

Illusion and Dazzle: Adversarial Optical Channel Exploits against Lidars for Automotive Applications

Hocheol Shin, Dohyun Kim, Yujin Kwon, and Yongdae Kim

Korea Advanced Institute of Science and Technology, Dajeon, Republic of Korea,
{h.c.shin, dohyunjk, dbwls8724, yongdaek}@kaist.ac.kr

Abstract. With the advancement in computing, sensing, and vehicle electronics, autonomous vehicles are being realized. For autonomous driving, environment perception sensors such as radars, lidars, and vision sensors play core roles as the eyes of a vehicle; therefore, their reliability cannot be compromised. In this work, we present a spoofing by relaying attack, which can not only induce illusions in the lidar output but can also cause the illusions to appear closer than the location of a spoofing device. In a recent work, the former attack is shown to be effective, but the latter one was never shown. Additionally, we present a novel saturation attack against lidars, which can completely incapacitate a lidar from sensing a certain direction. The effectiveness of both the approaches is experimentally verified against Velodyne’s VLP-16.

Keywords: attack, autonomous car, sensor, lidar, saturating, spoofing

1 Introduction

Of late, in the automotive industry, there is a trend shift towards autonomous vehicles. Most of the major automotive manufacturers have researched and/or invested in this technology and even companies outside the vehicular domain are considering autonomous vehicles as profitable future business ventures. In realizing autonomous vehicles, especially environment perception sensors such as radars, object-recognizing cameras, ultrasonic sensors, and lidars are critical; major sensor manufacturers (e.g. Velodyne, IBEO, and Mobileye) are attracting as much attention as the vehicle manufacturers.

Among the various environment perception sensors, the lidar, the target sensor in this work, has its own advantages that cannot be found in the other sensors. Compared to the current automotive radars and cameras, lidars have a considerably higher resolution and precision. Lidars can work both at daytime and nighttime unlike cameras, and can also recognize lanes, license plates, and street signs due to their retro-reflective surfaces [3]. These exclusive strengths render the lidar essential in autonomous driving platforms; they can be found on almost all autonomous vehicles except Tesla [17].

Although they are beneficial, lidars may be vulnerable to intentional external interferences, because they must be exposed to the outside. If the lidar in an autonomous vehicle is deceived by an attacker, it can lead to lethal outcomes, similar to a blind driver or a driver viewing illusions. Despite these risks, security against such threats are not being considered in the design of automotive




Fig. 1. Simplified illustration of a three-layer lidar operation.

lidars. In fact, during Black Hat Europe 2015, Petit et al. presented a work on remotely tempering a camera (Mobileye C2-270) and a lidar (IBEO LUX 3), with light [30]. Against the target lidar, they successfully induced multiple fake dots—sensed points that are not from real objects, but generated by the injected signal—in a wall-like shape by relaying and replaying the received lidar pulses with an intentionally added delay; they even induced multiple copies of the wall-like shape by repeating the waveform. However, they were only able to induce fake dots, further than the location of spoofer (this has even been specified as a limitation of their work). This is a critical limitation as an attack because the further the object is, the lesser is its effect on the victim vehicle. Therefore, at the time of spoofing, the most threatening object to the victim vehicle would not be the induced fake dots, but the attacker herself.

In this work, we have addressed such limitations. We demonstrated that it is possible to induce fake dots *closer* than the spoofer location. We also detail the *actual* attack process, which is considerably more complex than that of the previous work, such that the described process and parametric setup would be sufficient for other researchers to reproduce this work. Note that, inducing closer fake dots would not be possible without such detailed understanding of the process. Apart from the aforementioned contributions, we present a novel saturation attack against the lidars. By illuminating the lidar with a strong light of the same wavelength as that the lidar uses, we can actually erase the existing objects in the sensed output of the lidar. This approach was inspired by the work of Park et al., wherein they blinded a drop sensor in a medical infusion pump and rendered it unable to sense the fluid drops [29]. We also discovered that curved reception glass, which a number of off-the-shelve lidars adopt, can pose a severe threat to the lidar due to refraction/reflection. The target lidar we used to show the effectiveness of our attack was Velodyne’s VLP-16, which was never analyzed previously. In addition, we discuss practical aspects of the presented attacks along with several detailed scenarios. We also present multiple approaches to mitigate our attacks, and their limitations. Our contributions can be summarized as follows:

- We present the process of inducing fake dots closer than the spoofer location. This was considered to be impossible in the previous work.
- We introduce a saturation attack against the lidars, which can incapacitate a lidar from detecting objects.
- We present the attack process in considerable details for reproducibility.
- We discuss, in-depth, the resolution of problems pertaining to the deployment of attacks in reality, with detailed attack scenarios.

The remainder of this paper is organized as follows. Section 2 provides the required backgrounds for understanding this work. Section 3 presents the attack schemes for both attacks, and Section 4 the attack results. Sections 5 and 6 include the discussions and the related works, respectively. Finally, we conclude the study in Section 7.

2 Background

2.1 Lidar

Lidar is an *active remote sensing* method, or a sensor using this method to measure the distances to nearby objects. Here, *active sensing* is a way of analyzing the target of interest by exposing it to the energy (or signal) intentionally transmitted by the sensor itself. It is distinguished from the opposite, *passive sensing*, which examines the target of interest only by receiving energy from it. *Remote sensing* is a way of analyzing the target of interest without physical contact; examples include the telescope, radar, and seismometer.

The lidar was devised shortly after the advent of the laser, as a laser ranging device for the lunar laser ranging experiment [2]. Since then, it has been widely applied in fields such as meteorology [11], agriculture [40], topography [43], and altimetry [23]. Since the adoption of the lidar as one of the sensory systems for the test vehicle in the *DARPA*-funded Autonomous Land Vehicle project [31], its usage has expanded to advanced driver assistance systems [4, 8] and autonomous driving platforms [14, 10].

Limiting the scope of the environment perception sensors to automotive systems, there are roughly two types of lidars: *scanning* and *solid-state*. Scanning lidars are mainly composed of a/multiple laser transceiver(s) and a moving rotary system for scanning; they acquire an around-view by rotating the laser transceiver. However, the moving parts of scanning lidars contribute to its high cost and are limited in their reliability/durability. In contrast, solid-state lidars do not require moving parts for steering their laser beams. Although affordable solid-state lidars with acceptable performances are the ultimate goal of lidar manufacturers, currently, scanning lidars are dominant in the market due to lack of technical advancements, and solid-state lidars with equivalent performances are generally considered as the next-generation lidars [1, 12, 33]. Therefore, we confine our interest to scanning lidars only; in most cases, scanning lidars are denoted as lidars, for the rest of this work.

The working of a lidar is similar to that of a pulsed radar, and is quite simple. First, a lidar transmits a laser pulse, while spinning. When the transmitted pulse hits an object, a part of the transmitted energy reflects back to the lidar, as an echo. Note that, there can be multiple echoes, when the object does not fully block the transmitted pulse, possibly resulting in multiple echoes. Then, the echo(es) are received by the lidar, and the elapsed time (Δt) is measured. As light has a known constant speed (c) in air, the lidar can derive the distance (l) to the object using the following equation:

$$l = c\Delta t/2 \tag{1}$$

The lidar can also determine the direction in which the pulse is transmitted, from the rotation angle of its spin. Knowing both the direction and the distance, the lidar can *map* points. The lidar rotates to cover its field of view, resulting in a point cloud, i.e., the set of all the measured points. Multi-layer lidars either have multiple copies of this system with vertical slant angles between them or they also scan vertically. Fig. 1 illustrates the operation of a multi-layer lidar.

As the pulses are transmitted periodically, there are ambiguities in determining the elapsed time of the received echoes. Assuming that an echo was received, after the last pulse was transmitted, and that the elapsed time is Δt , the echo can either be that of the last transmitted pulse or of one of the previous pulses'. Denoting the Pulse Repetition Time (PRT) as T , the elapsed time can be any of $\Delta t + nT$. Therefore, to limit uncertainties, lidars and pulsed radars define the *receiving time* (Δt_{max}) and *dead time* (D). Whenever a pulse is transmitted, a lidar waits for its echoes, for the duration of the receiving time, and every echo received in that interval is considered as that of the last transmitted pulse. After the receiving time ends, the lidar ignores all the incoming pulses for the duration of the dead time; then, the next pulse is transmitted. This establishes the relationship, $\Delta t_{max} + D = T$; the *maximum distance* (l_{max}) of a lidar can be derived using Eq. (1) to be $l_{max} = c\Delta t_{max}/2$. Fig. 2 illustrates these relationships.

Additionally, for a lidar, a wide receiving angle (size of the receiver aperture) is not required, if it is precisely calibrated. Only the echoes falling into the receiving angle can effectively affect the sensing result. The receiving aperture needs to cover the direction of the pulse transmission only during the maximum round-trip time (Δt_{max}) of the light pulse. Thus, we can derive the minimum required receiving angle (Θ_R [°]) from the rotating speed (ω [°/s]) and the maximum distance of the lidar, as per the following equation:

$$\Theta_R = \Delta t_{max} \cdot \omega = \frac{2l_{max}}{c} \cdot \omega \text{ [°]} \quad (2)$$

Because the rotating speed of a lidar is numerically much smaller than that of light ($\omega \ll c$), and the maximum distance is in the range of several hundred meters, the minimum required receiving angle is very small. For example, this value is only 0.0048° ¹ for the Velodyne's VLP-16.

2.2 Sensor Attacks

Although it is not long since sensor attacks drew significant attention from the security academia, several researchers have studied various approaches in compromising the sensors and defending them. Given diverse types of attack channels for sensor attacks, Shin et al. [34] classified them into three types: *regular*, *transmission*, and *side* channel. Regular channel attacks target the sensing structure using the same type of physical quantity sensed by the victim sensor, e.g. sound wave for a microphone. Side channel attacks likewise target the sensing structure as in regular channel attacks, but use a physical quantity other than the one sensed by the target sensor, as in the case where Son et al. [37] affected gyroscope sensing results with acoustic stimuli. Lastly, transmission channel attacks

¹ $2 \cdot 100/3e8 \text{ [s]} \times 360 \cdot 20 \text{ [°/s]}$. Note that 20Hz is the maximum update rate of VLP-16.




Fig. 2. Relationship between the PRT, receiving time, and dead time.

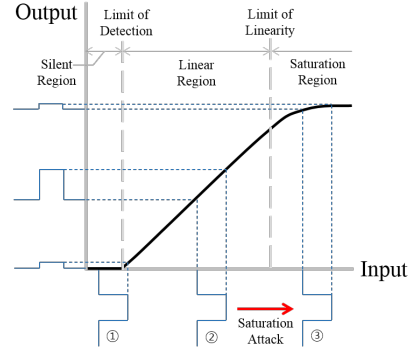


Fig. 3. Typical transition curve of a sensor and input-output relationships in the three regions of the curve: the *silent* (①), *linear* (②), and *saturation* (③) regions.

influence the channel connecting the sensing structure and the other parts of the system. For example, Foo Kune et al. intentionally induced electromagnetic interference (EMI) in the wire connecting an analog sensor and an amplifier to overwrite the sensor output [9]. For the rest of this paper, we focus on the regular channel attack, because the following two types of attacks against lidars all belong to that type.

Sensor Saturating

All sensors can be viewed as a form of transducers because they convert one type of inbound physical quantity into another type (mostly electric). Although it is ideal for transducers (particularly for sensors) to have linear transition curve, a certain degree of nonlinearity is inevitable. Fig. 3 depicts a typical sensor transition curve, and its input range can be divided into three regions. First, the *silent* region is an input range below the threshold of the sensor. The threshold also can be called the “Limit of Detection”, because input signals below the threshold will not be detected. Thus, the output of the sensor will be the same as that for a zero input signal, which is natural because every sensor has a limited sensitivity. Second is the *linear* region, which is the intended operation region or the dynamic range of the sensor. By design, all sensors should be guaranteed to work in this region, because the output is proportional to the input only in this region. As the input increases over the “Limit of Linearity”, the *saturation* region starts. In this region, the curve again becomes nonlinear, and the sensor cannot reflect the input changes well.

The principle of saturating is to push the overall level of the input signal (②) into the saturation region (③), in order to render the sensor unable to reflect the variations in the legitimate input signal. As shown in Fig. 3, an attacker can incapacitate a sensor by exposing it to excessive stimuli (②→③).

Sensor Spoofing

Different from saturating, whose goal is the denial-of-service (DoS), the goal of the sensor spoofing is to deceive the victim sensor. The attacker deceives

the victim sensor by exposing it to the attacking signal which simulates the circumstance that the attacker wants the sensor to believe. Simulating a fake circumstance exploits the *semantic gap* between what the circumstance really is and how the sensor perceives it to be. For example, an earthquake and a child shaking a seismometer are totally different, but it can seem similar to the sensor. Therefore, fabricating reality itself, e.g. spoofing a smoke detector by generating a real smoke, is not considered sensor spoofing.

For active sensors, in particular, sensor spoofing can be performed in more specialized forms. As mentioned in Section 2.1, active sensors expose the target of examination to their own energy; an active sensor can take a particular waveform (ping waveform) to differentiate its echoes from the other inbound signals. Therefore, the attacker should first acquire the ping waveform, and then *relay* it after an intentionally inserted delay to affect the victim sensor; this is called *sensor spoofing by relaying*. Besides, the received ping waveform can be duplicated during relaying, to amplify the effect.

The advantage of sensor spoofing is that it is not easy for the victim sensor to determine whether it is real or not. In many cases, it is almost impossible to detect the attack without external aids.

3 Attack Methods

3.1 Target System

We assume that the target is a scanning lidar system exposed to the exterior due to its role as an environment perception sensor. Although we focus on lidars for autonomous driving applications because attacking them leads to the most severe outcomes, the following attack schemes can also be applied to lidars for other types of applications, as long as they operate similarly.

For the case of inducing fake dots closer than the spoofer location, we assume one more condition: the ping waveform remains unchanged or at least changes predictably. We confirmed that most of the real-world lidar products for autonomous applications would meet this condition. We could not find any product with a random ping waveform as part of the specification. This can be cross-confirmed by measurements. We analyzed the Velodyne VLP-16 to confirm that it has a consistent ping waveform, and we could also infer that the IBEO LUX 3 had consistent ping waveform by examining the work of Petit et al. [30].

3.2 Attack Model

We list different models for the two types of attacks: saturating and spoofing. This is because the required attacker capabilities are different for each.

Saturating: The attacker can inject an attacking light into the target sensor remotely. The attacking equipment can transmit light, whose wavelength is the same as that used by the target, with sufficient intensity to saturate the target receiver. This includes the ability to aim and focus onto the target sensor.

Spoofing by Relaying: In addition to the ability to inject an attacking light into the target sensor, the attacker can receive a signal from the target. Thus, the attacker has both a receiver and transmitter to receive and inject.



Fig. 4. Lidars with curved reception glasses. Velodyne’s VLP-16, HDL-32E, IBEO’s LUX Mini, and Quanser’s M8 (from left).

3.3 Saturating

As described in Section 2.2, saturating renders the victim sensor unable to reflect the input signal changes. This line of attack is powerful, because saturation itself is unavoidable. The victim systems can easily detect the attack², but cannot prevent the sensor from saturating. As the size of the sensor output curve’s linear region is limited, irrespective of its size, its output will start to saturate at a certain input strength. This also applies to lidars, and by exploiting it, attackers can effectively perform DoS attacks. As the medium used for the attack is light, saturating against lidars can also be called *blinding*.

Lidars can be saturated by exposing the target lidar to an intensive light source with the same wavelength as that used by the lidar. We observed numerous induced fake dots with a weak light source, and the complete blinding of a certain direction with a strong light source. The effect of saturating will be described and illustrated in detail, in Section 4.2. The following points are characteristics common to the saturation attacks against lidars:

Stealthiness against Drivers and Pedestrians: In order to not hinder human driving and for eye safety, lidars use infrared (IR) lasers for their operation. The invisibility of the medium also assists stealthiness in saturating. Even if the target lidar is saturated by a high-intensity IR light source, human drivers and pedestrians would be unaware, rendering the attack effective.

Receiving Angle: As mentioned in Section 2.1, a wide receiving angle is not essential for lidars to sense objects in the field of view. Therefore, lidar receivers typically have much smaller receiving angles compared to the angle of view (360° for the case of VLP-16) of the lidar. This can limit the effect of saturating, because the attacking light comes from a certain direction, when the lidar is rotating. As a result, saturating cannot affect target’s field of view universally, but disturbs only a fan-shaped part of it; the angle of disturbance would be proportional to the receiving angle. Referring to Eq. (2), the minimum receiving angle for meeting the specification is sufficient to render saturating impractical. In reality, however, we found that the receiving angles of lidars are much larger than required, rendering them significantly more vulnerable to saturating, even without adopting multiple light sources to widen the angle of disturbance.

Curved Reception Glass: Due to the small receiving angles of scanning lidars, it can only affect the sectors in the direction of the attacker. However, we found that an oblique incidence of strong light onto the curved reception glass of VLP-16 can cause the appearance of fake dots in directions other than that of the

² However, we could not find any function alerting the occurrence of saturation.




Fig. 5. Speculations of how the oblique incidence of light onto a curved reception glass induces fake dots in a direction different from that of the actual light source.

attacking light source. In addition to VLP-16, there are several lidars with curved reception glasses e.g. the Velodyne HDL-32E, IBEO LUX mini, and Quarnergy M8 (Fig. 4). Although we are not 100% sure because we were only able to conduct a non-destructive analysis, the above-mentioned occurrence is most likely due to refraction or reflection on the curved glass surface. Fig. 5 illustrates these speculations. Fake dots in directions other than the direction of the attacker can be a severe threat to the victim, because the detected points have different significances according to their directions on roads. For example, an autonomous vehicle should not be hindered by vehicles on the other lanes, even if they are very close. Now assume the attacker vehicle is located slightly ahead of the victim’s vehicle in the lane next to the victim’s; exploiting the above effect, the attacker can generate fake dots in front of the victim, where nothing exists in fact. Fig. 6 depicts this attack scenario.

3.4 Spoofing by Relaying

Our approach for spoofing by relaying is basically the same as the principle used in the relaying attack method proposed by Petit et al. [30]. In this work, however, we also provide a method to generate fake dots *closer* than the attacker position. This was listed as one of the limitations of the Petit et al.’s work. We first start with the ideal process to understand how spoofing by relaying works in general, then discuss the actual process.

Ideal Attack Process

Lidars measure distances by measuring the round-trip time of the flight of light. A fired laser pulse flies until it meets an object, and is then reflected back to the lidar. Ideally, the procedure for spoofing by relaying is to mimic this process:

1. Prepare an attack tool composed of a receiver, an adjustable delay component, and a transmitter of the same wavelength as that used by the lidar.
2. Aim at the target lidar with the attack equipment.
3. Receive the target lidar pulse signal using the receiver.
4. Add the required delay using the delay component.
5. Fire a laser pulse back to the target lidar using the transmitter.

Theoretically, this process would induce only one fake dot, and the required delay (d_i) in step 4 to generate a fake dot at a distance (l) can be determined as follows. Let the distance between the spoofer and the victim lidar be l_s ; $l_s \leq l$




Fig. 6. Attack scenario exploiting a curved reception glass. The attacker and victim vehicles are heading the same direction, and the attacker obliquely illuminates the victim’s lidar with a strong light source.

because we cannot add a negative delay. Therefore, d_i should be the delay, which makes an echo appear $l - l_s$ further than the spoofer, i.e. the round-trip time of light for the distance, $l - l_s$. Using Eq. (1), it is derived as,

$$d_i = \frac{2(l - l_s)}{c} \quad (3)$$

Although the basic procedure is as mentioned above, there are two other points to be considered. One is the limited lidar receiving angle. Even if the attacker fires attacking pulses to the victim lidar, they cannot affect the victim, when the victim’s receiver is not facing the attack direction. Therefore, the attacking pulse should reach the victim lidar, while it is still within the receiving angle. The other is the lidar receiving time; as discussed in Section 2.1, lidars ignore echoes with delays larger than a certain threshold derived from their range, i.e. the maximum measurable distance. Only echoes that fall within the receiving time can affect the measurement. This applies to the attacker also; therefore, the attacker should fire back to the target lidar within the receiving time. For example, VLP-16 has a range of 100m, which results in a receiving time of $(2 \times 100\text{m}) / (3 \times 10^8\text{m/s}) = 667\text{ns}$. Therefore, in order to affect the measurement of the VLP-16, an attacker should fire back at least within 667ns.

Actual Attack Process

Although theoretically, the attack process is as discussed above, the actual process is quite different. First, the laser pulse from the lidar diverges. Accordingly, the attacker receiver obtains multiple adjacent laser pulses; however, only a part of these pulses exactly head in the direction of the receiver. This enables the attacker to detect the target’s laser pulse a few PRTs (T) in advance, compared to the case where the laser pulses do not diverge at all. Next, irrespective of how close the receiver and transmitter are placed in the attack tool, they are apart by a certain distance. Let us assume that they are arranged horizontally; as the horizontal resolution of scanning lidars are typically high, the laser pulse heading to the receiver and the pulse to the transmitter is not temporally the same. Consequently, there is a time difference (S) between the detection of a laser pulse by the receiver and the firing of a pulse toward the transmitter. Note that the round-trip time of light would have almost no effect, because the speed of light is much faster than the rotating speed of the lidar.

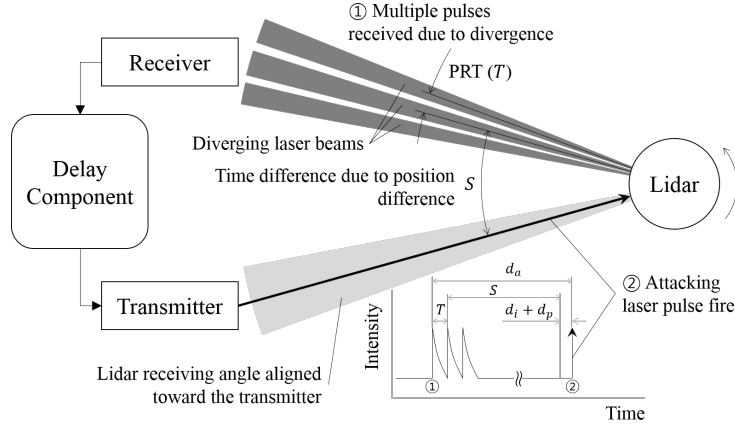


Fig. 7. Actual attack process: As the lidar rotates, multiple laser pulses, temporally separated by the PRT (T), are first captured by the attacker receiver (①). Then, after the actual required delay (d_a), the attacking laser pulse is fired (②). The graph below displays the temporal arrangement of events.

Owing to the above-mentioned phenomena, the required delay to induce a fake dot at a certain distance differs from Eq. (3) due to the time differences, T and S . Assuming that the receiver is illuminated by the lidar, before the transmitter and denoting the signal processing/propagation delay as d_p ,

$$d_a = d_i + nT + S + d_p \quad (4)$$

The time differences (T and S) are compensated by adding them to the ideal delay, because the delay component is triggered by the first received pulse. The delay d_p can be compensated likewise, because it is a constant delay which can be measured in advance. Note that, n multiplied by T is for the case, where the delay component is triggered multiple PRTs in advance. In addition, although $n, S \rightarrow 0$ as the distance between the lidar and the spoofer increases, attackers can enlarge n and S by increasing the receiver aperture size and the receiver-transmitter separation, respectively. Fig. 7 illustrates this process. This can be used for making a virtually *negative-valued delay* to generate fake dots closer than the attacker location. Assuming that $l < l_s$ in Eq. (3), d_i becomes negative. However, d_a will remain positive, because $T, S \gg |d_i|$.

A Notable Characteristics of Spoofing by Relaying Attack

- Stealthiness against Drivers and Pedestrians: As in saturating, spoofing attempts are invisible to human eyes.
- Inducing Multiple Fake Dots: If the lidar rotates at a constant speed, an attacker can generate multiple fake dots with one attack tool. This can be done by periodically firing back the attacking pulses, immediately after the first attacking pulse, with the same period as the PRT. The PRT of the target lidar can be approximately derived from the specification, and then, minutely adjusted by measurements. Let us denote the angular horizontal resolution of the target lidar, whose rotating speed is constant, as r_H [°], and the update rate as f [Hz]. Then,

Table 1. VLP-16 specification

# of Vert. Layers	16	Light Wavelength	903nm
Update Rate	5/10/20Hz	Angular Resolution	0.1/0.2/0.4° (hor.) 2° (ver.)
Range	100m	Field of View	-15° ~ 15° (ver.) 360° (hor.)

the theoretical interval between consecutive pulses can be derived as follows:

$$1 / \left(\frac{360}{r_H} \times f \right) = \frac{r_H}{360f} \text{ [s]} \quad (5)$$

Note that this is irrespective of the distance between the lidar and the attacker.

- Receiving Angle: Similar to saturating, a small receiving angle limits the maximum number of fake dots inducible by a fixed spoofer. Therefore, to increase the number of fake dots the attacker should utilize multiple transmitters.
- Curved Reception Glass: Although we did not experimentally confirm if spoofing attack using refraction/reflection on the curved glass is possible because we could not obtain a pulse laser source that was sufficiently strong, we expect the oblique incidence of a strong laser pulse to readily induce fake dots in sectors, other than the direction of the attacker. If this is possible, it will expose the victim vehicle to threats far more dangerous than that of saturating.

4 Experiments

In this section, we present equipment used and experimental setups for them. In addition, experimental results are provided with figures. Note that further details for the experiments, including videos and raw lidar packet capture for the attack, can be found in the appendices.

4.1 Experimental Setup

Target Lidar: We selected Velodyne’s VLP-16 [42] for verifying our attack methods. It is the lightest and the latest in the product lineup, and targeted for various mobile usages such as autonomous vehicles, UAVs, and robotics. Its specification related to this paper, is summarized in Table 1. Note that, the VLP-16 has an adjustable update rate and horizontal resolution, and they are in a trade-off relationship. For our case, they were set to lower values: 5Hz and 0.1°, respectively³. To check the effect of the attacks we required a visualizer for the sensing result. We used Velodyne’s official visualization software, VeloView [27], which visualizes the sensing result in real time by parsing the UDP packet stream from VLP-16, and supports recording into pcap files and replaying them.

Attack Tool for Saturating: For saturating, only a light source is required. We used a 30mW, 905nm laser module (\approx USD 40) as the weak light source, and a power-adjustable 800mW, 905nm laser module (\approx USD 350) as the strong one. Product names and pictures can be found in the appendices.

Attack Tool for Spoofing: The attack tool is as depicted in Fig. 7. We used an OSRAM SFH 213 FA (\approx USD 1) photodiode (PD) with additional comparator circuitry for the receiver⁴, and an OSRAM SPL PL90 (\approx USD 16) pulsed laser

³ Raw packet captures for 10Hz & 0.2° can also be found in the appendices.

⁴ Its detailed circuit diagram is given in the appendices.




Fig. 8. VeloView output during exposure to a weak light source. Fake dots are observable only in the direction of the light source. The maximum angle between the dots was measured to be 20° .




Fig. 9. VeloView output during oblique exposure to a strong light source. Fake dots are observable in a direction other than the light source.

diode (PLD) with a PCO-7110-40-4 (\approx USD 300) PLD driver from Directed Energy Inc. Note that, both of the PD and the PLD are not standalone; the PLD driver is required to generate the high-current pulses, essential for firing the laser pulses. For the delay component, we used an Agilent 33250A function generator with external-trigger mode in the burst n-cycle pulse output setup.

4.2 Saturating

For saturating, we illuminated the VLP-16 with the aforementioned light sources. As mentioned in Section 3.3, invisible light is one of the strengths of this attack. Thus, we used an IR viewer [32] to aim the light.

Weak Light Source: When the lidar was illuminated by a weak light source, we could observe numerous randomly-located fake dots, as depicted in Fig. 8. Because the experiment was conducted in a basement, every dot outside the room perimeter is apparently fake. As discussed in Section 3.3, induced fake dots were observed only in the direction of the light source. We suppose that the overall increase in the noise floor due to the injected light is the cause of the induced fake dots. The VLP-16 seems to have an absolute threshold for detecting echoes, and the raised noise floor might almost reach this threshold, causing the noise fluctuations lead to numerous fake dots.

Strong Light Source (Direct): We switched the light source to a strong one, and directly illuminated the lidar. We discovered that the lidar became completely blind in a sector, in the field of view (Fig. 10). We could also observe multiple fake dots as in the case of the weak light source and a severe degradation in the received signal strength in the direction of illumination.

Strong Light Source (Oblique): We obliquely illuminated the lidar, and observed fake dots in a direction other than that of the light source, as in Fig. 9. We also experimentally confirmed that curved glasses can change the incoming direction of the obliquely incident light. Details can be found in the appendices.




Fig. 10. VeloView output before (left) and after (right) exposure to a strong light source. We placed a metal plate ($41 \times 42\text{cm}^2$) in front of the lidar.

4.3 Spoofing by Relaying

We performed spoofing by relaying using the attack tool described in Section 4.1. We first aimed the attack tool on the lidar to receive its pulses. When the incoming pulses are captured by the PD, the comparator converts them into a series of 5V pulses. Then, these pulses are fed to the function generator, which is triggered by the first received pulse. The function generator waits for a predefined delay, and transmits a predefined number of copies of the output pulse to the PLD driver. Finally, the PLD driver lets the PLD fire laser pulses as signaled.

To induce multiple fake dots (Section 3.4), the intervals between the output pulses have to be matched to the PRT of the target lidar. Although the PRT can be derived using Eq. (5), the real value subtly varies. We analyzed the target lidar signal and found that the best approximation was $55.296\mu\text{s}$, whereas the theoretical value was $55.556\mu\text{s}$. We observed that the measured PRT remained the same over time and over various distances between the spoofer and the lidar. After determining the actual PRT, we encountered a problem in applying it as the output pulse interval. The smallest supported PRT of the PLD, OSRAM SPL PL90, was only $100\mu\text{s}$; therefore, to circumvent this problem, we set the output pulse interval as double of the actual PRT, $110.592 = 2 \cdot 55.296\mu\text{s}$. Then, we measured the delay d_a ; it was determined by setting the *cycle*—a function generator parameter to determine how many times the output pulses will be repeated after the inserted delay per a trigger—value to one, and gradually increasing the delay parameter of the function generator until a fake dot appeared. When the distance between the spoofer and lidar was approximately 5m, the delay was measured to be $663.3\mu\text{s}$. We could also conclude that the ping waveform of the VLP-16 was only a single laser pulse; else, we could not have observed any fake dot. Once we observed a fake dot by a single pulse, we gradually increased the cycle value. However, no matter how large the cycle was, no more than ten fake dots were observable. This may be because the receiving angle of the VLP-16 for the PLD used is approximately 2.0° ⁵, which corresponds to ten fake dots⁶. Fig. 11 shows the induced fake dots. Note that this scheme works outdoor under sunlight. Refer to the appendices for the details.

⁵ This is considerably smaller than the case in Fig. 8. The differences in the light source strength and beam diameter may be the cause.

⁶ As we fired attacking pulses for every two target lidar pulses, $10 \cdot 2 \cdot 0.1^\circ = 2.0^\circ$. Note that 0.1° was the horizontal resolution of VLP-16 then.




Fig. 11. VeloView output of the multiple induced fake dots.




Fig. 12. VeloView output of the fake dots closer than the spoofer. Note that the red dot, the closer it is to the lidar.

In Section 3.4, we present a method by which an attacker could induce fake dots closer than the spoofer. To confirm this, we gradually reduced the value of d_a until the induced fake dots were located between the spoofer and the lidar. Fig. 12 displays the induced fake dots located between the spoofer and the lidar. The lidar-to-spoofers distance and the delay were 12m and $1.959\mu s$, respectively.

We note that the exact value of d_a is not essential for inducing fake dots. In reality, a sufficiently large cycle would suffice. We observed multiple fake dots, when the cycle was set as 30, even with the delay parameter of the function generator set as zero. This is because whenever the cycle is increased by one, it is equivalent to adding a delay of $2\cdot\text{PRT}$. With the zero delay of the function generator, no delay other than d_p will be added. Therefore, the total delay for the m -th pulse will be just $2mT + d_p$ from Eq. (4). At a certain value among m 's (denote it m'), the relation, $2m'T + d_p \approx d_a = nT + S + d_p$, satisfies, which is equivalent to inducing a fake dot with $d_i = 0$ in Eq. (4). From that on, the pulses will start inducing fake dots.

5 Discussion

5.1 Practical Consideration for Attack Deployment

Aiming Problem: Aiming is one of the main obstacles in deploying attacks in practice. When the target vehicle moves, the attacker has to track the target lidar with the attack tool. However, advanced attackers may circumvent this difficulty by adopting the following approaches: Lidars are typically located at a fixed position on the vehicle, i.e. on the center of the roof or on the corners. Further, there are many cases on the road, when vehicles run straight with a constant speed. Therefore, an attacker may mount the attack tool on a vehicle with an accurate motorized mount, and deliberately follow/precede the target vehicle such that the relative speed becomes zero. This can render the situation almost similar to a stationary case. Next, attackers may adopt an optical system such as a *beam expander* [24] to widen the attacking beam width or spread the beam with an appropriate optical system such as concave lenses for a flashlight-like effect. Note that in this case, the decrease in light intensity due to expansion does not affect the effectiveness of the attack, because lidars are designed to mainly sense reflected lights, considerably weaker than direct illumination. Even

if a weak light intensity matters, attackers can utilize stronger light sources. Attackers can also install a *trap* on the road. With the attacking transceiver installed and calibrated in advance, the attacker can render the problem similar to a stationary case, because the speed of the victim vehicle is considerably slower than the rotating speed of the lidar and the speed of light.

Parameter Setting: Unlike in laboratory, attackers do not have access to the target sensor output, in reality. Therefore, the attacker cannot determine the best parameters for the attack tool. However, this would not be a serious issue because of the following reasons: First, most vehicles are mass produced, and are identical in terms of their sensors. Therefore, the attacker can obtain multiple types of vehicles, and analyze them to acquire the essential information for deploying the attacks, e.g. the PRT(s) and lidar position(s). Further, a precisely calibrated attack tool will work, regardless of the circumstances, and this calibration can be done in advance. Because real echoes and intentionally generated attack pulses are indistinguishable, spoofing by relaying will work as long as the transmitter and receiver are suitably aligned in the same direction. With such a calibrated attack tool aimed at the victim lidar, the only variable is the distance between the attacker and the target vehicle, which the attacker can measure by adopting additional sensors.

5.2 Potential Countermeasures

Redundancy and Fusion: If a vehicle is equipped with multiple lidars having an overlapping field of view, the effect of saturating and spoofing can be mitigated to a certain extent. However, this directly increases the cost, and is not a definitive solution because attackers can blind multiple lidars simultaneously. Besides, it is also not easy to detect spoofing, when fake dots are induced in non-overlapped zones. Likewise, the fusion of multiple types of sensors cannot be an ultimate solution either. Radars [44], cameras [30, 44], and ultrasonic sensors [44] have all been revealed to be vulnerable to either blinding/jamming or spoofing.

Saturation Detection: As discussed in Section 3.3, attempts to intentionally saturate a lidar can be easily detected, and the victim vehicle can adopt fail-safe mode. For example, it can abandon sensor outputs from the direction of the attack and move to the roadside, while slowing down. However, the victim will be unable to drive because saturation itself is inevitable. Further, on crowded roads, the fail-safe maneuver might rather endanger the victim vehicle.

Reducing the Receiving Angle: According to the calculation and measurement in Section 2.1 and 4.3 respectively, the receiving angle of VLP-16 (2.0°) is considerably larger than the minimum required size (0.0048°) for meeting the specifications. Therefore, reducing the receiving angle can mitigate the effect of saturating and spoofing. Both the angle of the region blinded by saturation and the maximum number of inducible fake points by spoofing can be reduced. However, reducing the receiving angle is not easy, because it is in a trade-off relationship with the lidar sensitivity [25]. Further, it would be difficult to reduce the receiving angle to the minimum required value due to the design margins.

Random-direction Pinging: Transmitting pulses in random directions can mitigate the effect of spoofing, because it is no longer possible to induce multiple fake dots by a single spoofer. However, it is practically difficult to apply this approach to current lidars with rotating scanners. Randomly rotating the scanners will severely degrade the reliability and durability of the lidar. Even current lidars have reliability issues due to their moving parts [1]. Further, the update rate, a key performance figure, will be reduced.

To avoid the problem of random rotation, lidars may maintain the current scan-by-spinning but transmit pulses at random instants. However, this will directly lead to update rate decreases. Lidars using this approach should spin faster to reach the required update rate, which may again lead to reliability issues. Currently, in our opinion, the best cost/performance effective mitigation against the induction of multiple (closer) fake dots is to electrically perturb PRTs while keeping the rotating speed constant. Such slightly perturbed PRTs will not severely degrade the performance/reliability, but will effectively prevent the attacker from predicting pulse-firing instants blocking aforementioned two types of threats.

Randomizing the Ping Waveform: Transmitting pulses with randomized waveforms and rejecting pulses different from the transmitted one can fundamentally prevent spoofing from inducing fake dots closer than the spoofer. Further, this also can help mitigate inter-lidar interference. Approaches of this type have been intensively studied for military radars [26]. However, this cannot prevent all spoofing attempts, because attackers can still induce fake dots further than the spoofer location.

Mitigating Curved Glass Effects: The best approach for removing unwanted effect of the curved reception glasses is to get rid of them. Indeed, several lidars (e.g. IBEO LUX 2010 and Velodyne HDL-64E) do not have them. Even if curved glasses are essential for the operation, designers may mitigate their adverse effect by carefully selecting glass materials or designing glass curvature so that obliquely incident attacking light cannot reach central receiving structures.

5.3 Other Points

Fatality of Induced Fake Dots: Unlike the case of the IBEO LUX 3 [30], where it was possible to generate many fake dots spanning 30° approximately, only up to ten fake dots were induced in the VLP-16. As previously noted, the ten fake dots correspond to an object 2.0° wide. This may not appear important initially, but its significance cannot be underestimated; for example, the size of an object spanning 2.0° , 55m away from the lidar would be 1.9m wide, which is almost as wide as most vehicles. As per the data from UK Department for Transport [39], 55m is the *braking distance* for a car driving at 60mph. Because the braking distance is the distance required solely for braking, even autonomous vehicles have no room for checking the authenticity of the observed dots, but need to immediately activate emergency braking or evasive maneuvers. Such sudden actions are sufficient to endanger the surrounding vehicles.

Increasing the Number of Induced Fake Dots: As revealed in the experiment, the number of fake dots by one attack tool is limited due to the size of

the receiving angle. However, by adopting multiple attack tools, they can be increased. Further, attackers can also induce a larger shape to the victim lidar by orchestrating multiple attack tools.

Comparison with the Previous Work: Although we have improved upon the previous work in many aspects, there are a certain issues that have not been dealt with or were inferior in the outcome. However, we emphasize that the target lidar was different; as noted before, the IBEO LUX 3 was used in the previous work, whereas the Velodyne VLP-16 was used in our case. We did not deal with the induction of multiple dots in a single direction. VLP-16 has three modes of operation: *last*, *strongest*, and *dual*. Among the three, only the dual mode allows up to two dots per direction; the other two modes permit only one dot. Therefore, for the last and strongest modes, inducing multiple dots in a single direction was fundamentally impossible. For the dual mode, to induce two dots in one angle, two attack pulses should not deviate more than 667ns. However, as discussed in Section 4.3, this small deviation was not possible under our single-PLD setup due to the smallest supported PRT of the PLD used. As Petit et al. used the same single PLD, the operation scheme of IBEO LUX 3 seems to differ from that of VLP-16. Further, we did not deal with the tracking/recognition of the induced fake dots. This was because Velodyne does not provide such a functionality for any of its products, whereas IBEO does, and there were no suitable alternatives. Finally, as previously mentioned, the difference in the spanning angle of the induced fake dots seems solely because of the difference between the receiving angles of the two lidars. If the receiving angle of LUX 3 had not been that large, it would not have been possible to observe such a wide span of the induced dots because the transmitter was also fixed in the previous work.

6 Related Work

Automotive Security: With the abrupt increase in the proportion of electronics in modern vehicles, vehicles are no longer safe zones against hacking threats. Since Koscher et al. first demonstrated the feasibility of vehicle hacking [16], numerous researchers have discovered vulnerabilities in vehicular networks and control units [18], demonstrated the feasibility of remote hacking [5, 19], and even the hacking of real vehicles [20]. To cope with these new threats, various approaches have been proposed as defensive measures [18, 6, 7, 13, 41, 22]. However, most works in this field focus on compromising and defending the structurally vulnerable control area network buses. In comparison, researches on vehicular sensor security are rare, despite its criticality for (semi-)autonomous vehicles. We have already discussed the contributions and limitations of Petit et al. [30] in Section 1; this work was the first in revealing that the vehicular sensors for autonomous driving can be easily tempered by external stimuli. Another notable work is that of Yan et al., who performed a comprehensive security analysis on environment perception sensors mounted on a real vehicle, the Tesla Model S [44]. They succeeded in jamming and spoofing the ultrasonic sensors, and only in jamming the mm-wave radar. They also demonstrated, like Petit et al., that cameras are extremely vulnerable to exposure to a strong light source.

However, the lidar was not dealt with, because the Model S does not have one. Finally, Shoukry et al. spoofed an anti-lock braking (ABS) sensor, another vehicular sensor that is a type of magnetic encoder [35]. They installed an attacking actuator next to the target sensor, and canceled the legitimate magnetic field from the sensor by emitting its reverse waveform. Then, they added the spoofing waveform, and it was injected without any disturbance. By simulation, they showed that by this attack, the ABS system would be unable to brake properly.

Sensor Attacks: Park et al. caused a medical infusion pump to over/under infuse fluids by injecting an IR laser to its drop sensor [29]. They illuminated the receiver of the drop sensor to render it unable to sense any fluid drops, which in turn led to over-infusion. To the best of our knowledge, this was the first attempt at inducing a critical high-level malfunction by saturating. With a *side* channel attack, Son et al. incapacitated a flying drone by inducing massive fluctuations in the gyroscope outputs with acoustic stimuli [37]. Trippel et al. further developed this idea over a DoS attack; they succeeded in controlling an RC car driven by a smartphone’s accelerometer output, only with the injection of acoustic stimuli to the MEMS-based accelerometer [38]. Finally, as an example of *transmission* channel attack, Foo Kune et al. injected fake sensor outputs by inducing EMI to the wire connecting an analog sensor and an amplifier [9]. They demonstrated that this can be exploited to induce malfunctions in implantable medical devices such as pacemakers and cardiac defibrillators.

Defenses against Sensor Attacks: To counter the aforementioned threats to sensors, several approaches have been proposed. Shoukry et al. proposed an active sensor spoofing defense scheme called PyCRA [36]. This is a spoofing detection scheme that detects spoofing attempts by turning off the active sensor transmitter at random instants such that the attacker cannot react to the sudden changes. When the sensor is attacked, the spoofing signal can be detected because no incoming signal is expected. However, the PyCRA cannot be applied to lidars or radars because it assumes the channel between the transmitter and receiver to be fixed, whereas lidars and radars have continuously changing channels because targets can be located anywhere. Further, Shin et al. pointed out that the PyCRA has a critical problem to be applied to analog-digital systems [34], because it can either lead to an arms race between the attacker and the defender or requires too many resources to be secure. For redundancy and fusion, most works in this field focus on sensor reliability/precision enhancements rather than on the security; relatively fewer works focus on security [21, 28, 15]. However, redundancy and fusion have limitations, as discussed in Section 5.2.

7 Conclusion

Lidars are undoubtedly one of the core sensors in autonomous vehicles. Being the eyes of safety-critical systems, such as cars, their reliability is critical and cannot be compromised, because it can endanger human lives. In this work, we have presented and experimentally verified two types of attacks that can severely degrade the reliability of lidars. Although we have listed many mitigative approaches in the discussion, they are either technically/economically infeasible

or are not definitive solutions to the presented attacks. We do not advocate the complete abandonment of the transition toward autonomous driving, because we believe that its advantages can outweigh the disadvantages, if realistic adversarial scenarios are appropriately mitigated. However, such considerations are currently absent; therefore, automakers and device manufacturers need to start considering these future threats before too late.

Acknowledgment

This work was supported by the Advanced Technology R&D Center of Hyundai AutoEver.

References

1. Ackerman, E.: Velodyne Says It's Got a "Breakthrough" in Solid State Lidar Design. <http://spectrum.ieee.org/cars-that-think/transportation/sensors/velodyne-announces-breakthrough-in-solid-state-lidar-design>, [Online; accessed: 24-Feb-2017]
2. Alley, C.O., Bender, P.L., Dicke, R.H., Faller, J.E., Franken, P.A., Plotkin, H.H., Wilkinson, D.T.: Optical radar using a corner reflector on the Moon. *Journal of Geophysical Research* 70(9), 2267–2269 (1965), <http://dx.doi.org/10.1029/JZ070i009p02267>
3. Beasley, E.: LiDAR and Autonomous Technology. <http://velodynelidar.com/blog/lidar-autonomous-technology/>, [Online; accessed: 9-March-2017]
4. Bhatia, P.: Vehicle Technologies to Improve Performance and Safety. Tech. rep., University of California Transportation Center (2003), <https://escholarship.org/uc/item/4zw4m05k>
5. Checkoway, S., McCoy, D., Kantor, B., Anderson, D., Shacham, H., Savage, S., Koscher, K., Czeskis, A., Roesner, F., Kohno, T., et al.: Comprehensive experimental analyses of automotive attack surfaces. In: *Proceedings of 20th USENIX Security Symposium*. USENIX Association (2011)
6. Cho, K.T., Shin, K.G.: Fingerprinting electronic control units for vehicle intrusion detection. In: *Proceedings of 25th USENIX Security Symposium*. pp. 911–927. USENIX Association (2016)
7. Dagan, T., Wool, A.: Parrot, a software-only anti-spoofing defense system for the can bus. *ESCAR EUROPE* (2016)
8. Distner, M., Bengtsson, M., Broberg, T., Jakobsson, L.: CITY SAFETY A SYSTEM ADDRESSING REAR-END COLLISIONS AT LOW SPEEDS. In: *Proc. 21st International Technical Conference on the Enhanced Safety of Vehicles* (2009)
9. Foo Kune, D., Backes, J., Clark, S., Kramer, D., Reynolds, M., Fu, K., Kim, Y., Xu, W.: Ghost Talk: Mitigating EMI Signal Injection Attacks against Analog Sensors. In: *IEEE Symposium on Security and Privacy*. IEEE (2013)
10. Ford Mediacy: Ford First Automaker to Test Autonomous Vehicle at Mcity, University of Michigans Simulated Urban Environment. <https://media.ford.com/content/fordmedia/fna/us/en/news/2015/11/13/ford-first-automaker-to-test-autonomous-vehicle-at-mcity.html>, [Online; accessed: 23-Feb-2017]
11. Goyer, G., Watson, R.: The laser and its application to meteorology. *Bulletin of the American Meteorological Society* 44(9), 564–575 (1963)

12. Higgins, S.: Solid-State LiDAR: A New Era of 3D Scanning. <http://www.spar3d.com/blogs/the-other-dimension/vol13no50-solid-state-lidar-a-new-era-of-3d-scanning/>, [Online; accessed: 24-Feb-2017]
13. Hoppe, T., Kiltz, S., Dittmann, J.: Security threats to automotive CAN networks—practical examples and selected short-term countermeasures. In: International Conference on Computer Safety, Reliability, and Security. pp. 235–248. Springer (2008)
14. Huynh, T.: Google self-driving car: everything you need to know. <http://www.techradar.com/news/car-tech/google-self-driving-car-everything-you-need-to-know-1321548>, [Online; accessed: 23-Feb-2017]
15. Ivanov, R., Pajic, M., Lee, I.: Attack-Resilient Sensor Fusion for Safety-Critical Cyber-Physical Systems. *ACM Transactions on Embedded Computing Systems* 15(1), 21 (2016)
16. Koscher, K., Czeskis, A., Roesner, F., Patel, S., Kohno, T., Checkoway, S., McCoy, D., Kantor, B., Anderson, D., Shacham, H., et al.: Experimental security analysis of a modern automobile. In: IEEE Symposium on Security and Privacy. pp. 447–462. IEEE (2010)
17. Lambert, F.: Tesla still has no plans to use LiDAR in consumer vehicles, but does use the tech for ‘ground truthing’. <https://electrek.co/2016/11/02/tesla-no-plan-for-lidar-self-driving-cars/>, [Online; accessed: 9-March-2017]
18. Miller, C., Valasek, C.: Adventures in automotive networks and control units. DEF CON 21 (2013)
19. Miller, C., Valasek, C.: A survey of remote automotive attack surfaces. Black Hat USA (2014)
20. Miller, C., Valasek, C.: Remote exploitation of an unaltered passenger vehicle. Black Hat USA (2015)
21. Montgomery, P.Y., Humphreys, T.E., Ledvina, B.M.: Receiver-Autonomous Spoofing Detection: Experimental Results of a Multi-antenna Receiver Defense Against a Portable Civil GPS Spoofer. In: Proceedings of the ION International Technical Meeting (2009)
22. Müter, M., Asaj, N.: Entropy-based anomaly detection for in-vehicle networks. In: IEEE Intelligent Vehicles Symposium. pp. 1110–1115. IEEE (2011)
23. NASA: Planetary Laser Altimetry. <https://tharsis.gsfc.nasa.gov/index.php>, [Online; accessed: 23-Feb-2017]
24. Newport Corp.: Optics: How to Build a Beam Expander, http://assets.newport.com/webdocuments-en/images/how_to_build_a_beam_expander_5.pdf
25. Osta, P.V.: The Basics of Microscopy, <http://www.vanosta.be/microscopy.htm>
26. Pace, P.E.: Detecting and classifying low probability of intercept radar. Artech House (2009)
27. ParaView: VeloView. <http://www.paraview.org/VeloView/> (2017), [Online; accessed 05-March-2017]
28. Park, J., Ivanov, R., Weimer, J., Pajic, M., Lee, I.: Sensor Attack Detection in the Presence of Transient Faults. In: Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems (2015)
29. Park, Y., Son, Y., Shin, H., Kim, D., Kim, Y.: This ain’t your dose: Sensor Spoofing Attack on Medical Infusion Pump. In: 10th USENIX Workshop on Offensive Technologies. USENIX Association (2016)
30. Petit, J., Stottelaar, B., Feiri, M., Kargl, F.: Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and LiDAR. Black Hat Europe (2015)

31. Pomerleau, D.A.: ALVINN, an autonomous land vehicle in a neural network. Tech. rep., Carnegie Mellon University, Computer Science Department (1989)
32. Public Lab: Near-Infrared Camera. <https://publiclab.org/wiki/near-infrared-camera> (2017), [Online; accessed 06-March-2017]
33. Quanergy Systems Inc.: Quanergy S3 Solid State LiDAR, the World's First Affordable Solid State LiDAR Sensor, to Begin Full Scale Manufacturing in 2017. <http://www.businesswire.com/news/home/20170103005387/en/Quanergy-S3-Solid-State-LiDAR-Worlds-Affordable>, [Online; accessed: 24-Feb-2017]
34. Shin, H., Son, Y., Park, Y., Kwon, Y., Kim, Y.: Sampling Race: Bypassing Timing-based Analog Active Sensor Spoofing Detection on Analog-digital Systems. In: 10 th USENIX Workshop on Offensive Technologies. USENIX Association (2016)
35. Shoukry, Y., Martin, P., Tabuada, P., Srivastava, M.: Non-invasive Spoofing Attacks for Anti-lock Braking Systems. In: International Workshop on Cryptographic Hardware and Embedded Systems. pp. 55–72. Springer (2013)
36. Shoukry, Y., Martin, P., Yona, Y., Diggavi, S., Srivastava, M.: PyCRA: Physical Challenge-Response Authentication For Active Sensors Under Spoofing Attacks. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. pp. 1004–1015. ACM (2015)
37. Son, Y., Shin, H., Kim, D., Park, Y., Noh, J., Choi, K., Choi, J., Kim, Y.: Rocking Drones with Intentional Sound Noise on Gyroscopic Sensors. In: Proceedings of 24th USENIX Security Symposium. pp. 881–896. USENIX Association (2015)
38. Trippel, T., Weisse, O., Xu, W., Honeyman, P., Fu, K.: WALNUT: Waging Doubt on the Integrity of MEMS Accelerometers with Acoustic Injection Attacks. In: IEEE European Symposium on Security and Privacy. IEEE (2017)
39. UK Dept. for Transport: The Highway Code - General rules, techniques and advice for all drivers and riders (103 to 158) - Rule 126, <https://www.gov.uk/guidance/the-highway-code/general-rules-techniques-and-advice-for-all-drivers-and-riders-103-to-158#rule126>
40. USDA: ARS study helps farmers make best use of fertilizers. <https://www.ars.usda.gov/news-events/news/research-news/2010/ars-study-helps-farmers-make-best-use-of-fertilizers/>, [Online; accessed: 23-Feb-2017]
41. Van Herrewege, A., Singelee, D., Verbauwhede, I.: CANAuth - A Simple, Backward Compatible Broadcast Authentication Protocol for CAN bus. In: ECRYPT Workshop on Lightweight Cryptography (2011)
42. Velodyne: Velodyne LiDAR Puck. http://velodynelidar.com/docs/datasheet/63-9229_Rev-C_VLP16_Datasheet_Web.pdf (2017), [Online; accessed 05-March-2017]
43. Vosselman, G., Maas, H.G.: Airborne and terrestrial laser scanning. Whittles Publishing (2010)
44. Yan, C., Xu, W., Liu, J.: Can You Trust Autonomous Vehicles: Contactless Attacks against Sensors of Self-driving Vehicle. DEF CON 24 (2016)

Appendices

Due to space limitation, appendices are posted to the website below:

<https://sites.google.com/view/ches17illusionanddazzle>