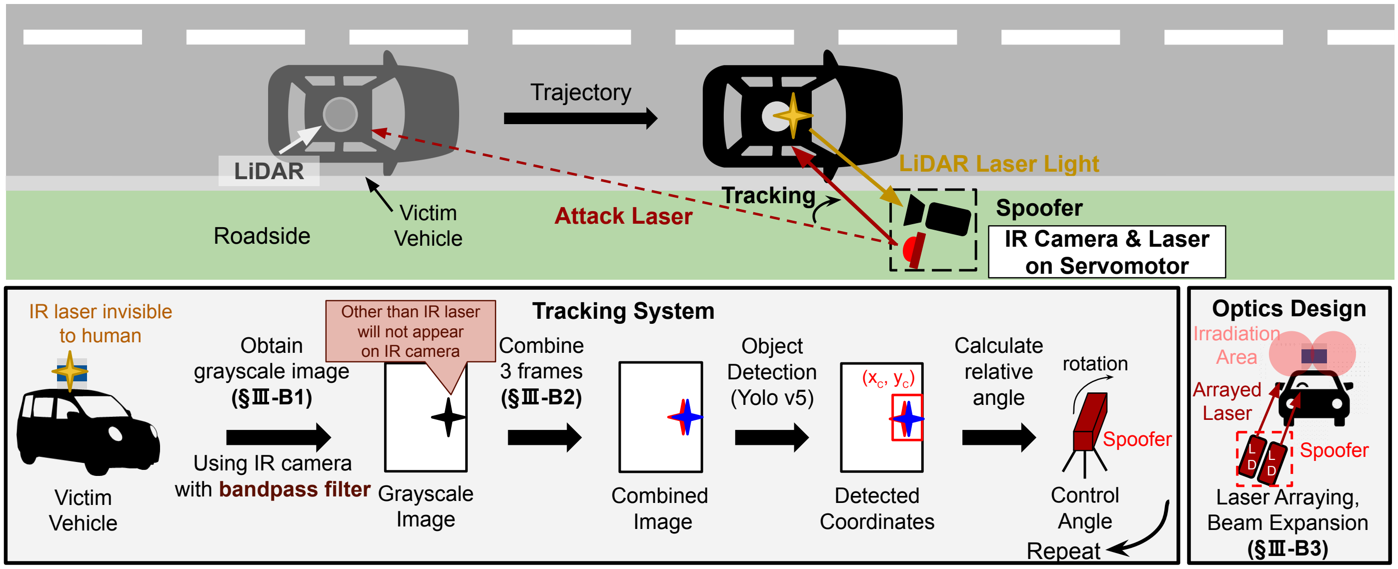

Fig. 1: Overview of our MVS system and improvements on optics design. By using an IR camera and a method of combining consecutive frames along the channel dimension, it is possible to stably track a LiDAR at long distances. In addition, the stability of laser irradiation is enhanced by arraying attack lasers.

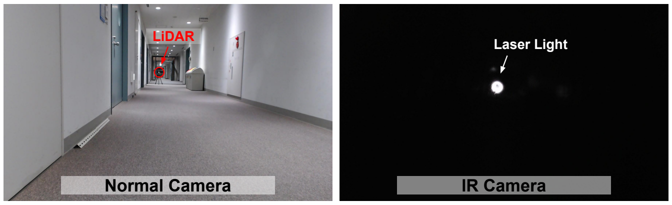

Fig. 2: Visibility difference at long distances (15m) when capturing LiDAR with a regular camera versus an IR camera. Using an IR camera makes it easy to detect LiDAR even at long distances. Moreover, since an IR camera captures only laser light, it is possible to detect it with the same object detection model regardless of the type of LiDAR.

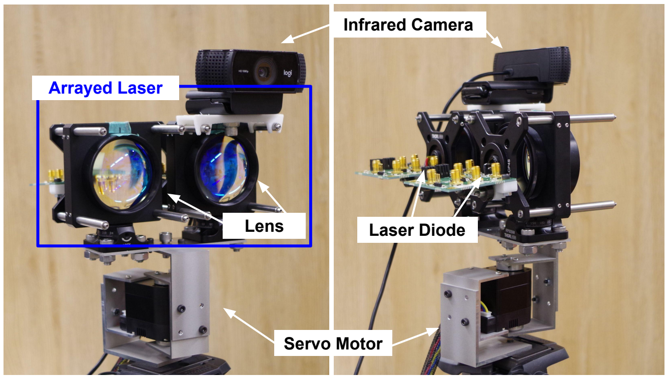

Fig. 3: Overview of our LiDAR spoofer setup.

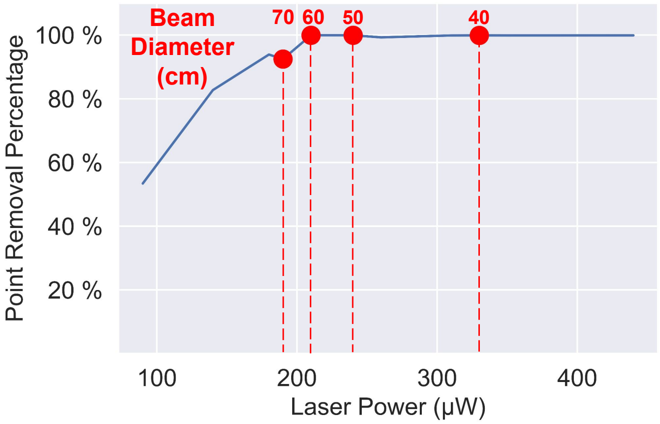

Fig. 4: Attack laser power vs. points removal rate. The red dashed lines show the power at each laser diameter.

located 2 m away, using varying laser powers. The outcomes of these tests are detailed in Figure 4, where we measured the laser power using a detector positioned on the wall. Our findings indicated that as the laser power diminishes, the attack success rate (ASR) also decreases, with a particularly sharp decline in ASR observed when the laser power falls below 200 μW. Since maintaining a high ASR is crucial for successfully deceiving object detectors, we decided to proceed with a laser diameter of 60 cm in subsequent experiments.

B. *Real-World MVS Attack Evaluations*

Experimental Environment. In accordance with the attack scenario outlined in §III-A, we carry out a series of experiments as illustrated in Figure 5. Assuming a roadside attack scenario, we position the spoofer 2 m from the edge of the driving lane. Since many self-driving cars [11], [21], [22] mount LiDAR on the roof of the car, we installed the victim LiDAR in a similar rooftop location. For the attack target, we selected a person standing 5 m in front of the spoofer (termed pedestrian). The vehicle, equipped with the victim LiDAR,

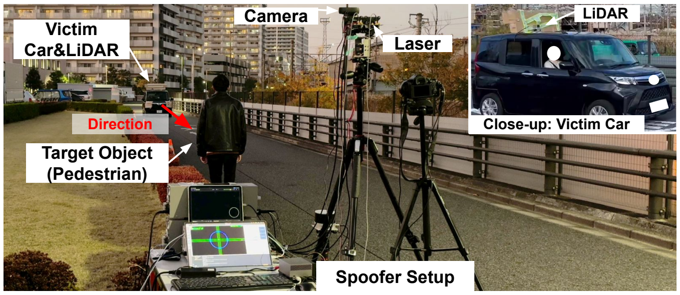

Fig. 5: The setup of our MVS attack experiment.

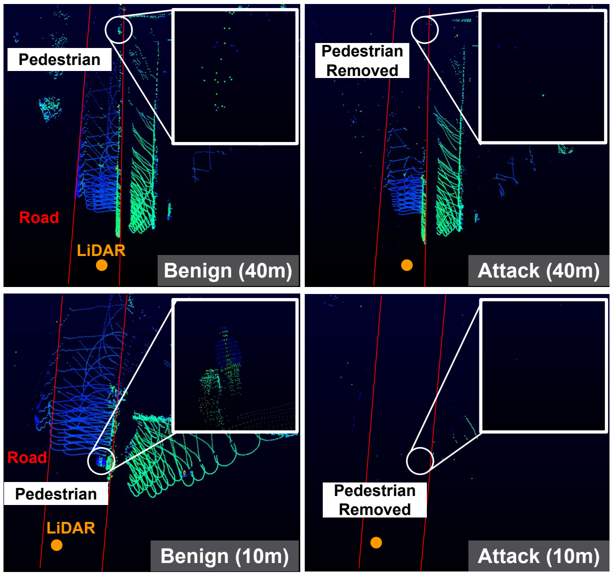

Fig. 6: An example of LiDAR point clouds visible from a vehicle moving at 35 $\mathrm{km/h}$. Our MVS attack successfully eliminated a pedestrian 40 m away, and it is possible to continue to remove it until it approaches 10 m.

commenced its approach from a distance of approximately 40 m from this pedestrian. To ensure safety, the target person was positioned 1 m away from the vehicle's driving path. Moreover, to explore how different driving speeds affect the MVS attack's difficulty, we conducted experiments with the car moving at two speeds: 10 $\mathrm{km/h}$ and 35 $\mathrm{km/h}$.

Attack Methodology. We conducted experiments employing our MVS system described in §III. For the method of attack, we specifically chose HFR attack [8]. HFR attack is notably potent and versatile, capable of targeting a wide range of LiDAR models, including NG-LiDARs. This broad applicability makes the HFR attack particularly threatening, which is why we chose to focus in this WIP paper. In our experiments, we set the HFR laser pulse frequency to 400kHz.

Point Cloud Based Results. Figure 6 shows the point cloud data captured by a victim LiDAR while driving at 35 $\mathrm{km/h}$, both with and without the attack. Notably, the disappearance of the pedestrian's point cloud, initially located 40 m away, evidences the efficacy of our MVS attack upon executing a removal attack. In Table III, we present the removal rate of the point cloud of the target person for each speed during the MVS attack. The percentage of frames in which the point is 100 %, 90 %, or 80 % or more removed during the attack is shown. Remarkably, for distances greater than 10 meters, the removal rate exceeded 90 % in more than 90 % of the frames, indicating a high level of consistency in our attack. Figure 7 shows the relationship between distance and removal rate for each frame during the attack. The stopping distance of cars at 35 $\mathrm{km/h}$ [23] is also shown in the figure. As shown in the figure, the removal rate decreases at short distances. (See Discussion B for details.)

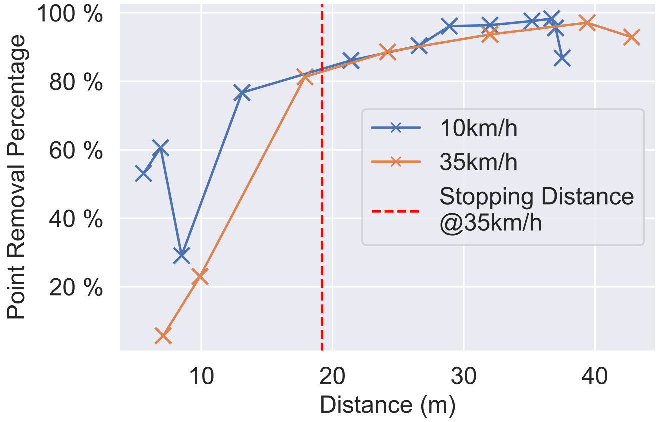

Fig. 7: The relationship between distance and the removed points of the pedestrian. Even at a speed of 35 $\mathrm{km/h}$, the success rate is over 80 % at distances beyond the stopping distance, and there is a risk of causing an accident.

TABLE III: Point-level success rate of our tracking removal attack with different threshold levels of removal percentage(RP). The percentage of frames in which the point is 100 %, 90 %, or 80 % or more removed during the attack is shown. At a distance of more than 10 meters, attacks are successful in more than 90 % of frames with more than 90 % removal rate.

	Success Rate(%)						
		All Distances			>10 m		
Speed	RP(%)	100%	>90%	>80%	100%	>90%	>80%
10 km/h		71.0	87.1	92.7	75.2	91.7	98.2
35 km/h		70.0	84.3	87.1	75.0	90.6	93.8

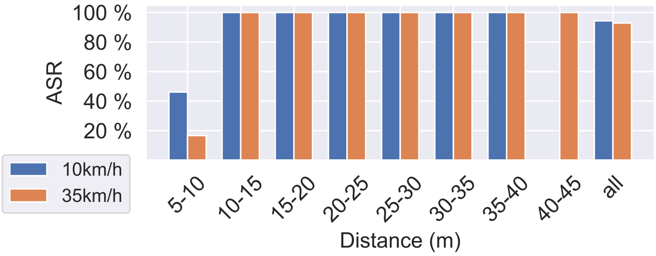

Fig. 8: Attack success rate (ASR) of MVS attack on livox detection. At 10 km/h, 40-45 m is blank as there are no results for this distance.

Object Detector Based Results. To assess the effectiveness of our MVS attack at the object detector level, we processed the

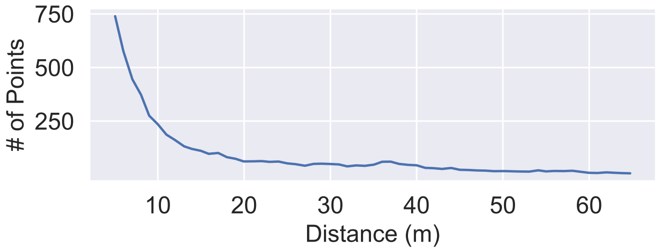

Fig. 9: Variation of distance and the number of point clouds constituting a person. At shorter distances, the number of person point clouds rapidly increases.

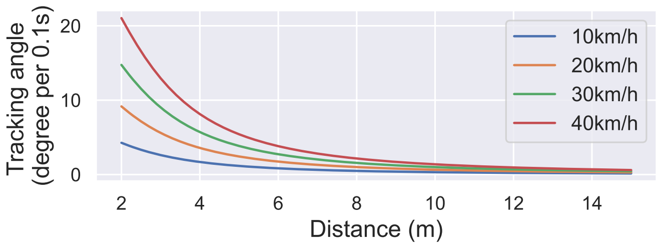

Fig. 10: The tracking angle between the spoofer and the target vehicle varies depending on the distance between them. The necessary tracking angle for each 0.1-second interval is plotted. As the distance decreases, the tracking angle increases sharply.

point cloud data obtained during the attack through the Livox Detection System [24]. Figure 8 shows the attack success rate (ASR) for each distance at each speed, defining success as instances where the object was not detected by the system. The object detector was successfully deceived in more than $90~\%$ of the frames during the MVS attack at both $10~\mathrm{km/h}$ ($94.4~\%$) and $35~\mathrm{km/h}$ ($92.9~\%$). Furthermore, analyzing the impact of distance reveals an important observation: at distances greater than 10 meters, even though the point cloud removal rate was not always 100%, the object detector consistently failed to recognize the presence of people in all frames. This indicates that even partial removal of a person's point cloud significantly impacts the object detector's ability to extract the necessary features for accurate detection.

When an AV is in motion, the distance it covers from the onset of braking upon detecting an obstacle until it halts completely, termed 'stopping distance', is directly proportional to the square of its speed (as annotated in Figure 7, the stopping distance for 35 km/h is approximately 19 m). This relationship makes AVs especially susceptible to attacks that deceive object detectors at medium to long ranges. Consequently, our preliminary results alarm that LiDAR spoofing attacks can pose a significant threat even in actual traffic scenarios.

V. Discussions

A. Limitations and Defense Discussions

1) Tracking at Higher Speeds: Our MVS system capably tracks vehicles up to 40 km/h, but struggles with higher velocities ($\geq$50 km/h) due to the current IR camera's 30 FPS frame rate and detection latency. To ensure accurate tracking of faster vehicles, rapid motor position updates are necessary. We plan to address this by integrating a new IR sensor with higher FPS and optimizing the detection algorithm to reduce latency.

2) Challenges in Short Ranges: While our MVS system achieves successful attacks at long distances, its attack capability is limited at short distances ($\leq$10 m). This limitation can be attributed to two factors. First, the increased density of person point clouds at closer distances complicates complete point elimination. As Figure 9 demonstrates, the number of point clouds constituting a person surges at distances under 10 m. Consequently, short-range attacks necessitate the removal of a significantly larger number of points to maintain the same effectiveness as long-range attacks, demanding enhanced attack capabilities. Secondly, as the distance decreases, the relative angular velocity between the vehicle and the spoofer increases, as shown in Figure 10 which illustrates the necessary tracking angle for a 0.1-second interval. While faster tracking as discussed in the previous section will also enhance close-range attack capabilities, attacks within 5 m, where the tracking angle grows exponentially, will remain challenging. We plan to address these limitations by improving our hardware. Meanwhile, we note that the short-range attack capability is not always necessary. For example, a collision is not avoidable if the distance and the victim and the attacked object is closer than the stopping distance, e.g., 19 m at 35 km/h.

3) Potential Countermeasures: Our current MVS systems are designed to track only one LiDAR sensor on a vehicle. As it is challenging for an attacker to simultaneously track multiple LiDAR sensors, deploying multiple redundant LiDAR sensors to sense the front area of the vehicle is an effective defensive strategy. However, increasing the number of LiDARs, which are among the more expensive vehicle sensors, leads to increased vehicle costs. Additionally, as the system relies on infrared detection to track distant LiDARs, randomly firing fake infrared lasers from locations away from the vehicle's LiDAR can effectively confuse the spoofer's detection. While this requires the installation of additional devices, it is less costly compared to increasing the number of LiDAR sensors.

B. Ethical and Safety Considerations

The experiments were safely carried out in controlled conditions on a private road with authorization. A human drove the experimental vehicle, and the area was surveilled to keep people off the road. Participants potentially exposed to the attack laser wore protective goggles for eye safety.

VI. Concluding Remarks and Future Plans

In this WIP paper, we presented an MVS system capable of launching an attack on cruising vehicles approaching $35~\mathrm{km/h}$. Our novel tracking and aiming methodology, employing IR cameras, archives attacks on long-range, high-speed vehicles from the roadside. We demonstrated that our attack can successfully remove a pedestrian from object detection results with $\geq$ 90% success rate in the real-world scenario.

In the future, we plan to expand our evaluation scope to dive deeper into the threat of MVS attacks. These include evaluating MVS attack on vehicles traveling at higher speeds, and verifying how the MVS attack robustness changes depending on the LiDAR scanning method and FOV specifications.

This will involve conducting tests on multiple LiDAR models, broadening our understanding of the vulnerabilities and defense mechanisms in ADAS systems utilizing LiDARs.

Acknowledgements

This research was supported in part by JST CREST JPMJCR23M4, JST PRESTO JPMJPR22PA, Amano Institute of Technology, NSF CNS-2145493, CNS-1929771, CNS-1932464, and USDOT UTC Grant 69A3552348327.

References

[1] "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," https://www.sae.org/standards/content/j3016_202104/, 2021.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] "Ride-Hailing App - Make the Most of Your Drive - Waymo One," https://waymo.com/waymo-one/.

[12] Y. Cao, J. Ma, K. Fu, R. Sara, and M. Mao, "Automated Tracking System for LiDAR Spoofing Attacks on Moving Targets," in *Proc. Workshop Automot. Auto. Vehicle Secur.(AutoSec)*, 2021, p. 1.

[13] 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.

[14] "PhantomX Robot Turret Kit,," https://www.trossenrobotics.com/.

[15] "VLP-16 User Manual," https://velodynelidar.com/wp-content/uploads/2019/12/63-9243-Rev-E-VLP-16-User-Manual.pdf.

[16] "Livox Horizon User Manual," https://www.livoxtech.com/3296f540ecf5458a8829e01cf429798e/assets/horizon/Livox\%20Horizon\%20user\%20manual\%20v1.0.pdf.

[17] "Logitech C922 Pro Stream 1080p Webcam + Capture Software," https://www.logitech.com/en-us/products/webcams/c922-pro-stream-webcam.960-001087.html.

[18] "Thorlabs - FBH905-10 Hard-Coated Bandpass Filter," https://www.thorlabs.co.jp/thorproduct.cfm?partnumber=FBH905-10.

[19] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," *Advances in neural information processing systems*, vol. 27, 2014.

[20] G. Jocher, "Ultralytics yolov5," 2020. [Online]. Available: https://github.com/ultralytics/yolov5

[21] "Baidu Apollo," https://github.com/ApolloAuto/apollo.

[22] "Cruise," https://www.getcruise.com/.

[23] "Vehicle Stopping Distance and Time ," https://nacto.org/docs/usdg/vehicle_stopping_distance_and_time_upenn.pdf.

[24] "Livox Detection V2.0," https://github.com/Livox-SDK/livox_detection.