

Who Is in Control? Practical Physical Layer Attack and Defense for mmWave-Based Sensing in Autonomous Vehicles

Zhi Sun¹, Senior Member, IEEE, Sarankumar Balakrishnan², Student Member, IEEE, Lu Su³, Member, IEEE, Arupjyoti Bhuyan⁴, Senior Member, IEEE, Pu Wang⁵, Member, IEEE, and Chunming Qiao

Abstract—With the wide bandwidths in millimeter wave (mmWave) frequency band that results in unprecedented accuracy, mmWave sensing has become vital for many applications, especially in autonomous vehicles (AVs). In addition, mmWave sensing has superior reliability compared to other sensing counterparts such as camera and LiDAR, which is essential for safety-critical driving. Therefore, it is critical to understand the security vulnerabilities and improve the security and reliability of mmWave sensing in AVs. To this end, we perform the end-to-end security analysis of a mmWave-based sensing system in AVs, by designing and implementing practical physical layer attack and defense strategies in a state-of-the-art mmWave testbed and an AV testbed in real-world settings. Various strategies are developed to take control of the victim AV by spoofing its mmWave sensing module, including adding fake obstacles at arbitrary locations and faking the locations of existing obstacles. Five real-world attack scenarios are constructed to spoof the victim AV and force it to make dangerous driving decisions leading to a fatal crash. Field experiments are conducted to study the impact of the various attack scenarios using a Lincoln MKZ-based AV testbed, which validate that the attacker can indeed assume control of the victim AV to compromise its security and safety. To defend the attacks, we design and implement a challenge-response authentication scheme and a RF fingerprinting scheme to reliably detect aforementioned spoofing attacks.

Index Terms—Autonomous vehicles (AVs), millimeter wave (mmWave) radar, cyber-physical security.

Manuscript received November 18, 2020; revised February 21, 2021 and April 2, 2021; accepted April 21, 2021. Date of publication April 28, 2021; date of current version May 21, 2021. This work was supported in part by the INL Laboratory Directed Research and Development (LDRD) Program through the DOE Idaho Operations Office under Contract DE-AC07-05ID14517 and in part by the U.S. National Science Foundation (NSF) under Grant CNS1626374 and Grant CNS1737590. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Kai Zeng. (Corresponding author: Zhi Sun.)

Zhi Sun and Sarankumar Balakrishnan are with the Department of Electrical Engineering, University at Buffalo, Buffalo, NY 14260 USA (e-mail: zhisun@buffalo.edu; sarankum@buffalo.edu).

Lu Su is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: lusu@purdue.edu).

Arupjyoti Bhuyan is with the Idaho National Laboratory (INL), Idaho Falls, ID 83402 USA (e-mail: arupjyoti.bhuyan@inl.gov).

Pu Wang is with the Department of Computer Science, University of North Carolina at Charlotte, Charlotte, NC 28223 USA (e-mail: pu.wang@uncc.edu).

Chunming Qiao is with the Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY 14260 USA (e-mail: qiao@buffalo.edu).

Digital Object Identifier 10.1109/TIFS.2021.3076287

I. INTRODUCTION

AUTONOMOUS vehicles (AVs) are envisioned to be the future of transportation. It is predicted that in the coming decades (by 2045) more than half of the new vehicles manufactured will be autonomous [1]. Fully autonomous vehicles need human-like cognition capabilities to safely navigate the environment and react to unforeseen circumstances. To achieve real-time human-like sensing capability, AVs rely on multitude of heterogeneous sensors such as camera, LiDAR, and radar. Millimeter wave (mmWave) radars have become an attractive choice in AVs due to their better spatial resolution thanks to the wide bandwidth available at mmWave spectrum. Currently, the mmWave sensor is a crucial component of state-of-the-art software systems for AVs such as Baidu Apollo [2] and Autoware [3]. Open-source software systems like *OpenPilot* uses mmWave radar as a major component for its lane keep assist and forward collision warning capabilities [4].

With increasing companies vying for the lucrative AV market and the predicted prevalence of AVs on the road, the safety-critical operation of AVs is important. AV safety has been a subject of intense scrutiny over the past few years. Even though AVs use state-of-the-art sensors and software systems to make reliable decisions and perform safe driving actions, multiple serious or fatal accidents involving AVs have been reported [5]–[7]. Those accidents show that sensors and software used in AVs are prone to vulnerabilities, which raises a serious question: *Can a powerful adversary take advantage of the vulnerabilities of the sensors used in AVs and cause an intentional safety-critical incident?* For example, the camera in the AV could be attacked by blinding it with laser beams thus derailing the AVs capabilities to detect lanes and traffic signs leading to severe consequences [8]. Similarly, LiDAR could be spoofed with fake obstacles which could be perceived as a vehicle thus spoofing the victim AV to make erroneous critical decisions which may lead to fatal crash [9]–[11].

Despite the important role of mmWave sensing in AVs, the understanding of its vulnerabilities is far from sufficient. In [12], attacks aiming to spoof chirp-based radar are investigated using software defined radios. However, the attack is highly impractical since the attacker uses a physical cable connecting to the victim's system to send the spoofing signals. The radar also only works at 2.4GHz, not the mmWave band. In [13], a preliminary attempt to spoof mmWave sensor in AVs

using waveform generators is reported. However, the designed spoof signals do not use the chirp waveform that is used in mmWave radar. Instead, the waveform generator only generates simple noise signals that can be easily detected by the victim. In summary, the attacks developed in existing works are too simple to cause meaningful security consequences in AVs.

To our best knowledge, no existing work performs comprehensive security analysis of the AV mmWave sensor module and its implications at driving decision level in real-world driving scenarios. On that front, this paper aims at performing practical physical layer security analysis of mmWave radar in AVs by designing and implementing practical attack and defense strategies. We ask the following questions: *Is it possible to reliably spoof an mmWave sensor used in the victim AV? And in doing so, will the attacker be able to cause meaningful security consequences to the victim AV?*

There are three key challenges in answering the above questions. (a) *Is it feasible to spoof a mmWave sensor to perceive a fake obstacle?* The attacker must be able to generate the replica of the signal used by the victim mmWave sensor, with a precisely controlled time delay. (b) *Can the attacker continuously spoof an obstacle to deceive the AV?* The attacker needs to continuously spoof the AV to influence the AV's decision making process. Hence, the attacker should continuously track the position of the victim AV to update the time delay with which the attacker must spoof. More importantly, in addition to (a) and (b), we ask a more ambitious question. (c) *Is it possible for the attacker to make the AV crash on to a lead obstacle by deceiving the victim AV to perceive the lead obstacle out of the danger zone?* To deceive the AV into recognizing the lead vehicle as out of the danger zone, for instance, perceiving it as an obstacle in adjacent lane, the attacker must deviate the obstacle from the real position to a strategically controlled position, which is a daunting task.

In this work, we solve all the above challenges and show that the mmWave sensor in the victim AV can be reliably spoofed using a multi-attacker approach, leading to severe security consequences. First, to address (a), we develop a mmWave radar adversary system using the state-of-the-art mmWave software defined radio [14], which is capable of transmitting a replica of the waveform used by the victim AV. Second, to address (b), we design a tracking module to continuously track the position and velocity of the victim AV and update the delay parameters of the spoofing signals. Finally, to address (c), we identify the challenges in faking the locations of existing obstacles and develop three sophisticated strategies: (1) random deviation of an obstacle by using Gaussian signal attack; (2) synchronous attacks that can fake the positions of obstacles in a controlled manner, and (3) since the perfect synchronization among distributed attackers is not practical, we develop a novel asynchronous attack strategy, which explores the signal correlation in the position estimation method used in the mmWave sensor. The attack strategy can reliably fake the locations of existing obstacles from their real positions to the attacker controlled strategic positions. We validate the developed strategies by conducting real-world

drive-by experiments and perform end-to-end security analysis using Baidu Apollo¹ based AV software [2].

To understand the severity of the security impact of our developed attacks, we construct 5 real-world attack scenarios: 1) *AV Stalling Attack*, which creates traffic chaos; 2) *Hard Braking Attack*, which endangers the safety of the passengers and the vehicle behind; 3) *Lane Changing Attack*, which forces the victim AV into making unintentional lane change; 4) *Multi-stage Attack*, which combines multiple attack stages with the goal of leading the AV to a fatal crash; and 5) *Cruise Control Attack*, which spoofs the cruise control function of the victim AV and leads to a fatal crash.

As defense mechanism, we propose and implement two physical layer solutions to detect the developed spoofing attacks. First, we propose a *challenge-response* defense where the victim AV detects the spoofing attacks from the response to the transmitted randomized challenge waveform. Second, we adopt an RF fingerprinting scheme to distinguish spoofed received signal from the legitimate received signal, thereby detecting the spoofing attacks.

Our major contributions are summarised as follows:

- We perform the first comprehensive physical layer security analysis of AV mmWave sensing. Practical spoofing attack strategies are developed and implemented, which are proved to be able to cause fatal consequences in real world. The proposed attacking strategies are not just heuristic but based on rigorous mathematical analysis on the state-of-the-art AV sensing algorithms, which has guaranteed spoofing performance.
- We develop a software defined radio based mmWave testbed to perform sophisticated single node and multi-node spoofing attacks on the mmWave sensor module in a Lincoln MKZ-based AV testbed.
- To study the severity of the developed attack on the AC security, we conduct extensive real-world experiments on the Lincoln MKZ AV fitted with state-of-the-art sensors. Five attack scenarios are designed to impart various levels of security impacts on the AV.
- We show that our spoofing attacks impact the end-to-end security of AVs and spoof the AV into making hazardous safety-critical driving decisions.
- As a defense against the spoofing attack, we propose and implement two spoofing attack mitigation strategies, including the *challenge-response* authentication and the RF fingerprinting mechanism.

II. RELATED WORK

A. Sensor Attacks

In [13], jamming and spoofing attacks are performed on ultrasonic sensors. Blinding attacks on camera and preliminary jamming attacks on mmWave radars are also investigated. In [8], physical sensor attack are performed on AV camera and LiDAR. In [15] and [10], adversarial attack are launched on LiDAR-based perception in AVs. In [16], an attack strategy is developed to spoof chirp-based ranging system. In [12]

¹Baidu Apollo is an open-source software that implements perception, planning, and control for AVs and is widely used by AV industry.

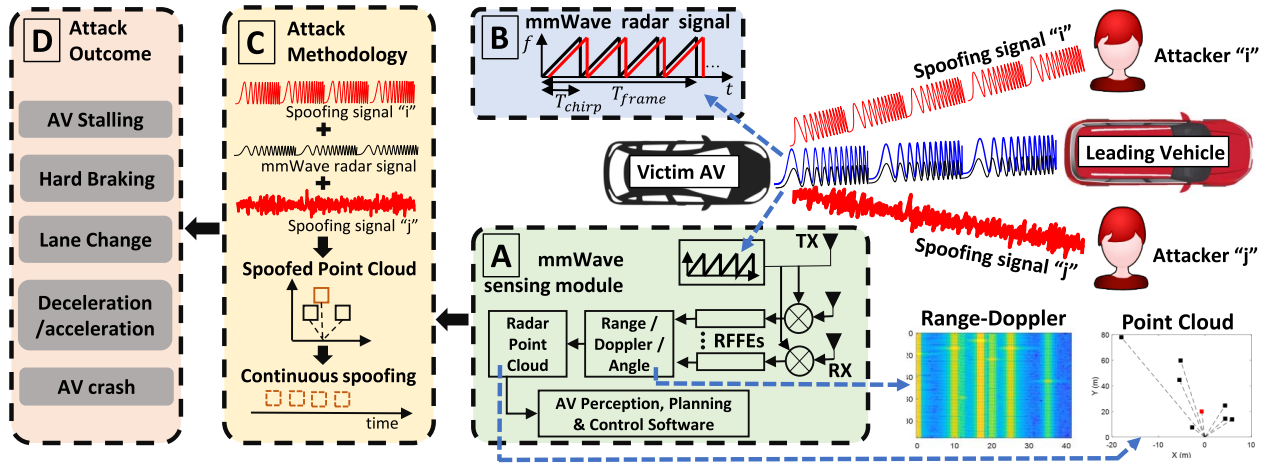


Fig. 1. Illustration of the adversarial setting and victim AV system overview. Block [A] shows the mmWave sensing module used in AVs. Block [B] shows the transmitted and received mmWave chirp waveform. Block [C] illustrates the attack methodology. Block [D] shows the driving decisions taken by the AV due to the attack described in block [C].

and [17], the feasibility of spoofing FMCW based radar is explored. However, the attacker has to use physical cables connected to the victim radar to launch the attack. In [18], a distance spoofing attack on mmWave FMCW radar is presented. Contrary to all those works, we perform the end-to-end security analysis of AV mmWave sensing system in real AV testbed in various practical real-world scenarios.

B. System Level Attacks

In [19] and [20], fault injection tools are proposed to study the safety and reliability of various AV hardware and software components. In [21], the resilience of *OpenPilot*, an open-source Adaptive cruise control (ACC), LKAS, and Assisted lane change system for AVs is evaluated, where radar is used as a primary sensor. Contrary to the security analysis using synthetic data, in this work, we evaluate the end-to-end security vulnerability of the AV by performing physical spoofing attack on the victim AV.

C. Adversarial Attacks

Recently, significant efforts have been made to investigate the security of AVs against adversarial sensor attacks [10], [22]–[25]. Those works target the machine learning model used in AVs, whereas this paper focuses on the physical attack of the radar used in AVs.

D. Defense

The integrity of sensor signals can be verified at the physical layer by considering the physics governing the sensor. For instance, the physical propagation characteristics of sensor signals can be exploited to design attack detection mechanisms. In [26], PyCRA, a physical challenge-response authentication scheme is proposed to defend magnetic and RFID sensors against malicious spoofing attacks. In this work, we propose to use (1) the *challenge-response* mechanism that randomizes the waveform parameters of the mmWave radar, as well as (2) the *RF fingerprinting* mechanism that exploits signal characteristics, to defend against malicious attacks.

III. SYSTEM OVERVIEW AND ADVERSARY MODEL

In this section, we overview the proposed attack system, introduce the background of the mmWave sensing module in AV, and discuss the threat model and attack goals. Fig.1 shows the overview of the attack system architecture and the AV mmWave sensing module. The proposed attack system consists of multiple distributed attackers (Attacker “i” and “j” in Fig.1) transmitting spoofing signals with the objective of deceiving the victim AV (black car in Fig.1) into making dangerous decisions. The mmWave sensing module of the victim AV is shown in block [A] of Fig. 1, which is introduced in Sec. III-A. The objective of the attackers is to spoof a reliable fake obstacle at an arbitrary location or fake the location of an existing obstacle (e.g., the leading vehicle), using the methods discussed in Sec. IV. The key methodology is to generate a spoofed point cloud (i.e., the positioning results derived by the mmWave radar) by combining the legitimate mmWave radar signal (block [B]) with the spoofing signals, as shown in block [C] in Fig. 1. Such spoofing attacks cause the victim AV to make dangerous safety-critical decisions as shown in block [D] in Fig. 1. The detailed attacker goals and assumptions in our threat model are discussed in Sec. III-B.

A. Background: mmWave Radar Module in AV

A mmWave radar has one transmitter and multiple receivers, as shown in block [A] in Fig. 1. The transmitter generates the chirp waveform (block [B] in Fig. 1), which is:

$$x(t) = \cos\left(2\pi f_{start}t + \frac{\pi B}{T_{chirp}}t^2\right), \quad (1)$$

where f_{start} is the transmission start frequency, B is the sweep bandwidth, i.e., $f_{end} - f_{start}$, and T_{chirp} is the chirp duration. Then the chirp signal is reflected back by the obstacle at distance d and received by the multiple receivers at the mmWave radar. The received signal is given by

$$r(t) = \alpha \cos\left(2\pi f_{start}(t - t_{delay}) + \frac{\pi B}{T_{chirp}}(t - t_{delay})^2\right), \quad (2)$$

where $t_{delay} = \frac{2(d+vt)}{\lambda}$ is the time delay due to an obstacle at distance d moving with velocity v with respect to the AV. This received signal is mixed with the transmitted signal in Eq. 1, which produces a cosine signal given by

$$y_{mixed}(t) = \cos\left(2\pi f_{start}t_{delay} - \pi \frac{B}{T}t_{delay}^2 - 2\pi \frac{B}{T}t_{delay}t\right), \quad (3)$$

where $\frac{B}{T}\tau$ is the beat frequency f_b , which is proportional to the distance d . The beat frequency can be calculated by $f_b = \frac{2RB}{cT_{chirp}} + f_{doppler}$ with $f_{doppler} = \frac{2v\cos\psi}{\lambda}$.

The mixed signal in Eq. 3 is further processed by the sensing module shown in block [A] in Fig. 1, which consists of three subsystems: *range & velocity estimation*, *direction estimation*, and *point cloud generation*.

1) *Range & Velocity Estimation*: The distance from the mmWave radar to the object is derived by taking the FFT of the signal in Eq. 3 and finding the f_b corresponding to the peak of the FFT. Then the distance is derived by $d = \frac{f_b T c}{2B}$. The mobility of the obstacle introduces a Doppler shift with the Doppler frequency $f_{doppler} = \frac{2vf_{start}\cos(\theta)}{c}$. Accordingly, in the mobile case, the received signal is given by

$$r(t) = \cos\left(2\pi f_{start}(t - t_{delay}) + \phi_{doppler} + \pi \frac{B}{T}(t - t_{delay})^2\right), \quad (4)$$

where $\phi_{doppler} = 2\pi f_{doppler} * n * T$ and n is the chirp index in a frame. The velocity of the obstacle is determined by taking second FFT across chirps in different time indices. The range FFT and Doppler FFT result in a Range-Doppler map of dimension $R_{bin} \times V_{bin}$ where R_{bin} is the number of range bins and V_{bin} is the number of velocity bins. An example of the Range-Doppler map is given in Fig. 1. The estimated range bin and velocity bin of each of the obstacle in the Range-Doppler map is denoted as a tuple (r_m^k, v_n^k) , where k is the index of the obstacle, m and n are the index of the range bin and velocity bin, respectively.

2) *Direction Estimation*: As shown in block [A] in Fig. 1, an AV mmWave radar uses multiple receiving antennas. The Range-Doppler maps of all receiving antennas result in a data cube of dimension $N \times R_{bin} \times V_{bin}$ where N is the number of receiving chains. The direction θ of each of the obstacle (r_m^k, v_n^k) can be accurately estimated using subspace-based beamforming methods, such as the Multiple Signal Identification and Classification (MUSIC) algorithm [27]. For a mmWave radar with N receiving antennas and arriving signals from K directions, MUSIC algorithm separates $N - K$ noise subspace Z from the K signal subspace $S(\theta_1, \theta_2, \dots, \theta_K)$ using subspace decomposition techniques such as Eigen decomposition. Then, the K arriving signal directions can be estimated by the K minima of function $\frac{1}{a(\theta)^* Z^* Z a(\theta)}$, where $a(\theta)$ is the steering vector corresponding to the direction θ .

3) *Point Cloud Generation*: The identified positions of the obstacles are presented as a Point Cloud for further processing by the AVs perception, planning, and control software modules. An example of the Point Cloud is given

in Fig. 1. The details of AV software modules are given in Appendix IX.

B. Threat Model and Attack Goal

1) *Threat Model*: We identify mmWave radar spoofing as a threat model to spoof the victim AV into making dangerous safety-critical decisions. A multiple-attacker spoofing model are considered, with the following assumptions:

a) *Prior knowledge*: The attacker has prior knowledge on the type of mmWave radar used in victim AV. Hence, the key parameters are available, including the waveform parameters (frequency sweep bandwidth and slope), the number of chirp signals used in a frame, and the duty cycle of the radar frame. The attackers do not need to know the perception algorithm used by the victim AV.

b) *Radar hardware physical access*: The attackers do not have physical access to the radar hardware, firmware, and software. The attackers cannot modify the firmware or inject false data in to the system.

c) *Position of the attackers*: The attackers are free to take any position. For stealthiness, the attackers in our experiments are stationary and are positioned on the side of the roadways. Also, due to the current hardware limitation, there are physical signal propagation constraints (discussed in Sec. IV-A4) that limits the attackers' positions. However, it should be noted that the proposed attacking strategies do not necessarily require the attackers to be stationary. It is possible to let mobile vehicles to carry the distributed attackers and launch the attacks in motion. However, it requires the attackers to acquire additional hardware and software to (1) localize themselves in real time and (2) synchronize with each other to coordinate the distributed attacks. It also require more powerful mmWave transmitter to overcome the physical signal propagation constraints.

d) *Attack equipment*: The attackers use a mmWave receiver for sensing the victim radar's signal and a mmWave transmitter for performing replay and spoofing attack. The transmitter and receiver are discussed in Sec. V-A.

2) *Attack Goal*: The attack goal is to spoof the victim AV into making potentially dangerous driving decisions or actions as shown in block [D] in Fig. 1, which endanger the occupants of the AV as well as the other vehicles in the vicinity.

IV. ATTACK DESIGN: PHYSICAL ATTACKS ON AV MMWAVE RADAR

In this section, we develop the two key spoofing strategies, including adding fake obstacles at arbitrary locations and faking the locations of existing obstacles, which form the more sophisticated attack scenarios described in Sec. VI.

A. Adding Fake Obstacles

The fundamental yet core attack capability of an attacker (for generality, assume "attacker i" in Fig. 1) is to reliably trick a victim AV's radar to perceive a fake obstacle spoofed by the attacker as a legitimate obstacle. An AV can be spoofed with a fake obstacle at a particular distance by transmitting a spoofing waveform identical to the waveform in Eq. 1 used




Fig. 2. Range FFT of a chirp showing the spoofed obstacle at 30m.




Fig. 3. Range FFT of a chirp showing the multiple spoofed obstacle.

by the victim AV with a carefully designed delay. An attacker could sense the transmitted radar waveform from the AV and either replay a previously recorded waveform or impersonate the transmitted waveform after a delay of $t_{delay} = \frac{2*d}{c}$, where d is the desired distance of the spoofed obstacle from the victim radar.

1) *Spoofing a Mobile Obstacle*: The received waveform reflected by the moving object with velocity v is given by Eq. 4. The velocity progresses linearly with chirp number, which can be estimated by the relationship $f_{doppler}nT_{chirp} = \frac{2v_{fstart}}{c}nT_{chirp}$. The attacker can use software defined radios to easily change the phase of its transmitted waveform. If the desired spoofed velocity is v , the manipulated phase should be $\frac{2v_{fstart}}{c}nT_{chirp}$ for each of the chirp n . By linearly increasing the chirp phase, the attacker can spoof an object with the fake velocity. As an example, Fig. 2 shows a Range FFT where a spoofed obstacle at 30m is perceived by the victim AV.

2) *Spoofing Multiple obstacles*: The transmitted spoofing signal is a superposition of signals corresponding to various delays τ_{delay} between the mmWave radar and the multiple spoofed obstacles. The spoofing signal is given by

$$x_{spoof}(t) = \sum_{n=1}^N \cos\left(2\pi f_c(t - \tau_{delay,n}) + \frac{\pi B}{T}(t - \tau_{delay,n})^2\right), \quad (5)$$

where N is the number of spoofed obstacles. For our testbed described in Sec. V-A, the delay τ_{delay} in terms of number of samples m can be calculated as $mT_s = \frac{d}{c}$, where $1/T_s$ is the sampling rate of the transmitter, and d is the distance between the two spoofed obstacles. As an example, Fig. 3 shows the Range FFT where multiple spoofed obstacles are created. The attacker simultaneously spoofs four obstacles with 10m distance between each of them.

3) *Influence of Attacker's Distance*: Since the attacker needs to sense the mmWave signal from the victim AV and transmit the spoofing waveform, the attacker should not be too far away.




Fig. 4. Received power of the spoofing signal as a function of attacker's distance.




Fig. 5. Received power of the spoofing signal as a function of attacker's angle.

Fig. 4 shows the received power of the attacker's spoofing signal for varying distances between the attacker and the victim radar, when using our mmWave testbed described in Sec. V-A. The maximum distance from the victim AV that the attacker can still reliably spoof is 26 m. This limitation is due to the limited gain and transmit power of the antenna module of the attacker [see Sec. V-B]. The gain of the antenna used for the attack is 23 dBi. The attack range can be significantly improved by using a high gain directional antenna [28].

4) *Influence of Attacker's Angle*: Since attackers are on road side, there is an angle between the victim AV's moving direction and the attacker antenna's direction. The attacker has restricted angular freedom in launching the spoofing attack since the mmWave radars in existing cars have a narrow field-of-view, typically $\pm 10^\circ$. Fig. 5 shows the received power of the spoofing signal for different angular positions of the attacker with respect to the victim radar, using our mmWave testbed (see Sec. V-A). When the attacker is within the field-of-view of the victim radar, the victim radar is reliably spoofed. However, as the attacker moves away from $\pm 10^\circ$, the spoofing attack fails since the attacker is in the side-lobe of the victim radar.

B. Faking Locations of Existing Obstacles

The goal of this attack is to fake the location of an obstacle so that the victim AV considers that obstacle is not in its region-of-interest (ROI) (see definition in Appendix. IX). For example, in a multi-lane scenario, the attacker can deviate the obstacle in front of the AV (e.g., the leading vehicle) to a position away from the current lane, which spoofs the AV into crashing onto the leading vehicle. To achieve such goal, we propose three location-faking attack strategies: *Random Signal Attack*, *Synchronous Attack*, and *Asynchronous Attack*.

1) *Random Signal Attack*: Considering the victim AV uses an optimal estimator to estimate the direction of detected




Fig. 6. Range-Doppler map and point cloud of the identified obstacles (T_i) and the spoofed obstacles ($s1$ and $s2$). The attackers are not synchronized leading to the two spoofing signals $s1$ and $s2$ resolved in Range-Doppler domain.

obstacles, the Cramer-Rao bound [29] for such estimator is:

$$\frac{1}{2} \left[\sum_{i=1}^K \text{Re} \left\{ x_{rx}^\dagger [i] \frac{\partial a}{\partial \theta} R_y^{-1} [I - a(a^\dagger a)^{-1} a^\dagger] R_y^{-\frac{1}{2}} \frac{\partial a}{\partial \theta} x_{rx} [i] \right\} \right]^{-1}, \quad (6)$$

where $R_y = \frac{1}{L} \sum_{i=1}^K E[H_{adv}[i]x_{adv}[i]x_{adv}^H[i]H_{adv}^H[i]] + \sigma_{noise}^2 I$ is the attacker signal covariance matrix with noise. It has been proved that the attacker's waveform that maximizes the variance of the estimation is drawn from a zero-mean Gaussian distribution [29]. Therefore, in this Random Signal Attack, we use our mmWave testbed (see Sec. V-A) as the attacker and let it transmit a Gaussian waveform with controlled power to influence the victim radar. From experiments, we prove that the attacker is able to deviate the angle of the detected objects by $\pm 5^\circ$. This attack can cause severe consequences on the AV's security as the true object's position is erroneously estimated. Then the AV could make the decision that the object does not exist in its ROI or the obstacle is in the adjacent lane, thus categorizing the obstacle as rather safe instead of flagging it.

a) *Drawback of random signal attack:* Since the attacker's waveform is Gaussian distributed, the angle deviation introduced by the attacker is random. As a result, the Random Signal Attack cannot cause guaranteed and controlled attacking consequence. Therefore, we next proceed to develop attack strategies to generate deviations with controlled angles.

2) *Synchronous Attack:* Two challenges need to be addressed to fake the obstacle's location in a controlled manner: First, the attacker need to spoof an obstacle at a controlled position, especially the angular position (since the distance position can be spoofed using the strategy developed in Sec. IV-A), thus creating an illusion the obstacles position has changed. Second, the attack need to simultaneously obliterate the reflection signal from the existing obstacle.

a) *Challenge 1:* We start with addressing the first challenge, i.e., spoofing the angular position of the obstacle. The key idea is to employ multiple attackers located at different angular positions, hoping the combined spoofing signals from those attackers can create an illusion that the new obstacle is located at an angular position in between those attackers.

However, the challenge is that existing mmWave radar can already distinguish the signals from different attackers. As discussed in Sec. III-A, the AV mmWave radar distinguishes object in the range r_m , Doppler v_m , and angle θ domain. Fig. 6

shows an example, the two attackers $s1$ and $s2$ are at the same distance $40m$ from the mmWave radar but with different angles (-10° and 7°). In the left Range-Doppler map of Fig. 6, the spoofed obstacles due to $s1$ and $s2$ have the same range bin m but appear in different velocity bins n and n' due to different phase offset, resulting in two obstacles at $(r_m^{(1)}, v_n^{(1)})$ and $(r_m^{(2)}, v_n^{(2)})$, which fails the spoofing attack.

To address the challenge, we introduce the synchronous attack, where the attackers are perfectly synchronised (in frequency, time, and phase), resulting in a correlated signal that make the spoofed obstacles appears in the same Range-Doppler bin $(r_m^{(1,2)}, v_n^{(1,2)})$. Here (1,2) refers to the obstacle spoofed by attacker 1 and 2. The next question is: *what is the angular position of this spoofed obstacle (1,2)? Can we control the angle?* As discussed in Sec. III-A, the subspace-based direction estimation algorithm calculates the direction of the obstacles by separating the signal subspace $S(\theta_1, \theta_2)$ from the noise subspace Z and finding the directions that are orthogonal to the noise subspace. The signal subspace $S(\theta_1, \theta_2)$ itself is a linear combination of the directions of all the received signals. We utilize such property and derive the following proposition for K synchronous signals.

Proposition 1: *For K synchronous signals in the same Range-Doppler bin $(r_m^{1,\dots,K}, v_n^{1,\dots,K})$ with directions $\theta_1, \theta_2, \dots, \theta_K$, the spoofed obstacles direction θ_{spoof} is the direction that maximizes $\|S(\theta_1, \theta_2, \dots, \theta_K)^H \mathbf{a}(\theta_{spoof})\|^2$ and is given by $\theta_{spoof} = \text{mean}(\theta_1, \theta_2, \dots, \theta_K)$.*

The synchronous attackers can launch the attack based on proposition 1. As a result, the victim AV identifies the combined spoofing signals from the multiple attackers as a single obstacle at $(r_m^{1,\dots,K}, v_n^{1,\dots,K})$ and direction θ_{spoof} . The position of the spoofed obstacle is given by $(x_{spoof}(t), y_{spoof}(t)) = \{x_{victim}(t) + d_{spoof} * \cos(\theta_{spoof}), y_{victim}(t) + d_{spoof} * \sin(\theta_{spoof})\}$, where d_{spoof} is the spoofed distance corresponding to $(r_m^{(1,2)}, v_n^{(1,2)})$ and (x_{victim}, y_{victim}) is the position of the victim AV.

b) *Challenge 2:* Then we address the second challenge, i.e., *How to obliterate the existing obstacle?* We employ an additional attacker to perform a jamming attack to overwhelm the genuine radar reflected signal. The power of the jamming signal P_j need to be higher than that of the reflected signal P_r received by the victim AV. $P_j = \frac{P_{attacker} G_{attacker} \lambda^2}{(4\pi d_{attacker})^2}$, where $P_{attacker}$, $G_{attacker}$, and $d_{attacker}$ are the transmit power, antenna gain, and distance of the attacker.

c) *Attacking results:* Fig. 7 shows the results of the synchronous attack. The scenario is the same as Fig. 6. Now the two attackers $s1$ and $s2$ at directions -10° and $+7^\circ$ are *synchronized*. The Range-Doppler map shows that signals from the two attackers fall into the same bin. The point cloud shows that the attackers' spoofed obstacle is identified as a single obstacle $s1, s2$ with direction 4° to the mmWave radar.

d) *Drawback of synchronous attack:* The attackers need to be perfectly synchronized in frequency, phase, and transmission starting time. Any mismatch in frequency or phase of the waveform between the attackers, or the transmission starting time could invalidate the spoofing waveform in frequency and velocity. For example, for a waveform with 300 MHz sweep




Fig. 7. Range-Doppler map and point cloud of the identified obstacles (T_i) along with the spoofed obstacle ($s1$ and $s2$). The attackers are perfectly synchronized leading to the two spoofing signals resolved as single obstacle $s1, s2$.

bandwidth and 30 μ s sweep time, the frequency bin width in the range profile is $\frac{f_{\text{sampling}} * c}{2 * \text{slope} * N_{\text{fft}}} = 681.8 \text{ KHz}$. If the carrier frequency offset (CFO) between the attackers is more than 681.8 KHz, the attack fails. Considering the mmWave radar works at 60 GHz or 77 GHz band, it is very difficult, if not impossible, to meet such synchronization requirement.

3) *Asynchronous Attack*: To avoid the stringent requirement of the synchronous attack, we propose a novel signal injection attack strategy that achieves the same performance only based on asynchronous attackers. This attack is designed based on the following three observations:

- The angle-of-arrival estimation method (see Sec. III-A2) separates the signal subspace from the noise subspace. Any correlation in the subspace results in incorrect angle estimation, as described in Proposition 1.
- The direction estimation itself is agnostic to the baseband representation of the transmitted waveform.
- The mmWave radar mixes the received waveform with the transmitted waveform to obtain the mixed signal in Eq. 3, which has a constant frequency of $2\pi \frac{B}{T} \tau$. Most signals other than a replica of the transmitted waveform (Eq. 1) result in mixed signals with high frequency that are filtered out by the low-pass filter of the victim AV's radar. However, if we use a filtered noise (correlated noise) to mix with the transmitted signal, a filtered signal can pass the low-pass filter as the output.

The key idea of the proposed signal injection attack is to use filtered noise to introduce undesired correlation in the received radar signal samples, so that we can use multiple attackers to fake the location of existing obstacles.

a) *Attack procedure*: Without loss of generality, we consider there are two attackers in the following explanation. The 1st attacker denoted as “attacker i” in Fig. 1 transmits a waveform with correlated noise $F(\text{noise})$ and with delay τ_{delay} to spoof an obstacle at distance d_{spoof} . The transmitted waveform is given by

$$TX_{\text{spoof},1} = \sqrt{P_1} \left(\cos \left(2\pi f_{\text{start}} t' + \pi \frac{B}{T} t'^2 \right) + F(\text{noise}) \right), \quad (7)$$

where $t' = (t - \tau_{\text{delay}})$, P_1 is the transmission power. The 2nd attacker denoted as “attacker j” in Fig. 1 transmits a correlated




Fig. 8. Range-Doppler map and point cloud showing the identified obstacles and the spoofed obstacle $s1, s2$. The asynchronous attackers perform distributed angle attack leading to the two spoofing signals resolved as single obstacle $s1, s2$.




Fig. 9. By varying transmission power, the asynchronous attackers can control the angular position of the spoofed obstacle.

filtered noise with power P_2 :

$$TX_{\text{spoof},2} = \sqrt{P_2} \cdot F(\text{noise}). \quad (8)$$

At the victim mmWave radar, after mixing with the radar waveform, the two spoofing signals become:

$$r_1 = a(\theta_1) \frac{g_1}{PL(d_1)} \sqrt{P_1} (\cos(\Phi) + F(\text{noise})); \quad (9)$$

$$r_2 = a(\theta_2) \frac{g_2}{PL(d_2)} \sqrt{P_2} \cdot F(\text{noise}), \quad (10)$$

where $a(\theta_1), a(\theta_2)$ are the direction vectors corresponding to the directions of two attackers; $PL(d_1), PL(d_2)$ are the two distance dependent path loss; $\Phi = 2\pi f_{\text{start}} \tau - \pi \frac{B}{T} \tau^2 - 2\pi \frac{B}{T} \tau t$.

Since the two spoofing signal r_1 and r_2 are correlated, the estimated direction θ_{spoof} of the spoofed obstacle is the direction that maximizes $\|S(\theta_1, \theta_2)^H a(\theta_{\text{spoof}})\|^2$. Attacker i's waveform determines the distance of the spoofed obstacle. Attacker j's waveform together with attacker i influence the angle of the spoofed obstacle. Meanwhile, attacker j also performs as a jammer to obliterate the existing obstacle. By this way, the two attackers do not need any synchronization but can simultaneously address the two challenges of the faking location attack described in Sec. IV-B2.

b) *Attacking results*: Fig. 8 shows the results of two attackers $s1$ and $s2$ performing *asynchronous* attack. The attackers spoofed obstacle is identified as a single obstacle $s1, s2$ at distance 40m and at angle 4° to the mmWave radar.

We further perform experiments to show how the attackers control the position of the spoofed obstacle. In this experiment, the two attackers are located at a direction of -10° and $+10^\circ$, respectively, as shown in Fig. 9. Fig. 9 shows the spoofed obstacle $s1$ at $+3^\circ$ with respect to the victim AV. By increasing the transmit power of attacker 1, we can move




Fig. 10. The mmWave software defined radio testbed.

the spoofed obstacle's position to s_2 , which is $+7^\circ$ with respect to the victim AV. In Sec. VI, through the case studies of different driving scenarios, we show how such distributed and asynchronous attack could potentially trick the AV into making dangerous driving decisions.

V. SYSTEM DESIGN

To cause meaningful security consequences, the attacker need to continuously spoof the victim AV based on the combination of the two attack strategies discussed in Sec. IV. Moreover, beside launching the attacks, the attackers also need to: (1) sense the signal from the victim AV; and (2) track the position of the victim AV to continuously spoof it.

In this section, we introduce our attacker system based on the state-of-the-art mmWave testbed. We first present the mmWave testbed that lays the foundation of our attacker system. Then we discuss the main components of the attacker system. Finally we discuss how the attacker system performs continuous spoofing attack based on the mmWave testbed and the spoofing strategies developed in Sec. IV.

A. mmWave Radar Testbed

Our attacker testbed is based on a state-of-the-art software defined radio (SDR) mmWave transceiver system from NI [14]. The testbed architecture is shown in Fig. 10. Each mmWave SDR can be configured either as a transmitter or a receiver. On the transmitter side, each mmWave SDR has a PXIe-3610 DAC digital-to-analog converter. The PXIe-3610 is a 14-bit DAC with 3.072 GS/s sampling rate. It provides interface to analog baseband signal with a maximum bandwidth of 2 GHz. Each receive mmWave SDR has a PXIe-3630 ADC with 12 bits resolution and supports 3.072 GS/s sampling rate. The antenna module is a SiBEAM 60 GHz phased array that supports 2 frequency bands (60.48 GHz and 62.64 GHz center frequency). The mmWave SDR has a PXIe-7902 FPGA module to support physical layer baseband signal processing design.

B. Attacker System

As shown in Fig. 11, each of the attacker systems consists of (1) **tracking system**, (2) **sensing system**, and (3) **transmitter**, working in tandem to form the attacker spoofing system.

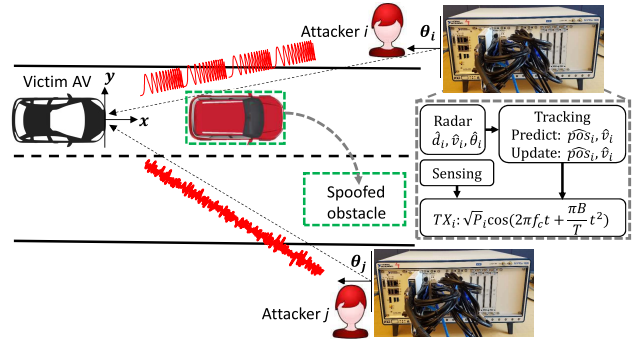


Fig. 11. Conceptual AV spoofing scenario showing distributed attackers and the attacker system modules.




Fig. 12. Tracking the distance, velocity, and angle of the victim AV using the Kalman filter.

1) **Tracking System:** The tracking system consists of radar and tracking modules. The radar is used to estimate the position and velocity of victim AV, based on Kalman filter (see Appendix. IX-2). To avoid interfering with the spoofing transmitter, the radar in tracking system operates in a different frequency band. This radar gives estimation results with a cycle period of 4.45 ms. Fig. 12 shows the distance, velocity, and angle predicted by the tracking system compared with the ground truth, which prove that the tracking system achieves accurate prediction of the victim's position and velocity.

2) **Sensing System:** The attacker needs to sense/detect the beginning of the transmitted mmWave signal (Eq. 1) from the victim AV. We use the receiver of our mmWave SDR testbed as the sensing system. According to the threat model, the attacker knows the frequency and waveform used by the victim mmWave radar (obtained from the radar's data sheet). To detect the signal from the victim AV, we use a correlator with the known victim chirp as a template [16]. For a 33 μ s chirp and with the receiver sampling rate of 3.072 GSamples/Sec, the number of digitized samples is 101,376. Through experiments, we prove that the attacker only needs 5000 samples of the stored template chirp to reliably detect the mmWave signal from the victim AV. Fig. 13 shows the received mmWave signal from the victim AV and the detected beginning of the signal using the correlator. With less than 5% of the samples, the signal from the victim AV is detected with high accuracy.

3) **Transmitter:** After detecting the signal from victim AV, the attacker's transmitter continuously spoofs the victim AV with the carefully designed signals described in Sec. IV. The spoofing waveform $x_{spoof}(t)$ is given by (1) Eq. 1 to add fake




Fig. 13. Detected mmWave signal from the victim AV.

obstacles, and (2) Eq. 7 and Eq. 8 to fake the locations of existing obstacles. Due to the inherent delay between sensing the victim's signal and triggering the transmitter (limited by our current testbed), the spoof signals are not sent in the same time period as the signal sensing. Instead, in our attacker system, we sense the victim AV in a time period t_1 and trigger the spoofing in a different time period t_2 as shown in Fig. 13.

C. Procedure of the Continuous Spoofing Attack

With the attacker system and its core modules described in Sec. V-B, we can explain the attack procedure. To reliably spoof an obstacle at a distance d_{spoof} , the adversary has to transmit the spoofing waveform under the delay constraints:

$$t_{AV-Attacker} + t_{sensing} + t_{delay} + t_{Attacker-AV} = \frac{2 * d_{spoof}}{c}, \quad (11)$$

where $t_{AV-Attacker}$ is the propagation delay from the victim AV to the attacker; $t_{sensing}$ is the time to detect the victim waveform; t_{delay} is the delay after which the attacker transmits the spoofing waveform; and $t_{Attacker-AV}$ is the propagation delay from the attacker to the AV. Based on Eq. 11, the attacker's distance $d_{Attacker}$ from victim AV is bounded by

$$d_{Attacker} \leq \frac{\left(\frac{2 * d_{spoof}}{c} - t_{sensing} - t_{delay} * c \right)}{2}. \quad (12)$$

The limiting factors for attacker's position include sensing time $t_{sensing}$ and delay t_{delay} . The sensing system (Sec. V-B) need anywhere between $1\mu s$ to $3\mu s$ on a state-of-the-art FPGA to sense the victim's signal. Assuming the attacker has a minimum switching speed of 10 ns , and the objective of the attacker is to spoof an obstacle at 50 meters, the distance from the attacker to the victim AV should be $d_{Attacker} \leq 38.75\text{ m}$. That means the attacker is only be able to spoof obstacles at a distance further than its physical distance from the victim AV. In order to overcome this limitation, in our system, we delay the attack by several signal frames. As a result, the attacker can be further away than the spoofed obstacle.

The attacker system works as follows. The multiple attackers are coordinated (e.g., through WiFi) but asynchronous. At the beginning of the attack, i.e., time t_0 , each of the attackers estimates the position of the victim AV using the tracking system (Sec. V-B1), and continuously tracks the position of the victim AV for subsequent time intervals. At time t_k , the attackers detect the mmWave signal from the victim AV using the sensing system (Sec. V-B2). Based on the estimated position of the victim AV and the attack goal (Sec. VI),

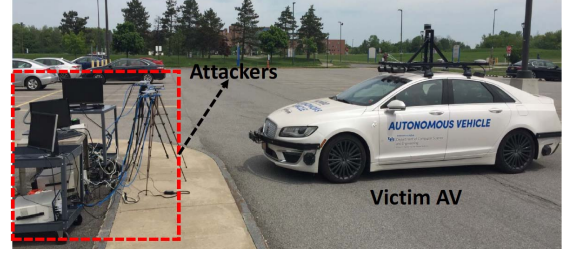


Fig. 14. Experiment set up showing the attacker mmWave testbed and the Lincoln MKZ-based AV testbed at University at Buffalo.

TABLE I
LONG RANGE RADAR PARAMETERS OF THE VICTIM AV

Parameter	Value
Sweep bandwidth	300 MHz
Ramp slope	10 MHz/us
Inter-chirp duration	3 us
Number of chirps	128
Chirp duration	30 us
Total frame time	4.224 ms
Active frame time	3.84 ms
Start frequency	62.61 GHz

the attackers transmit corresponding spoofing signals with delay τ_{delay} . The two-way propagation time due to sensing and spoofing needs to be considered when determining τ_{delay} . The attackers continuously calculate the spoofing delay τ_{delay} using the input from the tracking system. The spoofing strategy (transmitted waveform) depends on the goal of the attacker. For *AV stalling attack* (Sec. VI-1), *Hard Braking attack* (Sec. VI-2), and *Lane change attack* (Sec. VI-3), the attacker uses the strategy described in Sec. IV-A. For multiple stages attack scenarios *Multi-stage Attack* (Sec. VI-4) and *Cruise control attack* (Sec. VI-5), multiple distributed attackers employ combination of strategies described in Sec. IV-A and Sec. IV-B3.

VI. FIELD EXPERIMENTS AND SECURITY ANALYSIS

We construct five real-world driving scenarios to investigate the end-to-end security of AV under mmWave sensing spoofing attacks. We perform real-world experiments on a Lincoln MKZ autonomous vehicle testbed [30], [31], which was developed at University at Buffalo, as shown in Fig. 14. The AV is equipped with an TI IWR6843 radar [32] that uses the 60 GHz band for sensing. Table I gives the parameters of the radar. In addition to the mmWave radar, this AV testbed also has a LiDAR and several cameras. We use Baidu Apollo software [2] as victim AV's perception, planning, and navigation software. The parameters of the Baidu Apollo AV software are given in Table II in Appendix. IX-5. We perform the *software-in-the-loop* drive-by experiments using the Lincoln MKZ AV. We first spoof the Lincoln MKZ's mmWave radar. Then the detection results from the victim radar are input to the Baidu Apollo software to derive the decision of AV. Next, the driver moves the car according to the decision made by Apollo. This process is repeated for each of the time steps of the scenarios to evaluate the end-to-end security of the victim AV. Due to safety




Fig. 15. Experiment set up showing the radar mounted on the victim AV.




Fig. 16. AV stalling attack: Spoofed obstacle at 13m at multiple frames overlaid.

considerations, we performed our drive-by field experiments in an open parking lot. The AV is driven with a speed of approximately 10 *mph*. Fig. 14 shows the experiment set up with the mmWave testbed used as attackers. According to the threat model, we keep the attackers (i.e. the mmWave testbed) stationary and placed them on the roadside. The mmWave radar is mounted on the roof of the AV as shown in Fig. 15 due to cabling limitations.

1) *Scenario 1: AV Stalling Attack*: In this attack, the goal of the attacker is to spoof an obstacle in front of a stationary AV to stall it. For example, an AV waiting for the red light in a traffic junction might not move ahead when such an obstacle is detected by its radar, causing traffic confusion.

We performed experiments to emulate such a scenario. The victim Lincoln MKZ AV was at a distance of 20m from the attacker. After sensing and tracking the victim's signal as described in Sec. III-A at time t_k , the attacker transmits spoofing waveform in Eq. 1 continuously to spoof an obstacle at 13m in front of the victim AV. The location of the spoofed obstacle is given by $\{x_{spoof}(t_k), y_{spoof}(t_k)\} = (x_{victim}(t_k) + d_{spoof} * \cos(\theta)) + \epsilon_{var}, y_{victim}(t_k) + d_{spoof} * \sin(\theta) + \epsilon_{var}$. Fig. 16 shows the spoofed obstacle at 13m for consecutive frames. Fig. 17 shows the Baidu Apollo view with input from the victim radar. When the spoofed obstacle is within 13m distance, the AV chooses not to move ahead even when the traffic light turns green. In Baidu Apollo, if the spoofed obstacle is further than 15m and if there is an adjacent lane (two-lane road), the AV performs a side-pass and drive around the obstacle. In such scenario, the attacker could utilize the multiple obstacle spoofing method described in Sec. IV-A2 to simultaneously spoof multiple obstacles on both the lanes.

2) *Scenario 2: Hard Braking Attack*: In this attack, the objective of the attacker is to spoof an obstacle to force




Fig. 17. End-to-end demonstration of AV stalling attack. The AV is waiting at a red light traffic junction. The attacker spoofs an obstacle at 13m in front of the AV. The AV does not move ahead after the traffic light turns green, since it sees a spoofed obstacle.




Fig. 18. Hard braking attack: Spoofed obstacles at multiple time frames overlaid.



Fig. 19. End-to-end demonstration of AV hard braking attack. The AV is travelling at a speed of 50 km/hr. The attacker spoofs an obstacle at 27m in front of the AV. The AV's perception module detects the sudden obstacle and makes an emergency braking decision.

the victim AV make a hard braking decision thus endangering the safety of the occupants of the AV as well as other cars on the road. The attacker tracks the position pos_{victim} of the victim AV as in Sec. V-B1 and transmits a spoofing signal to spoof an obstacle near the victim AV. In a multi-lane scenario, the attacker spoofs multiple obstacles so that the victim AV's planning and control module does not initiate a sudden lane change action. We spoofed an obstacle at 27m, and the detected obstacle was fed to the AV software. To imitate an obstacle slowing down in front of the AV, the attacker gradually decreases the spoofed obstacle distance with respect to the victim AV. Fig. 18 shows the spoofed obstacle for consecutive frames and Fig. 19 shows the corresponding scenario where the victim AV is travelling at a speed of 50 km/hr. The attacker spoofs an obstacle at 27m in front of the AV. At a speed of 50 km/hr, the AV would cover 27m in 1.944 sec, which is less than the average driver reaction time of 2.3 sec. The AV initiates a hard braking decision in this case which could potentially lead to fatal situation.




Fig. 20. Illustration of the *planning* module (lane change) attack scenario. Victim AV is travelling at right lane before it changes lane due to spoofed obstacle.




Fig. 21. Lane change attack: Spoofed obstacles at multiple time frames are overlaid.




Fig. 22. The victim AV's planning algorithm charts an alternate path due to a spoofed obstacle.

3) *Scenario 3: Lane Change Attack*: The goal is to spoof the AV into making dangerous lane change decisions. We construct a scenario shown in Fig. 20. For this attack experiment, the Lincoln AV is initially at a distance of 25m from the attacker. The attacker detects the mmWave signal from the victim AV and tracks the position of the victim AV as described in Sec. V-B2 and Sec. V-B1, respectively. Upon estimating the position of the AV, the attacker spoofs an obstacle in the lane in which the victim AV is traveling. The attacker spoofs an obstacle at 40m at time t_i . To imitate the distance decreasing effect, the attacker spoofs the obstacle with decreasing distance at subsequent time frames as shown in Fig. 21. The victim AV was travelling at a speed of 30 mph. Fig. 22 shows the decision taken by the victim AV to change lane due to spoofed obstacle in its lane. As discussed in Appendix. IX-4, the cost function of the planning module decides the appropriate path decision based on various factors. Eq. 14 shows that when the spoofed obstacle greater than the collision distance threshold, the cost of the lane in which the AV is travelling increases compared to other lanes and the AVs planning module searches for other alternative path. The attacker could take advantage knowing the planning algorithm to deviate the victim AV from the planned path. Another observation we noticed is that, the AV decides to change lane only when it is within a threshold distance from the spoofed obstacle. Until it reaches the threshold distance, the AV slows down.

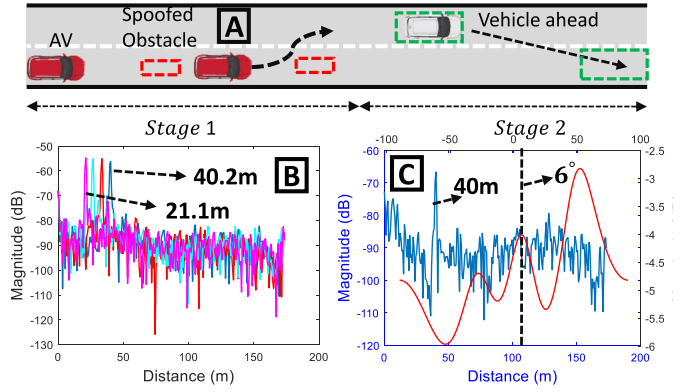


Fig. 23. Illustration of scenario 4. At stage 1, the victim AV changes lane due to a spoofed obstacle. At stage 2, the vehicle ahead is spoofed to a different location.

4) *Scenario 4: Multi-Stage Attack*: We construct a multiple-stage attack scenario by combining the intuitions and AV driving behaviour from the previous attack scenarios, which could lead the victim AV to a serious accident. We emulate a complex AV driving scenario shown in block [A] in Fig. 23, where the victim AV is initially cruising on a multi-lane road. We use distributed attack strategy discussed in Sec. IV-B3. The objective of the attacker is to force the victim AV to make a dangerous lane change (see Sec. VI-3) and subsequently crash on to a vehicle ahead of it.

As shown in Fig. 23 [A], the victim AV is cruising in lane 1 while another vehicle travels in lane 2. To force a lane change, the attacker spoofs an obstacle at distance d_{spoof} (40.2m) in front of the victim AV, as shown in block [B] in Fig. 23. The corresponding position of the spoofed obstacle at time t_i is $(x_{spoof}(t_i), y_{spoof}(t_i)) = \{x_{victim}(t_i) + d_{spoof} * \cos(\theta) + \epsilon_{var}, y_{victim}(t_i) + d_{spoof} * \sin(\theta) + \epsilon_{var}\}$. In order to consistently keep the cost of the current lane higher and to force the AV to think, the vehicle ahead is slowing down, the attacker spoofs obstacles at decreasing distance from 40.2m to 21.2m at subsequent time intervals, as shown in block [B] in Fig. 23. The position of the spoofed obstacles in subsequent time intervals are $(x_{spoof}(t_{i+1}), y_{spoof}(t_{i+1})) = \{x_{victim}(t_{i+1}) + (d_{spoof} - v_{rel} * c) * \cos(\theta) + \epsilon_{var}, y_{victim}(t_{i+1}) + (d_{spoof} - v_{rel} * c) * \sin(\theta) + \epsilon_{var}\}$, where v_{rel} is the relative velocity between the victim and the adversary.

At stage 2, two attackers perform the obstacle-deviation attack (Sec. IV-B3). The attackers spoof the location of the vehicle in lane 2 at 40m and at an angle of 6° to the victim AV, as shown in block [C] in Fig. 23. Such attack makes the vehicle ahead of the victim AV appear in a different lane than its true lane, which leads the victim AV to a potentially fatal accident. As discussed in Sec. VI-2, on the highway driving conditions, with a typical lane change speed of 45 to 50 mph, a lead vehicle at 30m that *disappears* from the victim AVs mmWave radar view could turn in to a fatal accident. In addition, as we show in the next attack scenario, when the leading vehicle *disappears* from the field-of-view, the victim AV could also accelerate to match the lane speed, causing a high speed crash.

5) *Scenario 5: Cruise Control Attack*: Traffic-Aware Cruise Control (TACC) systems can keep the AV cruising at a set

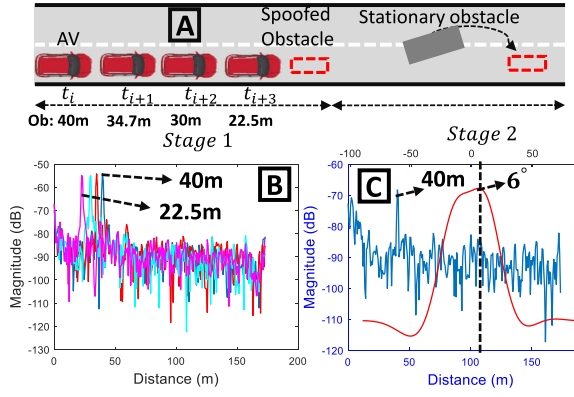


Fig. 24. Illustration of scenario 5. Victim AV decelerates and accelerates to the set speed in response to the spoofed obstacle ahead before crashing on to the parked truck.




Fig. 25. Distance/Angle of victim AV tracked by the attacker.

speed while at the same time adjusting its speed according to the traffic. Although such functionality offers comfort to the drivers, a powerful attacker can take advantage of the TACC techniques to manipulate the driving behaviour of the vehicle. For instance, the attacker could spoof the victim AV's mmWave radar into *seeing* a non-existent obstacle. Consequently, the victim AV changes its speed.

Inspired by the above discussion on TACC safety, we construct the following scenario. Instead of a real physical vehicle ahead of the victim AV, the attackers spoof a fake leading vehicle. The objective of the attack is to influence the victim AV to decelerate and accelerate in response to a spoofing attack, leading to a crash, as shown in Block [A] in Fig. 24. At the beginning, the attackers track the position of the victim AV (Sec. V-B1). Fig. 25 shows the tracked position of the victim AV. Before the attack, the victim AV *sees* no obstacle ahead of it in the decision making zone and maintains the set speed. At *stage 1*, the attacker spoofs the victim AV with an obstacle at 40m. In order to imitate a slow moving vehicle, the attacker continuously spoofs the victim AV with an obstacle at a decreasing distance. Block [B] in Fig. 24 shows the spoofed obstacles at various distances from 40m to 22.5m, which results in the victim AV decelerating as shown in Fig. 26. At *stage 2*, the attacker stops the spoofing attack, which makes the AV start accelerating as it *sees* no obstacle in front of it. In addition, the attackers employ the attack strategy provided in Sec. IV-B3 to fake the location of the stationary obstacle. The goal is to let the stationary obstacle




Fig. 26. Velocity of victim AV tracked by the attacker.

appear out of the way of the victim AV. The attackers spoof the fake location of the stationary obstacle at 40m with respect to the victim AV and at an angle of 6° , as shown in block [C] in Fig. 24, which results in the stationary obstacle deviating by 4m from the victim AV's lane. Fig. 26 shows the speed timeline of the victim AV in the experiment. The victim AV accelerates and reaches a steady velocity before it decelerates due to the spoofed obstacle. At around 15 sec, the victim AV accelerates again at which point the attackers spoof the stationary obstacle out of its way. The victim AV would have crashed on to the stationary obstacle at this point in time, leading to a fatal situation.

VII. DEFENDING STRATEGIES

To defend the above fatal attacks, we propose two defending strategies to reliably detect spoofing attacks on AVs.

A. Challenge-Response

Unlike LiDAR pulses that can be randomized to mitigate LiDAR spoofing attack, mmWave radar uses a well defined waveform designed to meet specific sensing capabilities [33]. Such predictability empowers attackers to perform the attacks developed in this paper. The attackers can eavesdrop on the victims waveform and deduce its parameters to perform the attacks. Alternatively the attacker could simply record the signal from the victim AV and replay it at a later time to launch the attacks. To address this problem, departing from the traditional fixed waveform design, we propose to use a radar frame consisting of chirps with varying parameters, yet meet the design criteria. The transmitted chirp in Eq. 1 is:

$$x(t) = \cos \left(2\pi f_{start}t + \pi \frac{B}{T_{chirp}}t^2 + \phi_{init}(n) \right), \quad (13)$$

where n is the index of the n^{th} chirp. The parameters f_{start} , $\frac{B}{T_{chirp}}$, and ϕ_{init} can be randomized across chirps or frames, thus mitigating the spoofing attacks. To that end, we introduce the *challenge-response* authentication for mmWave radar, where the AV radar transmits chirp with randomized parameter as a challenge. The response, i.e., the received signal at the radar, is verified for authenticity. Fig. 27 shows the AV transmitting a *challenge* with chirps having random parameters p_1 and p_2 . The parameter p could be phase ϕ_{init} or f_{start} .

In the experiment, we used ϕ_{init} and f_{start} as two parameters. For the method with ϕ_{init} as a random parameter, the AV radar transmits a sequence of chirps in Eq. 13 with a random phase ϕ_{init} for each transmitted chirp n . The attacker




Fig. 27. Illustration of *challenge-response* spoofing detection where the victim AV transmits a challenge with parameters p_1 and p_2 and the reflected signal is used as a response.




Fig. 28. Spoofed obstacle is *smeared* without distinct peak.

transmits a spoofing waveform with $\phi_{init} = 0$. Fig. 28 shows the radar detection of the victim AV. The attacker's spoofing signals fail to generate a peak. Instead, it appears as a smearing. In the second *challenge-response* method, the AV varies f_{start} randomly. Therefore, the attacker fails to register any meaningful spoofing effect on the victim AV.

B. Waveform Fingerprinting

Recently RF fingerprinting has gained attention in its ability to detect spoofing attacks. However, there are two challenges to utilize the RF fingerprinting techniques in radar. First, the uniqueness and non-linearity associated with the components of the transmitter hardware result in the RF features that are unique to each device. However, the radar transmitter shown in Fig. 1 is significantly simpler than communication devices, which leads to smaller feature set. Second, the received signal in radar is a superposition of the reflections from all the targets and the attackers. Those signals undergo two-way path loss in addition to reflection loss due to the reflecting object.

To address the challenges, we find that the statistical characteristics of the AVs mmWave signal is influenced by the attackers waveform, which give us the opportunity to detect the attacks. Hence, we propose to use the statistical features, including: *standard deviation*, *kurtosis*, *skewness* of magnitude and phase of the received signal [34]. The amplitude of the received signal at radar is influenced by the two-way propagation path of the transmitted signal. One of the key requirement of RF fingerprinting is that the wireless channel influences need to be eliminated. Therefore, the received signal in Eq. 3 is rms-normalized before feature extraction. The rms-normalization of Eq. 3 is $r_{rms}[n] = \frac{r[n]}{\sqrt{\sum_{n=1}^N r[n]^2}}$, where N is the total number of samples in a received chirp.




Fig. 29. 2-Dimensional representation of the features.

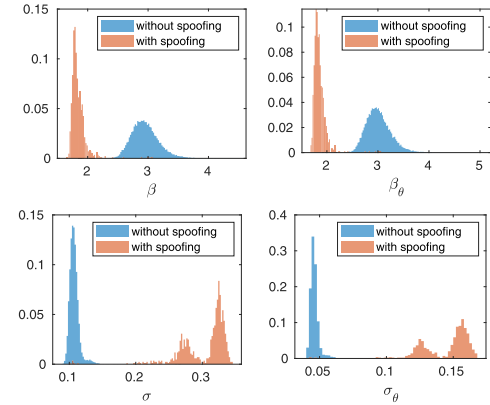


Fig. 30. Histogram of the received signal w/o spoofing.

The rms-normalized signal is then used to extract RF features. Since the AV mmWave signal does not use any specific preamble, we use the entire chirp of 256 samples to extract the features. Fig. 30 shows the distribution of the features when there is no spoofing attack and when the AVs mmWave radar is under spoofing attack. We can clearly see that the spoofing attacks change the distributions of the proposed features. As a result, the attacks can be reliably detected.

The AV software needs to detect spoofing attacks from the received signals. Only the training signals from the AV's own radar can be used. Therefore, we propose to use *one-class SVM* as the machine learning framework to identify spoofing attack [35]. We use 3000 received signals without spoofing attacks as training data. The testing data comprises 3000 signals under spoofing attack. Fig. 29 shows the 2-D embedding of the features. We can see distinct grouping between the features with (red points) and without (blue points) spoofing attack. The spoofing detection accuracy is 98.9%.

VIII. DISCUSSION

A. How Quick Can Attacker Spoof

In our mmWave testbed, the attacker senses the mmWave signal from the victim AV at frame t and spoofs at a future frame $t + k$. This is mainly due to the inherent delay of the testbed. The spoofing response time of the attacker depends on multiple factors, including the sampling rate of the DAC, the delay due to the cable connecting the DAC module and the baseband FPGA module, and sampling rate of the involved FPGAs. Better delay resolution can be obtained by using modules with higher sampling rate, and or implementing

the attack system as a single chip solution (e.g., system-on-chip) with guaranteed delay requirements. However, currently, the available state-of-the-art commercial mmWave testbed limits such implementations.

B. Effect of CFO

It should be noted that the testbed we used (see Sec. V-A) introduces CFO to the transmitted spoofing signals. CFO between the attacker and the victim AV radar manifests as an arbitrary Doppler in the received signal. The attacker testbed we currently use does not have any implicit control of CFO. In the experiments, we eliminate the impacts of CFO by calibration. In practical scenarios, to mitigate the impacts of CFO, the precise control of carrier frequency is needed; alternatively, the CFO between the attacker and victim must be estimated and compensated.

IX. CONCLUSION

In this paper, we perform the end-to-end security analysis of mmWave radars in AVs by designing and implementing practical physical layer attack and defense strategies in real-world scenarios. Two core attacking strategies are rigorously designed and validated, which can reliably add spoofing obstacles or fake the locations of existing obstacles. Based on the state-of-the-art mmWave and AV testbed, the attackers combine the core attacking strategies to continuously spoof the AV into making hazardous safety-critical driving decisions, leading to fatal accidents. To defend such attacks, we propose two defense strategies and prove the defending solutions can detect the sensor spoofing attacks with very high accuracy.

ACKNOWLEDGMENT

The authors would like to thank Roman Dmowski from the Department of Computer Science and Engineering at University at Buffalo for assisting in the autonomous vehicle field experiments.

APPENDIX A MODULES OF BAIDU APOLLO SYSTEM

1) *ROI Filter*: Region of interest (ROI) is the region that lies in the AV's driving path, which directly influences the AV's driving decisions. The detected objects should lie within the ROI, while objects outside ROI are discarded.

2) *Object Tracking*: Once the obstacles are filtered based on ROI, they are tracked using the Kalman filter.

3) *Track Matching*: Objects sensed by the sensors in the current sensing cycle are associated to the objects sensed in the previous cycle using data matching algorithms. Apollo Baidu matches the objects detected in a current radar frame to the existing tracks using gated Hungarian matching algorithm. When the object in the current radar frame does not match any existing tracks, new track is created for that object.

TABLE II
PARAMETERS USED BY THE APOLLO BAIDU'S PERCEPTION, PLANNING, AND NAVIGATION MODULES

Parameter	Value	Description
min-lane-change-length	5m	minimum distance to change a lane
min-lane-change-prepare-length	60m	minimum distance to prepare a lane change
follow-min-distance	3m	minimum distance to follow an obstacle
min-stop-distance-obstacle	6m	minimum stop distance from in-lane obstacle
max-stop-distance-obstacle	10m	maximum stop distance from in-lane obstacles
lane-change-prepare-length	80m	distance to prepare for lane change
min-lane-change-prepare-length	10m	minimum distance to prepare for lane change
min-nudge-distance	0.2m	minimum distance to nudge
max-nudge-distance	1.1m	maximum distance to nudge
min-yield-distance	5m	minimum distance to yield

4) *Path Planning*: AVs typically use a search algorithm with a certain cost function to decide the best path that an AV follows [36]. In Apollo Baidu, the cost function is $C_{total} = C_{env} + C_{obstacle}$ [36]. The obstacle cost $C_{obstacle}$ is given by

$$C_{obstacle} = \begin{cases} 0 & d > d_n, \\ C_{nudge}(d - d_c) & d_c \leq d \leq d_n, \\ C_{collision} & d < d_c. \end{cases} \quad (14)$$

5) *Baidu Apollo Parameters*: Important parameters used in Baidu Apollo AV software are summarized in Table II.

REFERENCES

- [1] T. Litman, *Autonomous Vehicle Implementation Predictions*. Victoria, BC, Canada: Victoria Transport Policy Institute Victoria, 2017.
- [2] [Online]. Available: <https://github.com/ApolloAuto/apollo>
- [3] S. Kato *et al.*, "Autoware on board: Enabling autonomous vehicles with embedded systems," in *Proc. ACM/IEEE 9th Int. Conf. Cyber-Phys. Syst. (ICCPs)*, Apr. 2018, pp. 287–296.
- [4] *openPilot: Open Source Driving Agent*. [Online]. Available: <https://github.com/commaai/openpilot>
- [5] A. Singhvi and K. Russell, *Inside the Self-Driving Tesla Fatal Accident*. Accessed: Jul. 12, 2016. [Online]. Available: <https://www.nytimes.com/interactive/2016/07/01/business/insidetesla-accident.html>
- [6] *Tesla in Fatal California Crash Was on Autopilot*. Accessed: Mar. 31, 2018. [Online]. Available: <https://www.bbc.com/news/world-us-canada-43604440>
- [7] *Rear-End Collision Between a Car Operating With Advanced Driver Assistance Systems and a Stationary Fire Truck*, Culver City, California, January 22, 2018. Accessed: 2015. [Online]. Available: <https://dms.ntsb.gov/public/62500-62999/62683/627686.pdf>
- [8] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote attacks on automated vehicles sensors: Experiments on camera and LiDAR," in *Proc. Black Hat Eur.*, vol. 11, 2015, p. 2015.
- [9] H. Shin, D. Kim, Y. Kwon, and Y. Kim, "Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications," in *Proc. Int. Conf. Cryptograph. Hardw. Embedded Syst.* Cham, Switzerland: Springer, 2017, pp. 445–467.
- [10] Y. Cao *et al.*, "Adversarial sensor attack on LiDAR-based perception in autonomous driving," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Nov. 2019, pp. 2267–2281.
- [11] J. Sun, Y. Cao, Q. Chen, and Z. Mao, "Towards robust lidar-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures," in *Proc. 29th USENIX Secur. Symp. (USENIX)*, 2020, pp. 877–894.
- [12] R. Chauhan, "A platform for false data injection in frequency modulated continuous wave radar," M.S. thesis, Utah State Univ., 2014. [Online]. Available: <https://digitalcommons.usu.edu/etd/3964>
- [13] C. Yan, W. Xu, and J. Liu, "Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle," *DEF CON*, vol. 24, p. 109, Aug. 2016.

- [14] *Introduction to the NI mmWave Transceiver System Hardware*. Accessed: Feb. 4, 2020. [Online]. Available: <https://www.ni.com/en-us/innovations/white-papers/16/introduction-to-the-ni-mmwave-transceiver-system-hardware.html>
- [15] Y. Cao *et al.*, "Adversarial objects against LiDAR-based autonomous driving systems," 2019, *arXiv:1907.05418*. [Online]. Available: <http://arxiv.org/abs/1907.05418>
- [16] A. Ranganathan, B. Danev, A. Francillon, and S. Capkun, "Physical-layer attacks on chirp-based ranging systems," in *Proc. 5th ACM Conf. Secur. Privacy Wireless Mobile Netw. (WISEC)*, 2012, pp. 15–26.
- [17] R. Chauhan, R. M. Gerdes, and K. Heaslip, "Demonstration of a false-data injection attack against an FMCW radar," in *Proc. 12th escar Eur.*, 2014, p. 135.
- [18] N. Miura, T. Machida, K. Matsuda, M. Nagata, S. Nashimoto, and D. Suzuki, "A low-cost replica-based distance-spoofing attack on mmWave FMCW radar," in *Proc. 3rd ACM Workshop Attacks Solutions Hardw. Secur. Workshop (ASHES)*, 2019, pp. 95–100.
- [19] S. Jha *et al.*, "Kayotee: A fault injection-based system to assess the safety and reliability of autonomous vehicles to faults and errors," 2019, *arXiv:1907.01024*. [Online]. Available: <http://arxiv.org/abs/1907.01024>
- [20] S. Jha *et al.*, "ML-based fault injection for autonomous vehicles: A case for Bayesian fault injection," in *Proc. 49th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, Jun. 2019, pp. 112–124.
- [21] A. H. M. Rubaiyat, Y. Qin, and H. Alemzadeh, "Experimental resilience assessment of an open-source driving agent," in *Proc. IEEE 23rd Pacific Rim Int. Symp. Dependable Comput. (PRDC)*, Dec. 2018, pp. 54–63.
- [22] C. Sitawarin, A. N. Bhagoji, A. Mosenia, M. Chiang, and P. Mittal, "DARTS: Deceiving autonomous cars with toxic signs," 2018, *arXiv:1802.06430*. [Online]. Available: <http://arxiv.org/abs/1802.06430>
- [23] A. Chernikova, A. Oprea, C. Nita-Rotaru, and B. Kim, "Are self-driving cars secure? Evasion attacks against deep neural networks for steering angle prediction," in *Proc. IEEE Secur. Privacy Workshops (SPW)*, May 2019, pp. 132–137.
- [24] Y. Jia *et al.*, "Fooling detection alone is not enough: Adversarial attack against multiple object tracking," in *Proc. Int. Conf. Learn. Represent.*, 2019, pp. 1–15.
- [25] A. Boloor, K. Garimella, X. He, C. Gill, Y. Vorobeychik, and X. Zhang, "Attacking vision-based perception in end-to-end autonomous driving models," *J. Syst. Archit.*, vol. 110, Nov. 2020, Art. no. 101766.
- [26] Y. Shoukry, P. Martin, Y. Yona, S. Diggavi, and M. Srivastava, "PyCRA: Physical challenge-response authentication for active sensors under spoofing attacks," in *Proc. 22nd ACM SIGSAC Conf. Comput. Commun. Secur.*, Oct. 2015, pp. 1004–1015.
- [27] M. Schneider, "Automotive radar-status and trends," in *Proc. German Microw. Conf.*, 2005, pp. 144–147.
- [28] *Wr-15 Waveguide Horn Antenna Operating From 50 GHz to 75 GHz With a Nominal 42 DBI Gain With UG-385/U Round Cover Flange*. [Online]. Available: <https://www.pasternack.com/images/ProductPDF/PE9881-42.pdf>
- [29] A. Abdelaziz, C. E. Koksai, and H. El Gamal, "On the security of angle of arrival estimation," in *Proc. IEEE Conf. Commun. Netw. Secur. (CNS)*, Oct. 2016, pp. 109–117.
- [30] *CAVAS: UB's Connected and Autonomous Vehicle Applications and Systems*. [Online]. Available: <https://ubwp.buffalo.edu/cavas/introduction-lincoln/>
- [31] R. Dmowski *et al.*, "Research at the intersection of big data and connected and autonomous vehicles," Univ. Buffalo, Tech. Rep., 2019. [Online]. Available: http://www.buffalo.edu/content/www/transinfo/Research/transportation-operations/research-at-the-intersection-of-big-data-and-connected-and-autonomous-vehicles/_jcr_content/par/download/file.res/TransInfo_CAV_Final_Report_lb1.pdf
- [32] *Ti IWR6843 Single-Chip 60-to 64-GHz mmWave Sensor (Rev. B)*. [Online]. Available: <https://www.ti.com/document-viewer/IWR6843/datasheet/description-x5103#%x5103>
- [33] V. Dham, "Programming chirp parameters in TI radar devices," Appl. Rep. SWRA553, Texas Instrum., Dallas, TX, USA, 2020.
- [34] K. Joo, W. Choi, and D. H. Lee, "Hold the door! Fingerprinting your car key to prevent keyless entry car theft," 2020, *arXiv:2003.13251*. [Online]. Available: <http://arxiv.org/abs/2003.13251>
- [35] R. Chalapathy, A. K. Menon, and S. Chawla, "Anomaly detection using one-class neural networks," 2018, *arXiv:1802.06360*. [Online]. Available: <http://arxiv.org/abs/1802.06360>
- [36] H. Fan *et al.*, "Baidu apollo EM motion planner," 2018, *arXiv:1807.08048*. [Online]. Available: <http://arxiv.org/abs/1807.08048>



Zhi Sun (Senior Member, IEEE) received the B.S. degree in telecommunication engineering from the Beijing University of Posts and Telecommunications (BUPT), Beijing, China, in 2004, the M.S. degree in electronic engineering from Tsinghua University, Beijing, in 2007, and the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, USA, in 2011. He was a Post-Doctoral Fellow with the Georgia Institute of Technology from 2011 to 2012. He joined the Department of Electrical Engineering, University at Buffalo, The State University of New York, Buffalo, NY, USA, as an Assistant Professor, in 2012. He is currently an Associate Professor with the University at Buffalo. His research interests include physical-layer security, wireless communication and networking in extreme environments, metamaterial enhanced communication and security, wireless intra-body networks, wireless underground networks, wireless underwater networks, and cyber physical systems. He was a recipient of the Outstanding Graduate Award at Tsinghua University in 2007, the BWN Researcher of the Year Award at the Georgia Institute of Technology in 2009, the Best Paper Award in IEEE Globecom in 2010, the NSF CAREER Award in 2017, the UB Exceptional Scholar–Young Investigator Award in 2017, and the Best Demo Award in IEEE Infocom 2017. He serves as an Editor for IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS and *Computer Networks* journal (Elsevier).



Sarankumar Balakrishnan (Student Member, IEEE) received the B.E. degree in electronics and communication engineering from Pondicherry University, India, in 2010, and the M.S. degree in electrical engineering from the National University of Singapore, Singapore, in 2013. He is currently pursuing the Ph.D. degree with the Department of Electrical Engineering, University at Buffalo, The State University of New York, under the supervision of Prof. Zhi Sun. He was the Associate Scientist of Temasek Laboratories, Singapore, focusing on digital beamforming and wireless system design using software defined radios. He also held an internship at the U.S. Research Center, Wireless Technology Group, Sony Corporation of America, focusing on mmWave wireless networks. His current research interest includes physical layer security for mmWave wireless networks.



Lu Su (Member, IEEE) received the M.S. degree in statistics and the Ph.D. degree in computer science from the University of Illinois at Urbana–Champaign in 2012 and 2013, respectively. He has worked with the IBM T. J. Watson Research Center and the National Center for Supercomputing Applications. He is currently an Associate Professor with the School of Electrical and Computer Engineering, Purdue University. He has published more than 100 articles in referred journals and conferences. His research interests include the Internet of Things and cyber-physical systems, with a current focus on wireless, mobile, and crowd sensing systems. He is a member of ACM. He was a recipient of the NSF CAREER Award, the University at Buffalo Young Investigator Award, the ICCPS 2017 Best Paper Award, and the ICDCS 2017 Best Student Paper Award. He serves as an Associate Editor for *ACM Transactions on Sensor Networks*.



Arupjyoti (Arup) Bhuyan (Senior Member, IEEE) received the Ph.D. degree in engineering and applied sciences from Yale University. He is currently a Wireless Researcher with the Idaho National Laboratory (INL) and the Technical Director of the INL Wireless Security Institute. His research interest includes the secure implementation of future generations of wireless communications with scientific exploration and engineering innovations across the fields of wireless technology, cybersecurity, and computational science. His specific goals are to assure communications among critical infrastructure systems supporting control of the electric grid, emergency response, and nationwide unmanned aerial systems. Before he joined INL in October 2015, he has an extensive industry experience in wireless communications.



Pu Wang (Member, IEEE) received the B.Eng. degree in electrical engineering from the Beijing Institute of Technology, China, in 2003, the M.Eng. degree in electrical and computer engineering from the Memorial University of Newfoundland, Canada, in 2008, and the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, in August 2013, under the guidance of Prof. Ian F. Akyildiz. He is currently an Assistant Professor with the Department of Computer Science, University of North Carolina

at Charlotte. Prior to joining UNCC, he was an Assistant Professor with the Department of Electrical Engineering and Computer Science, Wichita State University, from 2013 to 2017. His current research interests include modeling, analysis, and stochastic optimization of networked systems, with applications in software defined networking, cyber-physical systems, the Internet of Things, wireless sensor networks, cognitive radio networks, and electromagnetic nanonetworks.



Chunming Qiao is currently a SUNY Distinguished Professor and the Chair of the CSE Department, University at Buffalo. He has been leading the Laboratory for Advanced Network Design, Evaluation and Research (LANDR), UB, since 1993. His current research interests cover not only the safety and reliability of various cyber physical systems (such as transportation systems with connected and autonomous vehicles, critical infrastructures involving power grid and communications networks, and cloud services), but also algorithms and protocols for the Internet of Things, including smartphone-based systems and applications. He has published extensively with an H-index of more than 75 (according to Google Scholar). Several of his papers have received the Best Paper Awards from IEEE and Joint ACM/IEEE venues. He also has seven U.S. patents and has been serving as a Consultant for several IT and Telecommunications companies since 2000. His research has been featured in *BusinessWeek*, *Wireless Europe*, *CBC*, and *New Scientists*. He was elected to an IEEE Fellow for his contributions to optical and wireless network architectures and protocols.