Analyzing and Enhancing the Security of Ultrasonic Sensors for Autonomous Vehicles

Wenyuan Xu, *Senior Member, IEEE,* Chen Yan, Weibin Jia, Xiaoyu Ji, *Member, IEEE,* and Jianhao Liu

***Abstract*—Autonomous vehicles rely on sensors to measure road condition and make driving decisions, and their safety relies heavily on the reliability of these sensors. Out of all obstacle detection sensors, ultrasonic sensors have the largest market share and are expected to be increasingly installed on automobiles. Such sensors discover obstacles by emitting ultrasounds and analyzing their reflections. By exploiting the built-in vulnerabilities of sensors, we designed random spoofing, adaptive spoofing, and jamming attacks on ultrasonic sensors, and we managed to trick a vehicle to stop when it should keep moving, and let it fail to stop when it should. We validate our attacks on stand-alone sensors and moving vehicles, including a Tesla Model S with the 'Autopilot' system. The results show that the attacks cause blindness and malfunction of not only sensors but also autonomous vehicles, which can lead to collisions. To enhance the security of ultrasonic sensors and autonomous vehicles, we propose two defense strategies, single-sensor based Physical Shift Authentication (PSA) that verifies signals on the physical level, and Multiple Sensor Consistency Check (MSCC) that employs multiple sensors to verify signals on the system level. Our experiments on real sensors and MATLAB simulation reveal the validity of both schemes.**

***Index Terms*—Autonomous Vehicle, Ultrasonic Sensor, Security Analysis, Defense.**

I. Introduction

AUTOMOBILE is one of the most promising sectors for the Internet of Things (IoT). By converting vast amount of data into meaningful and actionable knowledge, the IoT can help solve many of modern society's challenges on automotive safety and transportation efficiency. Among them, autonomous (self-driving) vehicles is one notable achievement, and it holds the key to a widespread Internet of Vehicles. Self-driving technologies are built on *modern sensors* that enable vehicles to monitor the driving environment by themselves. Already, a preliminary stage of self-driving has been widely deployed as the Advanced Driver Assistance Systems (ADAS). Looking forward, vehicles will inevitably rely on these sensors and their measurements to become increasingly intelligent until they reach the full-fledged self-driving capability, i.e., require zero human interaction to make driving decisions. However, the safety of self-driving vehicles is determined by the reliability of sensors. Recent unfortunate fatal accidents [1], [2] of Tesla Model S with the Autopilot system [3] (i.e., the most advanced autonomous systems in the market) are caused by sensor failures, i.e., sensors cannot reliably detect neighboring cars in normal yet special road conditions. Such conditions are benign, and it is worth investigating the malicious scenarios to understand the following: (a) How will these sensors perform under intentional attacks in practice? (b) How will automobiles behave under such attacks? (c) How to enhance these sensors to defend against intentional attacks? This paper answers these questions with a case study on ultrasonic sensors.

Ultrasonic sensors detect obstacles and measure distance by probing the surroundings actively with pulses of ultrasound. They are widely used on IoT devices for ranging and occupancy detection, which correspond to two scenarios on a vehicle: (1) *parking* when a car is traveling at low speeds and (2) detecting blind spot at *high speeds*. In terms of parking sensors, more than half the new vehicles in Europe and Asia have rear parking sensors [4], and the Indian government will soon mandate all new vehicles to be equipped with such sensors to lower the risk of backover crashes [5]. Since NHTSA has reported 292 fatalities with 44% of them being children under five-year old and 18,000 injuries resulting from backover crashes every year [6], it is not surprising that the global automotive parking sensors market is predicted to grow steadily during the next few years with a compound annual growth rate of almost 24% by 2020 [7]. As a pioneer of autonomous vehicle in the consumer market, Tesla Motors has already employed ultrasonic sensors for its 'Autopark' and 'Summon' feature (self-driving with driver outside the vehicle) [8] and to monitor the blind spot at high speed [9]. Thus, it is critical to discover any vulnerabilities and implement remedies before billions of ultrasonic sensors are installed in vehicles and other IoT devices for various purposes.

Determining the right rules to make moral and ethical decision of self-driving cars is paramount, but it is far more complicated than one can imagine, because the driving decisions impact not only passengers' safety but also the safety of others. Thus, we focus on investigating the reliability of ultrasonic sensors and studying whether automobiles can make the right driving decisions that depend on these sensors. The goal is to design strategies that can avoid potential automobile collisions. Thus, sensors should detect all present obstacles and avoid false alarms, and vehicles should handle the following two scenarios correctly:

a) *Stop with obstacles:* a vehicle should stop moving towards obstacles on the driving path, and avoid active collision.
b) *Keep moving without obstacles:* a vehicle should keep moving when there is no obstacle on the driving path, and prevent passive collision with the unprepared traffic.

W. Xu, C. Yan, W. Jia and X. Ji are with the Department of Electrical Engineering, Zhejiang University, Hangzhou, China (emails: {xuwenyuan, yanchen, jwb8, xji}@zju.edu.cn). Dr. Xiaoyu Ji is the corresponding author.

J. Liu is with Qihoo 360 Inc., Beijing, China (email: liujianhao@360.cn).

We examine the feasibility of an attacker inducing incorrect driving decisions by exploiting sensor design vulnerabilities, and report vulnerabilities found on current sensor systems and real moving vehicles. To alleviate the threats, we design defense strategies that protect ultrasonic sensors against the attacks and facilitate reliable automated driving decisions.

A. Contributions

We summarize the attacks that lead to incorrect driving decisions in Table I, and list our contributions from security analysis and security enhancement aspects below.

Security Analysis. To analyze the security of ultrasonic sensors in automobiles, we perform black-box experiments, reverse engineer the sensors' printed circuit boards, and tap into the signal path. We verify the attacks on 11 stand-alone ultrasonic sensors in the laboratory and the on-board sensors of 7 vehicles (including a Tesla Model S) outdoors, using low-cost COTS hardware. To the best of our knowledge, we are the first to experimentally examine the feasibility of attacking ultrasonic sensors on real moving vehicles. We are able to induce the following incorrect driving decisions.

I. **A vehicle stops when it should keep moving.** We trick the on-board ultrasonic sensors to report non-existing obstacles anywhere within the detection range by designing two types of spoofing attacks—random spoofing and adaptive spoofing. We can force a moving Tesla in Summon mode to stop by creating imaginary obstacles.

II. **A vehicle keeps moving when it should stop.** We found that the design choices of ultrasonic sensors make it possible to hide obstacles. A jamming attack or an adaptive spoofing attack can both prevent sensors from reporting obstacles. Our experiments show that an attacker can cause a moving Tesla in hand-driving mode to collide into obstacles (e.g., students) from 10 meters way and cause collision in the Summon mode from 1 meter away.

Security Enhancement. Enhancing ultrasonic sensors is challenging because any defense strategy that requires a major modification of existing low-cost hardware will not be accepted by the automotive industry. The narrow operational frequency band and long physical delays of ultrasound make it difficult to utilize any traditional modulation-based schemes. We overcome all aforementioned challenges and design two security mechanisms: single-sensor based Physical Shift Authentication (PSA) and Multiple Sensor Consistency Check (MSCC). Both can be used alone or in combination.

I. **Physical Shift Authentication.** Despite the limitation of ultrasound, PSA allows a sensor to send random probing signals. Thus, it can detect obstacles reliably by checking whether the received echoes originate from the sensors.

II. **Multiple Sensor Consistency Check.** MSCC enables multiple sensors to collaboratively address more advanced attacks at a system level. It can detect spoofing attacks, measure distance resiliently, and localize obstacles (both real ones and attackers). Utilizing two *assistant sensors*, MSCC can achieve an improved detection rate.

We envision that the attack methodologies and enhancing technologies in this paper can provide insights for improving the security and reliability of autonomous vehicles, as well as ultrasonic sensors in other IoT applications, such as smart cities, smart home, medical diagnostics, SCADA platforms, and robot technologies [10], [11].

TABLE I: An overview of attack and defense goals.

Situation	Decision under Attack	Attacks
w/o obstacles	stop moving	random spoofing adaptive spoofing
w/ obstacles	keep moving	jamming adaptive spoofing

II. Background

In this section we briefly introduce state-of-the-art automated driving systems and the ultrasonic sensors.

A. Automated Driving System

An autonomous car (a.k.a., driverless car, self-driving car) is a vehicle that is capable of sensing its environment and navigating without human input. According to the SAE J3061 report [12], there are 6 levels of driving automation, from no automation (L0), to driver assistance (L1), partial automation (L2), conditional automation (L3), high automation (L4), and full automation (L5). Almost all cars that are at L2 and above are equipped with at least one type of active obstacle detection sensors, i.e., ultrasonic sensors, for parking assist. Tesla model S, as one of the most advanced autonomous cars in the market, is considered to be at L3 and has already implemented the 'Autopilot' system [3], which consists of functions like 'Autopark' and 'Autosteer' that monitor the surroundings and act accordingly. Research teams such as the ones at Google [13], Stanford [14], and Tesla Motors [15] are designing and experimenting fully autonomous prototype cars, but there is a long way to go before full application. Nevertheless, different levels of driving automation basically rely on the *same* types of sensor technologies, and the insights gained from analyzing sensors on Tesla can shed light on future automobiles.

B. Ultrasonic Sensors

Ultrasonic sensors were first introduced to automobiles as sensors of parking assistance systems in the early 1990s [16]. Ultrasonic sensors detect obstacles by transmitting and receiving ultrasound, which is one type of mechanical waves whose frequency is beyond the upper limit of human hearing (20 kHz). To measure the distance to an object, an ultrasonic sensor emits ultrasonic pulses (a.k.a. pings), and measures the time that it takes to receive the reflected pulses (a.k.a. echoes). Distance to the nearest obstacle is calculated based on the propagation time (time-of-flight, a.k.a. ToF) of the *first* received echo pulse according to the equation

$$d = 0.5 \cdot t_p \cdot v_s \tag{1}$$

where t_p is the propagation time of ultrasonic pulses, and v_s is the velocity of sound in air (343 m/s at 20 °C).

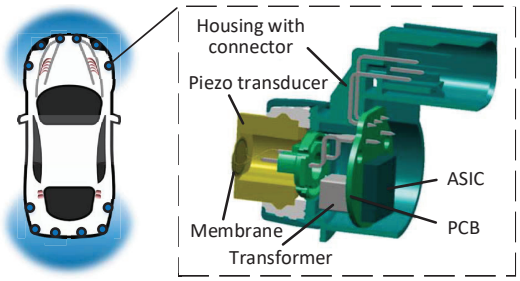

Fig. 1: The position and cross-section of ultrasonic sensors.

Fig. 1 shows an ultrasonic sensor consisting of a plastic housing, a piezoelectric transducer with an attached membrane, and a printed circuit board with the electronic circuitry and microcontroller to transmit, receive, and process the signals. We introduce the main sensing principles and our concerns as follows.

Piezoelectric Effect. Most ultrasonic sensors on automobiles utilize piezoelectric crystals as transducers [17], which can convert electric charges into mechanical vibrations and vice versa. For example, if a voltage is applied at the electrodes of a piezoelectric crystal, a mechanical deformation results and generates acoustic waves. On the contrary, an incoming acoustic wave creates oscillations of the crystal, which generate an alternating voltage at the electrodes. Note that it takes time for piezo transducers to emit stable mechanical vibrations, and we call this delay as *start-up time*.

Frequency. Ultrasonic sensors on vehicles typically operate within a frequency band between 40 and 50 kHz, which has been proved as the best trade-off between acoustical performance (sensitivity and range) and robustness against ambient noises. Frequencies higher than 50 kHz will lead to weaker echoes due to the attenuation of airborne sounds, whereas for frequencies lower than 40 kHz the proportion of interfering sound is larger [18]. Unlike speakers, ultrasonic sensors tend to work at their resonance frequencies and cannot efficiently transmit wide-band signals.

Distance Measurement. When a sensor receives a command from the electronic control unit (ECU) to transmit, its circuit excites the transducer with periodic waves at the resonance frequency for typically 300 μs, resulting in the membrane's vibration and emitting ultrasonic pings. Note that a transducer cannot listen while transmitting. Even after it stops transmitting, the sensor cannot receive echoes immediately until after a *ring-down time* (approx. 700 μs). Thus, ultrasonic sensors cannot detect objects in their close vicinity. Once rested, the membrane can be vibrated again by the echoes, which are converted to analog signals, then amplified, filtered, digitized, and compared to a threshold to determine the arrival of echoes.

III. Attack Overview

Before discussing the security vulnerabilities of ultrasonic sensors and their impact on automobiles, we specify our assumptions on the threat model, introduce the basic ideas of our attacks, and summarize the attack categories.

A. Threat Model

In this paper we focus on adversaries that attempt to attack a vehicle by falsifying the sensors' output *only* via the physical signal channels. We assume their capabilities as follows.

Sensor Assessment. We assume that an adversary is aware of the underlying principles of the sensor systems, and has budgets and access to obtain such sensors for assessment beforehand. The adversary can acquire the parameters of sensor designs, e.g., operational frequency, bandwidth, duty cycles, packet format, etc., and further explore sensor vulnerabilities. The adversary may be proficient with hardware design, and can exploit off-the-shelf hardware to accomplish the assessment.

Attack Scenario. The adversary can eavesdrop on the physical signals from on-board sensors, and actively generate forged echoes in an arbitrary form (frequency, amplitude, duration, phase, etc.), thereby corrupting or overpowering other concurrent physical signals in propagation.

Contactless. An adversary can be anywhere around the targeted vehicle and is free to move. However, she does *not* have control over the targeted vehicle or the on-board sensors. Moreover, the adversary must stay away from the targeted vehicle in order to remain stealthy during the attacks, and cannot make any physical alteration or damage to the sensors.

B. Physical Signal Level Attacks

In this work, we study *physical signal level attacks*, which take advantage of the physical sensing channels to disrupt or manipulate the sensor measurements.

Security Questions. Whether automobiles can make the right driving decisions in the two scenarios—with and without obstacles—depends on the reliability of sensors and the vehicles' pre-programmed logic to react to various situations. Thus, we would like to answer the following questions in the presence of physical signal level attacks.

- Will a sensor report the detection of obstacles when there is none?
- Will a sensor report no obstacle when one or multiple real obstacles exist?
- Will an automobile handle the output of sensors properly, especially abnormal sensory data?
- If sensors malfunction under attacks, what defense mechanisms can be adopted to cope with them?

Attack Basics. Since ultrasonic sensors emit ultrasounds to probe their surroundings, we utilize two types of well-known attacks as our building blocks to seek answers to the aforementioned questions: *spoofing*, whereby carefully crafted ultrasounds are injected so that they appear to come from non-existing sources and obfuscate real ones, and *jamming*, whereby noises are injected simply to interfere with sensors.

1) *Spoofing Attacks.* Spoofing attacks involve emitting carefully crafted signals (e.g., ultrasound pulses) that are identical to those transmitted by the sensors, i.e., with the same frequency, modulation, etc. As a result, the sensors may interpret the spoofing signals the same way as the authentic signals, and falsely detect non-existing obstacles. By carefully adjusting the timing of the

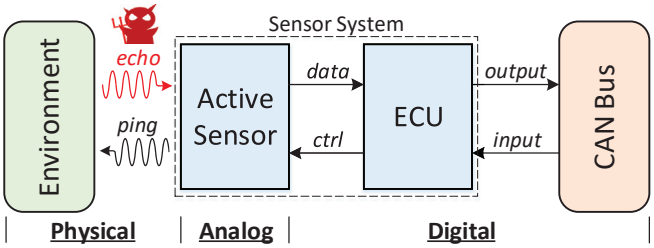

Fig. 2: A typical architecture of active sensor system on vehicles and the type of attack of interest in this paper. The 'echo' marked in red indicates being compromised by attack.

spoofing signals, an adversary may 'create' fake obstacles at various locations of her choice.

2) *Jamming Attacks.* Jamming attacks involve injecting similar but stronger signals to overpower the real ones. Sensors are typically designed to be robust against benign ambient noises, but they hardly expect strong interference. It is unclear whether sensors can detect objects in the presence of jamming attacks. In case the interference is so strong that it causes sensor denial-of-service, it is also unclear whether sensors and automobiles will fail gracefully and do not cause fatal accidents.

In the rest of the paper, we study both spoofing attacks and jamming attacks. Note that the longer the effective attack range is, the more practical the attacks will be. The effective range of the attacks relies on both the operational range of sensors and the transmission power of the attack equipment, which can be improved with budgets. The goal of this work is to validate the feasibility of the attacks, and we do not focus on intentionally maximizing the transmission power, thereby the reported attack range serves as a reference. In practice, a motivated attacker can increase the transmission power and boost the effective attack range.

C. Attack Categorization

Ultrasonic sensors are one type of active sensors that emit physical signals. To be general, we summarize attack classification in terms of active sensors. Fig. 2 illustrates the interaction of active sensors in an automobile. Inside an automobile, an electronic control unit (ECU) controls several active sensors, which emit probing pings to measure the environment and report distance to obstacles, if there are any, back to the ECU. Then, the ECU will transmit the measurement to other ECUs via the CAN bus to fulfill functions such as self-parking. We envision that as measurements are generated and transmitted from sensors to their ECU or other networked ECUs, three types of attacks that target at various levels of sensor systems are possible, which include the following.

1) Physical Signal Level Attacks: Physical level attacks take advantage of the physical sensing channels to disrupt or manipulate the analog sensor measurements. PS attacks only manipulate the surrounding physical environment or ambient signals, e.g., the echoes in Fig. 2, and do not affect the data processing path inside sensors.

TABLE II: Experimented stand-alone sensors and results.

Sensor	Frequency	Output under attacks		
		Spoofing		*Jamming*
		Random	Adaptive	
SRF01	42 kHz	Unsteady	Steady	Min
SRF05	40 kHz	Steady	Steady	Min
MB1200	42 kHz	Unsteady	Steady	Max
HC-SR04	40 kHz	Steady	Steady	Min
JSN-SR04T	40 kHz	Steady	Steady	Min
US-100	40 kHz	Steady	Steady	Min
RCW-0001	40 kHz	Steady	Steady	Min
URM04	40 kHz	Unsteady	Steady	Min
URM37	40 kHz	Steady	Steady	Min
Grove U. R.[1]	42 kHz	Steady	Steady	Min
Audi Q3	50 kHz	Unsteady	Steady	Max

[1] Grove Ultrasonic Ranger.

2) Sensor Hardware Level Attacks: Hardware level attacks manipulate sensor measurements by affecting how the sensory signals are collected and processed inside sensors. For example, Foo Kune et al. [19] demonstrate injecting voice signals into a Bluetooth headset and fake heart beats into an implantable pacemaker by blasting intentional EMI on the conducting wires inside sensors. Similarly, acoustic interference has been shown to be able to cause MEMS gyroscopes and accelerometer to malfunction [20], [21].

3) Digital Level Attacks: Digital level attacks are, by far, the most well-studied attacks, which include traditional cyber attacks that alter digital information or invade systems exploiting digital channels (e.g., network interfaces, file systems, memories). For instance, researchers have demonstrated attacking automobiles via cellular networks [22] and CAN bus [23].

Although active sensors can be vulnerable to all three types of attacks, we focus on physical signal level attacks because they are unique to active sensors and least studied.

IV. Attacking Ultrasonic Sensors

Analyzing the vulnerabilities of existing ultrasonic sensors begins with obtaining a thorough comprehension of their underlying principles. In particular, we investigate the frequency, period, and modulation schemes of the ultrasonic probing signals. Then, we design three types of attacks—random spoofing, adaptive spoofing, and jamming attacks—to understand whether sensors can detect obstacles reliably and whether automobiles will cope with abnormal situations properly. We validated all attacks on 11 models of stand-alone ultrasonic sensors (in Table II) in the laboratory. We also tested 7 models of vehicles (in Table III) outdoors.

A. Analyzing Sensors

To analyze the probing signals, we acquired 11 models of stand-alone sensors, and one of them is an OEM parking assistance system consisting of one ECU and four sensors, which is the same as the one on one of our tested vehicles. All of them report the distance to the closest obstacle if there is any within the detection range.

TABLE III: Experimented vehicles and results.

Model	Manufacturer	Sensor Number		Vulnerable to attacks[1]
		Front	Rear	
Q3	Audi	4	4	Yes
Model S	Tesla	6	6	Yes
Tiguan	Volkswagen	0	4	Yes
Polo	Volkswagen	0	4	Yes
Fiesta	Ford	0	3	Yes
Carnival	Kia	0	3	Yes
GLK 260	Mercedes-Benz	6	4	Yes

[1] All types of attacks include spoofing and jamming.

Methodology. From the public domain, we only learn that ultrasonic sensors operate in the range of 40 kHz to 50 kHz. To obtain details of the probing signals, we carried out two types of analysis. (a) Tapping into the signal pathway. The basic idea is as follows. Applying alternating voltages on piezoelectric crystals generates acoustic waves (i.e., mechanical oscillation), and the frequency and amplitude of the AC input signals determine the ones of the acoustic waves. Thus, analyzing the AC signals will reveal insights of the probing signals, and we manage to use oscilloscopes to intercept the periodic waves that drive the piezoelectric crystals. (b) Sampling over the air. Recording the emitted ultrasound directly will enable time and frequency domain analysis. However, generic microphones and ultrasonic sensors cannot record ultrasound that might spread tens of kHz. In the end, we used an off-the-shelf free field measurement microphone [24] (covering 4 Hz – 90 kHz) to sample the ultrasound, and it outputs electrical signals that can be fed into an oscilloscope, spectrum analyzer, or smartphone application for further analysis.

Probing Signals. Both methodologies reveal the same findings on sensors: The probing pings are of the form of square waves with a constant amplitude, frequency, and duration. Due to the short duration of each ping (e.g., 300 μs) and physical delays, ultrasonic sensors do not employ any modulation schemes.

Probing Periods. We discover that sensors emit probing signals periodically with two sensor algorithms (SA) illustrated in Fig. 3: (a) SA1, periodic according to the first echo and (b) SA2, periodic according to probing signals. Most of the tested stand-alone sensors work as SA1 and wait a fixed amount of time, T_1, either after the first echo or a timeout to transmit the next probes. All sensors on the tested automobiles are SA2 and transmit one ping every T_2, regardless of whether an echo occurs. This is because a vehicle ECU triggers multiple sensors in a predefined order.

B. Random Spoofing Attacks

Random spoofing attacks randomly replay the previously recorded sensor signals hopefully at the right timing to deceive sensors, and do not try to cancel the sensor signals received at this moment. Such attacks can at most delude the sensors to report a forged obstacle that is closer than any real obstacles. We formalize the signals received by a sensor under random spoofing attacks as

$$\Upsilon_i = \Psi_i(\tau_0) + \sum_{n=1}^{N} \Psi_j^*(\tau_n) \quad (2)$$

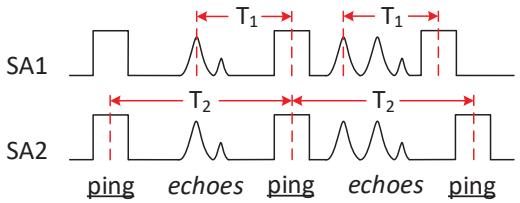

Fig. 3: Two probing algorithms: SA1 emits a probe after T_1 from the first echo. SA2 emits a probe periodically by T_2.

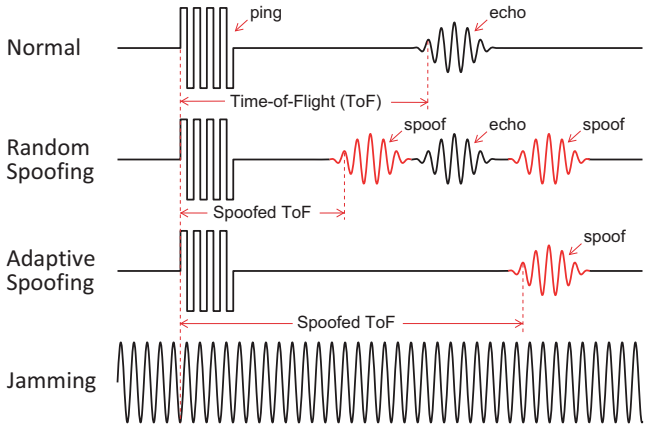

Fig. 4: An illustration of all ultrasonic attacks. The waveforms show signals a sensor receives in one sensing cycle, under no attacks, random spoofing, adaptive spoofing, and jamming attacks respectively.

where $\Psi_i(\tau)$ is the echo of the ith cycle received after time τ_0, $\Psi_j^*(\tau_n)$ is a spoof signal replayed based on a previous cycle j ($j < i$) and received after time τ_n, and N is the total number of spoof signals. An illustration of Υ_i is shown in Fig. 4.

Random Spoofing Timing. Since the nearest obstacle is the most important one to automobiles, only the *first* received echo signal is reported by sensors. If there is no obstacle nearby, the sensors will wait for a predefined timeout duration T_0, before starting the next cycle. T_0 is determined by the sensing range. Sensors only expect to receive echoes that are reflected from any obstacles within a sensing range. For a 2-meter range, T_0 is nearly 11.7 ms. Thus, for an effective attack, the spoof signals have to be received before the real echoes and within the timeout slot, whichever is sooner, i.e., $min(\{\tau_n\}_{n\in[1,N]}) < min(\tau_0, T_0)$ must be satisfied.

Building a Random Spoofer. To validate spoofing attacks, we acquired ultrasonic transducers whose working frequencies are the same as the ones of the target sensors. To drive the transducers, we utilize two types of off-the-shelf hardware: an Arduino [25] or a function waveform generator. Arduinos can output square waves of selected frequencies on the digital I/O pins using a built-in function named `Tone()`, which is mainly used to generate tones for speakers. Due to its low-cost nature, Arduino cannot generate a perfect periodic signal without any frequency jitters. Nevertheless, the generated square waves are sufficient for driving ultrasounds. In comparison, a function generator outputs signals with more stable frequencies and higher amplitudes.

Results. Random spoofing attacks can decrease the values

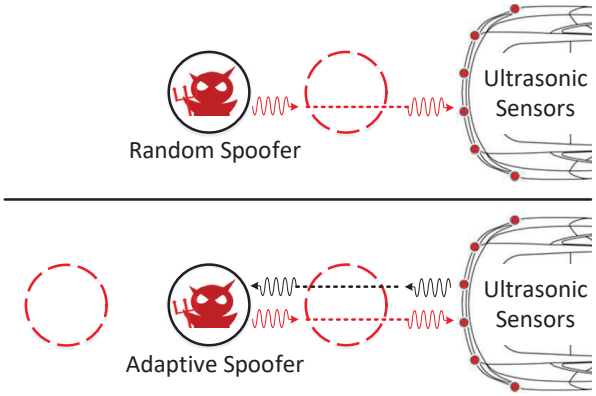

Fig. 5: Two types of spoofing attacks. The black circles are genuine obstacles and red dashed ones are fake obstacles that can be created by attackers.

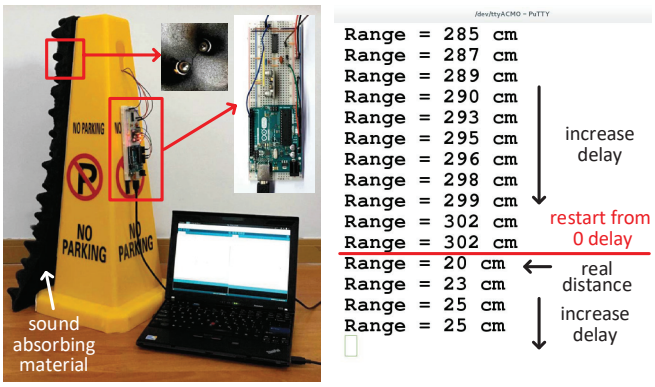

Adaptive Spoofer Spoofed Sensor Readings

Fig. 6: The disguised Arduino-based adaptive spoofer and incremental spoofing results on sensor SRF01. The adaptive spoofer mimics a moving-away obstacle by increasing the delay to transmit after receiving the sensor probe.

of sensor measurement. We validated the random spoofing attacks on stand-alone and on-board sensors. For both cases, we can create a forged obstacle that is closer than real obstacles, with slightly different observations.

1) Stand-Alone Sensors. For most stand-alone sensors that adopt SA1, by selecting a spoofing period that is larger than T_1 and smaller than $T_1 + T_0$, we were able to 'create' a non-existing obstacle that appears to be *stationary* on the sensor measurement. This is because the stand-alone sensor transmits the next probe after waiting for T_1 since the first received echo. If we emit a spoof signal at the period of $T_1 + 2d/v_s$, where v_s is the speed of sound, the sensors will output a constant measurement of distance d.

2) On-Board Sensors. By emitting spoof signals once every few milliseconds, we managed to make the automobiles report obstacles that did not exist, which kept causing alarms on all tested vehicles in Table III, and forced a moving Tesla Model S to stop in self-driving mode. The attacking distance was up to 2 meters when powered with an Arduino, and can be increased with higher transmission power. Since a vehicle ECU triggers each sensor periodically as SA2, random spoofing is unable to synchronize strictly with the pings, thus the imaginary obstacle appears to be *moving* on the sensor measurement. We observed the unsteadiness in various ways: an obstacle jumped back and forth between two locations; an obstacle appeared suddenly at random locations, as shown in Fig. 9(b). Nevertheless, none of the tested vehicles consider such abnormal results suspicious or warn drivers accordingly.

C. Adaptive Spoofing Attacks

A random spoofing attack cannot always precisely control the location of the spoofed obstacle and can at most create obstacles closer than the real ones. Thus, it can at most cause an automobile to stop unnecessarily. The goal of an adaptive spoofing attack, however, is to create a non-existing obstacle at an arbitrary yet stationary distance reliably, either closer (*subtractive*) or farther away (*incremental*) than the real one. Thus, it may cause an automobile to collide into a real obstacle, as illustrated in Fig. 5. To simplify the discussion, we assume that the adaptive spoofer is a real obstacle itself.

Adaptive Spoofing Timing. To forge a stationary obstacle at any distance d for even SA2 sensors, the adaptive spoofer has to transmit a spoof signal at the right timing, i.e., the sensor has to receive the spoof signal after a delay of $2d/v_s$ since the transmission of probing pings, where v_s is the speed of the sound. Thus, an adaptive spoofer has to listen and adjust to the concurrent sensor signals adaptively, and eliminate the existing echoes for incremental distance. As such, adaptive spoofing attacks will involve three phases: (a) receiving sensor signals, (b) eliminating the echoes, and (c) transmitting the spoof signals. Similarly, the signals received under adaptive spoofing are

$$\Upsilon_i = \Psi_i(\tau_0) - \alpha\Psi_i(\tau_0 + \delta) + \Psi_i^*(\tau_1) \tag{3}$$

where α is the index of signal cancellation and δ is the time delay introduced by real-time processing.

Since on-board sensors are typically triggered periodically, an adaptive spoofer can calculate and predict τ_1 by first measuring its distance to the sensor. Eliminating the echoes is not required when the imaginary obstacle is closer than the real obstacle, i.e., $\tau_1 < \tau_0$. However, in order for $\tau_1 > \tau_0$, echo elimination is necessary with proper α and δ. Acoustic quieting [26], [27] is a technique to cancel acoustic signals, and is well developed for stealth military submarines [28] as well as commercial noise cancelling headphones [29]. All these techniques can be utilized to eliminate echoes. In our implementation, we adopt a simple yet convenient method: we wrap the adaptive spoofer with sound absorbing materials (e.g., damping foams) and only expose the ultrasonic transducer in the air. The damping foams can passively absorb sound such that the reflected echoes are too weak to be detected by the sensors.

Building an Adaptive Spoofer. An adaptive spoofer has to be both a transmitter and receiver. We built an adaptive spoofer out of two ultrasonic transducers, amplification circuits, a buffer amplifier, an envelope detector, and an Arduino board. The adaptive spoofer listens to sensor signals, and controls the timing for emission. Meanwhile, we attached the adaptive

spoofer onto a traffic cone and wrap it with damping foams, as shown in Fig. 6.

Results. Adaptive spoofing attacks can decrease or increase the measured distances. We were able to create a *stationary* imaginary obstacle for both stand-alone and on-board sensors, and even manipulate its movement. The spoofer can be placed anywhere within the sensor's working range (normally a few meters). For demonstration, we place the adaptive spoofer 20 cm away from the sensor, and create an imaginary obstacle that is gradually moving away from the sensor. In particular, the spoofer will transmit a spoof signal with a delay of $2nTv_o/v_s$ $(n = 0, 1, 2, ...)$, where v_o is the speed of the imaginary obstacle, n is the echo sequence, and T is the sensing period. Fig. 6 shows the distance measured by a victim sensor under adaptive spoofing attacks—it illustrates the effectiveness of our attacks.

D. Jamming Attacks

Jamming attacks generate ultrasonic noises that induce continuous vibration on the sensor membrane, and render distance measurement impossible. The goal is to cause a sensor fail to detect real obstacles, which may cause collisions.

Jamming Parameters. A jamming attack continuously emits ultrasounds towards a sensor such that the jamming signals overwhelm the echoes, as shown in Fig. 4. The signals received under jamming attack are

$$\Upsilon_i = \Psi_i(\tau_0) + \int_0^{T_0} A\cos(\omega t) \tag{4}$$

where A is the jamming amplitude and ω is the frequency.

Resonant Frequency. From our measurement on several vehicles, we found that the operation frequency appears to be near 50 kHz. In practice we used off-the-shelf 40 kHz transducers for jamming because 50 kHz ones were unavailable. Since ultrasonic transducers operate around a narrow band, the 40 kHz transducer cannot emit 50 kHz ultrasounds efficiently. Nevertheless, 40 kHz turned out to be effective, and we believe the effective range could be expended with 50 kHz transducers.

Voltage Level. The amplitudes of sounds created by piezoelectric crystals rely on the voltage level of the signals that drive the crystals. Thus, the effective jamming distance is determined by the applied voltages. In our experiments, we use two types of equipment. Arduino can generate square waves with 5 volts maximum, and the function generator outputs up to 20 volts. The ultrasonic transducers that we obtained can take up to 70 volts, and we believe the effective attack range can go beyond what we observed.

Results. We have validated jamming attacks in the following three types of scenarios: (1) stand-alone ultrasonic sensors, (2) cars with parking assistance, and (3) a Tesla Model S with self parking and summon. In all experiments, a real obstacle exists, and it can be detected by the sensor when there is no attack.

1) Stand-Alone Sensors. Under jamming attacks, we observed two types of sensor outputs: *minimum distance*—the sensor reports an obstacle at its minimum detection distance (0 – 10 cm in our case), and *maximum distance*—the sensor detects no obstacle. We believe that two types of sensor design

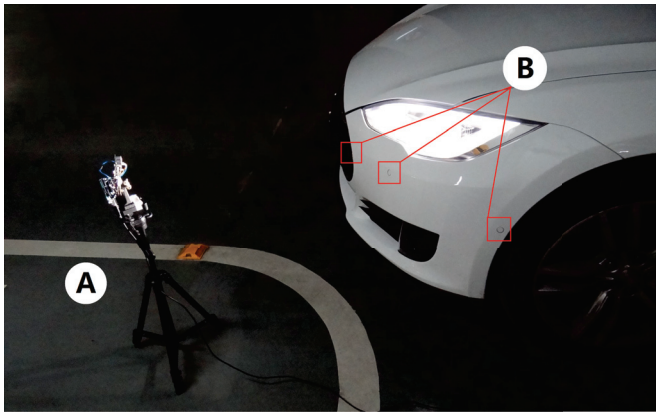

Fig. 7: Ultrasonic experiment setup on a Tesla Model S. A is the jammer, B is 3 sensors on the left-front bumper.

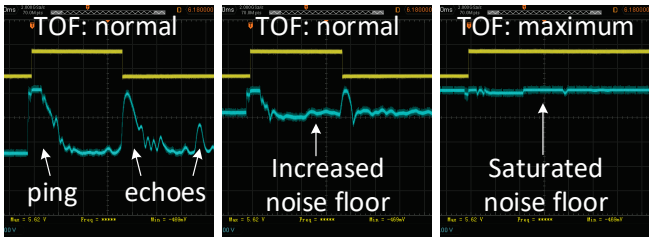

Fig. 8: Raw sensor signals showing noise suppression and the *Maximum Distance* result under strong jamming.

lead to these two results. For *minimum distance*, a sensor will consider the existence of an obstacle if the amplitude of received ultrasounds is larger than a predefined threshold. Under jamming attacks, once the sensor passes the ring-down period, it will receive the jamming signal and consider it as echoes from obstacles, resulting in *minimum distance*. For *maximum distance*, the sensors are designed to suppress ambient noises by adjusting their thresholds according to the noise level, as shown in Fig. 8. By tapping into the signal path of sensors, we realize that our jamming signal is recognized as noise, and to suppress the ambient noise, the sensor raises the threshold. As a result, the amplitude of the legitimate echoes is smaller than the threshold, thus the sensor cannot detect any echoes and reports maximum distance (i.e., no obstacle).

2) Vehicles with Parking Assistance. Next, we examined a few vehicles with driver assistance systems listed in Table III. The driver assistance systems on these cars all inform the driver about obstacles vocally or visually. As shown in Fig. 7, an ultrasonic jammer is placed in front of the car bumpers and can be detected. Once a jamming attack is launched, the vehicle can no longer detect the obstacle and no alarm is triggered (Fig. 9(c)). We believe that this maps to the *maximum distance* case and the design of these sensors aims at noise reduction. Using a function generator, we can effectively attack a moving Tesla from up to 10 meters away.

3) Tesla Model S with Self-Driving. We further tested jamming attacks on the *Autopark* and *Summon* features of Tesla Model S. We were wondering whether jamming attacks can prevent automatic parking systems from detecting the obstacles reliably. The results we observed turned out to be

(a) Real distance. (b) Spoofed distance. (c) Jammed distance.

Fig. 9: The dashboard on Tesla Model S showing distance to a nearby obstacle under (a) no attack (the distance is real), (b) spoofing attack (the distance is falsified), (c) jamming attack (the distance is maximum and not displayed).

prominent and worrisome.[1] When the Tesla is jammed in self-parking or summon mode, the car moving by itself will ignore obstacles and collide with them. The attack range is 1 meter when the jammer is driven by a function generator, and it could be increased excessively with power amplifiers.

E. Summary

In summary, we validated the following attacks.

- Random spoofing attacks can create imaginary obstacles that are closer to an ultrasonic sensor than real obstacles. Although the forged distance may change constantly and unrealistically, it can force a moving autonomous vehicle to stop when it should not.
- Adaptive spoofing attacks can create imaginary obstacles that are either closer or farther away than the real ones at predetermined distances. The attacks can force a moving vehicle to stop when it should not, or not to stop when it should.
- Jamming attacks can prevent ultrasonic sensors from detecting obstacles, and cause vehicle collisions.

V. Enhancing Ultrasonic Sensors

We design defense enhancement strategies that can cope with spoofing and jamming attacks against ultrasonic sensors, and aim at achieving the following levels of functions.

1) *Attack Detection.* At minimum, the defense strategies should be able to detect and report the attacks, thereby drivers or the self-driving systems can react to the attacks properly.
2) *Resilient Obstacle Detection.* Despite the spoofing attacks, the enhanced algorithm should be able to identify the spoofed echoes from real ones, and report the real distance.
3) *Attacker Localization.* The most challenging task is to localize the attackers. We believe the location information of an attacker can help a driver or the self-driving system to cope with the attacks, and can be used in forensics.

As illustrated in Fig. 2, an existing ultrasound-based obstacle detection system consists of a set of ultrasonic sensors (from 3 to more than 12) and an ECU. Thus, our enhancement includes two types of schemes:

[1]We demonstrate with a video on https://youtu.be/r4vS7YhT3DI.

TABLE IV: An overview of the defense scheme functions.

Attacks		PSA			MSCC		
		D [1]	R [1]	L [1]	D	R	L
Random Spoofing		√	√	×	√	√	√
Adaptive Spoofing [2]	–	√	√	×	√	√	√
	+	×	×	×	√	√	√
Jamming		√	√	×	×	×	×

[1] D: Attack Detection, R: Resilient Obstacle Detection, L: Attacker Localization.
[2] The '–' and '+' indicate subtractive and incremental attack results.

1) *Physical Shift Authentication (PSA)* based method that allows each individual sensor to detect attacks and to perform resilient obstacle detection.
2) *Multiple Sensor Consistency Check (MSCC)* based method that enables a set of ultrasonic sensors to collaboratively achieve the aforementioned functions.

We summarize the functions of each scheme against different attacks in Table IV. The schemes individually may not be able to protect the sensor system against all types of attacks, but their combination, as part of our systematic strategies, can enhance the reliability of the overall system, which we will discuss later.

A. Physical Shift Authentication

Physical Shift Authentication authenticates physical signals by shifting the waveform parameters. Essentially, physical signal level attacks are possible because ultrasonic sensors transmit pings of the *same* waveform throughout their lifetime and search for only the first echo via energy-based detection, i.e., detecting any ultrasonic signals whose amplitude is higher than a threshold. Thus, there is no bond between a ping and its echoes. To detect attacks and possibly reject spoofed echoes, it is important to bind them. As such, we propose a challenge-response scheme by customizing the ping waveform, then correlating the received echoes with the pings. A simple procedure of PSA can be summarized as:

Step 1. Randomize the ping waveform $\mathbf{X}$. Transmit.
Step 2. Receive. Measure the echo waveform $\mathbf{Y}$.
Step 3. If $\mathcal{C}(\mathbf{X}, \mathbf{Y}) < \alpha$, reject the echo.

where $\mathbf{X} = [x_1, x_2, ..., x_n]^T \in \mathbb{R}^n$ is a vector of n selected waveform features, $\mathbf{Y} = f(\mathbf{X}) = [y_1, y_2, ..., y_n]^T \in \mathbb{R}^n$ is a vector of n received waveform features, and $f \in \mathbb{R}^{n \times n}$ is a conversion function. $\mathcal{C}$ is a function that determines the correlation of the waveforms, and α is a threshold. Our hypothesis is that the waveform of real echoes should be correlative to the prior ping, however, a waveform $\mathbf{Y_A}$ from a passive attacker who does not know $\mathbf{X}$ will likely fail the challenge. Therefore, our approach in designing PSA is twofold: 1) examining the feasibility of physical signal authentication on ultrasound, and 2) detecting spoofed echoes and real ones.

A unique challenge we face, is that ultrasonic sensors receive real echo signals with *unknown parameters*, even when the transmitted signals are well-known and tuned. For example, the probing signal is typically

$$s(t) = \cos(\omega_c t), \quad t \in [0, T] \tag{5}$$

where ω_c is the carrier frequency, and T is the time duration. After reflection, the received signal becomes

$$r(t) = a\cos((\omega_c+\omega_D)(t-\tau)+\theta)+n(t), \quad t \in [\tau, \tau+T] \quad (6)$$

where a represents signal attenuation, ω_D is the Doppler shift, θ is a phase shift, $n(t)$ is the additive noise component, and τ is the time delay (proportional to the round-trip distance to the obstacle), all of which are unknown and dependent on the environment. There are sonar applications where the signal frequency, amplitude, phase, and the analog waveform are measured or estimated for advanced target measurement and identification [30]. However, it remains unknown whether these physical parameters can be used for signal authentication, especially when they include uncertainty after reflection.

To examine the feasibility of modulating these parameters, we choose amplitude, frequency, phase, and ping duration as candidates for the waveform feature x_k. Formally, we consider a sequence of customized pings, and let the waveform of the ith ultrasonic ping be

$$s_i(t) = A_i\cos(2\pi f_i t + \varphi_i), \quad t \in [\sum_{j=1}^{i}\Delta_j, \sum_{j=1}^{i}\Delta_j + T_i] \quad (7)$$

where A_i is the amplitude, f_i is the frequency, φ_i is the phase, and T_i is the duration of the ith ping. Δ_i is the period, i.e., the time period between the ith and $(i-1)$th ping (let $\Delta_1 = 0$). On existing sensors, all these waveform features are always constants, and pings are transmitted periodically with a fixed period Δ. In our experiments, we change each feature x_k independently, and seek for correlation in the reflected echoes.

Ultrasonic transducers inherently have physical delays as *start-up* and *ring-down* period [31]. Even if the signals that drive the sensors have constant amplitude, frequency and phase, during these special periods, the amplitude and frequency of the emitted ultrasound are unstable. Given that the duration of each ping is short and typically lasts for 8 to 20 cycles of sinusoid, it is impossible to modulate all these parameters efficiently within each ping. Thus, we prefer to modulate the waveform as *per ping* instead of *within each ping*, i.e., each ping will have a constant amplitude, frequency, duration and at most one phase shift, but they may be different between consecutive pings.

We define the process of transmitting the ith ping and receiving the following echoes as the ith **cycle**. We envision that the authentication of received waveform feature y_k can be done in two ways depending on the type of transmitted waveform feature x_k, as

a) **Per single cycle**, where $y_k(i)$ is received if and only if $x_k(i)$ has been transmitted, i.e., $y_k(i) \Leftrightarrow x_k(i)$.
b) **Per consecutive cycles**, where receiving a sequence of y_k implies that a sequence of x_k has been transmitted, i.e., $\{y_k(i), y_k(i+1), ..., y_k(i+m)\} \Rightarrow \{x_k(i), x_k(i+1), ..., x_k(i+m)\}$

The first case corresponds to those x_k that do not change dramatically during one cycle, possibly frequency and phase. The second case can be applied when x_k is highly dependent on the environment and reflecting surface, possibly amplitude and ping duration. Our experiments focus on examining the feasibility of employing the proposed x_k candidates for the above cases, i.e., whether they can be used for PSA and how to use them.

Although ideas similar to PSA have been tested on RF signals, its feasibility on ultrasonic sensors in automotive applications still remains unanswered. We ask the following questions: (a) Can we reliably modify the amplitude A_i, frequency f_i, phase φ_i, duration T_i, and period Δ_i of each ping? (b) Once the waveform of a ping is modified, will the corresponding echo change proportionally to the ping? (c) How reliably can we differentiate the modulated echoes with the spoofed echoes from a passive attacker?

To answer these questions, we use the following experiment setup. We set two ultrasonic transducers—one transmitter driven by a signal generator and one receiver connected to an oscilloscope—side by side toward an obstacle close by. The reflected signals can be observed and measured on the oscilloscope after amplification. We analyze the feasibility of each x_k candidate in the next few sections.

1) Frequency Shift: We ask the following questions for frequency shift: Can sensors create ultrasounds at various frequencies? Will reflection on an obstacle modify the frequency randomly?

Frequency Range. Since ultrasonic sensors can only emit ultrasounds in a narrow frequency band centered at their resonant frequencies determined by the diameters of the piezoceramics, we first measure the frequency response of a transducer. In this experiment, we place a wide-band microphone [24] 10 cm away from an ultrasonic transducer as we sweep the frequency of the stimulation signals from 35 kHz to 45 kHz. We plot Sound Pressure Level (SPL) of the received signals in Fig. 10(a). As the frequency of the stimulation signals deviate from the resonant frequency (40 ± 1 kHz), the SPL of the received ultrasounds reduces. To ensure the detection range of an ultrasonic sensor, it is reasonable to choose frequencies from 38.5 kHz to 41.5 kHz.

Obstacle Reflection. The start-up time varies for different transducers, but it is generally larger than the typical ping duration (8 to 20 cycles of sinusoid). We set the duration of each ping to 100 cycles (2.5 ms at 40 kHz) so that the frequency is stable, and unaffected by the start-up/ring-down period. Fig. 10(b) shows that when the obstacle is stationary, the frequencies of echoes are close to the ones of the stimulation signals, with a maximum frequency deviation of 0.212 kHz in the band from 39 kHz to 41 kHz. When the vehicle or obstacles are moving, the received echoes can be Doppler shifted. For example, a relative speed of 15 km/h can lead to a Doppler shift of 0.48 kHz. However, it can be solved by compensating the Doppler shift as a constant bias, due to the fact that it changes very little between two cycles (20 ms).

We envision that f_i can be shifted randomly between different pings, and the authentication can be done per ping, by checking if $y(i) = x(i)$.

Anti-Jamming. A similar idea of frequency hopping can be used to resist signal jamming, where the ultrasonic frequency is changed frequently according to a predetermined schedule. When the jammer signal strength isn't so strong that it saturates the receiver front end, advanced DSP techniques can

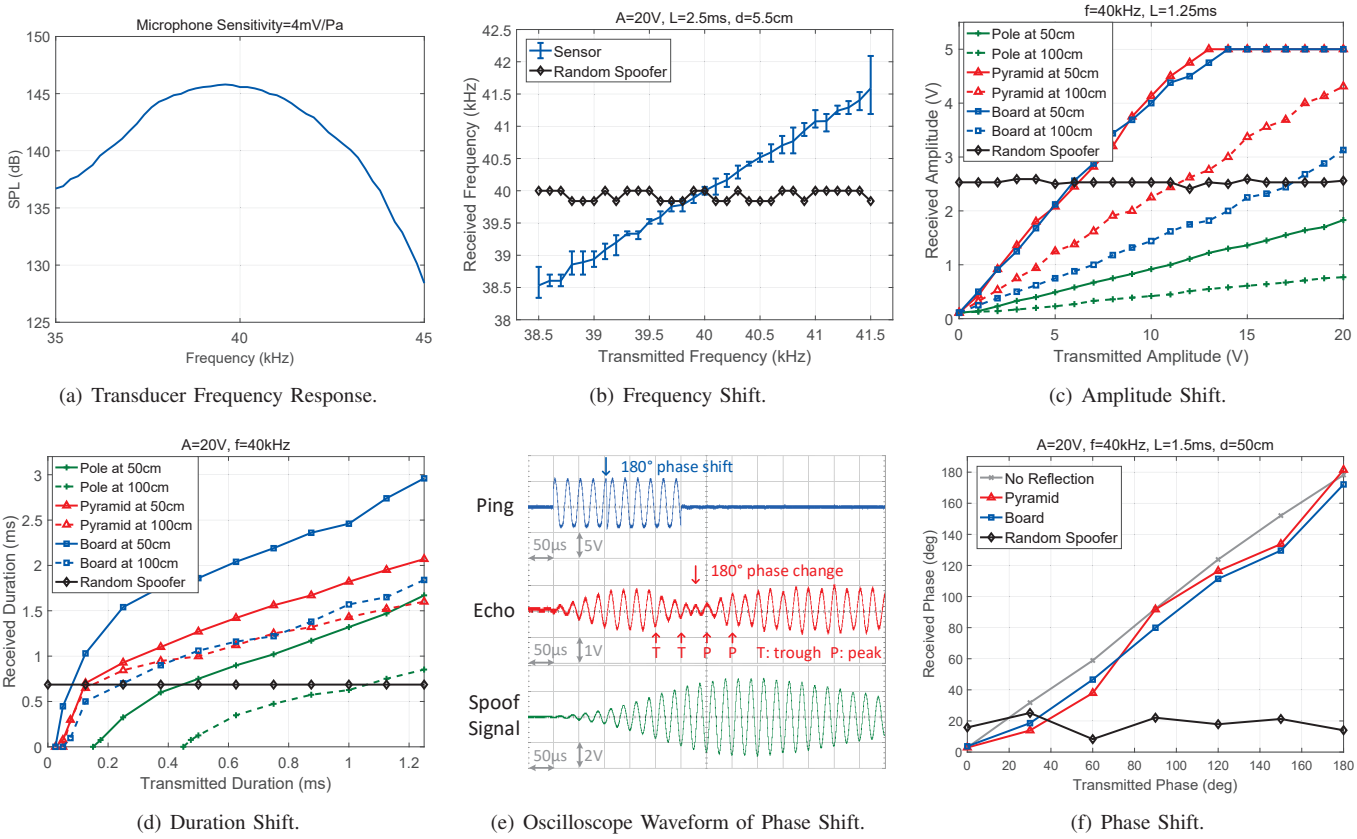

(a) Transducer Frequency Response. (b) Frequency Shift. (c) Amplitude Shift.

(d) Duration Shift. (e) Oscilloscope Waveform of Phase Shift. (f) Phase Shift.

Fig. 10: The experimental results of physical authentication schemes.

be used to estimate and cancel the effects of the jamming signal.

2) Amplitude Shift: The strength of received echoes depends on the transmission power, distance to, and surface of the obstacle. We place an obstacle 50 and 100 cm away from the emitter and the receiver, and experiment with three obstacles that are common on the road: a plastic bollard, a plastic pyramid cone floor sign, and a metal board (to mimic a car body), which are referred to as the *pole*, the *pyramid*, and the *board* in Fig. 10(c).

As shown in Fig. 10(c), the amplitudes of received signals increase with the transmitted signals (strong signals are clipped by the amplifiers at 5V), and different obstacles at different distances show various attenuation. From our observation, the conversion function of amplitude can be approximated as $y(i) = f(x(i)) = \lambda x(i)$, where λ is a constant for a certain stationary obstacle. For a moving obstacle, we assume λ remains constant in one cycle period (20 ms). Although it is difficult to verify echoes in per single cycle, we can transmit consecutive pings (e.g., two) with different amplitudes and measure their echoes. If the amplitude attenuation of all pings exhibits the same linearity, then all these echoes should be valid. In comparison, signals from a random spoofer will not obey this rule.

3) Duration Shift: We change the duration of each ping, i.e., T_i. Since ultrasound tends to attenuate quickly over the air, the minimum duration of each ping has to be long enough so that the corresponding echoes can be detected, e.g., more than 20 cycles (0.5 ms at 40 kHz) for the cases shown in Fig. 10(d). As we increase the duration beyond its minimum limit, the duration of received echoes are approximately linear to those of the corresponding pings. Empirically, it can be modeled as $y(i) = f(x(i)) = \lambda x(i) + \upsilon$, where λ and υ are constants for a certain obstacle. Similar to amplitude shift, it is reasonable to emit consecutive pings (e.g., three) with various duration and examine the consistency of linearity.

4) Phase Shift: Because ultrasounds are essentially mechanical vibration, we cannot change their phases immediately, i.e., not in the same way as how radio signals change. Physical laws require time to dampen a vibrating membrane and then drive it into the new phase. To understand this, consider a phase shift (e.g., 180°) as a sudden shift of force direction on the vibrating membrane of a transducer. When the force is shifted to a reverse direction, the membrane gradually decreases the amplitude and then increases following the new phase. Fig. 10(e) illustrates this phenomenon.

To validate whether we can perform an arbitrary phase shift, we performed the following experiments. We drive the ultrasonic transmitter with a vector signal generator [32] capable of phase modulation, and measure the phase shift of the amplified echoes on an oscilloscope. Suppose the phases before and after phase shift are φ_A and φ_B. To obtain the phase shift $|\varphi_A - \varphi_B|$, we introduce a reference signal with phase φ_R on another channel of the oscilloscope, measure two phase differences $\varphi_A - \varphi_R$ and $\varphi_B - \varphi_R$, and subtract the results. As shown in Fig. 10(f), the received phase shifts are

TABLE V: Comparison of different waveform parameters.

$\mathbf{x_k}$	N	PSA Features	Real	Attack
f	1	$\lvert y(i)-x(i)\rvert$	$\leq \varepsilon_f$	$> \varepsilon_f$
A	2	$\lvert y(i+1)-\frac{y(i)}{x(i)}x(i+1)\rvert$	$\leq \varepsilon_A$	$> \varepsilon_A$
T	3	$\lvert y(i+2)-y(i+1)-\lambda[x(i+2)-x(i+1)]\rvert$	$\leq \varepsilon_T$	$> \varepsilon_T$
φ	1	$\lvert y(i)-x(i)\rvert$	$\leq \varepsilon_\varphi$	$> \varepsilon_\varphi$

[1] $\lambda = \frac{y(i+1)-y(i)}{x(i+1)-x(i)}$

close to the modulated phase shifts, i.e., $y(i) \approx x(i)$. Similar to frequency shift, phase shift authentication can be achieved per single cycle.

5) Period Shift and Speed Filter: Period shift is designed as a supplement to the above schemes. Its motivation is to identify attacks by the instability of spoofing results. It requires the sensor to repeatedly probe for two or more consecutive cycles, and measure the difference in distance estimation between cycles. Intuitively, if a sensor and obstacle are relatively still, the distance measurement in different cycles should be almost the same. If either party moves, the relative speed can be calculated from the displacement per period. By shifting the probing period, we can increase the time jitter caused by spoof signals from random or subtractive adaptive spoofers, and filter out obstacles that appear to move at unrealistic speed. For example, consider a sensor that probes every 20 ms and covers a range of 3 meters. A spoofer that induces 1 ms jitter on the ToF will cause a distance offset of roughly 17 cm. However, moving 17 cm in 20 ms indicates a relative speed of 8.5 m/s (30.6 km/h), which is unlikely and should be rejected.

6) Performance Evaluation: Above experiments validate that shifting frequency, amplitude, duration, phase, and period can be used for distinguishing spoofed echoes that does not change in accordance with the modulated pings. To evaluate and compare the performance of each scheme, we propose a basic prototype that employs the features in Table V to validate echoes in N cycles. ε_k is the detection threshold for a specific feature x_k. Generally, we perform binary detection by considering feature values below ε_k as real echoes and those above as from attackers.

We calculate feature values with selected raw data[2] in Fig. 10, and plot the cumulative distribution function for each scheme in Fig. 11. In our case, with a confidence interval of 95% (True Negative Rate), the thresholds are: $\varepsilon_f = 0.18$, $\varepsilon_A = 0.25$, $\varepsilon_T = 0.19$, $\varepsilon_\varphi = 20.42$. Attack detection rates under these thresholds are: $\alpha_f = 87.1\%$, $\alpha_A = 90.0\%$, $\alpha_T = 87.3\%$, $\alpha_\varphi = 77.3\%$.

To better understand the performance, we plot the ROC curves in Fig. 12, and require the authentication to be finished in no more than three cycles. Under the same time efficiency, the performance ranks as: *Frequency Shift* > *Phase Shift* > *Amplitude Shift* > *Duration Shift*. Especially, two cycles of frequency shift will outperform two cycles (minimum requirement) of amplitude shift, with a detection rate of 98.35% under 5% false positive rate. The detection rate can be increased to 99.8% with three cycles of frequency shift, and even further with wider transducer frequency range. In addition,

[2] 38.6 kHz $\leq f \leq$ 41.4 kHz; $A_T \geq 5$ V, $A_R < 5$ V; $T \geq 0.5$ ms.

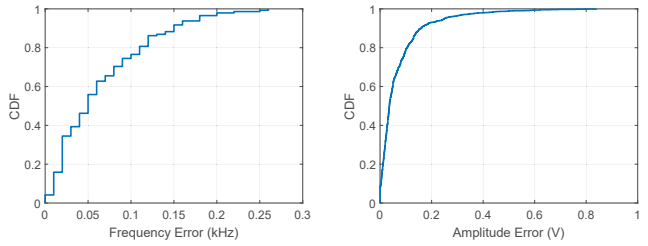

(a) Frequency Shift. (b) Amplitude Shift.

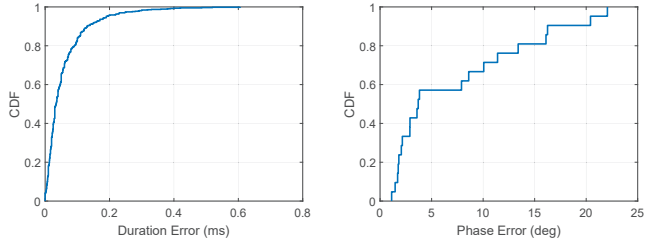

(c) Duration Shift. (d) Phase Shift.

Fig. 11: Cumulative Distribution Function (CDF) of PSA feature errors.

a sensor may employ multiple schemes interchangeably for higher security.

7) Likelihood Ratio Detector: The above basic prototype is based on N cycles of measurements, and it can be improved by considering more past measurements. Given the received signal, R, and the ultrasonic sensor, S, the task of attack detection is to determine whether R originates from S. We propose one general approach with likelihood ratio test, which exploits prior knowledge of statistical descriptions of data to choose amongst a candidate set of populations. The test is based on likelihood ratio, which expresses how many times more likely the data are under one model than the other. We restate the attack detection task as a basic hypothesis test between

$$H_0 : R \text{ originates from the sensor } S.$$

and $$H_1 : R \text{ does not originate from the sensor } S.$$

The optimum test to decide between these two hypotheses is a likelihood ratio test given by

$$\Lambda = \frac{p(R|H_0)}{p(R|H_1)} \begin{cases} \geq c & \text{accept } H_0 \\ < c & \text{reject } H_0 \end{cases} \tag{8}$$

where $p(R|H_i), i = 0, 1$ is the conditional probability density function of the observed signal R when the hypothesis H_i is true. c is the decision threshold for accepting or rejecting H_0. The task can be solved with a right probability model of feature value distribution, and we will not discuss here.

B. Multiple Sensor Consistency Check

We present an ECU algorithm that can achieve resilient obstacle detection and attacker localization based on a single-transmitter multiple-receiver sensor structure, where at a time only one sensor transmits and multiple sensors (three or more) receive, and check the consistency of measurement. Since

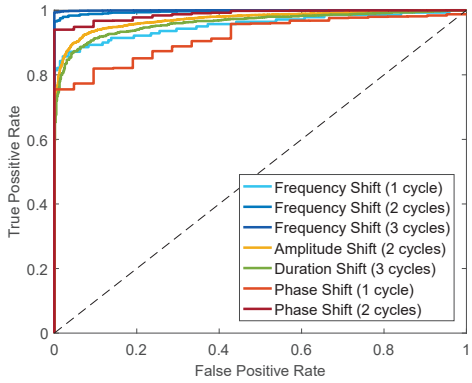

Fig. 12: Receiver operating characteristic (ROC) curves for different PSA parameters in no more than 3 cycles.

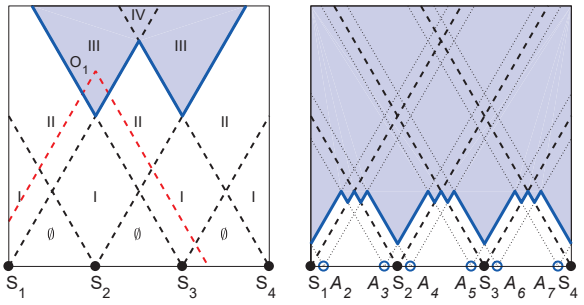

(a) Regular MSCC. (b) Enhanced MSCC.

Fig. 13: Simulation layout for two MSCC schemes. The detection area (shaded) is enlarged with enhanced MSCC.

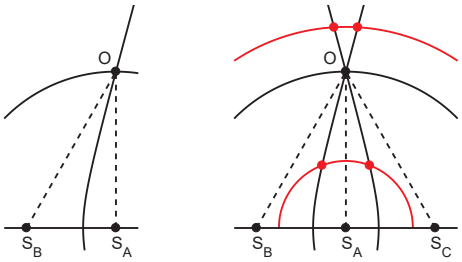

(a) Localization. (b) Resilient Detection.

Fig. 14: Illustrations of MSCC scheme principles.

modern vehicles are typically equipped with multiple sensors, they can support MSCC-based defense strategy. The basic idea of the scheme is to utilize the redundancy information from different sensor positions to detect the inconsistency caused by the spoofer. In particular, an ultrasonic transducer emits ultrasound beams in the shape of a torchlight, which exhibits an emission angle (e.g., 60° for [33]). Thus, as a spoofer emits a forged echo, multiple sensors may receive it, and combine their measurement to check consistency, thereby reject the spoofed echoes and accept the real ones.

In summary, depending on the number of sensors that can overhear the echoes, the ability of this algorithm is as follows:

- *Two sensors.* MSCC cannot detect attacks, but can localize obstacles no matter real or fake.
- *Three or more sensors.* MSCC can detect attacks, estimate the distance to obstacles resiliently under random and adaptive spoofing attacks, and localize the attacker.

1) Regular MSCC and Enhanced MSCC: Since the detection rate of MSCC depends on the available number of sensors, we design two sensor layouts to support MSCC algorithms: a regular MSCC and an enhanced MSCC. A regular MSCC contains N sensors that are evenly distributed on the front and rear of an automobile, e.g., a 200 cm line. The enhanced MSCC adds two more *Assistant Sensors* (AS) beside each of the N sensors. As shown in Fig. 13(b), A_3 and A_4 are placed 5 cm away from the original sensor S_2. Assistant sensors do not transmit, instead, they receive echoes and detect attacks based on the MSCC method.

2) Localizing Obstacles: The underlying principle of MSCC based scheme is to localize obstacles with a pair of sensors, as illustrated in Fig. 14(a). In one sensing cycle, only one sensor (e.g., S_A) transmits, while both sensors receive and measure ToFs. S_A's output t_A is the ToF from S_A to and from the obstacle, while S_B's output t_B is the ToF from S_A to the obstacle and then to S_B. Consider the case that the obstacle is farther to S_B than S_A by $\Delta D = (t_B - t_A) \cdot v_s$, where v_s is the speed of sound. According to the measurement of S_A, the possible location of the obstacle is a circle of a radius $t_A \cdot v_s/2$ centered at S_A. According to the difference between t_B and t_A, the possible location is on a branch of hyperbola with foci S_A, S_B, and vertex distance $2a = \Delta D$. Thus, the true location of the obstacle is the intersection of the two curves.

3) Resilient Obstacle Detection: Let two sensors each transmits for one cycle, we can obtain two location estimations of an obstacle. If the obstacle is real, the two estimation should be almost the same. In the presence of a random spoofer or a subtractive adaptive spoofer who do not adjust to the timing of probing signals, the difference between two estimations will be distinct. However, an incremental adaptive spoofing attacker adds a constant delay as soon as receiving the probing signals, and will create the same estimations in each cycle. Thus, two sensors are insufficient to detect incremental adaptive spoofing attacks. MSCC scheme aims to solve this problem by introducing a third sensor S_C in Fig. 14(b).

Consider a MSCC structure with three sensors. Since one sensor pair will provide one location, two pairs will provide two. When sensor S_A transmits, all three sensors receive and output ToFs: t_A, t_B, and t_C. Likewise, given t_A and t_C, the potential location of the obstacle is on the hyperbola with foci S_A, S_C, and vertex distance $2a = (t_C - t_A) \cdot v_s$. Now the circle centered at S_A with $r = t_A \cdot v_s/2$ will intersect with two hyperbolas at two points, i.e., two locations.

We then check the obstacle locations. If echoes are reflected from a real obstacle, the radius $t_A \cdot v_s/2$ will be the real obstacle distance d_{AO}. Thus, the two locations will coincide at obstacle O. However, if O is a spoofer emitting ultrasound actively, the two locations will be distinct as t_A is spoofed (as the red circles show). Given the ToFs from a basic triple-sensor MSCC structure, any attempt to spoof the sensors translates into the inconsistency of localizations in one sensing cycle. By filtering out the inconsistent obstacles, we can achieve resilient obstacle detection.

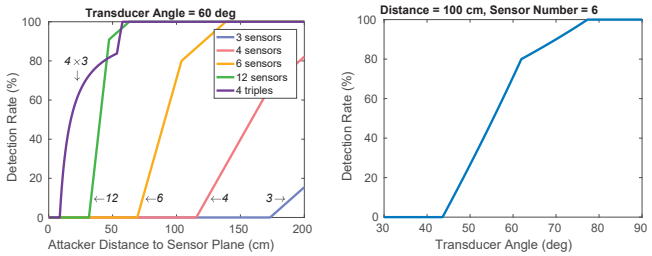

(a) Sensor Number and Structure. (b) Transducer Angle.

Fig. 15: Simulation results showing that detection rate can be increased with enhanced MSCC, more sensors and larger transducer angle.

4) Localizing Attacker: Notice that the spoof signals are emitted at the *same* time from the attacker but received by multiple sensors at different time. The ToF is determined by the geographical layout of sensors. Thus, for a stationary emitter (i.e., the spoofer), the two hyperbolas in Fig. 14(b) is fixed, and their intersection O is the location of the spoofer no matter what the spoofed t_A is. This method is also known as multilateration in radio navigation systems.

5) Evaluation: To validate the effectiveness of MSCC method, we implemented the algorithm in MATLAB, and studied both the regular MSCC and enhanced MSCC designs. We randomly set an attacker in a 200×200 cm region ahead of the sensors, as shown in Fig. 13. We assume the attacker is aimed perpendicularly to the sensor plane, and has the same transducer angle as the sensors for performance trade-off. The dashed lines indicate the limits of sensor transmission/receiving, and divide the region into areas by the number of sensor overlap. Since the attacker can only be detected and localized with three or more sensors, our scheme will detect attackers above the blue margin. The enhanced MSCC design enlarges the detection area (shaded blue) by 3 times. Moreover, as shown in Fig. 15(a), the enhanced MSCC (the 4 triples) dramatically improves the detection rate when the attacker is close, comparing with the same number (12) of evenly distributed sensors. In addition, Fig. 15(a) and 15(b) together show that the detection rate can be raised by enlarging the detection area, through either increasing the sensor number or transducer angle.

C. Systematic Strategies

Though we propose PSA and MSCC as effective mitigation methods to the attacks in this paper, systematic design strategies are necessary for enhancing the reliability of the overall sensing systems against future threats. We envision that the security principle in designing ultrasonic sensor systems should adopt the following two complimentary aspects simultaneously.

- *Individual sensors* should provide reliable measurements.
- *Multiple sensors* should collaborate to achieve reliability which is otherwise impossible for individual sensors.

1) Securing Individual Sensors: For individual sensors, each of their measurements can be enhanced independently or collectively. PSA is an effective method to secure each measurement independently, while Kalman Filters can be used to improve the reliability of a sequence of measurement. A Kalman Filter [34] fuses data measured in successive time intervals to provide a maximum likelihood estimation of a parameter, e.g., distance, and it is widely used to track obstacles by filtering and prediction. Since most attacks introduce abrupt changes to the sensor measurements, a Kalman Filter could reduce the impact of transient attacks and detect attacks by setting a threshold on the error correlation matrix.

2) Multi-Sensor Data Fusion: Three types of sensor fusion are applicable to ultrasonic systems: complementary, competitive, and cooperative. *Complementary fusion* refers to the configuration that sensors do not directly depend on each other but can be combined to provide a more complete image of observation. It is the current configuration on most vehicles—multiple sensors are installed around a vehicle in order to detect obstacles from different directions. *Competitive fusion* involves adding redundant sensors for the measurement of the same obstacle. Since it is difficult for an adversary to jam or spoof multiple sensors at the same time, a voting scheme or confidence tags [35] can be used to indicate the trustworthiness of an observation. *Cooperative fusion* derives information that is unavailable from individual sensors. MFCC is such a scheme that can estimate two-dimensional locations and achieve resilient obstacle detection. Measurements from ultrasonic sensors can also be fused at a higher level with other types of sensors, e.g., cameras and Lidars, if there are any.

VI. Discussion

Overhead of Defense. PSA requires to authenticate echoes by waveform processing, which is a mature technology in ultrasonic level measurement [31]. MSCC requires multiple sensors to perform trilateration, which is a software based solution and can be employed on vehicles to improve distance measurement [18].

Defense Awareness. Even if the attacker is aware of the defense mechanisms, she cannot predict the randomness of PSA, or hide her geographical properties imposed by physical laws from MSCC.

Dealing with Other Attacks. Other attacks that might bypass our defense may be available, e.g., wide-band jamming attack. However, we aim to raise the bar of attacks and let automobiles make the best driving decision, e.g., a vehicle could stop when it detects jamming attacks. In addition, we encourage integrating information from different types of sensors to improve the resilience against more attacks.

Dealing with Multiple Attackers. An attacker may want to disrupt the defense mechanisms with multiple coordinated transducers. However, PSA will be able to detect each attacker, and MSCC can detect subtractive adaptive spoofing attacks in the presence of multiple colluding attackers. Again, our intention is to raise the bar of attacks, instead of building a bulletproof system, and we invite researchers to improve the reliability of sensors.

VII. Related Work

Previous work on vehicle security mostly focused on the digital communications inside and outside an automobile.

The infrastructure of modern vehicles is designed in such a way that all components are networked with each other by the CAN-bus to exchange information. This structure facilitates the functionality and efficiency of modern vehicles, but poses a serious threat in addition to potential insecure components [36], [37]. Several studies [38], [39] have shown the feasibility of launching CAN-bus attacks, mainly through OBD-II port, to cause malfunction and even take control of the car. In addition, it is possible to launch attacks [23], [22] remotely, if it contains external attack surfaces [40].

Several studies have examined the security of passive sensors, e.g. microphones and medical devices [19], MEMS gyroscopes [20], and MEMS accelerometers [21]. As for active sensors, Petit *et al.* [41] and Shin *et al.* [42] examined the security of LiDAR. The security of FMCW radars has been studied in [43] without experiment on real cars. Shoukry *et al.* [44] proposed a physical challenge-response authentication scheme for magnetic encoders and RFID tags. We argue that their methods may not apply well to ultrasonic sensors due to the physical latency of acoustic vibration and piezoelectric transducers, i.e., acoustic waves cannot change fast enough to satisfy the scheme requirements as RF signals. Despite the valuable insights gained from previous studies, different types of sensors adopt distinct underlying physical principles, therefore pose unique security challenges. We generalize our methods as *Physical Signal Level Attacks*, validate this approach on real vehicles with autonomous driving system, and propose enhancement strategies. Some topics of this paper have been presented by the authors previously in a non-academic talk [45].

The differentiation of ultrasonic probing signals has been studied mainly for solving inter-sensor interference, by means of TDMA scheduling [46], [47] or multi-code modulations [48], [49]. The structure of multiple ultrasonic sensors has been used for identifying multiple objects [50]. Spatial-temporal phase dynamics has been utilized for RFID-based relative object localization [51]. These schemes are mainly developed for applications (e.g., target tracking and RFID) rather than automobiles, and do not consider malicious attacks on vehicular sensors.

VIII. Conclusion

The reliability of ultrasonic sensors remains a critical question, especially when it shapes the safety of autonomous vehicles. Our work validated the feasibility of three types of attacks—random spoofing, adaptive spoofing, and jamming attacks—on ultrasonic sensors, and show that they can cause incorrect driving decisions on moving vehicles in autonomous driving. To enhance the reliability of ultrasonic sensors, we proposed two types of defense strategies, single-sensor based Physical Shift Authentication (PSA) and Multi-Sensor based Consistency Check (MSCC).

Appendix A
Responsible Disclosure

We have informed Tesla Motors Inc. in January and June 2016 about the vulnerabilities reported in this paper, and discussed improvement with the product security team in July 2016. They have acknowledged our findings, and are in the process of improving these sensors and systems. A Tier-1 sensor supplier, Bosch, has also been informed.

Acknowledgment

This work has been funded in part by Qihoo 360 Technology Inc., NSFC 61472358, NSFC 61702451, NSF CNS-0845671, and the Fundamental Research Funds for the Central Universities 2017QNA4017.

References

[1] Tesla, "A tragic loss," https://www.tesla.com/blog/tragic-loss, 2016.
[2] The New York Times, "Autopilot cited in death of Chinese Tesla driver," https://www.nytimes.com/2016/09/15/business/fatal-tesla-crash-in-china-involved-autopilot-government-tv-says.html, 2016.
[3] Tesla, "Enhanced Autopilot," https://www.tesla.com/autopilot, 2017.
[4] M. Hikita, "An introduction to ultrasonic sensors for vehicle parking," http://www.newelectronics.co.uk/electronics-technology/an-introduction-to-ultrasonic-sensors-for-vehicle-parking/24966/, 2010.
[5] K. Sethi, "Reverse parking sensors will be soon made mandatory on new cars," https://auto.ndtv.com/news/reverse-parking-sensors-will-be-soon-made-mandatory-on-new-cars-1456634, 2016.
[6] NHTSA, "Federal motor vehicle safety standard," https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/rear_visibility_nprm_12032010.pdf, 2010.
[7] Technavio, "Global automotive parking sensors market 2016-2020," https://www.technavio.com/report/global-automotive-electronics-global-automotive-parking-sensors-market-2016-2020, 2016.
[8] Tesla, "Model S software release notes v7.1," https://www.teslamotors.com/sites/default/files/pdfs/release_notes/tesla_model_s_software_7_1.pdf, 2016.
[9] ——, *Model S Owner's Manual v8.0*, 2016.
[10] F. W. Kremkau, *Diagnostic Ultrasound: Principles and Instruments*. WB Saunders Company, 2001.
[11] H. Kuttruff, *Ultrasonics: Fundamentals and Applications*. Springer Science & Business Media, 2012.
[12] SAE On-Road Automated Vehicle Standards Committee, "Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems," 2014.
[13] Google, "Google self-driving car project," https://www.google.com/selfdrivingcar/, 2016.
[14] Stanford Autonomous Driving Team, "Welcome," http://driving.stanford.edu/, 2016.
[15] Tesla, "Full self-driving hardware on all cars," https://www.tesla.com/autopilot, 2017.
[16] R. Katzwinkel, R. Auer, S. Brosig, M. Rohlfs, V. Schöning, F. Schroven, F. Schwitters, and U. Wuttke, "Einparkassistenz," in *Handbuch Fahrerassistenzsysteme*. Springer, 2012, pp. 471–477.
[17] M. Noll and P. Rapps, "Ultraschallsensorik," in *Handbuch Fahrerassistenzsysteme*. Springer, 2012, pp. 110–122.
[18] H. Winner, S. Hakuli, F. Lotz, and C. Singer, *Handbook of Driver Assistance Systems: Basic Information, Components and Systems for Active Safety and Comfort*. Springer Publishing Company, Incorporated, 2015.
[19] D. Foo Kune, J. Backes, S. S. Clark, D. Kramer, M. Reynolds, K. Fu, Y. Kim, and W. Xu, "Ghost talk: Mitigating EMI signal injection attacks against analog sensors," in *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE, 2013.
[20] Y. Son, H. Shin, D. Kim, Y. Park, J. Noh, K. Choi, J. Choi, and Y. Kim, "Rocking drones with intentional sound noise on gyroscopic sensors," in *Proceedings of the 24th USENIX Security Symposium*, 2015.
[21] T. Trippel, O. Weisse, W. Xu, P. Honeyman, and K. Fu, "WALNUT: Waging doubt on the integrity of MEMS accelerometers with acoustic injection attacks," in *Proceedings of the IEEE European Symposium on Security and Privacy*, 2017.
[22] C. Miller and C. Valasek, "Remote exploitation of an unaltered passenger vehicle," *Black Hat USA*, 2015.
[23] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno, "Comprehensive experimental analyses of automotive attack surfaces," in *Proceedings of the USENIX Security Symposium*, 2011.

[24] Cry Sound, "Cry343 free field measurement microphone," http://www.crysound.com/product_info.php?4/35/63, 2017.
[25] Arduino, "Arduino and Genuino project," https://www.arduino.cc/, 2016.
[26] S. A. Cummer and D. Schurig, "One path to acoustic cloaking," *New Journal of Physics*, vol. 9, no. 3, p. 45, 2007.
[27] J. Li and J. Pendry, "Hiding under the carpet: A new strategy for cloaking," *Physical Review Letters*, vol. 101, no. 20, p. 203901, 2008.
[28] L. He, "Development of submarine acoustic stealth technology," *Ship Science and Technology*, vol. 28, no. s2, pp. 9–17, 2006.
[29] Bose, "Noise cancelling headphones," https://www.bose.com/en_us/products/headphones/noise_cancelling_headphones.html, 2017.
[30] R. D. Hippenstiel, *Detection Theory: Applications and Digital Signal Processing*. CRC Press, 2001.
[31] S. Milligan, H. Vandelinde, and M. Cavanagh, *Understanding Ultrasonic Level Measurement*. Momentum Press, 2013.
[32] Keysight Technologies, "N5172B EXG X-series RF vector signal generator, 9 kHz to 6 GHz," http://www.keysight.com/en/pdx-x201910-pn-N5172B, 2017.
[33] Jinci Technology, "NU40C16TR-1," http://www.jinci.cn/en/showgoods/857.html, 2017.
[34] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
[35] B. Parhami, "A data-driven dependability assurance scheme with applications to data and design diversity," in *Dependable Computing for Critical Applications*, 1991, pp. 257–282.
[36] M. Wolf, A. Weimerskirch, and C. Paar, "Security in automotive bus systems," *In Proceedings of the Workshop on Embedded Security in Cars*, 2004.
[37] M. Wolf, A. Weimerskirch, and T. Wollinger, "State of the art: Embedding security in vehicles," *EURASIP Journal on Embedded Systems*, vol. 2007, no. 1, pp. 1–16, 2007.
[38] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage, "Experimental security analysis of a modern automobile," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2010.
[39] C. Valasek and C. Miller, "Adventures in automotive networks and control units," *DEF CON*, 2013.
[40] C. Miller and C. Valasek, "A survey of remote automotive attack surfaces," *DEF CON*, 2014.
[41] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote attacks on automated vehicles sensors: Experiments on camera and lidar," *Black Hat Europe*, 2015.
[42] H. Shin, D. Kim, Y. Kwon, and Y. Kim, "Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications," in *Proceedings of the International Conference on Cryptographic Hardware and Embedded Systems*. Springer, 2017, pp. 445–467.
[43] R. Chauhan, "A platform for false data injection in frequency modulated continuous wave radar," Ph.D. dissertation, UTAH STATE UNIVERSITY, 2014.
[44] 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 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015.
[45] C. Yan, W. Xu, and J. Liu, "Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle," *DEF CON*, 2016.
[46] P. Cheng, F. Zhang, J. Chen, Y. Sun, and X. Shen, "A distributed TDMA scheduling algorithm for target tracking in ultrasonic sensor networks," *IEEE Transactions on Industrial Electronics*, vol. 60, no. 9, pp. 3836–3845, 2013.
[47] F. Zhang, J. Chen, H. Li, Y. Sun, and X. S. Shen, "Distributed active sensor scheduling for target tracking in ultrasonic sensor networks," *Mobile Networks and Applications*, vol. 17, no. 5, pp. 582–593, 2012.
[48] A. Heale and L. Kleeman, "A sonar sensor with random double pulse coding," in *Proceedings of the Australian Conference on Robotics & Automation*, 2000.
[49] K.-W. Jörg, M. Berg, and M. Müller, "Using pseudo-random codes for mobile robot sonar sensing," in *Proceedings of the IFAC Symposium on Intelligent and Autonomous Vehicles*, 1998.
[50] Z. S. Lim, S. T. Kwon, and M. G. Joo, "Multi-object identification for mobile robot using ultrasonic sensors," *International Journal of Control, Automation and Systems*, vol. 10, no. 3, pp. 589–593, 2012.
[51] L. Shangguan, Z. Yang, A. X. Liu, Z. Zhou, and Y. Liu, "STPP: Spatial-temporal phase profiling-based method for relative RFID tag localization," *IEEE/ACM Transactions on Networking*, vol. 25, no. 1, pp. 596–609, 2017.

Wenyuan Xu is currently a professor in the College of Electrical Engineering at Zhejiang University. She received her B.S. degree in Electrical Engineering from Zhejiang University in 1998, an M.S. degree in Computer Science and Engineering from Zhejiang University in 2001, and the Ph.D. degree in Electrical and Computer Engineering from Rutgers University in 2007. Her research interests include wireless networking, network security, and IoT security. Dr. Xu received the NSF Career Award in 2009, a CCS best paper award in 2017, and an ASIACCS best paper award in 2018. She was granted tenure (an associated professor) in the Department of Computer Science and Engineering at the University of South Carolina in the U.S. She has served on the technical program committees for several IEEE/ACM conferences on wireless networking and security, and she is an associated editor of TOSN.

Chen Yan received his B.E. degree in Electrical Engineering from Zhejiang University, Hangzhou, China in 2015. He is currently a Ph.D. student in the College of Electrical Engineering at Zhejiang University. His research interests includes wireless security, sensor security and IoT security. He received a *Best Paper Award* from the 2017 ACM CCS. He was acknowledged by Tesla Motors in the Security Researcher Hall of Fame in 2016.

Weibin Jia received his B.S. degree in Automation from Northeastern University in 2015. He received his M.S. degree in Control Theory and Engineering from Zhejiang University in 2018. His research interests include vehicle security, battery and electromagnetic wave. He won the first prize of automobile cracking competition in HackPWN.

Xiaoyu Ji received his B.S. degree in Electronic Information & Technology and Instrumentation Science from Zhejiang University, Hangzhou, China, in 2010. He received his Ph.D. degree in department of Computer Science from Hong Kong University of Science and Technology in 2015. From 2015 to 2016, he was a researcher at Huawei Future Networking Theory Lab in Hong Kong. He is now an assistant professor with the department of Electrical Engineering of Zhejiang University. His research interests include IoT security, wireless communication protocol design, especially with cross-layer techniques. He won the best paper awards of ACM CCS 2017, ACM ASIACCS 2018 and IEEE Trustcom 2014. He is a member of IEEE.

Jianhao Liu is the team leader of Qihoo 360 vehicle cybersecurity team. He specializes in the vehicle cybersecurity. He made a research on the remote control car and found a security vulnerability of Tesla Model S. As a security expert, he is well experienced in security service, security evaluation and penetration and has been employed by various information security organizations. In addition, he has delivered speeches on various conferences such as Blackhat, Defcon, CanSecWest.