EMI-LiDAR: Uncovering Vulnerabilities of LiDAR Sensors in Autonomous Driving Setting using Electromagnetic Interference

Sri Hrushikesh Varma Bhupathiraju
University of Florida
bhupathirajus@ufl.edu

Jennifer Sheldon
University of Florida
jsheldon@ufl.edu

Luke A. Bauer
University of Florida
lukedrebauer@ufl.edu

Vincent Bindschaedler
University of Florida
vbindsch@cise.ufl.edu

Takeshi Sugawara
The University of Electro-Communications
sugawara@uec.ac.jp

Sara Rampazzi
University of Florida
srampazzi@ufl.edu

ABSTRACT

Autonomous Vehicles (AVs) using LiDAR-based object detection systems are rapidly improving and becoming an increasingly viable method of transportation. While effective at perceiving the surrounding environment, these detection systems are shown to be vulnerable to attacks using lasers which can cause obstacle misclassifications or removal. These laser attacks, however, are challenging to perform, requiring precise aiming and accuracy. Our research exposes a new threat in the form of Intentional Electro-Magnetic-Interference (IEMI), which affects the time-of-flight (TOF) circuits that make up modern LiDARs. We show that these vulnerabilities can be exploited to force the AV Perception system to misdetect, misclassify objects, and perceive non-existent obstacles. We evaluate the vulnerability in three AV perception modules (PointPillars, PointRCNN, and Apollo) and show how the classification rate drops below 50%. We also analyze the impact of the IEMI injection on two fusion models (AVOD and Frustum-ConvNet) and in real-world scenarios. Finally, we discuss potential countermeasures and propose two strategies to detect signal injection.

CCS CONCEPTS

• **Security and privacy → Embedded systems security**.

KEYWORDS

Sensor attack, EMI injection, Autonomous driving

ACM Reference Format:
Sri Hrushikesh Varma Bhupathiraju, Jennifer Sheldon, Luke A. Bauer, Vincent Bindschaedler, Takeshi Sugawara, and Sara Rampazzi. 2023. EMI-LiDAR: Uncovering Vulnerabilities of LiDAR Sensors in Autonomous Driving Setting using Electromagnetic Interference. In *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '23), May 29-June 1, 2023, Guildford, United Kingdom.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3558482.3590192

WiSec '23, May 29-June 1, 2023, Guildford, United Kingdom

ACM ISBN 978-1-4503-9859-6/23/05…$15.00
https://doi.org/10.1145/3558482.3590192

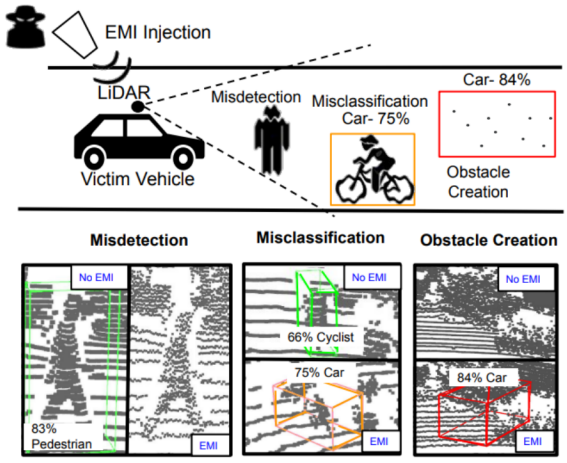

Figure 1: Overview of the IEMI injection on LiDAR. The perturbations by the IEMI signal injection can cause misdetection (bottom-left), misclassification (bottom-middle) or create new obstacles in AV perception models (bottom-right).

1 INTRODUCTION

Self-driving cars are being rapidly developed and are already achieving unprecedented levels of autonomy. Companies such as Tesla, Waymo, and Lucid are providing Autonomous Vehicle (AV) services for operation in real-world scenarios [3]. This autonomy is achieved through the Perception systems in AVs that analyze surrounding environments to facilitate object detection and safe-driving decisions. In addition, many AVs have adopted LiDAR (Light Detection and Ranging) sensor-based detection in their perception stacks to ensure reliable object detection. LiDAR sensors collect depth measurements from the vehicle's surroundings to construct a precise 3D point cloud of the environment.

Recently, researchers have demonstrated attacks that can compromise LiDAR-based detection systems [8, 12, 23, 57]. Generally, these works alter the sensing mechanisms to manipulate the raw sensor data. The attacks can further be designed so that the preceding machine-learning (ML) models behave as desired by the attacker [8, 12]. These attacks can induce fake obstacles [12] in or altogether remove obstacles from AV perception [8, 10]. At the

same time, other works use adversarial perturbations to induce misdetection and misclassifications [4, 10, 12, 23, 65, 76].

These previous works use light signal injection to spoof return signals to LiDAR sensors to achieve desired results from the detection model. The impact of such light injection attacks at the sensor and ML model levels has been extensively studied; this has thus revealed several accuracy and practicality constraints. In particular, light injection attacks require extreme precision for aiming and tracking the target sensor. This can be difficult to achieve when the target sensor is on a moving vehicle.

As a result of previous work's focus on light injection, the vulnerabilities of LiDAR sensors to other potential attack vectors have been overlooked. Motivated by this observation, we study the impact of Intentional Electro-Magnetic Interference (IEMI) on LiDAR sensors. In particular, we demonstrate that it is possible to impact AV perception with IEMI injection and the potential vulnerability of AV systems to such injections. This vulnerability lies in the susceptibility to EM waves of time-of-flight (TOF) circuits that compose modern 3D spinning LiDAR sensors used in autonomous driving systems. This work focuses on the following research questions: **(i)** *How does IEMI injection impact the sensing mechanisms of TOF circuits and LiDAR sensors and, consequently, 3D point cloud acquisition?* **(ii)** *How can the effect of IEMI injection on sensor output be characterized?* **(iii)** *How can these injections be leveraged to degrade the performance of object detection models, and what are potential ways to defend them?*

This work explores LiDAR sensor's susceptibility to intentional EMI injection. We observe that the sensor data perturbations caused by the injected signal can misdetect and misclassify objects and create non-existent objects in the LiDAR perception models, as shown in Figure 1. This paper makes the following key contributions.

LiDAR's Susceptibility to IEMI. This work identifies and presents a novel signal injection vulnerability in TOF circuits used in LiDARs, that can manipulate selected parts of the raw sensor data. We comprehensively analyze the injected signal's effect on sensor circuitry in a spinning LiDAR sensor to demonstrate how the vulnerability can be exploited in the real world.

Causality. The susceptibility of TOF sensors to EMI signal injection can stem from any region of the sensor. We study the causality of this phenomenon in TOF sensors. An extensive understanding of this effect is required to be able to secure future devices. We form a hypothesis regarding the causality of this vulnerability based on tests of how various components in the circuit are affected.

Vulnerability Characterization. To understand how EMI signal injection can be leveraged to impact object detection, we model the vulnerability to characterize the perturbations caused by the injection of sensor data. To do this, we first characterize the injection and quantize the effect in terms of various parameters (in Section 4.2). Afterward, we model the resulting perturbations in the 3D point cloud and their relation to the injected signal.

Impacts on AV Perception Models. We evaluate the impact of the injection on state-of-the-art object detection models (PointPillars [32] and PointRCNN [52]) and the industry-grade segmentation model Baidu Apollo [6], in both synthesized and real-world scenarios and estimate the IEMI signal required to impact the AV system's performance significantly. The injection drops the classification rate to less than 50% for all the models in the case of pedestrian and cyclist obstacles while dropping the detection rate to nearly 20% for each class. We further evaluate the effect of such perturbations on two state-of-the-art camera-LiDAR fusion models, AVOD [30] and Frustum-ConvNet [62], to demonstrate how sensor fusion models do not fully mitigate the vulnerability. The injection drops the object detection rates for pedestrian and cyclist classes below 50% in the case of AVOD and for pedestrian and car classes in the case of Frustum-ConvNet. Overall, the object detection rates of all the classes are reduced by over 10% for both models. Finally, we perform real-world analysis to demonstrate the practicality of exploiting the vulnerability, where we study the effect of EMI injection on object detection models in controlled outdoor scenarios. The injection reduced the Intersection Over Union (IOU) values of the detected objects by nearly 20%, dropping them below the required threshold for detection (0.5). We also discover and evaluate how the EMI perturbations can generate additional, not existent objects in the AV field-of-view (FOV).

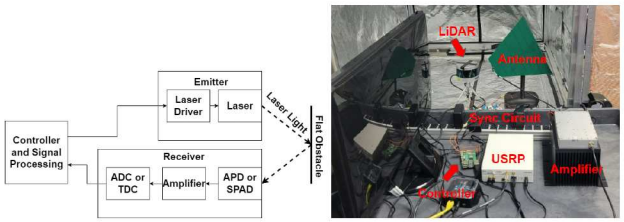

Figure 2: (Left) Block diagram of the fundamental components of TOF circuits. (Right) Faraday Cage setup.

Defense. Finally, we investigate existing defense methods against signal injection, and we demonstrate that state-of-the-art denoising techniques only improve model performance by 6%. Then, we propose and evaluate two novel IEMI injection detection methods: *Ground Point Based* and *Point Intensity Based*. The first method achieves a 98.74% True Positive Rate (TPR) and 94.35% True Negative Rate (TNR), while the latter achieves a 100% TPR and TNR.

2 BACKGROUND

2.1 Time-Of-Flight Sensors in AVs

Time-of-flight sensors fire and receive returning laser pulses to measure the distance between the sensor and nearby objects using the equation $d = t \cdot c/2$ where d is the distance to the target, t is the time taken for the fired pulse to return, and c is the speed of light [6]. Groups of TOF sensors are used in LiDAR sensors [26, 51] to create maps of obstacles' reflection distances called *point clouds*. Spinning 3D LiDAR sensors, in particular, construct the entire vehicle FOV by rotating stacks of TOF circuits horizontally 360° to capture the reflections.

In this paper, we focus on spinning LiDAR sensors that AV companies, including General Motors' Cruise [22], and Google's Waymo [5] use, due to their wide FOV for obstacle avoidance, higher resolution than RADARs, and more comprehensive 3D data than cameras [33, 70].

TOF Circuit Functioning.

A TOF circuit can be broadly divided into an emitter circuit that includes pulse timing, laser driving circuitry, and lasers and a receiver circuit that include photodiodes, amplifiers, and ADCs [17]. The laser and its driver circuitry transmit controlled light pulses of a given frequency and time. The reflections (or return echoes) from nearby obstacles are then received by a photodiode, e.g., avalanche photodiodes (APDs), which produce photocurrent. The photocurrent is then amplified by a transimpedence amplifier (TIA) and digitized by an analog-digital converter (ADC) before being processed, as shown in Figure 2 (left).

2.2 Intentional EMI Attacks

Fluctuations in an electromagnetic field or EMI surrounding a susceptible electrical component in a circuit (e.g., an amplifier, a diode, or a wire) will induce an electric current (or voltage) that can appear in the circuit output signal [29, 44].

Traditional circuit protection techniques against EMI usually include adding filters and shielding into their designs [29]. However, several works have shown how EMI can interfere with the correct functioning of hardware components, such as ADCs, GPIO pin readings [63], and amplifiers [15, 49]. In addition, researchers have demonstrated how malicious actors can leverage intentional EMI (IEMI) to actively manipulate the output signal of a circuit or sensor readings [28, 31, 60].

2.3 Perception in AVs and Adversarial Attacks

The AV navigation system typically comprises four main components: Perception, Localization, Planning, and Control [42]. Perception refers to the ability of the AV system to collect and process relevant information from the surrounding environment. The later modules then use the collected information to contextualize the current situation, decide how to achieve the vehicle's goals, and execute the chosen driving actions.

LiDAR gathers 3D depth measurements as point clouds (see Section 2.1). The perception module's challenge is identifying which points represent relevant objects and then classifying them. Identifying potential relevant objects in the vehicle's trajectory is done through the segmentation of the point cloud into easy-to-process chunks, such as 2D grids [30, 67], voxels [75], or pillars [32]. Alternatively, models such as Pointnet [45], Apollo [6], and PointR-CNN [52] feed the raw point cloud to a model, which then generates bounding boxes around potential clusters of interest.

Adversarial Attacks on Autonomous Systems. Adversarial example attacks add nearly imperceptible amounts of noise to a model's input to force a change in the output [16, 58].

Adversarial example attacks can be performed against LiDAR point clouds by perturbing or adding points [11, 21, 45, 45, 64, 65]. However, these attacks are simulated and not particularly realizable in the real world or focused on only one model type. Instead, physical attacks are able to manipulate the perception models to detect non-existent obstacles [9, 12, 57] or hide real obstacles with fake noise around legitimate objects [10, 23, 59, 76]. In addition, some attacks entirely remove the point cloud of existing obstacles [8]. While the above works represent threats to AVs, we present a novel AV vulnerability in the form of EMI injection, which not only decreases detection against many of the most popular object detection models but also affects the basic TOF circuits used in several LiDAR sensor families, as described in Section 4.

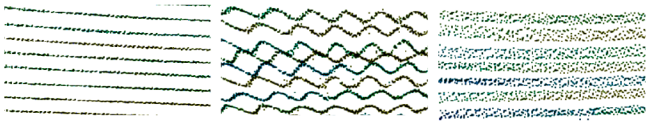

Figure 3: (Left) Point cloud of a flat vertical object with no injection (XY plane). (Middle) Resulting sinusoidal pattern during EMI injection at 960.9 MHz. (Right) Resulting random pattern during EMI injection at 977.4 MHz.

3 THREAT OVERVIEW AND MODEL

3.1 Vulnerability Overview

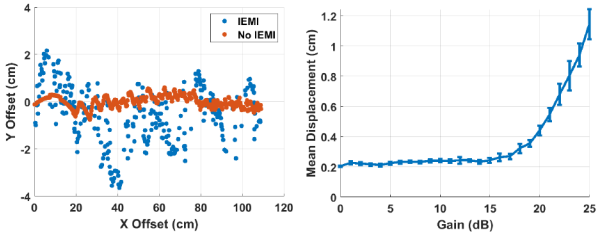

Figure 4: (Left) Slamtec 2D LiDAR perturbations with and without the 949.8 MHz sine injection. (Right) Slamtec S2 XY plane displacement vs injection power.

As described in Section 2.2, researchers have demonstrated how EMI can be used to maliciously alter the outputs of hardware components (e.g., ADC, amplifiers, etc.) and sensors [15, 31, 49, 60, 63, 66].

With this intuition, our goal is to discover whether TOF circuits can be affected by this external interference.

EMI Signal Injection. The circuit's susceptibility to EMI signals can differ significantly depending on the injected electromagnetic (EM) wave frequency. These frequencies are mainly determined by the circuit's design. Hence, we first identify the existence of an EM wave that can affect the output distance of LiDARs. To accomplish this, we perform, in a controlled environment (Figure 2 right), a frequency sweep of sine wave injections in 400–1000 MHz on two different types of spinning LiDARs based on TOF circuits: (i) the Slamtec 2D spinning LiDAR [51] with a single row of points and (ii) the 3D spinning LiDAR Velodyne VLP-16 with 16 rows of points to produce a 3D point cloud [26]. Both sensors are susceptible to specific frequencies, as shown in Figures 3 and 4.

In the 2D Slamtec LiDAR, the induced perturbation results in a displacement along the XY axis. While in the VLP-16, we identified two types of 3D patterns: a sinusoidal pattern (with the maximum amplitude measured with an IEMI sine wave injection at 960.9 MHz) and a random cloud point distribution along each row (with the maximum distribution with an IEMI sine wave injection at 977.4 MHz) along both XY and XZ axes and all the 360° VLP-16 LiDAR FOV. With changes in distances, we also observed a drop in the cloud point recorded reflection intensities. In the rest of the paper,

we refer to these perturbations as the *sinusoidal perturbation* and *random perturbation*. We use them as references for the adversarial capability evaluation in Section 4.2, and evaluate their impact on AV perception models and fusion systems in Section 5.

3.2 Threat Model

We consider an adversary that generates an electromagnetic field in the vicinity of an AV which uses data from a LiDAR sensor to detect and classify objects on the road. The generated electromagnetic field induces a voltage on the LiDAR's internal TOF circuits, creating perturbations in the resulting point cloud (see Figure 1). This perturbation may alter the AV perception system's behavior, e.g., detecting a pedestrian instead of a car or not detecting any obstacle in the affected area, which can result in dangerous automatic maneuvering of the AV.

Attacker Goals. We consider two potential goals for the adversary: (i) affect a specific region of space in the LiDAR FOV or (ii) target and track a specific obstacle in the LiDAR FOV. In both cases, the adversary aims to induce misclassification or misdetection of one or more obstacles already present in the FOV.

Attacker Capability and Assumptions. To exploit this vulnerability, an adversary can consider the following parameters: the distance between the adversary and the victim LiDAR (d_A), the distance between victim LiDAR and the target obstacle (d_O), affected horizontal FOV angle (perturbed angle) (θ), the injected IEMI frequency (f), and IEMI transmission gain (G). We assume the adversary can control the output gain G of the IEMI transmitter, which is linearly proportional to the output transmission power, and the affected region in the LiDAR FOV or simply a *horizontal angle* θ. This control can be achieved by synchronizing the transmitted EMI signal with the rotation of the LiDAR, as demonstrated in previous works [8, 12, 57] (see Appedix A for details). From our experiments on the 3D spinning LiDAR (VLP-16) we found that the adversary cannot perturb individual rows (meaning affecting a specific vertical angle of the LiDAR FOV). This is due to the fact that the interference affects all the TOF circuits stack of the 3D sensor in the field region covered by the transmitting antenna range. Thus we assume this same constraint in our analysis. We also assume that the attacker can discover the vulnerable frequencies by acquiring the same type of LiDAR sensor used by the victim AV. Based on these assumptions, we characterize the key parameter ranges in Section 4.2.

Previous Knowledge. Our analysis focuses on TOF spinning LiDARs. The adversary can acquire information about the target LiDAR from publicly available sources (e.g., manuals and datasheets). The adversary can track the approximate position of the target obstacle to aim the attack in the required FOV of the victim LiDAR using basic tracking algorithms [9]. We also assume that the attacker has knowledge of the type of machine learning model used by the victim AV Perception module. We consider this assumption reasonable since attackers can leverage information from publicly available sources released by companies [6] (e.g., open source code base) or obtain white-box access by additional engineering efforts to reverse engineer the model [39, 40, 73]. The attacker does not require access to the victim's LiDAR sensor or tamper with the AV.

Attack Scenario. As in previous works [8, 12], we assume the attacker can hide the attack equipment in a car or place it near the road, as depicted in Figure 1. This is possible because EMI injection does not require precise aiming and can pass through walls and glass. Note that in this work, we only characterize single-scene static scenarios as proof-of-concept analysis.

4 VULNERABILITY CHARACTERIZATION

4.1 TOF Circuit Susceptibility to IEMI

As a preliminary study to understand the cause of the interference effect, we place the VLP-16 LiDAR in a controlled area with no obstacles and no legitimate return echoes.

The rotation speed of the LiDAR is set at 300 RPM and dual mode setting to ensure the maximum horizontal resolution [26]. For the IEMI injection, we use an N210 USRP [47] with a ZHL-5W-202-S+ amplifier [37] and a trapezoidal directional antenna. A photodetector and a Raspberry Pi controller allow to control of the affected region (defined by the horizontal angle θ) in the LiDAR FOV, as shown in Figure 2. We use this setup for all the experiments in this work with the VLP-16 LiDAR. The safety considerations regarding the performed injection are described in Section 7.

As shown in Figure 10 in the Appendix, at increasing transmitting power, the IEMI generates cloud points in the targeted area. Since the transmitted laser pulses never return, we conclude that the receiver part of the TOF circuits of the LiDAR is likely to be susceptible to the EMI waves.

Off-the-shelf APD Module. To confirm the aforementioned hypothesis that the receiving circuit is susceptible to our IEMI, we analyze an off-the-shelf Hamamatsu C12703-01 APD module [20]. This high-sensitivity photodetector is typically employed in TOF receiving circuitry, as described in Section 2.1.

To achieve fine-grade precision in the injection, we built a flying probe [72] and performed a frequency sweep ranging from 400–1000 MHz to find potential susceptible frequencies. During this experiment, we also place a Faraday pouch on the back of the board to eliminate the IEMI effect on cables and power connection.

As per our hypothesis, a sine wave IEMI injection at 31 mW and 550 MHz produces an output voltage variation. We then placed the tip of the probe in four different areas of the circuit corresponding to specific hardware components within the circuit (a measurement junction (1), a voltage regulator (2), and two internal operational amplifiers (3-4)) to identify potential susceptible components, as shown in Figure 10 in the Appendix. The measured average output voltage for 300 captures has the highest value when the injected EMI wave is closer to the two amplifiers, suggesting they are most significantly impacted. These results also confirm findings of previous work on the susceptibility of operational amplifiers to EMI injection [60]. Furthermore, the low vulnerable frequency compared to the board's spectral response range (minimum 300 GHz [20]) suggests that electromagnetic induction, rather than the photoelectric effect (which requires a minimum frequency to generate current [53]), causes the output variation. Thus, the vulnerability is not dependent upon the wavelength the LiDAR uses.

Aliasing Effect. As shown in Figure 2, the amplifier circuit output is then fed to the ADC for digital processing. Since the APD module does not provide an ADC, we explore the presence of an aliasing effect in the VLP-16 LiDAR to understand whether the ADC contributes to the observed perturbation patterns.

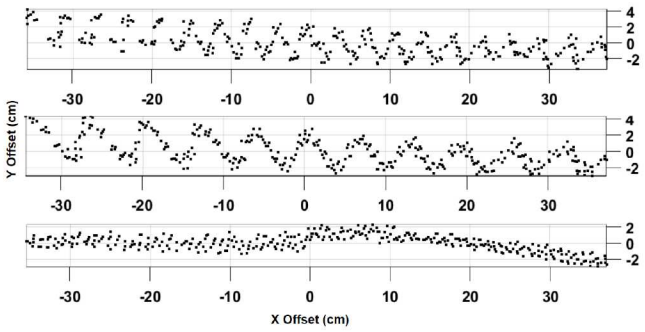

Figure 5: (Top) Injection of 960 MHz. (Middle) Increment of the injected sine wave frequency by 200 Hz (Bottom) Increment of 400 Hz.

The amplifier output voltage should be converted at a certain sampling rate by LiDAR internal ADCs. According to the Nyquist-Shannon sampling theorem, the sampling rate should be at least twice the signal's maximum frequency. If the system acquires data at an insufficient rate, the signal will be incorrectly processed at a lower frequency, generating the aliasing effect [27].

Using the same setup as the first experiment, we found that, by varying the injected frequency, the sinusoidal perturbation pattern in Figure 3 periodically repeats at 9 kHz frequency intervals. This behavior suggests that the induced voltage produced by IEMI injection is internally sampled with a fixed sample rate of 9 kHz as the injected sine wave aliases to the same sinusoidal perturbation after a complete cycle. Figure 5 shows the manipulation of the aliasing effect by shifting the injection to produce different sinusoidal patterns and, eventually, a near DC pattern. With a 3 kHz shift, we also produce the random pattern.Based on this analysis, we believe that the sinusoidal and random point cloud perturbations described in Section 3 are the results of the aliasing.

4.2 Adversary Capability

We characterize the adversary's capability based on the parameters listed in Section 3 and real-world experiments.

Minimum Horizontal Angle. As described in Section 3, the injected sine wave affects all horizontal lines (along the Z axis) and all the horizontal FOV angle (along the azimuth - XY plane) of the 3D LiDAR. An adversary can only control the firing timing to match with 3D LiDAR scans to affect a specific horizontal angle θ. This angle is proportional to the LiDAR horizontal revolution time ($\theta/360 = t/p$) where t is the time taken to scan a target region, and p is the revolution period (as determined by the RPM). To characterize the minimum horizontal angle (meaning the minimum region an adversary can influence with our setup), we measure the maximum time taken for the USRP to synchronize with the LiDAR and achieve a successful injection. After recording 600 LiDAR revolutions, we found a maximum latency of ~1.6 ms, corresponding to ~ 3° of the LiDAR FOV. This result depends on the hardware latency of the trigger circuit and USRP transmitter as well as the software latency of the controller. We thus consider this as the minimum attack angle achievable by the attacker in all the experiments in this work. We believe this value can be further reduced using more sophisticated equipment.

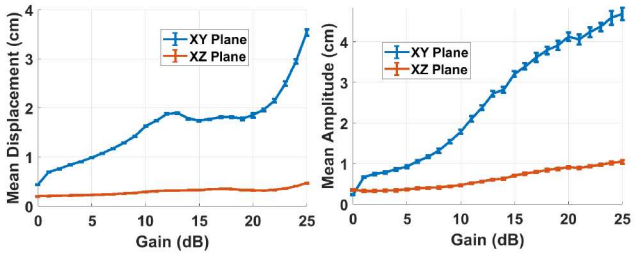

Figure 6: (Left) Mean distance displacements vs gain. (Right) Mean sinusoidal amplitude values vs gain.

Target Object Distance. To characterize the perturbation on the point cloud of existing obstacles at different distances d_O, we placed a flat surface as the target object at increasing distances (1 m, 2 m, 3 m, and 4 m) away from the target LiDAR in an indoor controlled environment. We recorded 150 frames each for the no-injection case (meaning the point cloud without the perturbation) and during the sine wave injection at the two relevant frequencies (960.9 and 977.4 MHz) that induce both the sinusoidal and random perturbation. We found the resulting sinusoidal wavelength (in the space domain) linearly increases with increasing d_O[1]. We also found that sinusoidal amplitude remains constant with d_O. We did not observe any change in the point cloud distribution for the random pattern, which remains constant. Based on our causality analysis, the wavelength increases with d_O because the induced output voltage remains constant with constant gain. However, the scale of the perturbation changes with distance (as the measured chord length is a function of angle and object distance).

EMI Transmitter Distance. To test the effect that parameter d_A has on the perturbations, we placed a flat object at a distance of 1 m from the LiDAR and performed 25 dB gain sine wave injections (the maximum allowed by our setup) with the antenna placed at various distances from the LiDAR. We found that the perturbation is still detectable (SNR above 0 dB) above 1 m from the LiDAR[1]. In this experiment, we were only able to test up to 160 cm away due to controlled space limitations. Leveraging this analysis and EMI transmission theory (particularly the effects of near-field and far-field injections) the attacker can estimate the EMI emission power required for achieving further distances.

IEMI Emission Power. We use the maximum gain G of 25 dB in our setup, corresponding to 2 W of EMI transmission power after amplification (0 dB USRP gain corresponds to 50 mW). We follow the same procedure: placing a flat object 1 m away and locating the transmitter at 0.5 cm from the LiDAR. From our causality analysis, we expect that increasing the gain (meaning increasing the transmitted power of the EM wave) increases the induced voltage, which causes a more substantial effect on the point cloud. This experiment shows that the sinusoidal pattern's peak amplitude increases with increasing gain while the wavelength remains constant. The mean displacement of the random distribution also increases. We also found that the intensity decreased (note that the intensity

[1]Supplementary data and graphs can be found at: https://cpseclab.github.io/EMILidar/

is identical in the XY and XZ planes as it is a property of each cloud point). We hypothesize that this occurs because the higher power injections are able to surpass the static noise threshold in the VLP-16 [18]. Because the induced voltage falls below the legitimate return pulse voltage, but above the noise threshold, it appears as a low-intensity point.

4.3 Perturbation Modeling

To study the effect of the perturbations caused by the IEMI on AV object detection models, we synthesize the perturbations measured for the VLP-16 LiDAR. We first conducted real-world experiments at increasing EMI signal power injected (until 25 dB gain, meaning the maximum EMI power achievable by our hardware setup) and increasing distances. We found a linear relationship between the power of the injected signal and the amplitude of the resulting interference in the point cloud (similar to previous works' findings on EMI injection on sensors [31]). Then, we used those data to model the sinusoidal and random pattern distribution up to 70 dB using Equation 1 (in Appendix B) and a uniform distribution model, to show the attack's consequences under more powerful injections (see Appendix B for details).

5 EVALUATION

5.1 Evaluation on Detection Models

We synthesize the perturbations modeled in Section 4.3, on the KITTI dataset [1], a widely used dataset for AD research that comprises diverse real-world driving scenes. Then we assess the effects of the EMI perturbations on two popular state-of-art LiDAR-based object detection models, PointPillars [32] and PointRCNN [52], to examine the impact on object detection and classification rates for different detection architectures. PointPillars adopts a voxel-based approach [46] that voxelized irregular point clouds into grids with multiple channels, forming a bird-eye view of the scene, and then processes it using convolutional neural networks. On the other hand, PointRCNN uses a point-based method with PointNet [45] as its backbone and a two-stage detection process that employs point feature vectors instead of voxelization to preserve the point cloud's geometric features.

5.1.1 Object Misdetection and Misclassification. In this evaluation, we consider target objects from three different classes: pedestrian, cyclist, and car, which are the most relevant classes for AVs in the KITTI dataset. We consider all the pedestrian (4458) and cyclist (1627) objects available in the dataset for our evaluation. In the case of car obstacles, we only select the car obstacles from the test set of the KITTI dataset (14386) to reduce the computational cost.

To evaluate the effect of the signal gain on detection models, we first set θ to be the horizontal span of the respective target object and synthesize the perturbation for each obstacle independently. We increment the gain by 1 dB increments until 70dB as we observe it to cause a significant effect on the performance of the tested models. We then study the effect of θ on model performance. For this, we set the gain to be 70 dB and increment the θ at 3° intervals corresponding to the minimum resolution of the perturbation in the real world until we reach 40°. We choose 40° as the maximum θ since we observe it to be the maximum perturbation angle required to cover any target object in the LiDAR FOV over the KITTI dataset.

Evaluation Metrics. We evaluate the effect of the synthesized perturbations on LiDAR-based models using Object Detection Rate (ODR) and Classification Rate (CLR) as metrics. We consider the object detection of an individual object successful if the corresponding prediction has an Intersection over Union (IOU) greater than the desired threshold with respect to the ground truth. Similarly, we consider object classification successful if the detected object is classified as the correct object class. We perform two different analyses based on the IOU threshold. In the first analysis, we set the IOU threshold to 0 (WIOU). Here, if the predicted object has an IOU with respect to the ground truth greater than 0, we consider it a successful prediction. In the second analysis, we evaluate based on the default IOU thresholds as proposed in the corresponding works of each model (DIOU). The DIOU values for PointRCNN are 0.7 for cars and 0.5 for cyclists and pedestrians classes [32]. The DIOU values for PointPillars are 0.6 for cars and 0.5 for cyclists and pedestrians classes [52].

Results and Observations. Table 1 shows object detection and classification results for PointPillars and PointRCNN models for WIOU and DIOU analyses. PointPillars drops at least 20% in ODR and 30% in CLR for cyclist and pedestrian object classes. Car object class shows only an 11% drop of ODR and CLR in WIOU analysis and a 19% drop in DIOU analysis. PointRCNN shows a 19%, 13%, and 1% drop in pedestrian, cyclist, and car objects classes, respectively. PointRCNN shows a poor performance in DIOU analysis relative to PointPillars, and the ODR and CLR are further depleted with IEMI perturbations. We believe that the robustness of the car obstacles is due to the significantly larger size of the obstacle compared to pedestrian and cyclist objects.

We also study the effect of the perturbation angle θ on ODR for both models. The ODR drops by 32%, 51%, and 31% in pedestrian, cyclist, and car object detection for PointPillars in DIOU analysis. For PointRCNN, we observe an ODR drop of 26%, 25%, and 59% in pedestrian, cyclist, and car, respectively, in DIOU analysis. The results visualized in Figure 7 indicate that increasing θ beyond the span of the target obstacle does not decrease ODR further. Figures 7 (a) and (b) show a sudden drop in the ODR with 3°, indicating that even small perturbation angles can significantly impact the model's object detection. Finally, PointRCNN-based architecture has been shown to be robust to geometric variations of the point cloud data in previous works [69]. We hypothesize that this is the reason for the smaller misclassification rate compared to PointPillars.

5.1.2 Object Creation Evaluation. We observe that the perturbations induce False Positives (FP) in the object detection results. To measure the effect, we synthesize both sinusoidal and random perturbations in the entire front view of the point cloud (80°) at 70 dB gain over the validation set of the KITTI dataset (3769 scenes). We then evaluate the number of fake obstacles induced by the perturbations in the form of FPs and report the recall and precision values of the model with the IEMI perturbations.

Results and Observations. We observe that the sinusoidal perturbations induce 8.76 FPs per frame and 7.95 FPs per sample over the validation set of the KITTI dataset for PointPillars. For

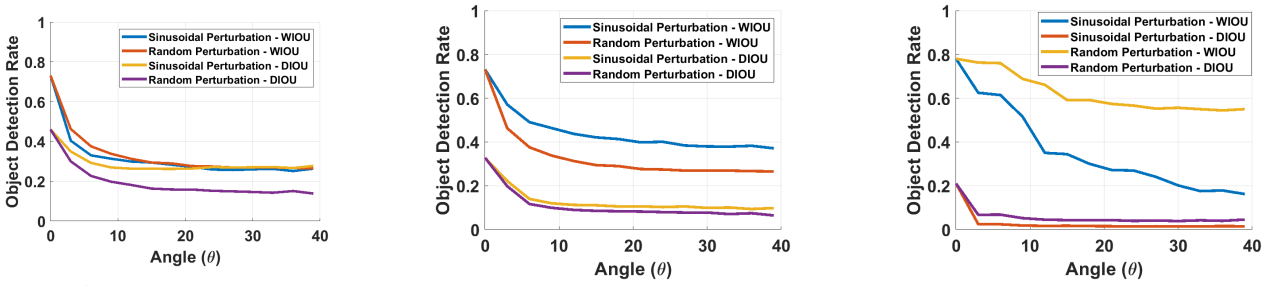

(a) PointPillars Pedestrian (b) PointRCNN Pedestrian (c) Apollo Pedestrian

Figure 7: Object Detection Rates (ODR) for pedestrian at increasing perturbation angles (θ) and 70 dB transmitted IEMI gain[1].

Table 1: ODR and CLR of LiDAR-based models on WIOU and DIOU for sinusoidal and random perturbations at 70dB gain.

Point Cloud Type	Pedestrian		Cyclist		Car	
	WIOU	DIOU	WIOU	DIOU	WIOU	DIOU
PointPillars ODR (%) / CLR (%)						
No Perturbations	81/73	46/45	88/75	74/69	89/88	77/77
Random	60/43	15/14	66/23	24/18	86/85	71/71
Sinusoidal	67/38	17/13	70/45	44/40	78/77	58/58
PointRCNN ODR (%) /CLR (%)						
No Perturbations	70/51	31/26	61/43	40/26	87/86	70/40
Random	65/40	21/15	55/41	30/20	86/85	68/37
Sinusoidal	51/31	12/8	48/36	18/13	86/85	46/19
Apollo ODR (%) / CLR (%)						
No Perturbations	78/57	21/21	71/44	20/15	72/66	40/38
Random	75/42	7/5	67/38	9/5	69/64	31/30
Sinusoidal	62/24	4/2	58/28	4/2	67/62	28/23

PointRCNN, the perturbations induce 4.77 and 4.97 FPs per sample. As a result, the precision (resp. recall) of PointRCNN over the entire KITTI validation set drops from 0.33 (resp. 0.35) to 0.32 (resp. 0.30) for random perturbations and to 0.23 (resp. 0.15) for sinusoidal perturbations. For Pointpillars, on the other hand, the precision drops from 0.447 to 0.444 and 0.408 for random and sinusoidal perturbations, respectively. Meanwhile, the recall drops from 0.776 to 0.753 and 0.719 for random and sinusoidal perturbations, respectively. Note that both models exhibit low precision even without any perturbations. Although the perturbation creates fake obstacles, precision drops only by a small amount because the fake obstacles do not correlate with the false positives perceived in the no perturbation case.

5.2 Evaluation on Segmentation Models

Baidu Apollo 5.0 [6], is an open-source framework for industry-grade AV systems adopted by several AV companies such as Chery, Hyundai, Volvo, and Ford [3]. Apollo's segmentation model filters the region of interest from the point cloud (typically 60 m), maps the information into a 512×512 grid, and extracts feature information for each cell. A deep neural network takes this feature matrix and gives the probability of each cell belonging to an obstacle. Candidate objects are then selected from these probabilities based on object scores, object height, and classification scores. For evaluating misclassification rates and object creation of Apollo, we use the same metrics of Section 5.1.

Impact on Object Segmentation and Classification. We first study the impact of the IEMI perturbations on the Apollo model for the segmentation and classification of the samples described in Section 5.1. For this analysis, we only consider a 70 dB gain to evaluate the maximum effect caused by the perturbations, and 0.5 as the DIOU values for cars, cyclists, and pedestrians classes [6] Similar to Section 5.1, we consider a perturbation angle wide enough to cover the corresponding target obstacle to studying the impact of gain variation. We then study the effect of the perturbation angle θ on segmentation with 3° increments until 40°.

Results and Observations. Table 1 shows the ODR and CLR results of Apollo segmentation with sinusoidal and random perturbations. The DIOU analysis shows a drop in ODR to less than 5% in pedestrian and cyclist obstacle classes and 30% in car obstacle classes. The WIOU shows a maximum drop of 16%. The CLR rate, however, drops below 30% for all the obstacle classes.

The results demonstrate that the ODR of the segmentation model declines gradually with increasing θ, unlike PointPillars and PointRCNN, as shown in Figure 7. While ODR for car, cyclist, and pedestrian obstacle classes drop to a minimum of 67%, 58%, and 62% in gain analysis, the ODR for the same classes drops to 39%, 19%, and 16%, respectively at 40° perturbation angle. Object detection models extract features from local sub-parts of the point cloud individually. Segmentation models, on the other hand, annotate the scene point by point, giving more granular information about the scene. We hypothesize that this makes the segmentation model more susceptible to our EMI perturbations in the region around the target obstacle.

Impact on Object Creation. We study the impact of the perturbations on inducing new objects into the segmentation model, measuring the number of false positives introduced. To do this, we synthesize the perturbations in the entire front view of the point cloud with a 70 dB gain for the KITTI validation set, similar to Section 5.1.2. The precision for Apollo segmentation drops from 0.28 to 0.18 for random perturbations and 0.09 for sinusoidal perturbations. On the other hand, the recall drops from 0.58 to 0.55 and 0.52 for sinusoidal and random perturbations. The drop in precision

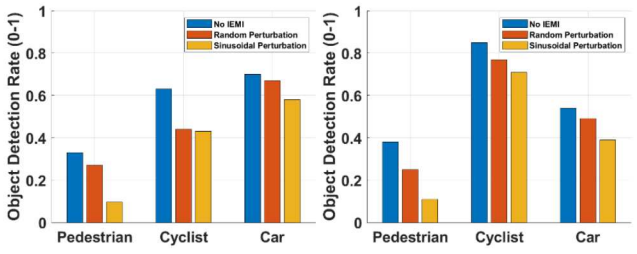

Figure 8: Object Detection Rates for AVOD (left) and Frustum-ConvNet (right). Results for 500 random samples per object class with synthesized perturbations at 70 dB gain.

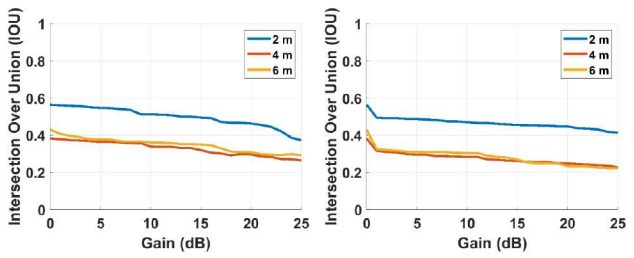

Figure 9: IOU of predicted pedestrian objects from PiFiNet at 2, 4, and 6 meters away from the VLP-16 LIDAR for sinusoidal (left) and random (right) perturbation injection.

is more significant than that in recall indicating that the rate of fake obstacle creation is high.

5.3 Sensor Fusion Evaluation

Camera-LiDAR sensor fusion models extract features from each of the sensors used and allow them to complement each other to improve the accuracy of 3D object detection. This analysis shows that the perturbations caused by the IEMI injection deteriorate the performance of state-of-the-art sensor fusion models. We evaluate two popular camera-LiDAR sensor fusion architectures: AVOD [30] and Frustum-ConvNet [62]. AVOD is a feature-level fusion model that extracts feature maps from RGB images from cameras and BEV images of the LiDAR individually and then combines them. Frustum-ConvNet is a cascaded semantic-level fusion model that creates frustum-level features on the LiDAR point cloud from each region proposal from the camera image. AVOD and Frustum-ConvNet are highly representative of state of the art and have been used in prior works to evaluate attacks [8, 19].

Methodology. For this analysis, we randomly select 500 objects for each of the cyclist, pedestrian, and car classes from the validation set of the KITTI dataset. We then synthesize perturbations on the LiDAR point cloud, as discussed in Section 4.3. Similar to the synthesized evaluation in Section 5.1, we set θ to be the horizontal span of the target object and synthesize the perturbations for each object corresponding to a 70 dB signal gain. We do not perturb or alter the input images to the fusion models.

Evaluation Metrics. We consider ODR and CLR as a metric to evaluate the effect of perturbations on sensor fusion models. Here, we consider the default IOU thresholds for each model (0.7 for car and 0.5 for pedestrian and cyclist obstacles for both models).

Results and Observations. Figure 8 shows the classification rates for the two sensor fusion models when injecting the sinusoidal and random perturbations for pedestrian, cyclist, and car classes. The perturbations drop the ODR below 50% for pedestrian and cyclist object classes in AVOD and pedestrian and cyclist object classes in Frustum-ConvNet. The ODR for the car object class in AVOD and the cyclist object class in Frustum-ConvNet show robustness against the perturbations.

The CLR for both models follows a similar trend as ODR with a drop below 50% for pedestrians and cyclists in AVOD. No other scenarios in Frustum-ConvNet show misclassifications, indicating that the CLR remains the same as ODR. This happens because the model's region proposals for object detection come from camera image data that is not perturbed by our injection. Our results show a significant deterioration of the model performance on pedestrian and cyclist obstacle classes in AVOD and pedestrian and car obstacles in Frustum-ConvNet. Furthermore, the perturbation signal generates new obstacles in the sensor fusion models. For AVOD, the precision drops by a maximum of 0.43 to 0.33, and the recall drops from 0.56 to 0.45. For Frustum-ConvNet, the precision drops from 0.69 to 0.52 while the recall from 0.46 to 0.32.

5.4 Real World Scenario

We further conduct a proof-of-concept experiment in a real-world scenario, by injecting the EMI signal into a Velodyne VLP-16 LiDAR. The antenna was placed at 0.5 cm from the victim LiDAR. Since the LiDAR has only 16 channels, the detection models of Section 5.1 cannot be used, as they have been trained on the KITTI dataset, which contains point cloud samples collected from a 64-channel Velodyne HDL-64E. Hence, we use PiFiNet [24], an attentive pillar network-based model trained on the JRDB dataset [36]. The JRDB Dataset is a large-scale multi-modal dataset collected from a social mobile manipulator JackRabbot [55]. This dataset contains predominantly indoor scenarios and it does not contain any vehicle classes which makes it suitable for our controlled scenario. We thus conduct our experiments using a pedestrian obstacle.

Methodology. We test the PiFiNet model under our IEMI injection with the target pedestrian obstacle at 2, 4, and 6 m away from the victim LiDAR in a static scenario. We observe that beyond 6 m, in our real-world setup, the PiFiNet model could not detect the target pedestrian obstacle even without any injection. This could be because, at far distances, the point cloud formed by VLP-16 becomes sparse. Considering this, we only analyze near LiDAR detection scenarios. We manually annotate the ground truth 3D bounding box around the pedestrian obstacle. We increment the transmitted signal gain in 1 dB intervals and measure the IOU metric of the predicted bounding box with respect to ground truth. We repeat the experiment for frequencies corresponding to sinusoidal and random perturbations. We limit the gain to 25 dB for safety.

Results and Observations. Figure 9 shows the results of PiFiNet for pedestrian object detection with EMI signal injection. We average the IOU values of 25 VLP-16 frames for each pedestrian position and signal gain. The random perturbations in the LiDAR frame correspond to a sharp drop in IOU for smaller gain values. The sinusoidal perturbations lead to a gradual drop in IOU with increasing gain, similar to what was observed for PointPillars. We

hypothesize this similarity is because both PointPillars and PiFiNet use a voxel-based method to form voxelized grids of multiple channels. The results indicate that at a 2 m target pedestrian distance, the IOU of the object drops below 0.5 for only 2 and 14 dB gains in random and sine signal injection, respectively. The average IOU at a 2 m pedestrian distance drops from 0.56 to 0.37 with a standard deviation (STD) of 0.04 over 20 frames. For 4 m and 6 m pedestrian obstacle distances, the average IOU drops from 0.38 and 0.43 to 0.22 with STD of 0.09 and 0.13. These results confirm that signal injection at higher gains leads to a greater IOU drop.

We also found that the injection induces fake obstacles into real-world data. The sinusoidal perturbations create 5.1 FPs per frame, while random perturbations create 6.3 FPs per collected frame. This reduces the precision (resp. recall) of the model from 0.33 (resp. 0.68) to 0.16 (resp. 0.56) for sinusoidal perturbations and to 0.13 (resp. 0.48) for random perturbations.

6 DEFENSES

6.1 Existing Defenses

6.1.1 Hardware Defenses. An obvious defense to prevent EMI signal injection is shielding. However, active shielding can cause inaccurate output in sensors due to parasitic interference [48, 54]. Shielding enclosures and EMI materials such as Mu-Metal [13] instead can impact the LiDAR rotation speed due to added weight. Enclosure-based shielding is also limited by the need for gaps near the photodiodes and laser diodes to allow emitting and receiving pulses. Finally, shielding can only attenuate the interference [50], and an EMI signal with sufficient gain can still be able to penetrate through the shielding and affect amplifiers, as demonstrated in previous work [31, 60].

Tu et al. [61] propose detecting EMI injection by using a dummy sensor. During IEMI injection, the offset in the output of the dummy sensor can be used to retrieve the legitimate sensor data from the original sensor. However, the use of a dummy sensor would increase the size and cost of the LiDARs. Moreover, since the EMI signal can affect differently different circuit parts, the effectiveness of this defense strategy might vary considerably based on the EMI injection directionality.

6.1.2 Point Cloud Denoising. The IEMI signal injection adds perturbations to the LiDAR point cloud data. Thus, a possible way to mitigate its effects is to reconstruct the original point cloud using denoising models [35, 43]. We study the Score Based (SB) denoising model [35], which uses score matching to estimate unconditional distributions in data. SB outperforms previous state-of-art denoising models [25, 43] under various noise signals (e.g., isotropic Gaussian noise, and simulated LiDAR noise).

Baseline Model Methodology. We use the training pipeline provided by SB to train the model on the KITTI dataset. In particular, following the model configuration, 1) we only consider the objects (e.g., cars, pedestrians, cyclists) in the 40° FOV in front of the vehicle, and 2) the cloud points were reduced to 10,000 points using the farthest point down sampling [74]. Finally, we include uniform random noise in the samples to train the model. We call this the reference Baseline Model.

Results and Observations. We evaluate the model's denoising capability on the same test samples used in Section 5.3. Since denoising aims to improve the object detection performance, we evaluate the denoising model using Object Detection Rate (ODR) and CLassification Rate (CLR) with WIOU. Table 2 shows the resulting ODR of PointPillars and PointRCNN on the denoised point clouds. The rates show that the baseline model on uniform random noise (Baseline + U) could reconstruct the original point cloud from the uniform noise perturbations. This demonstrates that the model can denoise point clouds from the KITTI dataset.

Table 2: Object Detection Rates (ODR) for No Noise, Baseline Model on Uniform Random Noise (Baseline + U), Baseline Model on IEMI perturbation (Baseline + IEMI), variation in ODR for Sinusoidal Model with sinusoidal IEMI perturbation (Sin. Mod. Var.), and variation in ODR for Random Model with random IEMI perturbation (Rand. Mod. Var.)

Class	No Noise	Baseline + U	Baseline + IEMI	Sin. Mod. Var.	Rand. Mod. Var.
PointPillars ODR (%) / CLR (%)					
Pedestrian	78/57	78/56	472	-31/-12	6/-3
Cyclist	84/77	82/77	46/7	-19/-40	1/-9
Car	98/96	97/96	67/66	-28/-30	-29/-28
PointRCNN ODR (%) / CLR (%)					
Pedestrian	81/44	81/44	36/5	6/-3	-15/-15
Cyclist	65/47	64/45	40/36	5/7	-15/-5
Car	88/87	88/87	42/41	6/4	-39/-40

We then evaluate the performance of the baseline model on our synthesized IEMI perturbations (Baseline + IEMI). The results in Table 2 show that the ODR of the baseline model significantly decreases with IEMI perturbations indicating that the model is insufficient to reconstruct the point clouds stably.

6.1.3 Adversarial Denoising. To improve the denoising results, we perform adversarial training of the SB model using our IEMI perturbations (random and sinusoidal) following the same methodology and metrics as the baseline model evaluation. We call the resulting models the Sinusoidal Model and the Random Model. Our evaluation considers the denoising effectiveness at the maximum synthesized adversarial capability (70 dB gain).

Results and Observations. Table 2 demonstrates how the ODR and CLR vary on Sine (Sin. Mod. Var.) and Random (Rand. Mod. Var.) denoising models. The evaluation shows that the best-case results increase the ODR only by 6% and the CLR by 7% out of all the scenarios. We hypothesize this is because current denoising models [25, 35, 43] are designed to consider 1-2 cm cloud point perturbations. While IEMI injection considers larger displacements.

6.1.4 Subsampling. 3D point subsampling methods might be used to improve the robustness of models under noisy data and adversarial attacks [34, 38, 68]. However, our EMI injection perturbs all the cloud points in the affected region, reducing the effectiveness of subsampling. We verify our hypothesis applying such algorithm [68] on 300 random samples (100 cars, 100 pedestrians, 100 cyclists from the KITTI dataset). For PointRCNN, the object detection rate increased by 8% at maximum, for Apollo the data show a degradation (maximum 21% drop). Finally, for PointPillars, the object detection

rate increased only for the pedestrians class by 3% maximum (we measure a drop for the remaining classes).

6.2 Proposed Defenses

6.2.1 Ground Point Based Detection. LiDAR ground filtering in AV Perception systems discards the ground cloud points from the LiDAR FOV outside the Region-of-Interest (ROI), in front of the vehicle's trajectory. Since the IEMI perturbations affect the entire point cloud region within the perturbation angle θ, it also perturbs the ground points in that region. Our methodology extracts the ground points of the LiDAR scene (e.g., using a ground point segmentation model integrated with Apollo-based Perception systems) and analyzes them to detect potential IEMI injection.

Evaluation Method and Results.

Our approach assumes that the minimum affected θ is 3°, consistent with the adversary's capability. We consider the ROI to be an 8 m distance in the front view of the LiDAR. Then, we calculate the SNR for the Z-coordinates of all the cloud points in the ROI. During the injection, we expect the SNR of the ground points to show a spike by scanning for increasing horizontal angles θ. We evaluated our methodology over the entire KITTI validation set and observed a 94.35% TNR and 98.74% TPR for binary IEMI detection. Furthermore, the methodology could predict the correct θ value with a 78.62% accuracy. The average runtime of this method over the KITTI dataset is 14.4 ms/scene with Intel Core i7-10870H CPU (2.2GHz) and 32GB RAM.

6.2.2 Point Intensity Based Detection. Our analysis shows that for high gain values (above 20 dB), the average intensity of the points in the perturbation region drops by almost five times in the case of the Velodyne LiDAR sensor family due to the static calibration settings of the sensor [26]. Thus a faster and simpler alternative defense approach can leverage this factor to detect potential injection.

Using the same methodology as the in-ground point-based detection, the potential IEMI injection is detected if a region's average point intensity is below a given threshold (which we set to 0.1). This methodology achieves a 100% TNR and TPR over the entire KITTI dataset. The strategy also predicted the perturbation angle θ with 100% accuracy. The average runtime of this method with the same Intel machine is 0.49 ms/scene.

7 OBSERVATIONS AND LIMITATIONS

Safety and Ethics Considerations. All the experiments in this work were conducted in a controlled environment using shields and Faraday Cages. In addition, all researchers were trained regarding EMI safety procedures and regulations [2]. We are in the process of notifying the LiDAR manufacturer companies about our findings.

Solid-State LiDARs. Although our proof-of-concept analysis is based on a spinning TOF 3D LiDAR (VLP-16), our experiments of Section 4.1 show this vulnerability affects the receiving part of ToF circuits. We thus believe that solid-state LiDARs that use the same basic functioning (TOF) as spinning LiDARs [7, 14] can be also affected by our IEMI. However, the adversarial control of the affected region (i.e., the horizontal angle) will depend on the scanning methodology used by the solid-state LiDAR. We leave this analysis as future work.

Point Cloud Ring. We observed that the VLP-16 LiDAR produces an external ring of cloud points at 50 m with the LiDAR location for high-gain injections. We limited our proof-of-concept evaluation to the linear behavior of the VLP-16 since the AV perception model's ROI is generally relatively closer (a max of 44 m in KITTI).

Limitations and Future Work. Our setup is limited to static scenarios evaluation, and we have not conducted end-to-end evaluations on real autonomous vehicles. We discuss the challenges and practicality of our attack under real-world constraints in Appendix A. Another limitation is that we only study and characterize the perturbation on one 3D and one 2D LiDAR. However, our preliminary analysis demonstrates the severity of the threat for TOF circuits, demanding future investigations.

8 RELATED WORK

Injection attacks on sensors have been performed using various out-of-band signals, including acoustic [71], optical [56], and EMI [31, 60]. These signals have been exploited also to affect embedded systems, including medical devices [41]. Several works [8, 11, 57] uses laser pulses to spoof fake cloud points on target LiDARs to create not-existent obstacles. Hallyburton et al. [19] use a similar methodology to attack camera-LiDAR-based sensor fusion models. Unlike the aforementioned works, EMI perturbations do not require high aiming precision.

Recent attack methodologies have synthesized adversarial objects to elude object detection in ML models [12, 59, 76]. Xiang et al. [65] initially showed that 3D point clouds can be vulnerable to projected adversarial perturbations. Tu et al. [59] create 3D-printed adversarial structures to avoid detection while, Cao et al. [10] and Abdelfattah et al. [4] show that such attacks can also evade camera-LiDAR fusion models.

9 CONCLUSIONS

Our work identifies a new class of LiDAR sensors and TOF circuit vulnerabilities resulting in manipulations of object detection algorithms used in AVs. Our evaluation results demonstrate the effectiveness of the induced perturbations using EM waves against two object detectors PointPillars and PointRCNN and one industry-grade detector Apollo. While it is clear that EMI injection can cause Perception and fusion models to fail, it is not clear the extent of this threat over AV technology. In this work, we begin to characterize this threat to understand how better address it.

ACKNOWLEDGMENTS

This research was supported in part by the National Science Foundation under CNS-1933208 and CNS-2055123, and JSPS KAKENHI Grant Number 22H00519. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

[1] 2017. KITTI Vision Benchmark: 3D Object Detection. http://www.cvlibs.net/datasets/kitti/eval_object.php?obj_benchmark=3d. Accessed: 2021-08-17.
[2] 2019. FCC - Radio Frequency Safety. https://www.fcc.gov/general/radio-frequency-safety-0.
[3] 2022. Navigant Research Names Waymo, Ford Autonomous Vehicles, Cruise, and Baidu the Leading Developers of Automated Driving Systems. https://ww

w.businesswire.com/news/home/20200407005119/en/Navigant-Research-Names-Waymo-Ford-Autonomous-Vehicles-Cruise-and-Baidu-the-Leading-Developers-of-Automated-Driving-Systems. Accessed: 2023-01-08.
[4] Mazen Abdelfattah, Kaiwen Yuan, Z. Jane Wang, and Rabab Ward. 2021. Adversarial Attacks on Camera-LiDAR Models for 3D Car Detection. In *IROS 2021*.
[5] Ron Amadeo. 2017. Google's Waymo invests in LIDAR technology, cuts costs by 90 percent. *Ars Technica, Jan* 10 (2017).
[6] Baidu Inc. 2017. Apollo. http://apollo.auto. Accessed: 2021-10-08.
[7] Benewake. 2018. CE30-A Solid State Array LiDAR Specification. http://statics3.seeedstudio.com/assets/file/bazaar/product/DE-LiDAR-CE30-A-Datasheet-V010-EN.pdf.
[8] Yulong Cao, S. Hrushikesh Bhupathiraju, Pirouz Naghavi, Takeshi Sugawara, Z. Morley Mao, and Sara Rampazzi. 2023. You Can't See Me: Physical Removal Attacks on LiDAR-based Autonomous Vehicles Driving Frameworks. In *USENIX Security 23*.
[9] Yulong Cao, Jiaxiang Ma, Kevin Fu, Sara Rampazzi, and Morley Mao. 2021. Automated Tracking System For LiDAR Spoofing Attacks On Moving Targets. In *AutoSec Workshop, 2021*.
[10] Y. Cao, N. Wang, C. Xiao, D. Yang, J. Fang, R. Yang, Q. Chen, M. Liu, and B. Li. 2021. Invisible for both Camera and LiDAR: Security of Multi-Sensor Fusion based Perception in Autonomous Driving Under Physical-World Attacks. In *IEEE S&P 2021*.
[11] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Zhuoqing Morley Mao. 2019. Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving. In *CCS '19*.
[12] Yulong Cao, Chaowei Xiao, Dawei Yang, Jing Fang, Ruigang Yang, Mingyan Liu, and Bo Li. 2019. Adversarial objects against LiDAR-based autonomous driving systems. *arXiv:1907.05418* (2019).
[13] DDL Chung. 2000. Materials for electromagnetic interference shielding. *Journal of Materials Engineering and performance* 9 (2000), 350–354.
[14] Cygbot. 2019. 2D 3D Dual Solid State TOF LiDAR. https://www.cygbot.com/2d-3d-dual-solid-state-tof-lidar.
[15] F. Fiori and P.S. Crovetti. 2002. Nonlinear effects of radio-frequency interference in operational amplifiers. *IEEE Trans. Circuits and Systems I* 49, 3 (2002), 367–372. https://doi.org/10.1109/81.989173
[16] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
[17] David S Hall. 2020. High definition lidar system. US Patent App. 15/700,543.
[18] David S Hall, Pieter J Kerstens, Yupeng Cui, Mathew Noel Rekow, and Stephen S Nestinger. 2018. LIDAR based 3-D imaging with varying pulse repetition. US Patent 10,048,374.
[19] R Spencer Hallyburton, Yupei Liu, Yulong Cao, Z Morley Mao, and Miroslav Pajic. 2022. Security analysis of camera-LiDAR fusion against black-box attacks on autonomous vehicles. In *USENIX Security 22*.
[20] Hamamatsu Photonics K.K. 2019. *APD Modules C12703 series*. Hamamatsu Photonics K.K.
[21] Abdullah Hamdi, Sara Rojas, Ali Thabet, and Bernard Ghanem. 2020. AdvPC: Transferable adversarial perturbations on 3D point clouds. In *ECCV 2020*.
[22] Mark Harris. 2017. GM Cruise Snaps Up Solid-State Lidar Pioneer Strobe Inc. *IEEE Spectrum* 11 (2017).
[23] Zhongyuan Hau, T Kenneth, Soteris Demetriou, and Emil C Lupu. 2021. Object Removal Attacks on LiDAR-based 3D Object Detectors. In *Workshop on Automotive and Autonomous Vehicle Security (AutoSec)*, Vol. 2021. 25.
[24] Yuhang He, Wentao Yu, Jie Han, Xing Wei, Xiaopeng Hong, and Yihong Gong. 2021. Know your surroundings: Panoramic multi-object tracking by multimodality collaboration. In *CVPR 2021*. 2969–2980.
[25] Chao Huang, Ruihui Li, Xianzhi Li, and Chi-Wing Fu. 2020. Non-local part-aware point cloud denoising. *arXiv preprint arXiv:2003.06631* (2020).
[26] Velodyne LiDAR Inc. 2019. USER'S MANUAL AND PROGRAMMING GUIDE: VLP-16 Velodyne LiDAR Puck Rev E. https://velodynelidar.com/wp-content/uploads/2019/12/63-9243-Rev-E-VLP-16-User-Manual.pdf.
[27] J.Schesser. 2017. Sampling and Aliasing. https://web.njit.edu/~joelsd/Fundamentals/coursework/BME310computingcw6.pdf.
[28] Chaouki Kasmi and Jose Lopes Esteves. 2015. IEMI threats for information security: Remote command injection on modern smartphones. *IEEE Trans. Electromagnetic Compatibility* 57, 6 (2015), 1752–1755.
[29] Mandeep Kaur, Shikha Kakar, and Danvir Mandal. 2011. Electromagnetic interference. In *2011 3rd International Conference on Electronics Computer Technology*.
[30] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. 2018. Joint 3d proposal generation and object detection from view aggregation. In *IROS 2018*. IEEE.
[31] Denis Foo Kune, John Backes, Shane S Clark, Daniel Kramer, Matthew Reynolds, Kevin Fu, Yongdae Kim, and Wenyuan Xu. 2013. Ghost Talk: Mitigating EMI signal injection attacks against analog sensors. In *IEEE S&P 2013*. 145–159.
[32] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. 2019. Pointpillars: Fast Encoders for Object Detection from Point Clouds. In *CVPR 2019*.
[33] Timothy B. Lee. 2018. Why Spinning Lidar Sensors Might be Around for Another Decade. Ars Technica.
[34] Xingyi Li, Wenxuan Wu, Xiaoli Z Fern, and Li Fuxin. 2021. The devils in the point clouds: Studying the robustness of point cloud convolutions. *arXiv preprint arXiv:2101.07832* (2021).
[35] Shitong Luo and Wei Hu. 2021. Score-based point cloud denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
[36] Roberto Martin-Martin, Mihir Patel, Hamid Rezatofighi, Abhijeet Shenoi, JunYoung Gwak, Eric Frankel, Amir Sadeghian, and Silvio Savarese. 2021. Jrdb: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *IEEE transactions on pattern analysis and machine intelligence* (2021).
[37] Mini-Circuits. 2019. High Power Amplifier ZHL-5W-202-S+. https://www.minicircuits.com/pdfs/ZHL-5W-202-S+.pdf.
[38] Ehsan Nezhadarya, Ehsan Taghavi, Ryan Razani, Bingbing Liu, and Jun Luo. 2020. Adaptive hierarchical down-sampling for point cloud classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
[39] Seong Joon Oh, Bernt Schiele, and Mario Fritz. 2019. Towards reverse-engineering black-box neural networks. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning* (2019), 121–144.
[40] Soham Pal, Yash Gupta, Aditya Shukla, Aditya Kanade, Shirish Shevade, and Vinod Ganapathy. 2019. A framework for the extraction of deep neural networks by leveraging public data. *arXiv preprint arXiv:1905.09165* (2019).
[41] Youngseok Park, Yunmok Son, Hocheol Shin, Dohyun Kim, and Yongdae Kim. 2016. This ain't your dose: Sensor spoofing attack on medical infusion pump. In *WOOT 16*.
[42] Scott Drew Pendleton, Hans Andersen, Xinxin Du, Xiaotong Shen, Malika Meghjani, You Hong Eng, Daniela Rus, and Marcelo H Ang Jr. 2017. Perception, planning, control, and coordination for autonomous vehicles. *Machines* 5, 1 (2017), 6.
[43] Francesca Pistilli, Giulia Fracastoro, Diego Valsesia, and Enrico Magli. 2020. Learning graph-convolutional representations for point cloud denoising. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*. Springer, 103–118.
[44] Milica Popovic, BD Popovic, and Z Popovic. 2006. Electromagnetic induction. *Fundamentals of Engineering Electromagnetics* (2006).
[45] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. 2017. Pointnet: Deep learning on point sets for 3D classification and segmentation. In *CVPR 2017*.
[46] Rui Qian, Xin Lai, and Xirong Li. 2022. 3d object detection for autonomous driving: a survey. *Pattern Recognition* 130 (2022), 108796.
[47] Ettus Research. 2012. USRP N200/N210 Networked Series. https://www.ettus.com/wp-content/uploads/2019/01/07495_Ettus_N200-210_DS_Flyer_HR_1.pdf.
[48] Ferran Reverter, Xiujun Li, and Gerard C M Meijer. 2006. Stability and accuracy of active shielding for grounded capacitive sensors. *Measurement Science and Technology* 17, 11 (sep 2006), 2884. https://doi.org/10.1088/0957-0233/17/11/004
[49] JM Roe. 1978. *Integrated Circuit Electromagnetic Susceptibility Handbook. Phase III*.
[50] Sowmya Sankaran, Kalim Deshmukh, M. Basheer Ahamed, and S.K. Khadheer Pasha. 2018. Recent advances in electromagnetic interference shielding properties of metal and carbon filler reinforced flexible polymer composites: A review. *Composites Part A: Applied Science and Manufacturing* (2018). https://doi.org/10.1016/j.compositesa.2018.08.006
[51] Shanghai Slamtec Co., Ltd. 2019. RPLiDAR S2. https://www.slamtec.com/en/S2.
[52] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. 2019. PointRCNN: 3D object proposal generation and detection from point cloud. In *CVPR 2019*.
[53] Giuseppe Schirripa Spagnolo, Adriana Postiglione, and Ilaria De Angelis. 2020. Simple equipment for teaching internal photoelectric effect. *Physics Education* (2020).
[54] Enrique Mario Spinelli and Ferran Reverter. 2009. On the stability of shield-driver circuits. *IEEE Transactions on Instrumentation and Measurement* (2009).
[55] Stanford Vision and Learning Lab. 2017. JackRabbot. https://svl.stanford.edu/projects/jackrabbot/.
[56] Takeshi Sugawara, Benjamin Cyr, Sara Rampazzi, Daniel Genkin, and Kevin Fu. 2020. Light Commands: Laser-Based Audio Injection Attacks on Voice-Controllable Systems. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 2631–2648. https://www.usenix.org/conference/usenixsecurity20/presentation/sugawara
[57] Jiachen Sun, Yulong Cao, Qi Alfred Chen, and Z Morley Mao. 2020. Towards robust LiDAR-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures. In *USENIX Security 20*. 877–894.
[58] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv:1312.6199* (2013).
[59] James Tu, Mengye Ren, Sivabalan Manivasagam, Ming Liang, Bin Yang, Richard Du, Frank Cheng, and Raquel Urtasun. 2020. Physically Realizable Adversarial Examples for LiDAR Object Detection. In *CVPR 2020*.
[60] Yazhou Tu, Sara Rampazzi, Bin Hao, Angel Rodriguez, Kevin Fu, and Xiali Hei. 2019. Trick or heat? Manipulating critical temperature-based control systems

using rectification attacks. In *CCS '19*. 2301–2315.
[61] Yazhou Tu, Vijay Srinivas Tida, Zhongqi Pan, and Xiali Hei. 2021. Transduction Shield: A Low-Complexity Method to Detect and Correct the Effects of EMI Injection Attacks on Sensors *(ASIA CCS '21)*. Association for Computing Machinery.
[62] Zhixin Wang and Kui Jia. 2019. Frustum ConvNet: Sliding Frustums to Aggregate Local Point-Wise Features for Amodal 3D Object Detection. *IROS 2019* (2019).
[63] David A Ware. 2017. *Effects of intentional electromagnetic interference on analog to digital converter measurements of sensor outputs and general purpose input output pins*. Ph. D. Dissertation. Utah State University.
[64] Yuxin Wen, Jiehong Lin, Ke Chen, CL Philip Chen, and Kui Jia. 2020. Geometry-aware generation of adversarial point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 6 (2020), 2984–2999.
[65] Chong Xiang, Charles R Qi, and Bo Li. 2019. Generating 3D adversarial point clouds. In *CVPR 2019*. 9136–9144.
[66] Zhifei Xu, Runbing Hua, Jack Juang, Shengxuan Xia, Jun Fan, and Chulsoon Hwang. 2021. Inaudible Attack on Smart Speakers With Intentional Electromagnetic Interference. *IEEE Transactions on Microwave Theory and Techniques* (2021).
[67] Bin Yang, Wenjie Luo, and Raquel Urtasun. 2018. Pixor: Real-time 3D object detection from point clouds. In *CVPR 2018*. 7652–7660.
[68] Jiancheng Yang, Qiang Zhang, Rongyao Fang, Bingbing Ni, Jinxian Liu, and Qi Tian. 2019. Adversarial attack and defense on point sets. *arXiv preprint arXiv:1902.10899* (2019).
[69] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 2020. 3dssd: Point-based 3d single stage object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11040–11048.
[70] Kentaro Yoshioka. 2022. A Tutorial and Review of Automobile Direct ToF LiDAR SoCs: Evolution of Next-Generation LiDARs. *IEICE Transactions on Electronics* (2022).
[71] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. 2017. Dolphinattack: Inaudible voice commands. In *CCS '17*.
[72] Min Zhang. 2021. A CAPACITIVELY COUPLED PIN INJECTION METHOD – AN ALTERNATIVE TO BCI TEST.
[73] Qifan Zhang, Junjie Shen, Mingtian Tan, Zhe Zhou, Zhou Li, Qi Alfred Chen, and Haipeng Zhang. 2022. Play the Imitation Game: Model Extraction Attack against Autonomous Driving Localization. In *Proceedings of the 38th Annual Computer Security Applications Conference*. 56–70.
[74] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. 2018. Open3D: A Modern Library for 3D Data Processing. *arXiv:1801.09847* (2018).
[75] Yin Zhou and Oncel Tuzel. 2018. Voxelnet: End-to-end learning for point cloud based 3D object detection. In *CVPR 2018*. 4490–4499.
[76] Yi Zhu, Chenglin Miao, Tianhang Zheng, Foad Hajiaghajani, Lu Su, and Chunming Qiao. 2021. Can We Use Arbitrary Objects to Attack LiDAR Perception in Autonomous Driving?. In *CCS '21*.

APPENDICES

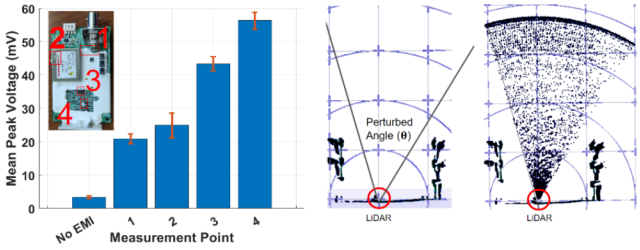

Figure 10: (Left) Remote signal injection at various measurement points (no injection, a PCB junction (1), a voltage regulator (2), and two internal operational amplifiers (3-4)). (Right) Before and during the IEMI injection at θ=45° (VLP-16).

A ADVERSARIAL SETUP DETAILS

To perturb the point cloud in a specific target region of the LiDAR's FOV, the attacker needs to synchronize with the spinning LiDAR transmitted pulses. In our methodology (used in previous works for laser injection attacks against LiDAR sensors [8, 12]), a photodiode with a wide FOV is used to capture LiDAR pulses and activate the IEMI injection in the window of time when the LiDAR scans a particular region (horizontal angle). To target a specific object in the LiDAR FOV, the attacker can use a tracking system to track the position of moving objects with respect to the victim AV [9].

Aiming to a Moving Vehicle. Attacking a moving vehicle from the roadside involves tracking the victim LiDAR and aiming the antenna toward the sensor location. Prior research [8, 9] has shown that laser-based spoofing attacks using our same synchronization methodology are feasible against moving vehicles, despite the LiDAR scan signal becoming sparser at greater distances. To overcome this, the team used multiple photodiodes and a robot turret. It's worth noting that the precision required for aiming a directional antenna is lower than that for a focused laser beam, as the antenna's transmission coverage area is wider. Additionally, a directional antenna with high coverage width can provide more stable injection at larger distances.

B PERTURBATION MODELING DETAILS

Modeling Sinusoidal Perturbations. We characterize the sinusoidal perturbation using the following system of equations of a typical sine wave:

$$\begin{cases} y = A_y \cdot \cos(2 \cdot \pi \cdot x/\lambda_y) \\ z = A_z \cdot \cos(2 \cdot \pi \cdot x/\lambda_z) \end{cases} \quad (1)$$

where x, y, and z represent Cartesian coordinates of a given point in the point cloud in the 3D space, the λs are the wavelengths of the sinusoidal pattern in a given plane, and the As are the sinusoidal amplitude measured in a given plane. Note that we model the perturbation as a sinusoidal wave in the time domain, while in our case, it represents a spatial displacement of the cloud points in each vertical line. We characterize the change in point intensity values based on the relation. The amplitude A is then modeled as a linear function of the gain G, as $A = \alpha \cdot G + \beta$ with $\alpha = 0.00032$ and $\beta = 0.00228$ for the XZ plane and $\alpha = 0.0019$ and $\beta = 0.0012$ for the XY plane as derived from our empirical experiments. In the same fashion, λ is defined as a linear function of the target object distance d_O, as $\lambda = \gamma \cdot d_O + \phi$, where $\gamma = 0.07$ and $\phi = 0.001$ for the XZ plane and $\gamma = 0.07$ and $\phi = 0.0004$ for the XY plane.

Modeling Random Perturbations. We characterize the random perturbations with respect to point cloud distributions within a certain range ΔR for each LiDAR horizontal point cloud line, as shown in Figure 3 (Right). Based on our empirical experiments in Section 4.2, the random point cloud displacement and range only depend on the injection gain G. We thus model such displacement ΔD as a standard deviation with respect to the original point cloud location before the EMI injection as $\Delta D = \alpha \cdot G + \beta$ with $\alpha = 0.0005$ and $\beta = 0.15$ for the XY plane and $\alpha = 0.0003$ and $\beta = 0.13$ for the XZ plane. $\Delta R = \gamma \cdot G + \phi$ models the affected range along each horizontal point cloud line. Where, for the XY plane, $\gamma = 0.558$ and $\phi = 1.847$ for maximum displacement and $\gamma = -0.298$ and $\phi = -0.421$ for minimum displacement. The maximum displacement in the XZ plane is given by $\gamma = 0.121$ and $\phi = 0.364$, and the minimum displacement by $\gamma = -0.082$ and $\phi = -0.758$. Similar to sinusoidal perturbations, we characterize the change in point intensity values based on our empirical experiments.