# SoundFence: Securing Ultrasonic Sensors in Vehicles Using Physical-Layer Defense

Jianzhi Lou[1], Qiben Yan[1], Qing Hui[2], Huacheng Zeng[1]
[1]SEIT Lab, Computer Science & Engineering, Michigan State University, East Lansing, MI, USA
[2]Electrical & Computer Engineering, University of Nebraska–Lincoln, Lincoln, NE, USA
{loujianz, qyan, hzeng}@msu.edu, qing.hui@unl.edu

***Abstract*—Autonomous vehicles (AVs), equipped with numerous sensors such as camera, LiDAR, radar, and ultrasonic sensor, are revolutionizing the transportation industry. These sensors are expected to sense reliable information from a physical environment, facilitating the critical decision-making process of the AVs. Ultrasonic sensors, which detect obstacles in a short distance, play an important role in assisted parking and blind spot detection events. However, due to their weak security level, ultrasonic sensors are particularly vulnerable to signal injection attacks, when the attackers inject malicious acoustic signals to create fake obstacles and intentionally mislead the vehicles to make wrong decisions with disastrous aftermath. In this paper, we systematically analyze the attack model of signal injection attacks toward moving vehicles. By considering the potential threats, we propose SoundFence, a physical-layer defense system which leverages the sensors' signal processing capability without requiring any additional equipment. SoundFence verifies the benign measurement results and detects signal injection attacks by analyzing sensor readings and the physical-layer signatures of ultrasonic signals. Our experiment with commercial sensors shows that SoundFence detects most (more than 95%) of the abnormal sensor readings with very few false alarms, and it can also accurately distinguish the real echo from injected signals to identify injection attacks.**

## I. Introduction

Recent decades have witnessed the rapid development of autonomous vehicles (AVs), which has an emerging trend of taking over the control authority inside vehicles. The autopilot system of Tesla is one of the most well-known assisted driving systems, which is expected to enable autonomous driving in the future [1]. Currently, there are more than 800,000 Tesla autopilot equipped cars on the road [2]. The Tesla autopilot system utilizes a radar, 12 ultrasonic sensors, as well as multiple cameras [1] to sense the road conditions. The sensors, which sense and provide the environmental data to the decision-making system, have become the core components of the AVs.

The heterogeneous sensors are complex functioning systems, and each type of sensor has its own corresponding working scope. For example, LiDARs and millimeter wave radars sense median-long ranges, while ultrasonic sensors detect obstacles nearby [3]. The readings from heterogeneous sensors are comprehended to assist the decision-making process by domain controllers. The absence or fault in any one of them is lethal to the security and safety of autonomous driving systems.

978-1-6654-4108-7/21/$31.00 

Fig. 1. Signal injection attacks: (Left) creating a fake obstacle; (Right) manipulating the distance to a real obstacle.

However, just as humans have perceptual problems under certain circumstances, such as the bad weather, stress, or fatigue, the cyber-physical sensors are also suffering from misperceptions. For instance, cameras are less effective during the nights [4] and require supplemental methods to sense the environments. Unfortunately, not only do these sensors suffer from the environmental interference, but they can expose victims to various cyber or physical attacks. One particularly popular type of attacks is to inject attacking signals to disable the sensors (i.e., jamming attack) or enforce faulty readings (i.e., signal injection attack). In this paper, we focus on the security of ultrasonic sensors, which play a crucial role in identifying nearby cars or objects during parking or driving.

Comparing with the jamming attack that disables the functionality of ultrasonic sensors, the signal injection attack is more subtle and harder to detect. The attacker creates fake obstacles or manipulates the distance of real obstacles by injecting a spoofing signal into the sensor. Traditionally, ultrasonic sensors play the role of parking assistant. As Fig. 1 shows, a fake obstacle can result in parking in a dangerous position. Also, if the attacker spoofs the sensor to detect a real obstacle at a wrong distance away, the vehicle will miscalculate the position and hit the obstacle to cause accidents. In AVs, the ultrasonic sensors may be responsible for object (such as human or animal) identification [5] when the vehicle is moving or switching lanes, in which case the signal injection attack may lead to more disastrous consequences. Specifically, if a fake obstacle is detected on a busy road, an accident is likely to occur because of the sudden stop or turning around.

Existing studies [6]–[8] have discussed the security issue of ultrasonic sensors, designed and demonstrated signal injection attacks in a laboratory environment or on stationary vehicles. However, it is rare to see real-world signal injection attacks on moving vehicles. One of the main reasons is the difficulty of injecting spoofing signal successfully and consistently. Yan et al. [8] show that the success rate of a random injection is usually below 10%. Moreover, since the sensor operates multiple rounds in one second, a single successful injection is not enough to spoof the driving system. Maintaining a

fake obstacle in a steady and reasonable distance requires the attacker to precisely emit the injection signals to ensure their correct arrival for a series of multiple consecutive rounds. However, in a dynamic environment, it is challenging to control the timing precisely in a millisecond level, especially when the distance between the attacker and the sensor is rapidly changing. As a result, *"how an attacker could improve the chance of successfully attacking a moving vehicle"* remains an open question.

Meanwhile, the mitigation strategies to the signal injection attack are also under investigation. For instance, Xu et al. [7] propose an authentication approach by shifting the physical waveform of the sensor's emitting signal and verifying the received echo. However, many commercial ultrasonic sensors cannot change their frequencies or amplitudes. The cost of involving highly customized devices will become an unnecessary burden to system modularization and standardization in industry. Therefore, another important arising question is: *does there exist any defense mechanism against signal injection attack that can be applied on a commercial ultrasonic sensor?*

To answer these two open questions, we plan to systematically analyze the attack models of ultrasonic signal injection attack. Specifically, we will discuss the possible signal injection attacks for different types of attackers, and we will further show that under certain circumstances, even a less-powerful attacker can perform a signal injection attack that creates fake obstacles with steady distances. Moreover, we propose SoundFence, a physical-layer defense system which can be applied on commercial ultrasonic sensors. SoundFence does not rely on specially designed equipment, but secures a sensor by randomly adapting the sensor's pulsing periods and extracting physical-layer features of the received signals. In particular, the pulsing period adaptation brings perturbation to the spoofed sensor reading, but does not interfere with the benign reading, for which we design detectors to locate the abrupt perturbations. We further explore the difference between the physical-layer signal characteristics of real echo signals and that of injected spoofing signals by analyzing their side echos.

In summary, this paper aims to mitigate the signal injection attacks towards ultrasonic sensors on AVs, and it makes the following contributions:

- We systematically analyze the attack models of ultrasonic signal injection attacks to moving vehicles, and specify the potential attack strategies from different types of attackers with different capabilities.
- We illustrate that even a less-powerful attacker can steadily perform signal injection attacks under certain circumstances, which indicates that the threat of signal injection attacks is real and must be taken into consideration when designing and assembling AV systems.
- We propose an acoustic physical-layer defense system that effectively and efficiently detects malicious signal injection attacks by pulsing period adaptation and side echo analysis. We experimentally demonstrate the effectiveness of the proposed defense system.

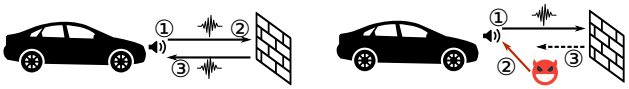

Fig. 2. Ultrasonic sensing and basic signal injection attack: (left) ultrasonic sensing; (right) signal injection attack.

## II. Background & Observations

In this section, we will briefly describe the operation mechanism of ultrasonic sensors and the basic signal injection attack. We will also share our preliminary testing and analysis results on the commodity sensors we purchased.

### A. Ultrasonic Sensors and Basic Signal Injection Attacks

Ultrasonic sensors generally detect obstacles and measure distances in the physical world by actively probing the surroundings with pulses of ultrasound. Most commercial ultrasonic sensors use piezoelectric transducers, which leverage the piezoelectric effect to convert electric charges into mechanical vibrations (and vice versa), and they typically operate within the frequency range between 40 and 50 kHz [7].

Fig. 2(a) is an illustration of how an ultrasonic sensor normally works. The sensor periodically generates pulses (①) and emits them in a waveform signal. The signal hits the surface of the obstacle and reflects back (②). After receiving the echo (③), the sensor computes the distance using the time of flight (ToF):

$$dist = v_{sound} \times (t_a - t_s)/2,$$

where $t_s$ and $t_a$ are the emitting and arriving time of the pulsing signal, and $v_{sound}$ is the speed of sound.

Obviously, the sensor will not wait for the echo signal forever. After waiting for a fixed *timeout duration* $T_0$ [7] without receiving any reflected echos, the sensor determines there are no obstacles nearby and moves on to the next round of measurement. Note that the existing sensors only accept the **first arrival echo** and its arriving time without any signal verification [7], which brings security loopholes — if an attacker injects a pulsing signal at the same frequency, which happens to arrive earlier than the real echo, the sensor will recognize a fake obstacle by computing the wrong distance.

Clulow et al. [9] first introduced the attacks to ToF-based protocols. Specifically, as shown in Fig. 2(b), an attacker injects a spoofing signal, and it arrives ($t'_a$,②) before the real echo ($t_a$,③) (and after the $t_s$,①). Then, the sensor will compute a spoofed distance using the erroneous arriving time:

$$dist_{spoofed} = v_{sound} \times (t'_a - t_s)/2,$$

which means that the attacker creates a fake obstacle in front of the real one.

The attack slot is defined as the time slot between the end of emitted signal and the start of return echo. If the injected signal arrives within the attack slot, the attack is successful, and otherwise, it will be ignored by the sensor. If the attackers aim to create a fake obstacle ahead of the real one (as shown in Fig. 1), they can use *Acoustic Quieting* [10], [11] techniques, such as using sound-absorbing foams, to cancel the real echo signal. When there are no obstacles or the obstacles are hidden by acoustic quieting, the attack slot increases to the timeout

TABLE I
FEATURES OF DIFFERENT ULTRASOUND SENSORS

| Sensor Type | Frequency | $T_0$ | Attack | Blanket Absorption |
|---|---|---|---|---|
| Rb-Dfr-720 [12] | 40 kHz | 38 ms | Succeed | Fail |
| Parallax PING [13] | 40 kHz | 18.5 ms | Succeed | Succeed |
| HC-SR04 [14] | 40 kHz | 60 ms | Succeed | Succeed |

duration $T_0$. If the injected signal falls outside $[t_s, t_s + T_0]$, the attack will be unsuccessful.

### B. Observations from Commercial Sensors

The ultrasonic sensors have two probing modes: *SA1* and *SA2* with the corresponding pulsing period $T_1, T_2$ [7]. In SA1 sensor, the next round measurement begins by waiting for a fixed pulsing period $T_1$ after receiving the echo or timeout; whereas in SA2 sensor, the pulsing period $T_2$ starts right after emitting the pulsing signal.

We purchased three models of commercial ultrasonic sensors, and explored their unique features. All the tested sensors work at 40 kHz. The sensors on vehicles may operate with different frequencies, beam pattern, or maximum distance, but the basic ToF-based design is the same. For each type of sensor, we let one sensor work regularly, and re-programmed another instance to play the role of an attacker, such that the attacker can generate exactly the same signal waveform as the benign sensor. Real attackers are also expected to purchase and utilize these ultrasonic sensors to generate the attack waveforms. As Table I depicts, all types of sensors can be successfully attacked by a spoofer, i.e., a spoofer can deceive the sensor to recognize a fake obstacle, which indicates that the ultrasonic signal injection attack could be prevalent.

As mentioned earlier, the injected signal needs to be falling within the attack slot to launch a successful attack: the larger the attack slot is, the higher chance the attack succeeds. Here, we further compare their timeout duration $T_0$, and the value of $T_0$ is derived by checking the documentation or observing the measurement result when a timeout happens. The results show that the sensor HC-SR04 has the longest $T_0$, which indicates a higher attack success rate. We also test if the ultrasonic signals from different sensors will be absorbed by a blanket. Table I shows that a blanket in a close distance (about 20 cm) is sufficient to absorb the signal of Parallax PING and HC-SR04 sensors, but it fails with Rb-Dfr-720. According to the reading from the Rb-Dfr-720 sensor, the signal either reflects or passes through the blanket. We conjecture that the failure of absorption is caused by the weatherproof feature of Rb-Dfr-720 sensor, which makes its signal more robust against disturbance. In the following experiments, we use HC-SR04 mainly due to its longer timeout duration. Although the attacks against other sensors require the attacker to inject the signals more precisely, they achieve a similar attack impact.

## III. ATTACK ANALYSIS

### A. Threat Model

We assume that the attackers have the resources such as signal analyzers and generators to simulate the sensor's signals, and they can use ultrasonic transducers to generate ultrasonic signals. Their goal is to inject these signals into sensors on an AV to trigger abnormal behaviors. In general, the attacks towards ultrasonic sensors on an automobile will cause two dangerous scenarios (see Fig. 1): i) a vehicle stops moving when there are no obstacles on the path, while other nearby vehicles may crash into it; ii) a vehicle moves toward obstacles on the driving path to cause collision. For ease of exposition, we also assume that the attacker is able to deploy acoustic quieting materials over the surrounding in advance, e.g., the attacker can cover the roadside trees or walls with the materials. For the attackers who cannot perform acoustic quieting, they can still generate a non-existing obstacle or fake a closer distance of real obstacle, which is harmful as we mentioned before.

TABLE II
PARAMETERS OF ATTACK MODEL

| Parameter | Description |
|---|---|
| $i$ | The round number of a sensor's measurement |
| $t_s[i]$ | The time that sensor emits pulsing signal in round $i$ |
| $t_a[i]$ | The time that echo signal returns in round $i$ |
| $t'_a[i]$ | The time that spoofing signal arrives in round $i$ |
| $dist_{fake}[i]$ | The desirable fake distance by the attacker in round $i$ |
| $t_e[i]$ | The time that attacker should start injecting in round $i$ |
| $D(t)$ | The (predicted) distance between the attacker and sensor at $t$ |

We categorize the attackers according to their capabilities:

- **Level 0**: the attacker has **no knowledge** about the target.
- **Level 1**: the attacker has some information obtained **offline**. For instance, the attacker can obtain the target sensor's pulsing period by reading documentations, conducting firmware reverse engineering or measurements.
- **Level 2**: the attacker can obtain more information, besides the offline knowledge, by **initial observations** (e.g., signal emitting time, speed) on each target vehicle.
- **Level 3**: the attacker is able to obtain **real-time information**, and it has abundant resources to perform eavesdropping, replaying, precise position prediction based on the trajectory and real-time speed, etc. Moreover, it can use sophisticated control-theoretic techniques such as Kalman filtering [15] for fairly accurate estimation and prediction.

### B. Generic Attack Model

In general, what the attackers intend to achieve is to inject a signal at a proper timing. A properly injected signal results in a spoofed ToF, which in turn corresponds to an object's distance that the attacker desires. Moreover, considering the temporal consistency, the attacker also tries to spoof several subsequent rounds of measurements to control the sequence of spoofed distances, or at least make them appear steady.

Here, we build a generic attack model using parameters listed in Table II. Since we assume the attackers can perform acoustic quieting to cancel out the real echo, its arriving time $t_a[i]$ in the table is actually not used in our model. Suppose that in the $i$th round, the attacker aims to mislead the sensor in detecting fake obstacle(s) at the distance $dist_{fake}[i]$. First, the attacker computes the desired arriving time of the inject signal:

$$t'_a[i] = t_s[i] + 2dist_{fake}[i]/v_{sound}.$$

Next, since the vehicle is moving, the attacker needs to predict the distance $D(t'_a[i])$ at $t'_a[i]$, based on which the attacker computes the emitting time $t_e[i]$ as follows:

$$\begin{aligned} t_e[i] &= t'_a[i] - D(t'_a[i])/v_{sound} \\ &= t_s[i] + (2 \cdot dist_{fake}[i] - D(t'_a[i]))/v_{sound}. \end{aligned}$$

According to our attack categorization, only the Level 3 (most knowledgeable and powerful) attackers can achieve the most generalized version of signal injection attack by real-time eavesdropping (to obtain $t_s[i]$) and precise position prediction (to obtain $D(t'_a[i])$).

However, there still remain some physical limitations, such as the propagation and processing delays of injected spoofing signals, and the processing delays of computation and position estimation. These limitations prevent $dist_{fake}[i]$ from being too close. We let $dist_{timeout} = v_{sound} \times T_0/2$ be the maximum achievable distance, which is calculated using timeout duration $T_0$ as the ToF. Since $dist_{fake}[i]$ cannot exceed $dist_{timeout}$ (otherwise the injected signal will go beyond the attack slot), the possible choices of fake distance are quite limited.

### C. Practical Attack Strategy

Although the Level 3 attacker is powerful, such a real-time, eavesdropping-based attack is still hard to perform because of the physical limitations mentioned above. To develop a practical attack strategy, an intuitive idea is to utilize some pre-obtained information to simplify the computation and allow the attacker to prepare in advance.

First, the attacker does not have to detect the sensor's emitting time $t_s[i]$ in every round. As we have described in Section II-B, the sensors have fixed pulsing periods ($T_1$ or $T_2$) according to their probing mode (SA1, SA2), which can be accessed by the attacker in advance. For SA1 sensor, the next pulse depends on the received signal. If the spoofing is successful at $t'_a[i]$ in current round $i$, the sensor will emit pulse for the next round at $t_s[i+1] = t'_a[i] + T_1$. In other words, once there is a successful attack on the current round, the sensor's transmission on the subsequent rounds will be taken over by the attacker. For SA2 sensor, the next pulse follows the current one with a fixed period, i.e., $t_s[i+1] = t_s[i] + T_2$. The attacker can observe several initial pulses to estimate the emitting time of the subsequent rounds.

Second, in most cases, the attacker does not need a precise prediction of the victim position. The attacker can elaborately choose a region (e.g., a straight lane), where the vehicles are supposed to move on a simple trajectory with a constant speed. Moreover, the attacker's desired spoofing distance $dist_{fake}[i]$ is not necessary to change in every round. In practice, the attacker prefers to choose a constant spoofing distance $dist_{fake}[i] = d$ for all rounds, or to increase or reduce the distance linearly (for simulating a fake approaching/leaving obstacle), which lowers the difficulty of the attack. After aforementioned simplifications, a Level 2 attacker will be able to perform signal injection attacks based on the initial observations and inferences.


(a) Short-period "flooding" (b) Attacking (static) SA1 sensors

Fig. 3. (a) Random, short ToF; (b) steady ToF (① attacker's pulsing (and also arriving) period $T_1 + 2dist_{fake}[i]/v_{sound}$; ② SA1 sensor's pulsing period $T_1$; ③ spoofed ToF $2dist_{fake}[i]/v_{sound}$.)

### D. Zero-effort Spoofing with Short Pulsing Period

Although the aforementioned practical considerations decrease the difficulty of attack, the attacker needs to earn prior knowledge of victims via initial observations. For a zero-effort attacker (or called *Ignorant Attacker* in [16]), although it is hard to control the fake distance, there are plausible solutions to reach similar outcomes.

Short-period spoofing approach has been adopted by [7], [8] to increase the success rate of attack. As Fig. 3(a) shows, the attacker can flood the target sensor by a great quantity of spoofing signals with a short pulsing period. When the attacker's pulsing period is short enough, the sensor readings will be fluctuating within a small range, which causes a fake obstacle to be very close with relatively steady sensor readings.

### E. Limited-effort Spoofing Using Sensor's Pulsing Period

As mentioned before, the major difficulty that prevents steady attacks in multiple rounds is the uncertainty of spoofing signal's propagation delay, which is caused by the motion of the vehicle and the changing distance. However, with a relatively fixed distance, it is possible to perform successful, steady attacks with a limited effort. Here, we adopt a signal injection attack strategy by periodically inject spoofing signals with a fixed pulsing period $T$ ($T_1 < T < T_1 + T_0$).

To attack SA1 sensors with pulsing period $T_1$, the attacker should set pulsing period as $T_1 + 2dist_{fake}[i]/v_{sound}$. When the distance is relatively fixed, the propagation time of attack signal in each round becomes constant. Meanwhile, for SA1 sensor, a successful injection in the previous round (that arrives at $t'_a[i-1]$) allows the attacker to control the emitting time of next sensing pulse ($t_s[i] = t'_a[i-1] + T_1$). Since the next injection signal arrives at $t'_a[i] = t'_a[i-1] + (T_1 + 2dist_{fake}[i]/v_{sound})$, the computed distance $dist[i]$ in the next round will be:

$$dist[i] = v_{sound} \cdot (t'_a[i] - t_s[i])/2 = dist_{fake}[i],$$

which is exactly the attacker's desired fake distance, implying that the fake distance can be manipulated and injected into the subsequent rounds by controlling the attack sensor's emitting time.

It is worth noting that: this attack does not depend on sensor's initial emitting time and spoofing signal's propagation time. As shown in Fig. 3(b), the only and most important factor is the difference between the sensor and attacker's pulsing periods (③). It requires a first successful signal injection, which helps taking over the subsequent rounds. In practice, the attacker can predetermine its pulsing period to be fixed

TABLE III
LIST OF POSSIBLE INJECTION ATTACKS WITH POTENTIAL ATTACK RESULTS.

| Attacker | Scenario | Required information/action | Potential injection result |
|---|---|---|---|
| Level 3 | General | Real time eavesdropping; Position predication | Arbitrary desired $dist_{fake}$ sequence |
| Level 2 | Simple environment (e.g., straight lane) | Prior knowledge; Initial observation; Simpler predication | Arbitrary desired $dist_{fake}$ sequence |
| Level 2 | (Relatively) stopped vehicle; SA1 or SA2 sensor | Prior knowledge; Initial observation | Predetermined $dist_{fake}$ sequence |
| Level 1 | (Relatively) stopped vehicle; SA1 sensor | Prior knowledge; Repeating random attempts | Predetermined $dist_{fake}$ sequence |
| Level 1 | (Relatively) stopped vehicle; SA2 sensor | Prior knowledge; Repeating random attempts | Non-predetermined $dist_{fake}$ sequence |
| Level 0 | General | Attack with very short period | Small $dist_{fake}$ sequence with small deviation |

(i.e., $dist_{fake} = d$) or slowly change the period to simulate a moving fake obstacle.

As for SA2 sensors which have a fixed pulsing period $T_2$ between transmissions, the intuitive attack approach is to let the attacker inject signals with the same pulsing period $T_2$, ensuring a fake obstacle with fixed distance in each round. If the attacker uses a slightly different pulsing period, the difference will be accumulated in the following rounds, resulting in steadily increasing/decreasing fake sensor readings.

Here, we further apply the aforementioned attacks in a dynamic environment where the vehicles are moving, based on the following observations:

i) The attacks can be achieved by **less-powerful attackers, with little effort.** For SA1 sensor, the attacker only needs to learn $T_1$ in prior. When the attacker does not observe the emitting time of the sensor, it can still launch attack successfully if it randomly attempts multiple times and injects several rounds in every attempt. If the success rate is $r$ in a random condition (less than 10% in [8]), it is still very possible (probability of $1-(1-r)^k$) to have a successful attack when the number of attempts $k$ is large. For SA2 sensor, the outcome is similar. A random attempt cannot predetermine the spoofed distance, but the sensor readings in the successful attempts will remain steady. The required functionality, timing and emitting, can be achieved by a Level 1 attacker with pre-programmed, battery-driven devices.

ii) The attacker can be hidden in **unnoticeable places**. According to our experiments, the attacker can be placed toward the sensor but with some angles within the sensor's detection coverage. In this case, although the attacker cannot receive the sensor's signal, the attacker's signal can still be injected into the sensor. As a result, if the attacker does not eavesdrop the sensor's signal, it may not have to be placed face-to-face with the sensor. In other words, the attacker can be placed not only on the obstacle, but also elsewhere like roadside, another vehicle moving along with the target, or even on the target vehicle.

The last two attack scenarios make the distance between the attacker and the sensor (almost) fixed, which indicates that the less-powerful attacker, like Level 1, can perform attacks with steady spoofed sensor readings on moving targets: an attacker can i) drive through the city and keep injecting signals to the victims nearby, or ii) adhere the attack device on the target in prior and trigger it at proper timing. Both of these attack strategies could cause disastrous damages.

Being concerned with the safety and hardware limitation, we did not validate the moving attack on a real vehicle. Instead, we performed a simulation of the last scenario (i.e., with attached attacker on the target vehicle) by attaching the sensor and the attacker on a suitcase. As Fig. 4(a) shows, the (SA1) sensor has a fixed pulsing period $T_1 = 100$ ms, and the attacker randomly attempts spoofing attacks. In each attempt, the attacker repeatedly injects with a fixed pulsing period, and we make the pulsing period increases with attempts. We test both (i) *static scenario* (with acoustic quieting) and (ii) *outdoor, moving scenario* with normal human walking speed.

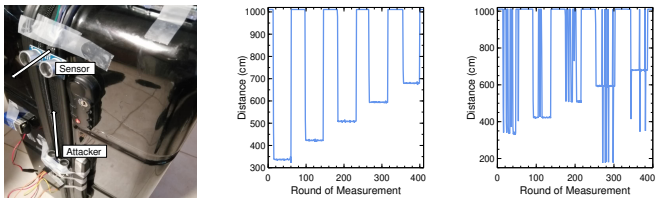


(a) Attack setup (b) Attacking static object (c) Attacking moving object (human-speed)

Fig. 4. Attack experiments (when the attacker sleeps, the sensor readings are beyond 10 m because of the acoustic quieting or outdoor environment).

As Fig. 4(b) and 4(c) shows, in the static scenario, the spoofing attack is successful, resulting in steady, predetermined distances. When attacking a moving object, although there are some failed injections (caused by the shaking during the walk on a non-flat road), the attacker can still inject most of the faked distances successfully. The experimental result shows that a well-prepared attacker could have a high success rate even on a moving object. We consider to safely verify the attacks on real moving vehicles, as well as implement the attacks with other capabilities/scenarios in our future work.

### F. Summary

To summarize, we categorize the attackers by their capabilities, and analyze the possible attack strategies for each type of attackers, which are listed in Table III. In particular, *we demonstrate that the attack strategies with attackers in Level 0 or 1 need to be taken into consideration by sensor designers, since they can be easily achieved by less-powerful attackers with limited efforts.*

## IV. DEFENSE AGAINST SIGNAL INJECTION ATTACKS

In this section, we will propose *SoundFence*, a physical-layer defense system which detects and rejects malicious signals or spoofed sensor readings. SoundFence can be applied on commercial ultrasonic sensors without professional equipment.

### A. Phase I: Pulsing Period Adaptation

As we have discussed in Section III, the majority of attack strategies require the sensor's pulsing period as key information. Only most powerful (Level 3 attacker) and zero-effort (Level 0) attackers can launch attack in general scenarios.

TABLE IV
ATTACK TOWARDS SENSORS WITH RANDOMIZED PULSING PERIODS.

| Attacker's pulsing period (ms) | Success rate | STD (cm) |
|---|---|---|
| 200 | 0.29 | 270.04 |
| 100 | 0.52 | 306.75 |
| 50 | 1.00 | 245.87 |
| 10 | 1.00 | 49.26 |

**Randomizing Pulsing Periods.** The first defense approach of SoundFence focuses on the randomization of sensor's pulsing period. The randomization does not change anything on the ToF of normal echo but builds huge barriers to the attackers. As we have mentioned in Section II-A, the spoofing signal must arrive before timeout (or the real echo), but the length of attack slot is small comparing with the entire pulsing period. In most of the attacks, the attacker utilizes sensor's pulsing period to schedule the spoofing signal to ensure them arriving properly. However, when the sensor's pulsing period is randomized, it becomes difficult to inject a signal within the attack slot, which causes plenty of failed attacks.

We measure the attack success rate after uniformly randomizing the sensor's pulsing period between 100 to 300 ms, which provides randomness without considerably affecting the functionality. In fact, the pulsing period of vehicle sensors are approximately 100 ms [8]: after our randomization, there are still at least 3 readings per second). We set attackers with different pulsing periods, and compute the success rate by counting the successes/failures from the first success to the last success. As Table IV illustrates, regardless of the attacker choosing 200 ms (average of sensor's period) or 100 ms (minimum), plenty of spoofing signals cannot be injected successfully, and the sensor reading (with acoustic quieting) shows the maximum result $dist_{timeout}$. Using a shorter period like 10 ms or 50 ms (zero-effort attack in Section III-D), the attacker is able to "flood" the sensor and increases the success rate. It is worth noting that a victim sensor with a smaller timeout duration, or an attacker without acoustic quieting, will further lower the success rate.

**Abrupt Change Detectors.** Moreover, even when the spoofing signal falls inside the attack slot, the spoofed distance will jump back and forth, since the spoofed ToF changes abruptly because of the sensor's randomization. On the other hand, the randomization does not affect the benign sensor readings. In most cases, the real distance changes smoothly, with only very few abrupt change happens when the obstacle appears/disappears. As a result, the abrupt change in sensor reading can be utilized by SoundFence to detect the signal injection attack. We design a detector $\mathcal{D}(w, \tau, \gamma)$, which counts the abrupt sensor reading changes in the past $w$ measurements. If there are more than $\gamma$ reading changes $|dist[i] - dist[i-1]|$ exceeding threshold $\tau$, it is highly likely to be a signal injection attack.

The detector $\mathcal{D}(w, \tau, \gamma)$ has to satisfy two requirements: i) it should detect as many attacks as possible; ii) it should avoid triggering false alarms on benign readings. These requirements elaborate the selection of parameters $(w, \tau, \gamma)$ of the detector. However, after manually selecting several sets of parameters, our experimental results show that there are no such universal parameters which can satisfy all the situations. Some parameters work well for attacks on a certain pulsing period without triggering false alarms, but the performance is poor on other pulsing periods.

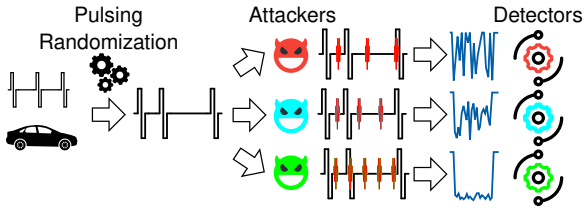

Fig. 5. Multiple detectors for different spoofed sensor reading patterns.

This phenomenon is caused by the difference in the patterns of spoofed sensor readings. When attacking with small pulsing period, the success rate is high, which implies that the sensor reading keeps being spoofed, and the deviation is relatively small (but still larger than benign readings). Consequently, a small distance threshold $\tau$ with a large count threshold $\gamma$ is more helpful to detect the attacks with short pulsing periods and to distinguish them from benign cases. On the other hand, when the attacking pulsing period is larger, the extent of change increases, but the success rate becomes very low, which implies that the spoofed readings are mixed up with some steady, benign readings ($dist_{timeout}$ in our experiments), impairing the counting of abrupt changes. In this case, we infer that a larger $\tau$ with a smaller $\gamma$ is more helpful.

Considering the difference of the spoofed sensor reading patterns, it is infeasible to select an optimal set of parameters. A sound solution is to deploy multiple detectors $\mathcal{D}_1(w_1, \tau_1, \gamma_1), \mathcal{D}_2(w_2, \tau_2, \gamma_2)$..., which work in parallel. Each of them is responsible for detecting and reporting the attacks in a certain pulsing period (or a subset range of periods) as shown in Fig. 5.

### *B. Phase II: Side Echo Signal Analysis*

While pulsing period adaptation prevents most of the prior-knowledge-based attacks (as well as some zero-effort attacks), the most-knowledgeable attackers with the highest capability may still succeed. Although there are several limitations, the real-time eavesdropping-replaying attack can bypass the first phase of SoundFence. To further improve the performance and reliability of SoundFence under the threat of remaining attacks, we design another approach, which is based on the physical-layer signatures of the received signal.

We extract features by leveraging two observations on signal reflection: i) in a real-world environment, acoustic signals are more likely reflected by multiple surfaces, through multiple directions at different distances away; and ii) the reflected signals attenuate more than the direct signals. The first observation results in the prevalence of weaker side echo signals after the main echo (or after the direct injection), while the second observation indicates the difference of signal strength between real echo signal and directly injected signal. Although the main echos are the most powerful one among all the echos, they are

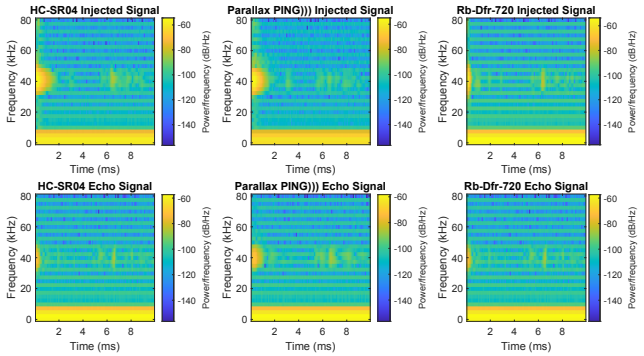


Fig. 6. The spectrogram of injected signals versus real echo signals (window size: 32; FFT length: 64). Side echos can be observed between 2-10 ms.

still less noticeable than the attack signal, which is directly injected (and usually at a close distance) without reflection.

We conducted an experiment in a laboratory environment, and the spectrograms of injected and real echo signals are displayed in Fig. 6, which shows a noticeable difference. Both of the real and injected signals have similar side echos which are caused by the environments (2-10 ms), but the injected signal has more power (brighter) in 0-2 ms, comparing with the real signal. While the numerical value of signal strength are impacted by many environmental factors, the ratio of side echo can help distinguishing the real and injected signals.

SoundFence utilizes such characteristics of side echo signals for physical-layer defense. It runs on the commercial sensor's microcontroller, and allows the sensors to keep receiving echo signals after the main echo arrives. The main echo and side signals are separated by a short idle period following the length of the original pulsing signal. Our proposed scheme utilizes *spectral power ratio*. First, after signal separation, we measure the spectral powers of these two signals by computing the Short-Time Fourier Transform (STFT) of the echo signals. The obtained signal spectrogram contains frequencies and corresponding power over time. Then, the cumulative spectral power $\vec{S}_{power}$ is computed, which is a vector that contains the total accumulated power per frequency:

$$\vec{S}_{power}(f) = \sum_{t=0}^{L} S_{STFT}(f,t), \tag{1}$$

where $f$ is the index of the selected frequency that falls in the range of 40∼50 kHz, $S_{STFT}(f,t)$ represents the STFT feature at a specific time and frequency index, and $L$ is the time duration of the signal of interest. The *spectral power ratio* $\lambda$ of the main and side echo signals will be computed as:

$$\lambda = ||\vec{S}^{side}_{power}||/||\vec{S}^{main}_{power}||, \tag{2}$$

where $\vec{S}^{side}_{power}$ and $\vec{S}^{side}_{power}$ are the spectral power vectors of side and main signals, respectively. If $\lambda$ is less than a threshold, we consider the main echo signal as an injected signal.

## V. Experimental Evaluation

In this section, we will evaluate the performance of the two defense phases of SoundFence. For the first phase, we assume the attacker has level 1 or less capability and can only inject signals with a fixed pulsing period.

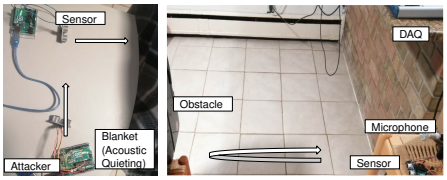


Fig. 7. Experiment settings for phase I (left) and II (right) defense evaluation.

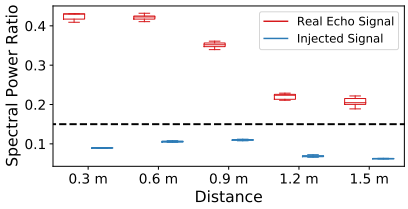


Fig. 8. Distribution of spectral power ratio with different distances.

### A. Pulsing Period Adaptation

We use Arduino UNO Rev3 [17] to control and obtain reading from HC-SR04 ultrasonic sensor. As mentioned in Section II-B, we use another ultrasonic sensor for attack, and perform acoustic quieting using a blanket with close distance, which is enough to quiet HC-SR04 sensor. For simplicity, the sensor and the attacker are placed on a table (Fig. 7(a)) to simulate the fixed relative position which we have discussed in Section III-E. For the attacker, we try different pulsing periods: 10 ms, 50 ms, and 100 ms. We choose not to use 200 ms pulsing period since the attack success rate is too low (see Table IV). The attacker is automatically switching between *sleep* and *attack* modes. During the sleep phase, it stops working for 5 seconds; during the attack phase, it keeps sending spoofing signals for about 5 seconds. For instance, if the attacking pulsing period is 50 ms, the attacker repeats emitting 100 spoofing signals. Considering the processing delay, the total attacking time may be slightly longer than 5 seconds. Here, we generate a sequence of distances which contain benign and fake readings. The detector $\mathcal{D}$ is required to detect the attacks as soon as they happen. The detection rate is defined as the ratio between the number of detected attack phases and that of total attack phases.

We also evaluate the false alarm rates of the proposed scheme. However, if we do not perform acoustic quieting in the attacking sequences, the arrived echo will considerably affect the success rate of spoofing attack. As a matter of expediency, we decide to collect real readings separately, and we run the detectors on the real reading to check the false alarms. Since there are no attack/sleep phases in the real reading, the false alarm rate can be computed by:

$$\frac{\text{\# of readings exceeds threshold of } \mathcal{D}}{\text{\# of total readings}}.$$

After collecting the malicious and benign data, we exhaustively search for the values of parameters $(w, \tau, \gamma)$ from a set of plausible values for each attacking pulsing period. Table V

TABLE V
OPTIMAL PARAMETERS $(w, \tau, \gamma)$ AND THEIR PERFORMANCE

| Attacker's pulsing period (ms) | $(w, \tau, \gamma)$ | Detection rate | False alarm |
|---|---|---|---|
| 10 | (20, 10cm, 15) | 96.7% | 0.12% |
| 50 | (13, 80cm, 8) | 97.3% | 0% |
| 100 | (7, 250cm, 2) | 95.3% | 3.3% |

lists the optimal parameters with the corresponding detection rate and false alarm rate. For all the three attack periods, the detection rates exceed 95%.

### B. Side Echo Signal Analysis

To evaluate the performance of side echo analysis, we collect real echo and injected signals, and test their distinguishability. Since HC-SR04 sensors do not have analog output, we use an ultrasonic microphone (Avisoft-Bioacoustics CM16/CMPA [18]) to receive the signal, and adopt a data acquisition system (NI USB-6341 [19]) to capture the data with sampling rate 160kHz (Fig. 7(b)). In a real-world defense implementation, SoundFence could probe the sensor and access the analog signals directly from the sensor [20].

When collecting real echo signals, we place the sensor and microphone together at the same direction, and the microphone receives the echo reflected from an obstacle; during the attack, we place the microphone and attacker face to face, so that the microphone receives injected signal directly. We collect 50 real echo signals and 50 injected signals respectively, with distance ranging from 0.3 m to 1.5 m. Different from the outdoor experiment, in an indoor environment, the reflection from the walls interferes with our result. As a mitigation, we stay away from the walls, and only clip 8 ms-long pieces from the beginning of main signal, to ensure all the side echos are from the reflections in close distances.

We apply STFT (window size=32, nfft=64, noverlap=16) on the collected signals to compute their power density, and then we choose index $f$=40 kHz to compute the spectral power ratio $\lambda$. According to our observation, the duration of the main signal is about 1 to 1.5 ms. We test different $L_{main}$ and $L_{side}$ to see whether the computed $\lambda$ can be utilized to distinguish among all the 50 injected signals and 50 real signals. Table VI depicts the detailed attack detection performance with different parameters. In a nutshell, we found that when using $L_{main}$=1 ms to compute $\vec{S}^{main}_{power}$, the selection of $L_{side}$ is limited, which affects the detection performance. However, with $L_{main}$=1.5 ms, it is much easier to distinguish between real and injected signals using a certain threshold (e.g., $\theta$=0.15 when $L_{side}$=6.5 ms) as shown in Fig. 8. The results demonstrate that, with a properly configured threshold, the proposed scheme can reach 100% detection rate with zero false alarm.

Note that in this evaluation, we adopt the optimized parameters to achieve the best defense performance. In a real-world environment with real vehicles, deriving the optimized parameters is challenging. However, due to the lightweight design of the defense mechanisms, the vehicles can iteratively attempt different parameters to increase the chance of detecting signal injection attacks.

TABLE VI
DETECTION RESULTS WITH DIFFERENT LENGTHS OF MAIN/SIDE SIGNALS

| Main signal | Side echo | Best $\theta$ | Detection rate | False alarm |
|---|---|---|---|---|
| 1 ms | 4 ms | 0.14 | 80% | 0% |
| 1 ms | 5 ms | 0.18 | 100% | 14% |
| 1 ms | 6 ms | 0.19 | 100% | 0% |
| 1.5 ms | 4 ms | 0.10 | 100% | 0% |
| 1.5 ms | 5 ms | 0.12 | 100% | 0% |
| 1.5 ms | 6 ms | 0.14 | 100% | 0% |
| 1.5 ms | 6.5 ms | 0.15 | 100% | 0% |

## VI. RELATED WORK

**Attacks towards Ultrasonic Sensors.** Recent studies have discussed the security of ultrasonic sensors. A series of contact-less attacking experiments, which include jamming, spoofing (with short-period attack), and acoustic quieting are performed in [8]. Gluck et al. [20] create fake obstacles by injecting noises into the victims. Lim et al. [21] also assess the vulnerabilities of ultrasonic sensors by attacking HC-SR04 Sensor. Moreover, Xu et al. [7] utilize the acoustic quieting technique to create a fake obstacle ahead of the real one. They successfully generate steady fake sensor readings in a laboratory environment based on the analysis of the pulsing period. In this paper, different from the prior work, we illustrate that the pulsing-period based attack is a realistic attack against moving vehicles, and it can be launched by pre-programmed attacking devices with very limited functionalities.

**Defenses for Protecting Ultrasonic Sensors.** Besides attacking experiments, Xu et al. [7] propose a defense approach to protect and authenticate emitted signals by shifting physical parameters or checking the data consistency from multiple sensors. Another defense approach [16] models the received signal strength based on pulse length, distance and multi-path, and rejects the abnormal signal by measuring the length of received signal above a threshold. There are also general defense mechanisms to secure different types of sensors. For instance, PyCRA [22] secures active sensors against injections by randomly turning off the sensor and observing the readings. However, PyCRA cannot be applied for securing ultrasonic sensor due to the additional oscillation (i.e., ringing time) of the ultrasonic sensor after signal transmission. Different from the existing approaches, our proposed defense, SoundFence, does not require any additional equipment to generate special waveform, nor does SoundFence require the cooperation of multiple sensors. Table VII provides a comparison of different defense methods.

## VII. DISCUSSION & FUTURE WORK

Although SoundFence achieves good performance in detecting injection attacks, there are also some limitations. First, if the attacker's pulsing period is extremely small (such as 1 to 2 ms), it may bypass the detection of SoundFence Phase I, since the perturbation of sensor readings will not be large enough to be detected. On the other hand, the spoofing signals will flood the time domain of acoustic channel, and it becomes difficult to extract side echo features for SoundFence Phase II. Due to the hardware limitation, we did not further investigate the attacks in a very short period, and we will explore it in future.

TABLE VII
COMPARISON WITH OTHER DEFENCE METHODS

| Defence Method | Key Idea | Applicable Sensors | Requirement / Limitation |
|---|---|---|---|
| Xu et al. [7] | Shifting physical signals / Multiple sensor triangulation | Ultrasonic sensor | Requires to change the waveform (frequency, amplitude, etc.) / Requires multiple sensors |
| Lee et al. [16] | Receiving signal modeling | Ultrasonic sensor | Requires to change the waveform (duration) |
| Shoukry et al. [22] | On-Off challenge | Active sensors | Ultrasonic sensors continue oscillating after being turned off |
| SoundFence (Ours) | Pulsing period randomization / Side echo analysis | Ultrasonic sensor | Requires to change the pulsing period / Requires to access the analog signal |

Another promising future topic is the real-world implementation of attacks toward vehicles. In this paper, we theoretically analyze the feasibility of attacking moving vehicles with a simplified human walking speed validation. It remains an open question on how to launch attack effectively and secretly on a fast-moving vehicle. For Level 3 attackers who eavesdrop and replay in real-time, relaying the acoustic signal through wireless channels (which is used in cancelling acoustic noise by [23]) may help eliminate the physical limitation of the attacker. For less-powerful attackers who should adhere to the vehicle, delivering spoofing signal through surface materials [24] may help hide the attacker and improve the performance.

## VIII. CONCLUSION

In this paper, we systematically analyzed the attack models of spoofing the ultrasonic sensors on a moving vehicle based on different knowledge levels of attackers. We discovered that less-powerful attackers can also perform signal injection attacks to create steady fake obstacles with minimum efforts, which should be seriously considered by security researchers and communities. We further proposed SoundFence, a physical-layer defense system that works on commercial ultrasonic sensors to reject malicious signals or sensor readings. SoundFence detects most of the abnormal sensor readings with minor false alarms, and it can also distinguish real echo from injected signals. In the future, we will work on further investigation of the other potential attacks and design of corresponding countermeasures. This study corroborates that it remains a long journey to enter a securely connected cyber and physical world for safe, reliable autonomous and semi-autonomous AV systems, and to benefit AV sensing system design and its prototype security evaluation.

## ACKNOWLEDGEMENT

We thank the reviewers and their constructive comments. This work was supported in part by NSF CNS-2113618, CNS-1950171 and CNS-1949753.

## REFERENCES

[1] "Tesla autopilot," https://www.tesla.com/autopilot, Accessed on Jan. 12, 2021.

[2] D. Z. Morris, "Tesla could deliver 'full self-driving' within weeks," https://fortune.com/2019/11/20/tesla-full-self-driving-car-tsla-stock, Accessed on Jan. 12, 2021.

[3] K. Ren, Q. Wang, C. Wang, Z. Qin, and X. Lin, "The security of autonomous driving: Threats, defenses, and future directions," *Proceedings of the IEEE*, vol. 108, no. 2, pp. 357–372, 2019.

[4] M. Venables, "Sensing an autonomous vehicle future," https://spie.org/news/photonics-focus/marapr-2020/sensing-an-autonomous-vehicle-future?SSO=1, Accessed on Jan. 12, 2021.

[5] N. Heath, "Tesla's autopilot: Cheat sheet," https://www.techrepublic.com/article/teslas-autopilot-cheat-sheet/, Accessed on Jan. 7, 2021.

[6] S. Sedihpour, S. Capkun, S. Ganeriwal, and M. Srivastava, "Implementation of attacks on ultrasonic ranging systems," in *Demo at the ACM Conference on Networked Sensor Systems (SenSys)*, vol. 10, 2005.

[7] W. Xu, C. Yan, W. Jia, X. Ji, and J. Liu, "Analyzing and enhancing the security of ultrasonic sensors for autonomous vehicles," *IEEE Internet of Things Journal*, vol. 5, no. 6, pp. 5015–5029, 2018.

[8] C. Yan, W. Xu, and J. Liu, "Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle," *DEF CON*, vol. 24, no. 8, p. 109, 2016.

[9] J. Clulow, G. P. Hancke, M. G. Kuhn, and T. Moore, "So near and yet so far: Distance-bounding attacks in wireless networks," in *European Workshop on Security in Ad-hoc and Sensor Networks*. Springer, 2006, pp. 83–97.

[10] S. A. Cummer and D. Schurig, "One path to acoustic cloaking," *New Journal of Physics*, vol. 9, no. 3, p. 45, 2007.

[11] J. Li and J. B. Pendry, "Hiding under the carpet: A new strategy for cloaking," *Physical review letters*, vol. 101, no. 20, p. 203901, 2008.

[12] RobotShop, "Weatherproof ultrasonic sensor w/ separate probe," https://www.robotshop.com/en/weatherproof-ultrasonic-sensor-separate-probe.html, Accessed on Jan. 7, 2021.

[13] ——, "Parallax ping ultrasonic sensor," https://www.robotshop.com/en/parallax-ping-ultrasonic-sensor.html, Accessed on Jan. 7, 2021.

[14] Mouser Electronics, "Ultrasonic Ranging Module HC – SR04," https://www.mouser.com/datasheet/2/813/HCSR04-1022824.pdf, Accessed on Jan. 12, 2021.

[15] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[16] S. Lee, W. Choi, and D. H. Lee, "Securing ultrasonic sensors against signal injection attacks based on a mathematical model," *IEEE Access*, vol. 7, pp. 107 716–107 729, 2019.

[17] Arduino, "Arduino Uno Rev3," https://store.arduino.cc/usa/arduino-uno-rev3, Accessed on Jan. 12, 2021.

[18] Avisoft, "Condenser ultrasound microphone Avisoft-Bioacoustics CM16/CMPA," http://www.avisoft.com/ultrasound-microphones/cm16-cmpa/, Accessed on Jan. 12, 2021.

[19] National Instruments, "USB-6341 Multifunction I/O Device," https://www.ni.com/en-us/support/model.usb-6341.html, Accessed on Jan. 12, 2021.

[20] T. Gluck, M. Kravchik, S. Chocron, Y. Elovici, and A. Shabtai, "Spoofing attack on ultrasonic distance sensors using a continuous signal," *Sensors*, vol. 20, no. 21, p. 6157, 2020.

[21] B. S. Lim, S. L. Keoh, and V. L. Thing, "Autonomous vehicle ultrasonic sensor vulnerability and impact assessment," in *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*. IEEE, 2018, pp. 231–236.

[22] Y. Shoukry, P. Martin, Y. Yona, S. Diggavi, and M. Srivastava, "Pycra: Physical challenge-response authentication for active sensors under spoofing attacks," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 2015, pp. 1004–1015.

[23] S. Shen, N. Roy, J. Guan, H. Hassanieh, and R. R. Choudhury, "MUTE: Bringing IoT to noise cancellation," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018, pp. 282–296.

[24] Q. Yan, K. Liu, Q. Zhou, H. Guo, and N. Zhang, "SurfingAttack: Interactive hidden attack on voice assistants using ultrasonic guided waves," in *Network and Distributed Systems Security (NDSS) Symposium*, 2020.