LiDAR Sensor Security of a Driverless Pod

Daniel S. Fowler
Secure Cyber Systems RG
WMG, University of Warwick
Coventry, UK
dan.fowler@warwick.ac.uk

Anh Tuan Le
Secure Cyber Systems RG
WMG, University of Warwick
Coventry, UK
a.le.1@warwick.ac.uk

Carsten Maple
Secure Cyber Systems RG
WMG, University of Warwick
Coventry, UK
cm@warwick.ac.uk

***Abstract*—Light Detection and Ranging (LiDAR) sensors have evolved from use within Advanced Driver-Assistance Systems (ADAS) to become critical sensors for the safe operation of Connected and Autonomous Vehicles (CAVs). Whilst it is common for engineers to design for the operational failure of a sensor system, if a sensor can be manipulated due to malicious acts, those acts must be taken into consideration. The LiDAR sensor on a prototype driverless *pod* was examined to determine potential external security weaknesses, included is a summary of previous LiDAR security work on attacks and mitigation. Any issues with LiDAR sensors should contribute to a CAV's security and health and safety risk assessments. A discussion on the security implications of using LiDAR as part of a CAV's sensor system is provided, with indicators for further research.**

***Index Terms*—LiDAR, CPS, CAV, security, sensors**

I. Introduction

LiDAR technology is well-established in automotive engineering, having been used for ADAS applications since the late 1990s [1]. Adaptive cruise control, collision avoidance, and lane-departure warnings are ADAS systems that increase vehicle complexity in return for benefits [2]:

- reducing traffic accidents;
- improving vehicle efficiency (saving fuel, reducing traffic jams through improved flow, and, thus, reducing a vehicle's impact on the environment);
- improving information to the driver, reducing mental stress and fatigue.

Building upon ADAS, LiDAR sensors have developed in complexity to become external sensory perception (exteroception) components in the self-driving systems [3] of CAVs. LiDAR sensors (alongside ultrasonic, radar, and camera) are used for detecting objects as part of a CAV's control system.

Fig. 1. A LiDAR sensor for a prototype last-mile driverless pod, infrared laser light is emitted from the small cylinder and collected in the larger lens.

CAVs can be used as a self-driving public transport vehicle, or autonomous pod, for use in *last-mile* [4], or *microtransit* [5], applications. In this work, a LiDAR sensor, Fig. 1, used on a prototype last-mile pod, has been examined to determine possible external security issues. Intelligent adversaries, i.e. *threat agents* [6], may consider sensor weaknesses as an attack vector to disrupt the pod's normal operation.

Section II revisits the motivation for attacks, how security impacts vehicle safety, how LiDAR operates, and its use in the pod. The limited existing research in LiDAR security is discussed in III. Potential weaknesses in the LiDAR sensor is examined in IV to VI and discussed in VII. VIII has mitigation techniques from literature before concluding.

II. Background and Motivation

Vehicles have become cyber-physical systems (CPS), i.e. robots, that are susceptible to attacks beyond those seen in traditional systems [7]. CAV complexity provides several attack vectors, including the ADAS and self-driving systems. Motivations for attacking vehicles and transportation systems have been established. The European E-safety vehicle intrusion protected applications (EVITA) project investigated the protection of vehicle systems, categorising attack aims and motivations [8]. Any attacks by threat agents are an attempt to disrupt a pod's operation or gain some advantage, therefore, any weaknesses in a system or its components should be understood to aid with risk mitigation.

A. Security as an input to vehicle safety

Sensors feed data to the pod's autonomous control system (ACS). The accurate perception of the environment is critical to the ACS. This means a sensor's ultimate function is one of safety. The safety of the pod, the passengers in it, and anything in the pod's vicinity, i.e., people, vehicles (including other pods) and local infrastructure. The pod is a large and heavy CPS that has the potential to cause injury and damage. Physical safety is a critical goal in any vehicle design, however, attacks against a vehicle's control system can compromise that safety design due to the modern vehicle's CPS nature [9]. Attacks may target the sensors, in this case, LiDAR, to manipulate the data sent to the ACS and cause the required disruption.

The possible attacks against the LiDAR sensors contribute to the pod's Threat Analysis and Risk Assessment (TARA),

supplementing the normal functional Hazard Analysis and Risk Assessment (HARA). This allows for suitable risk mitigation, possibly resulting in vehicle design changes if deemed necessary.

B. LiDAR's principle of operation

LiDAR devices use Time-of-Flight (ToF) to measure the distance from the sensor to an object in its Field of View (FoV). A pulse of light is emitted from the sensor and the time for the pulse to be reflected is measured. LiDAR will emit multiple points of light over a large FoV to provide information on any objects that may be present. The emitted light is typically from low power (class one) infrared laser so that it is not normally visible or harmful to humans.

For LiDAR the speed of light is used for the ToF calculations, for ultrasonic sensors it would be the speed of sound. Using the speed of light gives rise to challenging timing requirements. For example, the speed of light, c, takes 16.7ns to travel 5 metres to an object, and another 16.7ns to travel back to the sensor, a total of 33.4ns for the ToF. An equation to derive the distance, D, from ToF, t, is:

$$D = (c \cdot t)/2 \tag{1}$$

Taking the value for c, the speed of light, through the air of 299,702,547m/s and with a ToF of 33.4ns the calculation using (1) is $(299702547 \cdot 33.4e^{-9})/2 = 5.005m$. The time-sensitive nature of ToF sensors can limit their performance.

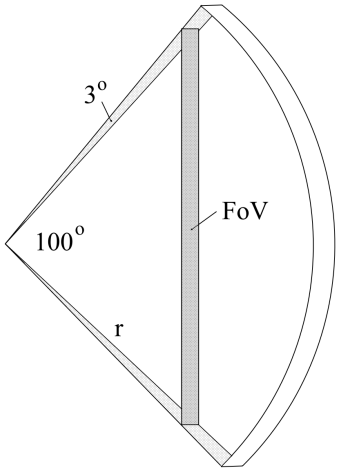

Fig. 2. The field-of-view for the LiDAR sensor, from above

C. Autonomous vehicle's LiDAR sensor

The pod's LiDAR sensor specified range measurement is from 6 to 30 metres with a 5cm accuracy. This is due to the different reflectivities of targets. Pulsed light at the 905nm infrared (IR) wavelength is emitted from the sensor and forms a Field-of-View (FoV). The light covers an angle of 100° in the horizontal plane (HFoV) and a narrow 3° in the vertical (VFoV), see Fig. 2. The total width, w, and height, h, of the FoV illuminated by the sensor at a distance r varies (Table I), calculated from (2) where θ is either the HFoV or VFoV angle. The LiDAR's distance readings are transmitted to the ACS via a data network.

$$w \text{ or } h = 2 \cdot r \cdot sin(\theta/2) \tag{2}$$

TABLE I
FoV HORIZONTAL WIDTH (W) AND VERTICAL HEIGHT (H) INCREASE WITH DISTANCE (R)

r	w	h
0.25	0.383	0.013
0.5	0.766	0.026
1	1.532	0.052
5	7.66	0.262
10	15.321	0.524
30	45.963	1.571

in metres (m)

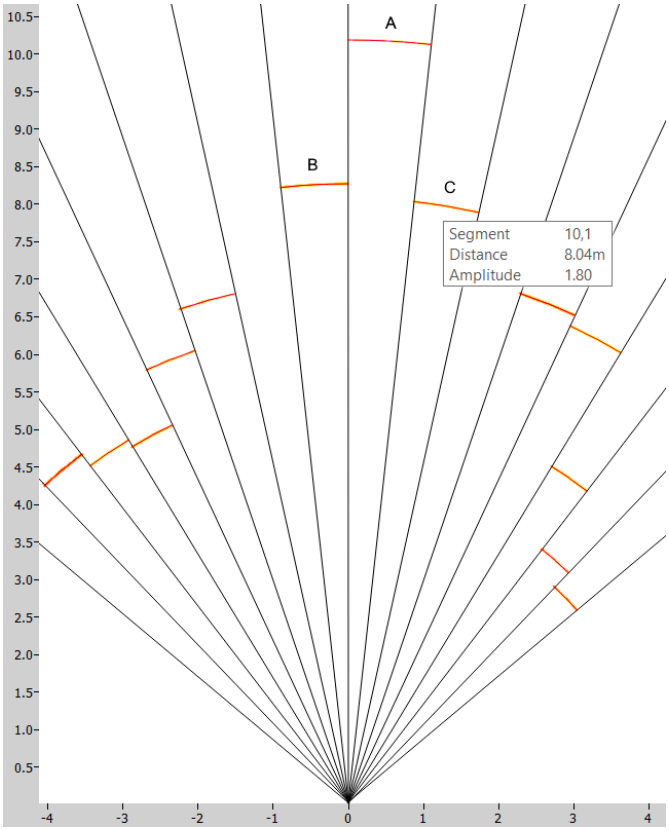

Fig. 3. Distance to object in the LiDAR's test software, metre scales, lines B and C are two objects at 8m, and A is an object at 10m

The LiDAR is constructed from 16 independent solid-state elements, dividing the FoV into 16 segments. Each of the segments, which can be disabled to reduce the active FoV, cover 6.25° of the 100° HFoV. In Fig. 3 the sensor manufacturer's software is showing the segmented FoV. Two large objects have been placed within the LiDAR's FoV at 8m from the sensor, positioned to be detected in segments 8 and 10 (B and C), indicated with lines. A third large object,

positioned at 10m (A), is in segment 9. Other environmental objects at the sides of the FoV are being detected.

The driverless pod needs to be aware of stationary and moving objects that enter the space directly in front of it. Furthermore, objects that approach from the sides or rear are a potential safety hazard. Therefore, multiple LiDAR sensors are used to form a 360° FoV, six at full 100° HFoV, and one set at 50° HFoV the direction of travel, see Fig. 4.

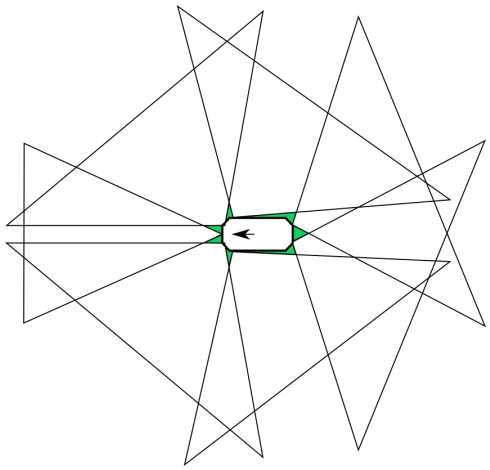

Fig. 4. Multiple LiDARs (six with 100° and one with 50° HFoV) on a self-driving pod to provide all-round distance to object detecting, the green area shows the resultant small blind spot, the arrow is the direction of travel

To support the LiDARs, twelve ultrasonic sensors are deployed around the pod for additional close range proximity detection, plus, a forward-facing and rear-facing radar to aid with multiple moving target detection. The pod's ACS performs sensor fusion to understand where objects are located in the pod's vicinity and determine subsequent driving actions. The final safety devices are contact strips on the pod's bumpers to detect objects that have evaded all the sensors. All of the pod's sensors are located in the lower third of the vehicle's height, therefore, the pod has a VFoV limitation.

III. Previous attacks on LiDAR

Sensor systems are usually designed for functional objectives and without consideration for malicious attacks, i.e. there is an assumption of the integrity of the sensed data [7], [9]. However, a threat agent will consider how to use normal functionality to gain an advantage. The pod cannot operate as intended without the LiDAR signals, this allows an agent the opportunity to manipulate the pod via those signals. Understanding the physical nature of LiDAR can help agents in achieving signal manipulation.

There are well established [10] challenges in measuring the distance to an object with a vehicle-mounted LiDAR. These include object surface reflectivity (the different reflection coefficients of materials determined by Fresnel equations), surface textures, the size and shape of an object, environmental conditions (varying weather), the movement of the vehicle, whether an object is moving or stationary, and undulating landscapes. These physical constraints may aid potential attacks.

An agent can disrupt the normal LiDAR signals using techniques that include eavesdropping, relaying, replaying, tampering, spoofing and blinding (saturation) [7], [9], [11], [12]. Depending on the technique and objectives, the chosen technique may produce noise, fake echoes, fake objects, or mask objects, thus creating a hazard that does not physically exist. However, although agents can use physical characteristics to perform attacks, those characteristics can limit the effectiveness of the attacks, in terms of the number and type of fake objects, and the range of the attack [11]. In [13] it was noted that despite the ability to fake physical signals, the vehicle system may ignore data because it is not reflective of the real world. To increase the attack success rate sophisticated transformations had to be applied to ensure that spoofed signals were decoded as false physical objects.

IV. Pod LiDAR sensor limitations

Reported attacks against LiDAR have concentrated on injecting modified pulsed IR signals to fool a LiDAR-based system. However, having reviewed the capabilities of LiDAR and the solid-state sensor deployed on the pod, possible limitations on the pod's sensing could be deduced. The deployed LiDAR's 100° HFoV and 3° VFoV are limited compared to a human's FoV, which is 180° and 150° respectively [14]. Although the LiDARs cover 360° around the pod, Fig. 4, that coverage is restricted to the lower section of the vehicle. Thus, the pod is blind to any obstruction located in its top section, Fig. 5.

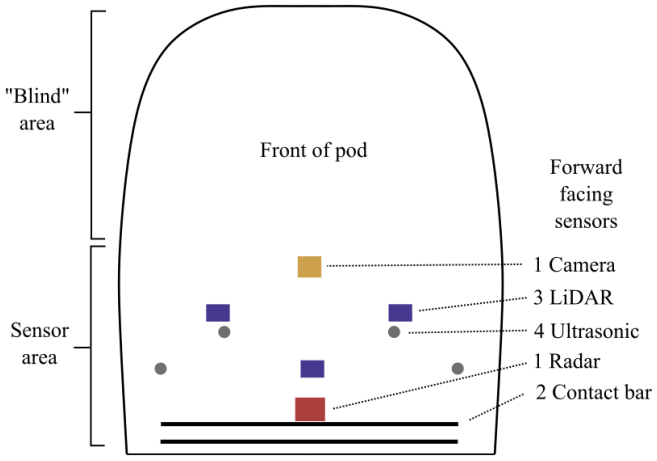

Fig. 5. Forward-facing sensors on the self-driving pod

Another limitation relates to the accuracy, precision and resolution of the LiDAR. The manufacturer specifies it as having an accuracy of ± 5cm for distance measurements, with successive measurements varying by ± 0.6cm and the measurement values varying by ± 1cm. These specifications suggest that macro objects that present a small surface to the sensor may not be reliably detected. Finally, the construction or surface of an object, see III, affects the IR beam's reflectivity and can affect reliable distance to object sensing.

A threat agent with knowledge of sensor limitations may be able to disrupt the normal operation of a CAV. Disruption

is easily achieved by simply walking in front of the pod, or placing a large object in front of it, to trigger the object detection mechanism of the ACS. However, a threat agent may want to avoid detection and disrupt the pod's operation stealthily. Experiments on the capabilities of the LiDAR were performed to investigate possible limitations.

V. Experimental methods

A single LiDAR sensor was mounted on a tripod directed at a controlled reference area, Fig. 3. Items placed between the control objects B and C at 8m are detected, i.e. showing A at 8m. No line shown between B and C indicates no object detected. Macro objects that presented a narrow surface to the sensor were placed within the FoV. This was to explore the limitations of the LiDAR's accuracy, precision and resolution. Furthermore, mesh type materials were used to test the disruption of the IR beam. This allows for the sensor to be tested for its handling of confusing reflections. Voids in mesh objects allow the IR light to travel beyond the object, whilst the solid material of the mesh reflects the light.

The objects tested consisted of cable and rope of 1cm diameter, narrow battens 2.5cm and 4cm in width, a mesh chair, wire fencing, and two types of plastic panels with voids. For the plastic panels, one had a pattern of 3mm holes creating 38% of open space, the second has 8mm holes which creates 68% open space. In a second experiment, narrow and mesh objects were tested against the functioning self-driving pod at a public trial site. The tests were limited due to operational constraints, and to ensure the safety of the operating crew and the general public.

VI. Results

The macro objects chosen do challenge the LiDAR's detection capability. Objects placed at the 8m control area are listed in Table II. The results show that some of the objects were not detected, i.e., were not displayed in the software. Some objects were detected intermittently, with several seconds between detection. Further work on quantifying the intermittent detection is warranted.

TABLE II
LiDAR solid-state sensor item detection deficiencies

Object	Detection result
Wire mesh fence	not detected
1cm diameter cable	not detected
4 loops of 1cm diameter cable	intermittent detection
1cm diameter rope	not detected
4 loops of 1cm diameter rope	intermittent detection
2.5cm wooden batten	not detected
4cm wooden batten	intermittent detection
Black mesh chair	intermittent detection
38% perforated sheet	not detected/intermittent
68% perforated sheet	intermittent detection

For one object, the plastic panel with 3mm holes, the LiDAR would not detect it, however, by changing the angle of the face presented to the sensor's FoV, effectively reducing the aperture size of the voids, the panel could be made to become intermittently detected.

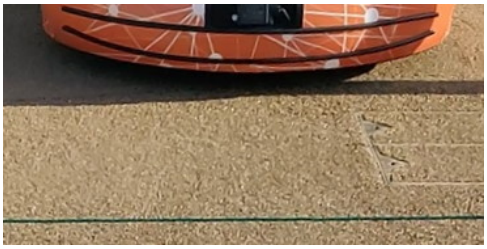

Fig. 6. Testing the capability to detect thin objects, the rope is detected in the FoV of the radar but is not seen when only within the LiDAR's FoV

Static tests using objects were performed with a functioning pod, along with limited tests with the pod operating at a low speed. The pod's radar sensor, Fig. 5, in combination with the LiDAR sensors, does provide a high degree of distance to object sensing. However, the previously highlighted sensor limitations, Table II, can affect the pod's operation. Fig. 6 shows a rope in front of the pod, which is not detectable by the LiDAR. However, the rope is detectable by the radar, but, if held above the radar's FoV, it is no longer seen by the pod, though being within the LiDAR's FoV. Similarly, the plastic panels with voids can be placed in front of the pod and not be detected by the ACS. This results in the pod impacting the test objects when it is in motion.

VII. Discussion on findings

Factors to consider in choosing a sensor for use on a CAV include functional aims, sensor abilities, weight, power consumption, interfacing, data processing requirements, and cost. However, attack considerations should be a factor in sensor choice. Results here show that knowledge of a sensor's capabilities allows for weaknesses in a design to be revealed, weaknesses are a vector for attack.

Detecting thin objects is one weakness. A physical attack would be to secure a cable/rope/wire at a height that would evade detection by the pod. This could be done on a known route of the pod, or by observing the direction of travel and fixing the cable at some point ahead of the pod. This aids the attackers in evading detection, particularly in areas with a low-density pedestrian population and no security cameras. The damage to the pod, and its occupants, from hitting a securely fixed cable could be significant. Furthermore, the thin object weakness could allow a threat agent to craft an object that evades detection, e.g. an Improvised Explosive Device (IED) constructed with a thin profile. The IED would be undetected until impacted by the pod and removes the need for an attacker to be in the vicinity for triggering.

The variability in the results from testing mesh objects, Table II, needs further examination. A mesh fabric chair is a large object but was intermittently detected by the LiDAR. A plastic panel with 38% open space could be placed to evade detection. A study of how the beam from LiDAR sensors is disrupted by mesh materials is warranted. Mesh objects are

not uncommon in the environment, and the fencing sample tested could not be seen by the LiDAR. Future work could examine how mesh materials may be used to disguise objects, and a wider range of meshes and fences with different size and shape voids tested. Furthermore, the results from this work, and other research, needs to be repeated with different makes of solid-state LiDAR sensors.

The pod's vertical blind spot is another weakness to exploit, not only from the risk from a fixed cable but any large and heavy object that could be suspended over the pod's route. Likewise, a drone could evade detection and fly close to, or land on, the pod. This presents the opportunity to deliver an airborne threat, from an innocuous listening or videoing device to an IED.

The types of weaknesses and threats demonstrated in this work will not be considered important by some. Indeed, a direct physical attack on the pod would be simpler. The advantages to the threat agent in attacking a vehicle via its sensors are stealth and, possibly, distance. Even if sensor attacks are regarded as unlikely, low risk and low probability, they must be considered as part of any system's TARA procedure. Summarising this pod's LiDAR sensor, some issues aid threat agents in masking attacks:

1) Previous LiDAR research has shown (see III) that an attack does not require physical contact with the vehicle, and with IR light, attacks cannot be seen.
2) There is a blind spot on the driverless pod and the LiDAR and sensor systems can be evaded entirely.
3) Difficult to detect objects could be placed ahead of the CAV.

Other undiscovered sensor issues and sensor attacks may increase risks to CPSs. As noted in [12], research into attacks on vehicular sensor systems is not common, and further studies on improving the resilience of vehicle sensor systems to attack are required. For this work, restrictions on access to a functioning pod limited investigations into the sensing systems. To overcome such limitations, and accelerate sensor testing, a simulated virtual environment [15] can be used.

VIII. Sensor attack mitigation

Mitigation strategies [7], [9], [11]–[13] for sensors have been proposed, these include:

- Signal emitting sensors, e.g., ultrasonic or LiDAR, could use some form of signal coding. For example, linear coding techniques, use of error detection codes, encryption, or signal randomisation. In [7] a technique called Randomised Pulse Redundancy is proposed, using random variation between successive pulses to obfuscate signals. In [11] randomly skipping pulses is suggested to detect fake pulse injection.
- System and sensor redundancy (including diversity [16]), or sensor fusion, could be used. This increases the difficulty of attack as multiple sensors and systems need to be compromised at the same time. In [11] multiple IR wavelengths are suggested for redundancy, plus, information fusion from vehicle-to-vehicle (V2V) communications could verify object presence (though that raises the question of verifying the authenticity of V2V data).
- Software filtering can remove unusual/outlying values, or compare unusual values from consecutive sensor samples. Such values may indicate falsified signals. For example, if the LiDAR is indicating improbable values, e.g. sequences of 30m, 10m, 25m, 5m in quick succession, it could indicate an attack. Example filters include a rolling mean or a Kalman Filter. However, filtering may reduce system performance and responsiveness.
- Sensors can deploy countermeasure techniques by transmitting a false signal to confuse the threat agent, switching between true and false signals, or redundant sensors can send out false signals.
- Temporal accuracy can be used, ensuring that signals meet exacting time constraints, any delays or mistiming can indicate potential attacks. Higher frequency light pulses would increase the difficulty for attackers, but also increase the sophistication and the cost of the LiDAR design, and may impact upon performance (for example reducing range). In [12], they suggest reducing the tolerances on the LiDAR's receiving angle to increase attack difficulty, however, that impacts the FoV.
- An intrusion detection system (IDS) for a CPS may be able to flag a suspicious signal. Suspicious or out-of-bounds signals could invoke a separate watchdog system that prevents the pod from entering an unsafe position. A *runtime monitoring* system can operate independently of the physical sensor signals by examining internal systems values [17]. For example, do the incoming signals correspond, within a tolerance threshold, with a worldview predicted by the watchdog. Once an attack is detected self-correcting mechanisms can be invoked, to maintain system operation despite the presence of an attack [18].

Natural variations in the physical world are a consideration in determining if an attack is real. The design of mitigation techniques must minimise the number of erroneously flagged (false positive) signals. Other considerations for mitigation design, often overlooked in the literature, are the vehicle engineering factors of component cost and component power consumption. Additional functionality increases both of these factors and will impact the total vehicle cost and energy consumption. Increasing energy consumption increases the power drain on the pod's battery, impacting operational time and the environmental footprint. Therefore, additional security functionality has a business implication, which may determine if a mitigation technique gets implemented.

A TARA and HARA process will raise questions to be addressed in a pod deployment. How should the pod respond in the presence of an attack? Will it still operate within its normal parameters and safety limits or does it need to enter a *limp mode* or stop altogether? Would a separate watchdog system that can halt the pod if it detects an unsafe decision be

beneficial? How are the passengers handled in such situations? Can remote control from an operations centre be invoked? Such questions may need to be considered for the overall system design.

For the driverless pod the risks can be mitigated by improving the sensor systems:

- autonomous vehicles often use sensor fusion to form a safety bounding box for the braking function (the box size based on the driving conditions, e.g., road size, traffic density, vehicle speed, etc.), objects moving into this safety box will invoke braking, this research implies the bounding box configuration should consider types of objects, as they can affect sensor accuracy for distance detection, e.g. for a road flanked with wire fencing the safety distances may need to increase to provide more time for distance to object processing;
- improve sensors to reduce blind spots;
- deploy the underutilised front-facing camera, for example, by implementing a camera-based collision warning system as a backup object detector;
- strengthen cross-sensor verification techniques, i.e. improved sensor fusion between LiDAR, radar, ultrasonic and camera;
- for attacks highlighted in other works, sensor components with built-in encryption could help with resilience by protecting from attacks targeted at physical sensor characteristics, at-the-edge encryption can ease the workload of the ACS by offloading encryption signalling techniques.

Whilst this work examined the external characteristics of a LiDAR sensor, future work can examine the sensor in other ways, for example, via the data or commands sent through the internal network, or side-channel attacks using electromagnetic radiation (EMR) aimed at disrupting the sensor's electronics. Furthermore, EMR is a consideration for information leakages. The mitigation of EMR attacks is generally via shielding techniques.

There exist few autonomous transportation systems, thus malicious risks are currently low. However, as future deployments of mobile CPSs increase, the need for research into the security aspects of sensor systems is required to mitigate risks that may be found by threat agents. For this LiDAR sensor, several aspects of attacks were examined, with practical experiments examining the physical limitations of its solid-state design. As a result, it raised some considerations with the security and safety of a CAV's sensor system, considerations to be discussed amongst engineers and researchers who are designing sensor systems for future CAVs.

IX. Conclusion

CAVs are complex CPSs and they present multiple attack surfaces to threat agents. The interpretation of the environment by a CAV is through its sensors. The sensors need to provide timely and accurate data for a vehicle to make a correct control decision. This is to provide a vehicle motion that ensures passenger and pedestrian safety. The prototype driverless pod operates correctly when sensing a wide range of objects. However, the results from examining its solid-state LiDAR sensor demonstrates the feasibility of adversaries using sensor design limitations to overcome the pods' normal safety mechanisms. An intelligent threat agent, with knowledge of the pod's sensor weaknesses, may gain an advantage towards disrupting the pod's operation, using it as a proxy for attacking the pod's passengers and impacting the wider transport system.

Acknowledgment

This work was supported by the Innovate UK project 103288 (CAPRI), grant TS/P012264/1.

References

[1] A. Bastian, P. Andreas, R. Holze, and et al., "Autonomous Cruise Control: A First Step Towards Automated Driving," in *Future Transportation Technology Conference & Exposition*. SAE International, 1998.

[2] K. Bengler, K. Dietmayer, B. Farber, and et al., "Three Decades of Driver Assistance Systems: Review and Future Perspectives," *IEEE Intelligent Transportation Systems Magazine*, vol. 6, no. 4, 2014.

[3] J. Hecht, "Lidar for self-driving cars," *Optics and Photonics News*, vol. 29, no. 1, 2018.

[4] M. Tight, F. Rajé, and P. Timms, "Car-free urban areas: A radical solution to the last mile problem or a step too far?" *Built Environment*, vol. 42, no. 4, 2016.

[5] A. Ongel, E. Loewer, F. Roemer, and et al., "Economic Assessment of Autonomous Electric Microtransit Vehicles," *Sustainability*, vol. 11, no. 3, 1 2019.

[6] International Organization for Standardization, "ISO/IEC 15408-1:2009(E) Information technology - Security techniques - Evaluation criteria for IT Security," ISO, Geneva, Tech. Rep., 2014.

[7] K. D. Akdemir, D. Karakoyunlu, T. Padir, and et al., "An Emerging Threat: Eve Meets a Robot," in *Trusted Systems*, L. Chen and M. Yung, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011.

[8] A. Ruddle, D. Ward, B. Weyl, and et al., "Security requirements for automotive on-board networks based on dark-side scenarios," European Commission, Tech. Rep. 1.1, 2009.

[9] A. M. Wyglinski, X. Huang, T. Padir, and et al., "Security of Autonomous Systems Employing Embedded Computing and Sensors," *IEEE Micro*, vol. 33, no. 1, 2013.

[10] F. Sandt and L. Pampagnin, "Perception for a transport robot in public environment," in *Proceedings of the 1997 IEEE/RSJ International Conference on Intelligent Robot and Systems. Innovative Robotics for Real-World Applications. IROS '97*, vol. 1, Grenoble, 9 1997.

[11] J. Petit, B. Stottelaar, M. Feiri, and et al., "Remote attacks on automated vehicles sensors: Experiments on camera and lidar," *Black Hat Europe*, vol. 11, 2015.

[12] H. Shin, D. Kim, Y. Kwon, and et al., "Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications," in *International Conference on Cryptographic Hardware and Embedded Systems*. Springer, 2017.

[13] Y. Cao, C. Xiao, B. Cyr, and et al., "Adversarial sensor attack on lidar-based perception in autonomous driving," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '19. New York, NY, USA: Association for Computing Machinery, 2019.

[14] M. L. Heilig, "El cine del futuro: The cinema of the future," *Presence: Teleoperators & Virtual Environments*, vol. 1, no. 3, 1992.

[15] J. P. Espineira, J. Robinson, J. Groenewald, and et al., "Realistic LiDAR with Noise Model for Real-Time Testing of Automated Vehicles in a Virtual Environment," *IEEE Sensors Journal*, vol. 8, 2021.

[16] B. Littlewood and L. Strigini, "Redundancy and Diversity in Security," in *Computer Security – ESORICS 2004*. Sophia Antipolis: Springer, 2004.

[17] A. Kane, O. Chowdhury, A. Datta, and et al., "A Case Study on Runtime Monitoring of an Autonomous Research Vehicle (ARV) System," in *Runtime Verification*, E. Bartocci and R. Majumdar, Eds. Cham: Springer International Publishing, 2015.

[18] M. Segovia, A. R. Cavalli, N. Cuppens, and et al., "Reflective Attenuation of Cyber-Physical Attacks," in *Computer Security, ESORICS 2019 International Workshops, CyberICPS, SECPRE, SPOSE, and ADIoT*. Luxembourg City: Springer International Publishing, 2020.