

SoK: A Minimalist Approach to Formalizing Analog Sensor Security

Chen Yan^{*§}, Hocheol Shin^{†§}, Connor Bolton^{‡§}, Wenyuan Xu^{*¶}, Yongdae Kim^{†¶}, and Kevin Fu^{‡¶}

^{*}Zhejiang University, {yanchen, wyxu}@zju.edu.cn

[†]KAIST, {h.c.shin, yongdaek}@kaist.ac.kr

[‡]University of Michigan, {mcbolto, kevinfu}@umich.edu

Abstract—Over the last six years, several papers demonstrated how intentional analog interference based on acoustics, RF, lasers, and other physical modalities could induce faults, influence, or even control the output of sensors. Damage to the availability and integrity of sensor output carries significant risks to safety-critical systems that make automated decisions based on trusted sensor measurement. Established signal processing models use transfer functions to express reliability and dependability characteristics of sensors, but existing models do not provide a deliberate way to express and capture security properties meaningfully.

Our work begins to fill this gap by systematizing knowledge of analog attacks against sensor circuitry and defenses. Our primary contribution is a simple sensor security model such that sensor engineers can better express analog security properties of sensor circuitry without needing to learn significantly new notation. Our model introduces transfer functions and a vector of adversarial noise to represent adversarial capabilities at each stage of a sensor’s signal conditioning chain. The primary goals of the systematization are (1) to enable more meaningful quantification of risk for the design and evaluation of past and future sensors, (2) to better predict new attack vectors, and (3) to establish defensive design patterns that make sensors more resistant to analog attacks.

I. INTRODUCTION

Cyber-physical systems such as airplanes [1], [2], autonomous vehicles [3], [4], and medical devices depend on the trustworthiness of measurement from trillions of embedded sensors for critical decision making [5]. A key research problem is how to build security into the analog circuitry of sensors to protect against denial of service and damage to the integrity of sensor outputs.

Today, manufacturers operate in a more reactive stance, closing sensor security holes one by one after each is discovered. We believe that a more effective approach is to mitigate security risks by design rather than reaction. Our work systematizes past papers on analog sensor security by providing a simple sensor security model based on transfer functions that map to the components within a sensor’s signal conditioning chain. The model offers a way for sensor engineers to more deliberately and concisely write down security requirements and limitations concerning analog security risks to sensors. This process helps an engineer to more quickly identify the security limitations of sensor design and to have a way to debate the effectiveness of various defenses.

[§]Co-first authors.

[¶]Corresponding authors.

Since our goal is to improve sensor design, we build our work on a wide community of publications yet focus on investigating how adversaries may intentionally induce *untrustworthy sensor output* (measurement) via analog attacks, i.e., those that employ analog signals to manipulate sensor output. Notice that attacks within this scope exploit a sensor’s implementation rather than the overall system incorporating that sensor. Thus, in this paper we do not consider the following attacks: (1) attacking the digital transmission of sensor data, e.g., on a vehicular CAN bus [6]–[9] and in sensor networks [10]–[14]; (2) utilizing sensors to fingerprint a device [15]–[18], invade privacy [19]–[26], or trigger malicious behaviors [27]–[30]; (3) modifying the measurand (the target of measurement), e.g., using a dummy finger, to fool sensors or the systems [31]–[40].

We adopt the term *transduction attacks* [41] to refer to analog attacks in our scope, i.e., those where victim sensor circuitry transduces an attacker’s malicious physical signals to analog ones. Transduction attacks are diverse in adversarial signal modality (e.g., light or acoustic waves), application, and vulnerable hardware with examples including: acoustic waves manipulating a drone’s gyroscope [42]–[44], lasers creating false points in an autonomous vehicle’s lidar [45], [46], and radio frequency waves injecting false audio into smart-device microphones [47], [48]. Existing work categorizes a set of attacks that partially includes transduction attacks [49], [50] or uses transfer functions to describe the security properties of analog-to-digital converters (ADCs) at high-frequency inputs [51], but engineers still lack a simple language to compare how well sensors defend against transduction attacks.

We propose a model to simplify transduction attack analysis based in (1) mathematical models of sensor physics and (2) abstracted methods to exploit these mathematical sensor models. Sensors may be modeled as a series of transfer functions, with a separate function for each sensor component in the signal processing chain. Each transduction attack exploits at least one transfer function. We find that one can abstract and categorize these methods to exploit the physics of sensor circuits based on two categories of basic steps. Each step describes a method to exploit a transfer function and is an abstraction above the physical component level. Thus, different transduction attacks may share some basic steps that exploit the same vulnerability in transfer functions regardless of sensor components or sensor type. This mathematical abstraction simplifies comparison of attacks across different sensors and signal modalities, as many

attacks may be described as a chain of five or fewer steps.

Additionally, the model enables defense abstraction across different types of sensors. Since a successful attack depends on the full chain of steps, the key to defending an attack is to mitigate at least one step in the attack chain. Within this model, each step and the mitigation to the step are all mathematically abstracted over sensor types and components.

Lastly, the cross-sensor analysis provided by the model enables rudimentary prediction of transduction attacks and defenses. A new attack is theoretically possible if one can construct a new series of steps. More importantly, a sensor designer may implement defenses for each known step at the design time. As such, even if an attacker discovers a new step, she could not easily construct a successful attack as existing defenses against other known steps may mitigate the attack.

Contributions. Our primary contributions pertain to a simple collection of formalisms to better express the analog security of sensor circuitry:

- **A simple sensor security model.** Building upon established notation and transfer functions from the signal processing community, our model introduces a vector of intentional, malicious noise to express how well a sensor design can resist various transduction attacks. The model enables more meaningful quantification of risk for the design and evaluation of past and future sensors.
- **Formalisms to help predict new attack vectors.** The small number of formalisms help a sensor engineer to capture the security properties of sensor circuitry. Our hope is that by empowering the engineer with a few simple mathematical tools and engineering steps, engineers will make fewer mistakes against known analog vulnerabilities as well as future unknown vulnerabilities.
- **Defensive design patterns.** By systematizing the knowledge of attacks against sensors, our work helps to establish defensive design patterns that make sensors more resistant to analog attacks.

II. SENSOR BACKGROUND

Existing studies on sensor security focus on analyzing the security of one or a few sensor types. Systematizing sensor attacks and defenses is challenging due to the heterogeneity of sensors. For example, there are more than 370 types of sensors on record [52] that rely on dozens of conversion phenomena for measurement [53]. In this section, we show that despite great diversity, sensors share components and properties that facilitate security exploits in general.

A. Sensors

A sensor is a device that outputs usable measurement in response to a specific measurand [54]. We present the signal conditioning chain of typical sensors in Fig. 1, where a sensor is represented as an interconnection of essential electronic components. Sensors transform a physical stimulus (input) to an analog intermediate and finally to a digital representation (output). We introduce the signal conditioning chain in the following.




Fig. 1: A general signal conditioning chain of sensors. Signals flow from left to right through each component and transform from the physical stimulus (input) to an analog intermediate and finally to a digital representation (output). Depending on the specific design, variations to this schematic may include multiple amplifiers or filters, no filters, filters before the transducer (e.g., CMOS) or amplifier, other circuits (e.g., comparators), etc.

Stimulus and Measurand. The measurand is a quantity that a sensor intends to measure, and a stimulus is a physical signal involved in measuring the measurand. For example, an accelerometer’s measurand and stimulus are acceleration and force, respectively. A thermocouple’s measurand and stimulus are temperature and heat. Despite the wide variety of stimuli, we have two measurement methods:

- Passive sensors* passively accept physical stimuli and do not emit external stimuli. For example, microphones are passive sensors that capture sound from the environment. The stimuli of passive sensors, e.g., light, sound, force, and chemicals, are generated by other objects in the environment or already exist.
- Active sensors* emit physical stimuli to an environment and actively measure the response after the stimuli’s interaction with the environment. For example, ultrasonic sensors/lidars measure the distance to objects by emitting ultrasound/lasers and receive the reflection. Active sensors are often used to measure quantities of tangible objects, such as obstacle distance, rotation speed, and liquid drop. For simplicity, we do not show the emitter in Fig. 1.

Transducer. Commonly a sensor’s first component, a transducer produces an analog electrical representation of the measurand by measuring a physical stimulus. Transducer heterogeneity, even for the same measurand or physical stimuli, is the primary source for sensor diversity. For example, dynamic, condenser, and piezoelectric microphones all can capture sound, but they rely on entirely different conversion phenomena [55].

Analog Signal Processing Circuits. Typically, a sensor must process a transducer’s analog signals to reduce noise while amplifying useful information. Standard components include amplifiers to increase the signal amplitude, filters to remove noise, envelop detectors, comparators, etc.

ADC. An analog-to-digital converter (ADC) digitizes analog signals for digital processing, storage, etc.

Note that a sensor may not contain all components shown in Fig. 1. For instance, some sensors do not have filters by design. Nevertheless, Fig. 1 represents a simplified yet functionally comprehensive structure of a modern sensor.

B. Common Properties for Sensor Exploits

Sensitive to Physical Signals. By design, sensors are sensitive to the target physical stimulus, even if the stimulus is unintended for measurement. This effect guarantees that at least one type of physical signal will affect the sensor.

Similar Analog Signal Processing. Analog signal processing circuits often remain similar despite transducer heterogeneity. For example, sensors commonly use amplifiers and filters, even if designed to measure different phenomena. Thus, exploits on similar signal processing circuits may remain similar even on different sensors, e.g., microphones and thermocouples.

Same Signal Modalities. There are three signal modalities for the signal conditioning chain: physical, analog, and digital, as shown in Fig. 1. The shared signal modalities show that the same signal properties in each modality may be exploited for attacks across various sensors.

Chain of Blind Trust. Sensors are essentially proxies of reality. Most sensor designs use a series of electric components to approximate the measurand. Typically, each component blindly assumes its input is valid. However, this blind trust can allow malicious signals to exploit components in the signal conditioning chain without detection.

III. ATTACK OVERVIEW AND SCOPE OF THE STUDY

This research lies on analog sensor security, which focuses on the trustworthiness of sensor measurement under the threat of analog attacks. Trustworthiness refers to whether the sensor measurement reflects reality. Within this scope, we use the term transduction attacks [41] to indicate sensor-related attacks that fall into the scope.

A. Transduction Attack

A transduction attack exploits vulnerabilities in the physics of a sensor to manipulate its output. In particular, an attacker generates malicious physical signals that are transduced to malicious analog signals in the sensor circuitry, either explicitly through transducers or implicitly through other components in the sensor. The types of malicious physical signals include but are not limited to the following items.

- **Electromagnetic radiation** refers to the waves of an electromagnetic (EM) field that propagate through space. It includes radio frequency (RF) waves, infrared, (visible) light, ultraviolet, X-rays, and gamma rays.
- **Sound** is a vibration that propagates as a wave of pressure through mediums, including gas, liquid, and solid. It includes audible sound, ultrasound, and infrasound.
- **A magnetic field** is created by magnetized materials and by moving electric charges (currents) such as those used in electromagnets.
- **An electric field** is generated by particles that bear electric charges in any form.

An attacker may generate malicious physical signals of any modality for transduction attacks, regardless of whether the signal is of the same type as the intended stimulus by design.

For example, attackers can use infrared (IR) to attack a lidar (which measures obstacles with IR) or RF signals to attack a microphone (which measures sound).

Out-of-scope: Sensor design vulnerabilities enable transduction attacks to produce untrustworthy sensor measurements. For example, attacks can cause a lidar to detect non-existing obstacles or a microphone to report non-existing sounds. Our intention is to improve the security of sensor design. Thus, we do not consider the attacks that intentionally modify the measurand. For instance, attackers can use a hairdryer to heat a thermocouple such that the temperature measurements are higher than the ones of the distant environment. However, the measurements, though manipulated, still reflects the temperature near the sensor. Similarly, we do not consider fake fingers for fingerprint sensors [56], IR decoy flares for infrared homing missiles [57], and melamine adulteration of milk for measurement of nitrogen content [58], since all involve measurand manipulation. Mitigating these attacks normally requires alternative defense mechanisms of the systems rather than sensor design, such as extra functionality (e.g., liveness detection, object recognition, and protein measurement) or fusing multiple sensors for the intended application. In short, we focus on systematizing academic knowledge on transduction attacks from the view of sensor designers.

B. Attacker Objectives

We consider two adversarial objectives.

Denial-of-service (DoS). The goal is to prevent a sensor from acquiring usable measurements. For instance, a strong acoustic signal can cause unreliable, seemingly random gyroscope output if the signal's frequency is close to the gyroscope's resonant frequency. Thus, such a signal may prevent proper flight for drones relying on gyroscope output [42].

Spoofing. The goal is to trick a sensor into providing seemingly legitimate but erroneous measurements. For example, for an RC car controlled by a phone's gravitational orientation, malicious acoustic signals can induce false acceleration measurements and control the RC car's movement while the phone remains stationary [59].

C. Threat Model

In this paper, we consider adversaries with the following assumptions.

Analog Attacks. An adversary focuses on affecting the analog signals in a sensor and does not interfere with digital measurement processing or transmission.

Sensor Assessment. We assume an adversary cannot tamper with victim sensors but can obtain similar sensors for assessment. The adversary may reverse engineer sensor design parameters, such as operational frequencies, bandwidth, signal format, etc., and explore vulnerabilities.

Attack Range. Transmission power generally bounds effective attack range, but an adversary may extend the range by emitting stronger physical signals at extra costs. Thus, we focus on other aspects contributing to attack feasibility rather than strictly range.

IV. TRANSDUCTION ATTACK SYSTEMATIZATION

Existing approaches to defending sensor security tend to follow an endless cat and mouse game—security researchers find a physics-based exploit, then manufacturers deploy an exploit-specific patch rather than create an overarching and measurable security goal to address the root causes of that specific exploit. The lack of a measurable, goal-oriented approach to analog sensor security makes it difficult to apply science to defensive design.

Thus, our model seeks a balance between the salient security properties of sensors while requiring minimal additional cognitive effort by established sensor experts. To achieve this, our model consists of a simple adaptation to a well-established language from the signal processing community.

A. Simple Sensor Security Model Overview

To avoid creating overly complicated notation and terminology, our systematization of knowledge on sensor security seeks to find a minimal amount of new formalisms sufficiently powerful for a descriptive model to characterize how secure existing and future sensor systems are to transduction attacks. An underlying goal is to ensure the model does not simply describe past attacks and defenses, but is predictive of the future and helps sensor engineers to better build measurable security into sensor circuitry.

The key to identifying transduction attacks is how to articulate the signal conditioning chain in sensors. Given that a transfer function [60] is a well-established concept in signal processing community to model system input and output relationships [61]–[63], our model extends this approach by encapsulating the notion of malicious interference as an additional vector of noise that maps to each stage of a sensor’s signal conditioning chain and related support circuitry. Our systematization of knowledge addresses the challenge of how to capture security characteristics in the existing representation of transfer functions by modeling transduction attacks as a vector of malicious noise.

Our simple sensor security model for transduction attacks consists of a series of cascading transfer functions followed by a formulaic set of basic steps that are sufficient to describe the predictive space of imaginable future transduction attacks against sensor security.

Specifically, two categories of basic steps comprise the main ingredients of present and future transduction attacks:

- (i) *Signal injection steps* describe key methods to inject malicious analog signals into sensors by emitting malicious physical signals and are vital to all transduction attacks.
- (ii) *Measurement shaping steps* describe key methods to shape injected analog signals to benefit the adversary and may be optional depending on the target sensor and attacker goal.

B. Simple Sensor Security Model

In this section, we introduce the traditional sensor model and base the formalization of transduction attacks on the traditional model (Fig. 2). Table I summarizes established notation and our minimal additions to capture security properties.

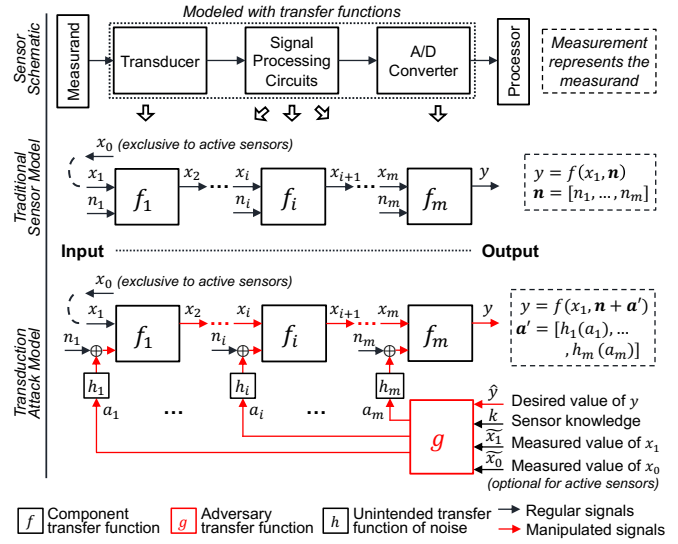


Fig. 2: A generalized sensor schematic (top), the traditional and established way of modeling a sensor (middle), and our adaptation to model a transduction attack (bottom). The signal processing community models a sensor as cascading transfer functions. Each transfer function (e.g., f_i) has two inputs—the legitimate signal x_i and the noise signal n_i . The output signal x_{i+1} becomes the input to the next transfer function f_{i+1} . x_1 is the stimulus to the transducer (f_1), and it is either pre-existing or transformed from a probing signal x_0 generated by the sensor if it is an active one. The noise n_i may come from the electrical circuit or the external environment. The transduction attack model consists of a simple adaptation of established practices by representing analog security exploits as a vector of intentional noise $[a_1, a_2, \dots, a_m]$ that leads the output y to an attacker-desired value of \hat{y} .

1) *Sensor Model:* Our approach models each sensor component as a transfer function, which characterizes its input-output relationship. Transfer functions are normally expressed in the Laplace domain for analysis of the dynamic response [60], [64] and in the time domain for analysis of the static response [61], [65], [66]. Here, we use the time-domain representation of transfer functions, because transduction attacks usually damage or exploit the static characteristics of sensors, e.g., accuracy and nonlinearity.

As shown in the middle of Fig. 2, traditional approaches characterize the i th sensor component as a mathematical transfer function f_i :

$$x_{i+1} = f_i(x_i, n_i) \quad (1)$$

where the inputs consist of a legitimate signal x_i and a noise signal n_i , and the output signal is x_{i+1} . Note that x_{i+1} is also the input to the $(i + 1)$ th component. All signals are time-varying, and here we omit the time variable t in $x_i(t)$ for simplicity. The noise n_i is an unwanted disturbance to x_i and may come from the electrical circuit itself (electronic noise) or the external environment (coupled noise). In practice, the noise signal n_i and legitimate signal x_i are mixed and cannot be separated easily [67]. Our conceptual model serves to emphasize that noise inputs can affect the outputs of the components and introduce errors to the measurement.

Examples. We first review a few examples of established transfer functions when the noise is additive, which is the most widely used noise model.

TABLE I: NOTATION.

Established notation from the traditional sensor modeling	
f_i, f	The transfer function of the i th component and the sensor
m	The number of components modeled as transfer functions
x_i, \mathbf{x}	The legitimate input to the i th component and their vector
x_0	The physical stimulus generated by active sensors
n_i, \mathbf{n}	The noise input to the i th component and their vector
y	The sensor measurement (at the ADC)
h_i	The unintended transfer function of external noise
Notation introduced by transduction attacks	
a_i, \mathbf{a}	The malicious physical signals emitted for the i th component and their vector
a'_i, \mathbf{a}'	The malicious noise signals injected into the i th component and their vector
g	The adversary transfer function
\hat{y}	The sensor measurement desired by the attacker
k	Knowledge of the sensor model
$\widetilde{x}_1, \widetilde{x}_0$	The real measurement of x_1 and x_0 (optional for active sensors)

- (i) A linear transfer function is defined as $x_{i+1} = c_0 + c_1(x_i + n_i)$, where c_0 is the intercept and c_1 is the slope. The linear relationship is expected for many components in their ideal states, such as position transducers and amplifiers [61].
- (ii) Nonlinear relationships apply to most components in practice. Common transfer functions include [61]: logarithmic function $x_{i+1} = c_0 + c_1 \ln(x_i + n_i)$ (e.g., photodiodes), exponential function $x_{i+1} = c_1 e^{k(x_i + n_i)}$ (e.g., thermistors), and power function $x_{i+1} = c_0 + c_1(x_i + n_i) + c_2(x_i + n_i)^2$ (e.g., silicon resistive sensors). Higher-order polynomial functions may be employed for other cases.

As shown in the middle of Fig. 2, traditional signal processing formalizes the transfer function of a sensor as the cascading components' functions:

$$y = f_m(\cdots f_2(f_1(x_1, n_1), n_2) \cdots, n_m) \quad (2)$$

where the output y is the sensor measurement and m is the number of components. The inputs to a sensor include the physical stimulus x_1 at the transducer and the noise signals at each component. Let a vector of the noise signals be $\mathbf{n} = [n_1, n_2, \dots, n_m]$ and the overall transfer function of the sensor be f , Eq. (2) simplifies to:

$$y = f(x_1, \mathbf{n}) \quad (3)$$

2) *Transduction Attack Formalization:* Although noise reduction with signal processing circuits is a common practice, sensor designers typically do not expect the intentional noise injected by transduction attacks. The goal of a transduction attack is to modify a sensor's measurement y by injecting intentional noise to the legitimate physical or analog input x_i . Formally, we define transduction attacks as follows:

Definition 1. A transduction attack alters the measurement of a sensor to approach an attacker-desired value \hat{y} by injecting intentional interference to the sensor. The measurement under a transduction attack can be represented as

$$y = f(x_1, \mathbf{n} + \mathbf{a}') \quad (4)$$

where $\mathbf{a}' = [a'_1, a'_2, \dots, a'_m] = [h_1(a_1), h_2(a_2), \dots, h_m(a_m)]$ is a vector of malicious noise signals injected by the attacker

at each of the m components, and $\mathbf{n} + \mathbf{a}' = [n_1 + a'_1, n_2 + a'_2, \dots, n_m + a'_m]$. h_i is the unintended transfer function of noise that is hidden before the i th component and models the propagation and coupling of external noise. We omit them from the traditional sensor model for simplicity.

We demonstrate the attack at the bottom of Fig. 2. In our model, the injected signal a'_i is of the same type as the original noise n_i . The combination of the injected signal a'_i and noise n_i may affect the output of components and the sensor measurement eventually, because subsequent signal conditioning components blindly trust inputs. A vector of the malicious physical signals emitted by the attacker can be formalized by an *adversary transfer function* g :

$$\mathbf{a} = [a_1, a_2, \dots, a_m] = g(\hat{y}, k, \widetilde{x}_1, \widetilde{x}_0) \quad (5)$$

where \hat{y} is the desired sensor measurement, k is the knowledge of the sensor model, \widetilde{x}_1 is the measured stimulus from the measurand, and \widetilde{x}_0 is the measured stimulus from the victim sensor if it is active. Independently, attackers may measure both x_1 and x_0 . Knowledge of the existing stimulus x_1 is necessary to achieve accurate control over the sensor's measurement, and measuring the stimulus x_0 emitted by active sensors is optional, e.g., in spoofing sensors such as lidars [45] and ultrasonic sensors [68].

3) *Case Study—Microphone Audio Spoofing:* We demonstrate how to model transduction attacks with transfer functions using a high-level example from previous work [47]. In this simplified example, an adversary's goal is to cause a spoofed sound measurement \hat{y} in the audible frequency range (20 Hz to 20 kHz) at a microphone's output in a quiet environment ($x_1 \approx 0$, $\mathbf{n} \approx \mathbf{0}$). To achieve this goal, the adversary exploits the conducting wire at the input of an amplifier (f_2) as an unintended transducer (h_2), which can unintentionally convert the adversary's electromagnetic signal a_2 to an electrical noise a'_2 via electromagnetic coupling modeled by h_2 (acting as an antenna). h_2 determines the optimal frequency (f_{a_2}) of the electromagnetic signal a_2 . Given the nonlinear property of an amplifier's transfer function, the adversary modulates the desired output \hat{y} on a carrier of f_{a_2} and recovers \hat{y} at the amplifier's output through intermodulation distortion [69]. The demodulated signal may survive in the signal conditioning chain and spoof the microphone's output as desired.

4) *Discussion:* Without transfer functions, designing a viable physical signal for transduction attacks is empirical and requires trial and error. With the transfer functions, we envision that the simple sensor security model facilitates the design of malicious physical signals \mathbf{a} that can modify the sensor measurements to a given value and quantifying the effectiveness. In particular, given the desired sensor measurement \hat{y} , an attacker can derive \mathbf{a} by mathematically solving an optimization problem¹. We can quantify the effectiveness

¹An attacker may acquire the transfer functions from public datasheets or by conducting sensor assessment. If accurate modeling of a transfer function is challenging, e.g., the unintended EM coupling effect, approximate solutions are applicable. Though inaccurate transfer functions may degrade our model, they seldom impair the effectiveness of attacks, especially when an attacker does not demand a highly accurate control over the sensor output.

of an attack by the expected error, i.e., $e = |\tilde{y} - \hat{y}|$, where \tilde{y} is the real sensor measurement. An attacker’s goal is to minimize the expected error e by optimizing the emitted physical signals.

a. Conversely, sensor designers tasked with threat mitigation may consider methods to attenuate undesirable signals injected through each transducer, including unintended transducers. In addition, designers should validate that real components match their theoretical behavior, even for abnormal input.

To construct a transduction attack, the simple sensor security model indicates two categories of basic steps.

- (i) The signal injection steps involve determining the necessary characteristics of the malicious physical signals that can inject noise to the signal conditioning path efficiently, including the injection point and signal parameters.
- (ii) In the measurement shaping steps, we consider how an attacker can construct physical signals such that the injected analog signals (“injected signals” hereafter) can shape the sensor measurement to desired values.

In the following, we elaborate on the two categories of steps.

C. Signal Injection Steps

To successfully and efficiently inject malicious analog signals into a sensor, an attacker must consider (1) the injection point, i.e., before or after the intended transducer, and (2) which signal types, amplitudes, and frequencies of the emitted physical signals are most efficient.

1) *Injection Point and Signal Type*: Since transducers are by design the only component that intends to accept external physical inputs, we divide the injection point into pre-transducer (at or before the intended transducer) and post-transducer (after the intended transducer). The injection point determines the required signal type.

Pre-Transducer. Attackers may exploit how transducers are designed to be sensitive to at least one type of physical signal (Section II-B). Most existing transduction attacks exploit injection points at or before the intended transducers. Intuitively, the malicious physical signal a_1 can be the same type of signal as the legitimate stimulus x_1 so that they are inherently noise inputs without conversion. For example, Petit et al. [46] and Shin et al. [45] succeeded in attacking lidars with light. Yan et al. [70] managed to jam and spoof ultrasonic sensors with ultrasound. Attackers can also use acoustic waves to attack sensors that measure movement vibrations, e.g., MEMS gyroscopes [42]–[44] and accelerometers [43], [59].

Post-Transducer. Components after the intended transducer may act as unintended transducers, converting noise or an attacker’s physical signal a_i into an electrical signal inside the sensor despite not being designed to do so. For example, wires in a circuit may act as unintended antennas by converting electromagnetic waves into analog electrical signals in the sensor via inductive or capacitive coupling [71]. Since the injected signal a'_i is usually weak after conversion, injection points before the pre-amplifier are typically preferred so that the injected signal may have a more dominant effect on the output than the legitimate input x_i .

Rasmussen et al. [72] first suggested adversarial usage of wires being unintended antennas, but as a possible pathway for wormhole attack [73] rather than in the context of transduction attacks. Foo Kune et al. [47] first exploited this effect for a transduction attack by injecting a voice signal into a Bluetooth headset. This attack employed intentional electromagnetic interference (EMI) on the conducting wire between the microphone and amplifier. Later work employed electromagnetic coupling from outside [48] and inside [74] a smartphone to inject malicious signals to its microphone, and compromised actuator control signals [75]–[77].

Attackers must consider how injection using a wave or field may affect several sensor components simultaneously, as it may complicate the attack in terms of the needed measurement shaping steps. For example, sound can vibrate all sensor components besides the transducer, and electromagnetic injection may induce currents at many parts of a circuit. However, as we will discuss later, an attacker with knowledge of the sensor may choose the malicious physical signal’s frequency to affect targeted components more than others.

2) *Efficient Injection and Signal Frequency*: The effectiveness and efficiency of signal injection depend on the physical signal’s form, especially the amplitude and frequency.

Amplitude. An injection is likely to be more effective with physical signals of higher amplitudes. Shorter attack distances or increased transmission power may result in higher physical signal amplitudes at the target sensor due to the power-attenuation law for physical signals, e.g., sound and electromagnetic waves [78]. However, in practice, one cannot easily increase the power indefinitely. For example, a high-power radio frequency signal, laser, or sound may cause damage to human bodies [79]. Inaudible ultrasound at high power may create audible byproducts [80], [81] due to nonlinear acoustics [82], [83], which makes inaudible attacks audible.

Frequency. Resonant frequencies widely exist for many mechanical and electrical components, e.g., MEMS transducers [84] and wires (antennas) [85]. At these frequencies, the injected signal can reach relative maximum intensities. An attacker may increase the efficiency of injection at the resonant frequencies of the target injection point. A standard method to find the resonant frequency is to conduct a frequency sweep and search for the highest amplitude response [42], [43], [47], [59], [86]. In addition to using a resonant frequency, an attacker may choose specific frequencies to satisfy the special requirement of attacks, e.g., to be hidden. For example, attacks that employ ultrasound (> 20 kHz) are inaudible to humans. Similarly, injecting infrared to cameras makes the attack invisible. Depending on the frequency, signals can be divided into two categories: in-band and out-of-band.

3) *In-Band and Out-of-Band Signals*: Signals with frequencies within the intended frequency band are in-band signals, and otherwise are out-of-band signals. For example, ultrasound is out-of-band for microphones that are expected to record audible sounds between 20 Hz and 20 kHz. Most MEMS accelerometers are expected to capture movement (mechanical vibration) under 750 Hz [59]; therefore, frequencies in the

kHz range may be out-of-band. Typically, injecting an out-of-band signal can have two problems. (1) Out-of-band signals do not directly interfere with ordinary inputs because they are at different frequencies. For example, one cannot discern ultrasonic noises in audible sounds. (2) Out-of-band signals are often attenuated by filters designed to remove signals outside the expected frequency band. For an injected out-of-band signal to affect measurement, an attacker will require measurement shaping steps to convert the injected signals to the expected frequency range, i.e., in-band.

D. Measurement Shaping Steps

Signal injection alone may be insufficient to meet all attacker goals; thus, an attacker may have to take additional steps to shape the injected analog signals to DoS or spoof the sensor. A common example being how an attacker may use step(s) to shape out-of-band signals to in-band for sensor spoofing. To shape the signal between transduction and digitization, an attacker may exploit specific properties of sensor components by properly designing the malicious physical signals. We define measurement shaping steps as these exploits of sensor component properties that shape the injected signals. In the following, we elaborate on how one can group several common measurement shaping steps across a wide variety of sensor types and existing studies, and show how these may be mathematically abstracted to fit into our model.

1) *Saturation (Sat.)*: Saturation is a common phenomenon in analog electronics referring to how a quantity cannot exceed an upper or lower bound [61]. For example, as shown in Eq. (6), an amplifier may become saturated when the input is beyond a threshold. In this case, the input increment no longer proportionally increases the output, which leads to clipping:

$$f_i(x_i, n_i + a'_i) = \begin{cases} c_1 \mathcal{A}(x_i, n_i + a'_i) & \text{if } \mathcal{A}(x_i, n_i + a'_i) \leq k \\ \text{const} & \text{if } \mathcal{A}(x_i, n_i + a'_i) > k \end{cases} \quad (6)$$

where $\mathcal{A}(x_i, n_i + a'_i)$ denotes the intensity of combined x_i and $n_i + a'_i$, c_1 is the amplification factor, and k is the saturation point. For amplifiers, the clipping voltage is normally determined by the power supply. Similar effects can happen to other components such as transducers, ADCs, etc. Saturation is undesirable, and sensors are designed to operate below the saturation point. However, an attacker can intentionally saturate a component by injecting a strong interference a'_i . In this way, the adversary may mask legitimate input and DoS a sensor [45], [46], [70], [87], or let in a DC signal component for spoofing when there was none [59].

2) *Intermodulation Distortion (IMD)*: IMD [69] can occur when a signal with two or more frequencies passes through a nonlinear component. For example, amplifiers, diodes, and transducers are generally known to be nonlinear; even ADCs show some level of inherent nonlinearity due to internal amplifiers [51]. IMD forms cross-products at new frequencies that are not present in the input signals. Specifically, the output signals include the sum and difference of the input frequencies.

For example, consider a nonlinear transfer function in a simple 2nd-order power series:

$$x_{i+1} = c_0 + c_1(x_i + n_i + a'_i) + c_2(x_i + n_i + a'_i)^2 \quad (7)$$

Suppose the mixed signals of $x_i + n_i + a'_i$ contain two frequencies, f_1 and f_2 ($f_1 > f_2$). The output x_{i+1} of this nonlinear transfer function contains frequencies at f_1 , f_2 , $f_1 - f_2$, $f_1 + f_2$, $2f_1$, $2f_2$, and a constant offset. Note that $f_1 - f_2$ may be below the original frequencies. An attacker can exploit IMD to convert malicious out-of-band signals to in-band, e.g., demodulating amplitude modulated (AM) signals [88]. Note that in radio receivers, IMD is a desired effect by design for down-converting signals to intermediate frequencies, such as in frequency mixers [89].

An attacker may exploit any sensor component that is characterized by a nonlinear transfer function for IMD. For example, Foo Kune et al. [47] utilized amplifier nonlinearity to recover a baseband voice from the injected electrical signals coupled from an RF carrier. Similarly, other studies [90]–[92] managed to recover voice commands from ultrasound by exploiting nonlinear microphones.

3) *Envelope Detection (Env.)*: Diodes and capacitors are essential in many circuits, especially for electrostatic discharge protection [93], [94]. However, they can also act as simple envelope detectors that demodulate AM signals. Foo Kune et al. [47] found several capacitor-diode pairs before a microphone's amplifier that could demodulate the injected signals.

4) *Aliasing (Ali.)*: According to the Nyquist-Shannon sampling theorem [104], if the frequency of a sampled signal is higher than half of the sampling rate, the signal will be indistinguishable from signals of other frequencies [105]. For example, if the sample rate of an ADC is F_s , then a signal at frequency f will have the same sampling result as a signal at frequency $F_s - f$. This effect is known as aliasing and typically should be avoided in signal processing. However, an adversary may exploit aliasing to convert the malicious out-of-band signals to in-band frequencies after the ADC. For example, Trippel et al. [59] and Tu et al. [43] managed to control the output of ADCs in MEMS accelerometers or gyroscopes by tuning the amplitude, frequency, or phase of the injected signals. Foo Kune et al. [47] managed to demodulate the injected signals after ADC by setting the carrier frequency equal to the sample rate.

5) *Filtering (Fil.)*: Ideally, filters before an ADC should remove all out-of-band signals and prevent aliasing. However, in practice it is difficult to manufacture a filter that can remove all out-of-band frequencies based on the designed cut-off frequencies while passing all in-band frequencies. Instead, there is a range of frequencies around the cut-off frequency that is attenuated but not completely removed. For example, lower-order filters have a wider transition range where signals remain only partially attenuated [89]. An attacker could exploit this property to design signals that pass the filter, but following components cannot handle properly. Trippel et al. [59] found many low-pass filters in MEMS accelerometers that show large transition ranges, and thus these filters do not sufficiently

TABLE II: SYSTEMATIZATION OF TRANSDUCTION ATTACKS WITH THE SIMPLE SENSOR SECURITY MODEL.

Application	Sensor		Exploited Component					Signal Injection			Measurement Shaping					Outcome		Paper	
	Type	C.	Trans.	Wire	Amp.	Filter	ADC	Point	Type	Freq.	Sat.	IMD	Fil.	Env.	Ali.	DoS	Spoof		
Automobile	Lidar	A	●	○	○	○	○	Pre	☀	In	●	○	○	○	○	●	○	[45] [45], [46]	
	Camera	P	●	○	○	○	○	Pre	☀	In	●	○	○	○	○	●	○	[46], [70]	
	Radar	A	●	○	○	○	○	Pre	📶	In	○	○	○	○	○	○	○	[70] [70], [95]	
	Ultrasonic Sensor	A	●	○	○	○	○	Pre	🔊	In	○	○	○	○	○	○	○	[68], [70] [68], [70]	
	Magnetic Encoder	A	●	○	○	○	○	Pre	🧲	In	○	○	○	○	○	○	○	[96], [97]	
Drones or Smart Devices	Optical Flow Sensor	P	●	○	○	○	○	Pre	☀	In	○	○	○	○	○	○	○	[98]	
	MEMS Gyroscope	P	●	○	○	●	●	Pre	🔊	Out	○	○	●	○	○	○	○	[42], [43] [43], [44], [99]	
	MEMS Accelerometer	P	●	○	●	●	●	Pre	🔊	Out	○	○	●	○	○	○	○	[59], [43] [59], [43], [99] [59]	
	Microphone	P	●	●	●	●	●	Post	📶	Out	○	○	●	○	○	○	○	○	[47] [47], [48]
								Pre	🔊	Out	○	○	○	○	○	○	○	○	○
	Touchscreen	A	●	○	○	○	○	Pre	⚡	N/A	○	○	○	○	○	○	○	[103]	
Hard Disk	MEMS Shock Sensor	P	●	○	○	○	○	Pre	🔊	Out	○	○	○	○	○	○	○	[86]	
Energy	Infrared Sensor	P	○	●	○	○	○	Post	📶	Out	●	○	○	○	○	○	○	[75], [76]	
	Pacemaker Lead	P	○	●	○	○	○	Post	📶	In	○	○	○	○	○	○	○	[47]	
Defibrillator Lead	Pre							☀	In	●	○	○	○	○	○	○	○	○	○

☀ Visible light or infrared 📶 RF waves 🔊 Audible sound or ultrasound 🧲 Magnetic field ⚡ Electric field ● Applicable ○ Probable ○ Not applicable
 C. Category A Active sensor P Passive sensor Pre Pre-transducer Post Post-transducer In In-band Out Out-of-band N/A Not available

attenuate higher-frequency signals that affect sensor output. From the defense point of view, note that though high-order filters may hinder the exploit of aliasing, they can also mask the trace of high-frequency malicious signals if the exploit happens before these filters, e.g., IMD or saturation.

E. Constructing a Transduction Attack

An attacker can build a transduction attack based on the chaining explorations of various signal injection steps and measurement shaping steps, which are determined by the target sensor and the desired attack outcome. We summarize the basic steps that have been exploited by existing transduction attacks in Table II. In the design of a transduction attack, an attacker can construct malicious physical signals by examining the possible steps as malicious signals go through the signal conditioning chain of a sensor. To this end, explicit knowledge of the sensor components and their transfer functions, i.e., the simple sensor security model, is necessary.

V. DEFENSE SYSTEMATIZATION

We divide defense mechanisms against transduction attacks into two broad categories: detection and prevention. Simply, detection mechanisms detect the presence of an attack while prevention mechanisms ensure proper or trustworthy sensor output even in the presence of an attack. With these two broad categories, we analyze and systematize existing defenses in terms of the aforementioned transfer functions and injection/measurement shaping steps. Table III summarizes our defense systematization with examples from previous work.

A. Detection Methods

Detection methods only detect attacks and do not ensure proper sensor output if attacked. However, they can be a starting point of more robust system-wide defenses, e.g., a trigger for activating other preventive measures or a fail-safe mode. Detection methods do not modify any existing transfer functions, but can be modeled as an additional binary function, χ , as shown in Eq. (8), which takes $\mathbf{x}' = [x_1, x_2, \dots, x_m, y(=x_{m+1})]$, \mathbf{n} , and \mathbf{a}' as inputs (See Table I).

$$\chi(\mathbf{x}', \mathbf{n} + \mathbf{a}') = \begin{cases} 1 & \text{if it is an attack} \\ 0 & \text{if it is not an attack} \end{cases} \quad (8)$$

Some detection methods may also include an active sensor's generated stimulus x_0 in addition to the basic parameters of \mathbf{x}' , \mathbf{n} , and \mathbf{a}' . This simple model represents a detection scheme based on the status (intermediate in/output signals and noise levels) of each component. Detection methods may be further categorized by which of an attack's step(s) the method defends.

1) *Detecting Signal Injection Steps:* Currently, three types of detection methods detect the exploitation of injection steps.

Detection by TX Randomization. Adopting randomness to an active sensor's transmission (TX) can often assist in detecting transduction attacks. Essentially, sensors with random transmission schemes may detect attacks when an adversary should be unable to predict the random pattern. Often, the sensor transmits a randomized physical stimulus and checks that the randomness is intact in the received stimulus. Any major changes to the signal indicate the presence of an attack. To include signal transmitters, Eq. (8) has to be expanded to embrace the transmitted stimulus, such that $\chi(\mathbf{x}', \mathbf{n} + \mathbf{a}'; r(x_0))$.

TABLE III: SYSTEMATIZATION OF TRANSDUCTION ATTACK DEFENSES

Goal	Cat.	Subcat.	Related Component					Injection and Shaping Steps					Xfer-Func.	Paper					
			TX	Trans.	Wire	Amp.	Fil.	ADC	Dig. ²	Point	Freq.	Sat.			IMD	Fil.	Env.	Ali.	
Detection	Inject.	TX randomiz.	●	○	○	○	○	○	●	Both	Both	○	○	○	○	○	N/A	[68], [87], [106]	
		Verif. Actuation	●	○	○	○	○	○	●	Both	Both	○	○	○	○	○		[47], [107]	
	Detect OOB Sig.	○	●	○	○	○	○	●	Pre	Out	○	○	○	○	○	[44], [86]			
	Saturation	○	●	○	○	○	○	●	N/A	N/A	○	○	○	○	○	[45], [70], [87]			
Prevention	Rand.Shap.	IMD Features	○	○	○	○	○	○	●	N/A	N/A	○	●	○	○	○	P1,P2	[90], [91]	
		TX Randomiz.	●	○	○	○	○	○	●	Both	Both	○	○	○	○	○		[45], [46], [68]	
	Shielding	RX Randomiz.	○	○	○	○	○	○	●	●	N/A	N/A	○	○	○	○	●	P1	[43], [59], [70], [98]
		Physical Barrier	○	●	●	●	●	●	●	Both	Both	○	○	○	○	○	P2	[42], [44], [47], [86], [87]	
		Spatial ASR ³	○	●	○	○	○	○	○	Pre	In	○	○	○	○	○	P1	[45]	
		Temporal ASR	○	●	○	○	○	○	○	Pre	In	○	○	○	○	○	P1	[46]	
		Spectral ASR	○	●	○	○	○	○	○	Pre	Out	○	○	○	○	○	P1,P2	[46]	
		LPF/BPF/HPF	○	○	○	○	○	○	○	N/A	N/A	○	○	●	○	○	●	P1,P2	[47], [59], [90]
	Filt.	Adaptive Filt.	○	○	○	○	○	○	○	Both	Both	○	○	○	○	○	P2	[42], [44], [47], [86], [90]	
		Out-of-phase Samp.	○	○	○	○	○	○	○	N/A	N/A	○	○	○	○	○	P1	[59]	
		Fusion	Spatial Fusion	●	●	●	●	●	●	●	Steps differ case by case					P3	[44], [45], [68], [70], [86]		
	Spectral Fusion		●	●	○	○	○	○	○	P1,P3						[46]			
	Temporal Fusion		○	○	○	○	○	○	○	P1						[46], [68], [98]			
	Comp. Quality Improv.		●	●	●	●	●	○						P1	[42], [44], [59]				

¹ Denotes the three xfer func. models of Section V-B. ² Digital Backend ³ Attack Surface Reduction ● Applicable ○ Not applicable

Here, x_0 represents the physical stimulus generated by active sensors originally, and $r(\cdot)$ denotes the randomizing function. A prime example of a TX Randomization Detection scheme is a time-based randomization scheme first shown by Shoukry et al. [106]. In the scheme, the sensor randomly ceases all stimulus transmission. Then, any received stimulus during this pause indicates the presence of an attack. While attackers with higher capabilities may still overcome this scheme [108], it would greatly increase attack difficulty. Other work has proposed the same time-based randomization scheme [87] for medical infusion pumps. In the same line, Xu et al. [68] proposed physical shift authentication to detect transduction attacks on automotive ultrasonic sensors by randomizing several waveform parameters.

Verifying Actuation. A system with both sensor(s) and actuator(s) may detect attacks by probing the surroundings periodically. Essentially, the actuator delivers a probing signal to the surrounding environment and compares the measured response with an expected response. A vast difference between expected and measured signals indicates the presence of an attack. The transfer-function notation is similar to that of detection by TX randomization, in $\chi(\mathbf{x}', \mathbf{n} + \mathbf{a}'; t(x_0))$ form, where x_0 indicates the preset probing signal and $t(\cdot)$ denotes the function that converts the probing signal to the environmental response. Prior work has suggested various forms of verifying actuation detection schemes. Foo Kune et al. [47] proposed the adoption of a cardiac probe that cross-checks whether cardiac signal readings coincide with the expected values after some sort of investigative actuation to the cardiac tissue. In the same line, Muniraj et al. [107] suggested an active detection method, which detects the attack based on “judiciously designed excitation signals” superimposed on the control signal.

Detecting Out-of-band Signals. Defenders may come up with an additional receiver that detects targeted out-of-band signals. For example, as suggested by prior work [44], [86], adopting an additional microphone can detect resonating sound

against MEMS sensors. Since sensors require only in-band stimuli to function, detecting out-of-band signals do not affect sensors’ functionality, and there can be many other variations according to the targeted out-of-band signals.

2) *Detecting Measurement Shaping Steps:* Previous work describes how to detect certain measurement shaping steps.

Saturation Detection. Several previous studies suggest saturation detection as a defense [45], [70], [87]. Saturating a component leaves the component in an abnormal state, that may be easily detectable with hardware or software support. The saturation detection function, χ_{sat} , monitors if the input of a vulnerable component exceeds a threshold, e.g., a voltage level. Assuming the i th component is saturated, χ_{sat} can be modeled as a logical inequality as shown in Eq. (9), where $\mathcal{A}(x_i, n_i + a'_i)$ denotes the intensity of combined x_i and $n_i + a'_i$, and ε is the saturation threshold.

$$\chi_{sat}(\mathbf{x}', \mathbf{n} + \mathbf{a}') = \mathcal{A}(x_i, n_i + a'_i) > \varepsilon \quad (9)$$

Detecting IMD Features. Studies have shown that attacks exploiting intermodulation distortion (IMD) for signal demodulation could leave identifying features in analog signals. Zhang et al. [90] proposed to search for features of IMD demodulated voice signals by the intensity at high frequencies (500 Hz to 1 kHz), and Roy et al. [91] suggested to search for signal correlation in the sub-50 Hz band. These features are introduced by IMD and may not be easily erased.

B. Prevention Methods

Prevention methods ensure proper sensor output even in the presence of a transduction attack, generally by attenuating malicious signals either inside or outside of the sensor. Prevention methods can roughly be modeled as three types:

P1: Component Modification. A defender modifies an existing component to reduce an attacker’s ability to exploit that function. Assuming the vulnerable i th component (f_i) is improved (f'_i), these defenses can be expressed as below.

$$|f'_i(x_i^{adv}, n_i + a'_i) - x_{i+1}| \ll |f_i(x_i^{adv}, n_i + a'_i) - x_{i+1}| \quad (10)$$




Fig. 3: Example of fused transfer-function chains. N chains with different numbers of internal components (m_1 through m_N) are synthesized together by the combining function \mathcal{C} to produce an intermediary output x_c . This intermediary output then undergoes M more transfer functions (f_{c1} through f_{cM}) to yield the final combined output y_c .

where x_i^{adv} denotes the input to f_i under attack and x_{i+1} is the output of f_i without an attack. Thus, modification ($f_i \rightarrow f'_i$) reduces an attack's effect on output (Eq. (10)).

P2: Component Addition. A defender inserts a new component to reduce the effect of the attack on subsequent components. In terms of transfer function representation, this type of defenses can be represented as below.

$$|f'(x_1, \mathbf{n} + \mathbf{a}') - y| \ll |f(x_1, \mathbf{n} + \mathbf{a}') - y|, \quad (11)$$

where f' is a new transfer function of the sensor obtained by adopting the new component, and y is the sensor output without attack in Fig. 2.

P3: Component Fusion. A defender fuses multiple chains of components, either full chains or parts of chains, to produce a new combined output. Fig. 3 illustrates an example of combining chains. N chains of transfer functions are synthesized together to produce an intermediary output x_c , which is then processed by additional M transfer functions to yield the final output y_c . Throughout the combining process, the effect of the attack becomes suppressed in the final output, y_c .

In the remainder of this section, we introduce various prevention methods built on Component Modification, Component Addition, and Component Fusion.

1) *Shielding*: Shielding mitigates injection steps by reducing exposure to external physical signals. It can be a simple solution to mitigate transduction attacks, but occasionally required hardware changes may be inadequate [86] or too costly to implement. Shielding typically adds additional hardware and thus corresponds to Component Addition.

Physical Barriers. A defender may add situation-specific barriers to attenuate external physical signals, e.g., using a Faraday cage to block electromagnetic signals. However, by design some sensors must be exposed to the outer environment, and thus physical barriers may not always be applicable. For example, a defender cannot block a lidar from lasers as the lidar needs to sense echoes of its own transmitted lasers to function. Previously discussed physical barriers include shielding conducting wires [47], acoustic damping [42], [44], [86], and optical shielding [87].

Attack Surface Reduction. Reducing the attack surface area selectively limits transducer exposure to external physical

stimuli, thus increasing injection difficulty while allowing the transducer to remain exposed to intended physical stimuli. Examination of previous work shows that these reduction methods can be classified further into spatial, temporal, and spectral categories. Spatial attack surface reduction attempts to confine transducer exposure only to the direction of the physical stimuli to measure. Spatial methods are especially relevant to sensors that take readings from their field of view piece by piece. For example, Shin et al. [45] suggested increasing the directivity of internal receivers/transmitters and removing curved reception glass in lidars. Similarly, temporal reduction limits the duration of transducer exposure, and spectral reduction limits the bandwidth of stimulus the transducer is exposed to. Petit et al. [46] proposed limiting lidar reception time (temporal reduction) and filtering out unwanted light frequencies (spectral reduction). Unlike spatial reduction, temporal reduction would correspond to Component Modification because reducing the effective duration can be implemented without additional hardware. Likewise, Spectral reduction may also be implemented by Component Modification when reducing the bandwidth to which the transducer is exposed is feasible without adopting new hardware.

2) *Filtering*: Filtering aims to attenuate malicious analog signals within the sensor without attenuating the legitimate signals. These defenses are suitable for sensors that focus only on a specific part of the analog signal, e.g., a specific frequency band. In terms of transfer-function representation, filtering defenses correspond to Component Modification (when improving existing filters) and Component Addition (when an additional filter is adopted).

LPF/HPF/BPF. A defender can use low/high/band-pass filters to attenuate frequency bands containing only noise or malicious signals. In addition, they may also mitigate IMD and aliasing by blocking frequencies that possibly induce such phenomena. Previous work suggests a variety of relevant applications for these filters [47], [59], [90].

Adaptive Filtering. A defender may be able to use adaptive filtering to attenuate injected signals when simple low/high/band-pass filters are inapplicable. In the context of transduction attacks, adaptive filtering methods typically find some reference of a malicious signal and then use this reference to filter it out. For example, a defender may augment a microphone with an additional wire to clearly receive an attacker's electromagnetic wave signal, and then use this reference to filter the malicious analog signal on the sensing path [47]. Alternatively, differential signaling can be employed to cancel out the injected signal [42], [47]. Other work has also employed adaptive filtering in additional contexts [44], [86].

Out-of-phase Sampling. A defender may make an analog-to-digital converter (ADC) adopt a special sampling pattern related to the frequency to which the injected signals are confined. This strategy can mitigate attacks that exploit ADC aliasing for output control (Section IV-D4) as shown by Trippel et al. [59].

3) *Randomization*: Adding randomness can often mitigate attacker influence on sensor output. Randomness can be ap-

plied to various components: transducers, ADCs, and even digital backends such as microcontrollers. This defense can be subdivided into receiver-chain (RX) and transmitter-chain (TX) randomization based on where the randomness is applied.

RX Randomization. RX randomization applies randomness to RX control parameters and can effectively deal with attacks where the injected signals can only be partially controlled, e.g., the raw injected signal is fluctuated randomly instead of a controllable bias [59] or only a fraction of the injected signal is controllable [98]. Under such partially controllable cases, attackers often exploit the predictable property of a certain component, e.g., sampling points of an ADC. Thus, randomizing the exploitable property may prevent full sensor output control. These defenses correspond to Component Modification because randomness can generally be adopted without any additional components. RX randomization can mitigate ADC aliasing as suggested by prior work on accelerometers and gyroscopes [43], [59]. Randomizing the ADC's sampling intervals prevents an attacker from predicting the exact quantization timing, increasing the difficulty of inducing controllable bias to the output. However, RX randomization may have other applications. Davidson et al. [98] suggested using a random sample consensus (RANSAC) [109] to defend a transduction attack on optical flow sensors based on the Lucas-Kanade method [110]. To extract true optical flow from the corrupted transducer output, RANSAC randomly picks a subset of features to form hypotheses, then makes all other features to vote for them.

TX Randomization. Prevention and detection by TX randomization operate on similar principles, but prevention by TX randomization (TX randomization in short) focuses more on enhancing the resiliency of sensors against attacks rather than detecting them. When randomness is added to various parameters (e.g., direction, waveform, and frequency) of the transmitted signal, the sensor itself, aware of the random pattern, can selectively concentrate only on meaningful information. In contrast, attackers unable to identify the random pattern embedded in the signal may not be as effective. TX randomization often corresponds to Component Modification, but may also correspond to Component Addition (e.g., adding actuators for spatial randomization). It also should be noted that the modified transfer function lies not in the receiver but in the transmitter chain. TX randomization has long been utilized for military applications, especially for radars [111], and for additional applications. The physical shift authentication [68] can recover real echoes in the received signal of an ultrasonic sensor. Petit et al. and Shin et al. [45], [46] suggested random-probing lidars, which correspond to spatial randomization, to defend spoofing attacks. Additionally, Shin et al. proposed randomizing lidars' ping waveform.

4) *Improving the Quality of Components:* Sensor designers may choose to improve the performance of certain components to mitigate attacker signals, though typically this increases production costs. The details of a specific component improvement, including the sensor and attack to defend, allows this class of defense to mitigate a variety of injection and measure-

ment shaping steps. For example, Trippel et al. [59] suggested using secure amplifiers whose dynamic range is large enough to cope with the exploited saturation. Son et al. [42] proposed to redesign MEMS gyroscopes to have resonance frequencies in non-critical frequency bands. Although specific approaches were not given, designing acoustic-resonance resilient MEMS gyroscopes were proposed as a defense by both Son et al. [42] and Wang et al. [44]. All three cases belong to Component Modification because no additional component is adopted.

5) *Sensor Fusion:* Defense by sensor fusion enhances resiliency against transduction attacks by utilizing output from multiple sensors. They can be divided further into (1) spatial fusion which utilizes multiple sensors, sometimes different types of sensors, (2) spectral fusion that adopts multiple redundant frequencies/wavelengths for measurement, and (3) temporal fusion which utilizes the history of measurement. Spatial fusion corresponds to Component Fusion as combines output of multiple transducers, i.e., multiple chains of transfer functions. It was commonly suggested as a defense by prior work [44], [45], [68], [70], [86], and can be applied to systems which can bear the cost of adopting multiple sensors. Spectral fusion can be thought as both Component Modification (when a single transducer suffices for such multi-band operation) and Component Fusion (when multiple transducers are required). Petit et al. [46] suggested utilizing multiple wavelengths to enhance lidars' resiliency against transduction attacks. Temporal fusion is typically more affordable than spatial as it does not typically require extensive hardware modification, and would be close to Component Modification. Previous work has suggested the use of sensor fusion as a mitigation. Xu et al. [68] proposed a special filter to remove maliciously injected echoes in ultrasonic sensors by examining echo consistency over multiple pulse cycles. Davidson et al. [98] suggested weighted RANSAC with momentum to increase resiliency against spoofing attacks. It not only utilizes RANSAC but also gives weights to each feature according to how consistent it was in earlier frames. Petit et al. [46] also discussed probing objects multiple times to limit the effectiveness of attacks. Additionally, it may be possible to combine spatial, spectral, and temporal fusion schemes to further enhance security.

VI. PREDICTING ATTACKS AND DEFENSES

High-level analysis of the signal injection and measurement shaping steps shown in Section IV may aid in the prediction of new attacks or defenses. A series of injection and shaping steps behave similarly across various sensor types. We briefly discuss how these similarities may allow one to synthesize new attacks through careful analysis. Lastly, we show how someone using our model may have been able to predict previous work and follow with a few predictions of future work.

A. Prediction Concepts

1) *Attack Chains:* Signal injection and measurement shaping steps may be grouped together to provide a quick, human-readable description of the mathematics behind an attack or part of an attack. For these series of steps, which we term

TABLE IV: ATTACK CHAINS FOR PRIOR AND PREDICTED TRANSDUCTION ATTACKS. USING OUR METHODOLOGY, OLDER WORK CAN SERVE AS A BASIS (BASIS) TO RETROSPECTIVELY PREDICT SUCCESSFUL ATTACKS (RETRO.). LAST, WE USE PRIOR WORK, INCLUDING THOSE RETROSPECTIVELY PREDICTED, TO PREDICT FUTURE WORK (FUTURE).

	Attack Target	Attack Chain	Paper
Basis	Microphone	$inject(EM) \rightarrow IMD \rightarrow filtering$ $inject(EM) \rightarrow filtering \rightarrow aliasing$	[47]
	Gyroscope	$inject(acoustic) \rightarrow filtering$	[42]
Retro.	Accelerometer	$inject(acoustic) \rightarrow filtering \rightarrow aliasing$	[59]
	Microphone	$inject(acoustic) \rightarrow IMD \rightarrow filtering$	[90]
Future	Magnetometer	$inject(acoustic) \rightarrow filtering \rightarrow aliasing$	—
	Accelerometer	$inject(acoustic) \rightarrow IMD \rightarrow filtering$	—

IMD: Intermodulation Distortion EM: Electromagnetic

attack chains, noting the order of steps is crucial. Different arrangements of the same steps describe separate attack chains. Additionally, attack chains share the abstraction of physical components from the steps they are based on, enabling conceptual comparison between transduction attacks even across different sensor types and components. Lastly, most attacks may be described as a chain of fewer than five steps, providing a quick, human-readable attack notation based in mathematics.

2) *Cross-Sensor Transferability*: The same attack chain that employs the same steps (Table II) in the same order should behave similarly despite sensor type or components exploited. This property results from how injection and shaping steps are a mathematical abstraction above the physical layer.

Defenses defined by which injection or shaping steps they mitigate (Table III) should also behave similarly despite sensor type or exploited components. Once again, this is because of steps being mathematical abstractions above the physical layer.

3) *Attack Chain Synthesis*: One may be able to combine two or more known sub-chains into new attack chains. One common method to create a new attack chain is to substitute a sub-chain in an existing chain with another sub-chain that accomplishes the same goal. The new sub-chain may accomplish the same mathematical goal as the original, but utilize different physical components or principles to do so. In doing so, the new sub-chain may succeed on a sensor or device the original chain did not. In particular, simply exchanging a known attack chain’s injection step with a new one may enable new attacks.

B. Attack Prediction: Case Studies

Analysis of signal injection and measurement shaping steps may allow an attacker to predict new transduction attacks via the cross-sensor transferability and attack chain synthesis concepts discussed above. We first present a case study to show how such an analysis may have aided in predicting previous attacks across microphones, accelerometers, and gyroscopes. Secondly, we provide examples of possible future attacks as an example of our prediction methodology.

1) *Retrospective Attack Prediction*: Attack chain analysis of prior work can elucidate conceptual similarities in attacks across different sensors. These similarities, along with the ideas of cross-sensor transferability and attack chain synthesis (Section VI-A), may have allowed someone to predict these

attacks. Shown in Table IV are several attack chains from prior work. To demonstrate our prediction methodology, we show how the first two attacks may have assisted in predicting the second two.

Basis 1: In 2013, Foo Kune et al. showed how to use electromagnetic waves to inject audio signals into microphones [47]. For both listed attacks, electromagnetic injection creates a malicious, amplitude modulated (AM) signal in the sensor. Then the $[IMD \rightarrow filtering]$ or $[filtering \rightarrow aliasing]$ sub-chains demodulate the injected signal using different principles. The demodulation results in the malicious audio signal remaining in the system.

Basis 2: In 2015, Son et al. [42] described how to use acoustic waves to disrupt drones. First, acoustic waves inject a signal via vibration of the MEMS gyroscope’s sensing mass. This signal is not filtered correctly, resulting in errant drone control.

Prediction 1: In 2017, Trippel et al. [59] demonstrated how to use acoustic waves to control MEMS accelerometer output. Step analysis may have assisted in predicting one of the attacks. First, analysis of an accelerometer reveals that the steps of injection via acoustic waves, improper filtering, and ADC aliasing are possible for MEMS accelerometers. Secondly, analysis of the acoustic injection step shows how the step may inject signals similarly to the second basis listed above, but also that these signals can be amplitude modulated like with the first basis. Third, the concept of cross-sensor transferability implies that the sub-chain $[filtering \rightarrow aliasing]$ should act similarly to the first basis. Last, the concept of attack chain synthesis allows the combination of the injection step and the $[filtering \rightarrow aliasing]$ sub-chain to predict this attack.

Prediction 2: In 2017, Zhang et al. [90] showed how ultrasonic waves could control microphone output. Yet this attack may also have been predicted via step analysis. First, step analysis reveals that attackers may use ultrasonic waves to inject an amplitude modulated signal into the microphone. While employing different physical principles, the injected signal is an AM signal similar to the first basis. Secondly, the AM demodulation sub-chain, $[IMD \rightarrow filtering]$, is also the same as in the first basis. Last, attack chain synthesis allows for these sub-chains to combine into a new attack chain, predicting this attack. Despite different signal types (electromagnetic vs. acoustic), different components for the IMD step (amplifier vs. transducer), and different frequencies (MHz vs. kHz), our analytical methods reveal the conceptual similarities of these attacks.

2) *Possible Future Attacks*: To further demonstrate our prediction methodology, we predict possible future attacks using analysis of prior work similarly to the previous retrospective prediction section.

Prediction 1: Step analysis predicts that acoustic waves may be able to control a MEMS magnetometer’s output. Nashimoto et al. [99], briefly mentioned that acoustic waves did affect MEMS magnetometer output. This seems similar to the acoustic injection steps for other MEMS sensors in prior work [42], [59], [86]. Additionally, the improper filtering step

may be available on some magnetometers similarly to other MEMS sensors [42]–[44], [59], [86], [99]. The aliasing step should also be available as magnetometer have analog-to-digital converters. Thus, one may be able to construct the [*inject (acoustic)* → *filtering* → *aliasing*] attack chain on the magnetometer, which is the same as found in previous work [59]. The ability to construct an attack chain known to work on another sensor suggests that the attack is possible, despite the different sensor types.

Prediction 2: Step analysis also predicts that an additional attack on MEMS accelerometers may be possible by employing the [*injection (acoustic)* → *IMD* → *filtering*] attack chain demonstrated in previous work [90]. The IMD step may be available on the amplifiers of some MEMS accelerometers, just as it is on amplifiers on other devices [47], [80], [100]. Additionally, attackers have already utilized acoustic injection and filtering steps on MEMS accelerometers [43], [59]. Thus, one could construct the full attack chain listed previously.

C. Defense Prediction: Case Studies

The property of cross-sensor transferability for injection and shaping mitigations (Section VI-A) can aid defenders in predicting successful mitigations. To demonstrate this predictive ability, we analyze how one may have predicted mitigations presented in previous work. Specifically, we use TX and RX randomization as case studies.

1) *RX Randomization:* Defenders may use RX randomization to mitigate an attacker’s ability to demodulate amplitude modulated signals via aliasing on an ADC. Trippel et al. [59] proposed this scheme in 2017 for use on MEMS accelerometers, and later Tu et al. also proposed a similar scheme on MEMS gyroscopes [43]. However, step analysis suggests that RX randomization may mitigate other attacks using ADC aliasing. For example, in an attack that uses a microphone’s ADC to demodulate an AM signal [47], an RX randomization scheme may mitigate the ADC aliasing step and defend against such an attack.

2) *TX Randomization:* Defenders may use TX randomization on active sensors to detect or prevent a variety of attacks. Researchers have suggested or demonstrated several cases of TX randomization such as with magnetic encoders [96], [106] and radars [95]. Step analysis also indicates that a TX Randomization would apply to other active sensors, such as ultrasonic sensors [68], [70] and lidars [45], [46].

VII. DISCUSSION

A. Improving the Simple Sensor Security Model

Our model merely serves as an initial step towards formalizing analog sensor security. Looking forward, we will continue to improve the model to more thoroughly cover existing and forthcoming studies. For example, it may be desirable to abstract new measurement shaping steps exclusive to in-band attacks on active sensors that manipulate measurement by adjusting the injection time [45], [46], [70] or signal frequency [96]. Incorporating emerging and potential attacks that modify the transfer functions, e.g., by injecting EMI to

the power lines of amplifiers [112] or heating temperature-sensitive components, may also be a direction for future work.

B. Improving Research Methodology

We hope and believe our model and related analytical methods will improve transduction attack research methodology. This model highlights the importance of discovering new signal injection and measurement shaping steps, as well as defenses for these steps. Our methodology demonstrates how these steps can apply across a wide range of sensors (Section VI-A). Thus, research based on the discovery or analysis of steps may have a broader impact on sensor design as a whole rather than only target a single sensor or system. Additionally, using the provided terminology can assist researchers and sensor designers in understanding new transduction attacks and how the attack may apply across sensors as a whole.

C. Predictive Defense Schemes

We discuss how the model enables two predictive defense schemes that enhance sensor resiliency to transduction attacks.

- **Predictive Attack Defense.** A sensor designer could employ the strategy of implementing a defense for every theoretical attack on a sensor. The model allows a designer to predict theoretical attacks on a sensor they are designing (Section VI). From this, the designer can adopt a simple strategy of ensuring there is at least one defense for each theoretical attack. In addition to the benefit mitigating possible future attacks, this approach may reduce the loss of time and money associated with redesigning, manufacturing, and distributing a new sensor each time a new attack is demonstrated against a sensor.
- **Predictive Step Defense.** Predictive step defense employs a different approach of designing a mitigation for every known signal injection or measurement shaping step in a sensor. In the model, a successful transduction attack requires all steps in its attack chain. Mitigating any step in the chain will mitigate the entire attack. So, a strategy that mitigates every known injection or shaping step will prevent all attacks from exploiting those mitigated steps, including attacks that have not yet been theoretically constructed. Thus, after this defense scheme is employed, an attacker would need to construct an attack chain of entirely comprising newly discovered steps. Therefore, predictive step defense provides a scheme for designers to protect their devices against unknown theoretical attacks at design time.

VIII. CONCLUSION

Security researchers and practitioners can use our simple sensor security model to better express and understand attacks employing physical signals to manipulate sensor output, and defenses against them. This model employs transfer functions and a vector of adversarial noise to allow comparison of attacks across sensors of different types. The model allows some predictive capability and enables new defense schemes to make sensors more resilient against future attacks.

ACKNOWLEDGMENTS

We thank our shepherd Prof. Brendan Dolan-Gavit and the anonymous reviewers for their constructive feedback. This work was supported in part by the ZJU-OPPO-OnePlus Joint Innovation Center, NSF CNS-1330142, and by an award from Mcity at University of Michigan. The views and conclusions contained in this paper are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the ZJU-OPPO-OnePlus Joint Innovation Center, NSF, or Mcity.

REFERENCES

- [1] AirSafe.com, "COPA airlines plane crashes," <http://www.airsafe.com/events/airlines/copa.html>, 1992.
- [2] T. C. Frankel, "Sensor cited as potential factor in Boeing crashing draws scrutiny," https://www.washingtonpost.com/business/economy/sensor-cited-as-potential-factor-in-boeing-crashes-draws-scrutiny/2019/03/17/5ecf0b0e-4682-11e9-aaf8-4512a6fe3439_story.html, 2018.
- [3] Tesla, "A tragic loss," <https://www.tesla.com/blog/tragic-loss>, 2016.
- [4] J. Stewart, "Why Tesla's autopilot can't see a stopped firetruck," <https://www.wired.com/story/tesla-autopilot-why-crash-radar>, 2018.
- [5] B. Cole, "The design challenges of a trillion sensor world," <https://www.embedded.com/electronics-blogs/cole-bin/4433743/The-design-challenges-of-a-trillion-sensor-world>, 2014.
- [6] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham *et al.*, "Experimental security analysis of a modern automobile," in *Proceedings of the 31st IEEE Symposium on Security and Privacy (SP)*. IEEE, 2010, pp. 447–462.
- [7] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno *et al.*, "Comprehensive experimental analyses of automotive attack surfaces," in *Proceedings of the 20th USENIX Security Symposium*. USENIX Association, 2011, pp. 447–462.
- [8] K.-T. Cho and K. G. Shin, "Fingerprinting electronic control units for vehicle intrusion detection," in *Proceedings of the 25th USENIX Security Symposium*. USENIX Association, 2016, pp. 911–927.
- [9] —, "Viden: Attacker identification on in-vehicle networks," in *Proceedings of the 2017 ACM Conference on Computer and Communications Security (CCS)*. ACM, 2017, pp. 1109–1123.
- [10] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*. ACM, 2005, pp. 46–57.
- [11] W. Xu, K. Ma, W. Trappe, and Y. Zhang, "Jamming sensor networks: attack and defense strategies," *IEEE network*, vol. 20, no. 3, pp. 41–47, 2006.
- [12] Y. W. Law, M. Palaniswami, L. V. Hoesel, J. Doumen, P. Hartel, and P. Havinga, "Energy-efficient link-layer jamming attacks against wireless sensor network MAC protocols," *ACM Transactions on Sensor Networks (TOSN)*, vol. 5, no. 1, p. 6, 2009.
- [13] I. Roufa, R. Miller, M. Hossen, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, and I. Seskar, "Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study," in *Proceedings of the 19th USENIX Security Symposium*, 2010.
- [14] I. Rouf, H. Mustafa, M. Xu, W. Xu, R. Miller, and M. Gruteser, "Neighborhood watch: Security and privacy analysis of automatic meter reading systems," in *Proceedings of the 2012 ACM Conference on Computer and Communications Security (CCS)*. ACM, 2012, pp. 462–473.
- [15] H. Bojinov, Y. Michalevsky, G. Nakibly, and D. Boneh, "Mobile device identification via sensor fingerprinting," *arXiv preprint arXiv:1408.1416*, 2014.
- [16] S. Dey, N. Roy, W. Xu, R. R. Choudhury, and S. Nelakuditi, "Accelprint: Imperfections of accelerometers make smartphones trackable," in *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2014.
- [17] X. Zhou, X. Ji, C. Yan, J. Deng, and W. Xu, "Nauth: Secure face-to-face device authentication via nonlinearity," in *Proceedings of the 2019 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2019, pp. 2080–2088.
- [18] J. Zhang, A. R. Beresford, and I. Sheret, "Sensorid: Sensor calibration fingerprinting for smartphones," in *Proceedings of the 40th IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019.
- [19] P. Marquardt, A. Verma, H. Carter, and P. Traynor, "(sp)iphone: Decoding vibrations from nearby keyboards using mobile phone accelerometers," in *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS)*. ACM, 2011, pp. 551–562.
- [20] L. Cai and H. Chen, "Touchlogger: Inferring keystrokes on touch screen from smartphone motion," in *Proceedings of the 6th USENIX Workshop on Hot Topics in Security (HotSec)*, 2011, pp. 1–9.
- [21] E. Miluzzo, A. Varshavsky, S. Balakrishnan, and R. R. Choudhury, "Tapprints: Your finger taps have fingerprints," in *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services*. ACM, 2012, pp. 323–336.
- [22] Y. Michalevsky, D. Boneh, and G. Nakibly, "Gyrophone: Recognizing speech from gyroscope signals," in *Proceedings of the 23rd USENIX Security Symposium*, 2014, pp. 1053–1067.
- [23] R. Spreitzer, "Pin skimming: Exploiting the ambient-light sensor in mobile devices," in *Proceedings of the 4th ACM Workshop on Security and Privacy in Smartphones & Mobile Devices*. ACM, 2014, pp. 51–62.
- [24] A. K. Sikder, H. Aksu, and A. S. Uluagac, "6thsense: A context-aware sensor-based attack detector for smart devices," in *Proceedings of the 26th USENIX Security Symposium*, 2017, pp. 397–414.
- [25] G. Petracca, A.-A. Reineh, Y. Sun, J. Grossklags, and T. Jaeger, "Aware: Preventing abuse of privacy-sensitive sensors via operation bindings," in *Proceedings of the 26th USENIX Security Symposium*, 2017, pp. 379–396.
- [26] A. Kwong, W. Xu, and K. Fu, "Hard drive of hearing: Disks that eavesdrop with a synthesized microphone," in *Proceedings of the 40th IEEE Symposium on Security and Privacy (SP)*. IEEE, May 2019, p. 15.
- [27] A. S. Uluagac, V. Subramanian, and R. Beyah, "Sensory channel threats to cyber physical systems: A wake-up call," in *Proceedings of the 2014 IEEE Conference on Communications and Network Security*, Oct 2014, pp. 301–309.
- [28] Y. J. Jia, Q. A. Chen, S. Wang, A. Rahmati, E. Fernandes, Z. M. Mao, A. Prakash, and S. J. University, "Contextlot: Towards providing contextual integrity to appified IoT platforms," in *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2017.
- [29] D. Kumar, R. Paccagnella, P. Murley, E. Hennenfent, J. Mason, A. Bates, and M. Bailey, "Skill squatting attacks on amazon alexa," in *Proceedings of the 27th USENIX Security Symposium*, 2018, pp. 33–47.
- [30] N. Zhang, X. Mi, X. Feng, X. Wang, Y. Tian, and F. Qian, "Dangerous skills: Understanding and mitigating security risks of voice-controlled third-party functions on virtual personal assistant systems," in *Proceedings of the 40th IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 1–16.
- [31] C. Sousedik and C. Busch, "Presentation attack detection methods for fingerprint recognition systems: a survey," *IET Biometrics*, vol. 3, no. 4, pp. 219–233, 2014.
- [32] A. Roy, N. Memon, and A. Ross, "Masterprint: Exploring the vulnerability of partial fingerprint-based authentication systems," *IEEE Transactions on Information Forensics and Security (TIFS)*, vol. 12, no. 9, pp. 2013–2025, 2017.
- [33] Z. Zhang, D. Yi, Z. Lei, S. Z. Li *et al.*, "Face liveness detection by learning multispectral reflectance distributions," in *Proceedings of the 2011 International Conference on Automatic Face and Gesture Recognition*, 2011, pp. 436–441.
- [34] B. Coporation, "How bkav tricked iphone x's face id with a mask," <https://youtu.be/i4YQRLQVixM>, 2017.
- [35] C. C. Congress, "Hacking the samsung galaxy s8 irisscanner," <https://media.ccc.de/v/biometrie-s8-iris-en#t=0>, 2017.
- [36] N. Carlini, P. Mishra, T. Vaidya, Y. Zhang, M. Sherr, C. Shields, D. Wagner, and W. Zhou, "Hidden voice commands," in *Proceedings of the 25th USENIX Security Symposium*, 2016, pp. 513–530.

- [37] N. Carlini and D. Wagner, "Audio adversarial examples: Targeted attacks on speech-to-text," in *Proceedings of the 2018 IEEE Security and Privacy Workshops*, May 2018, pp. 1–7.
- [38] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter, "Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition," in *Proceedings of the 2016 ACM Conference on Computer and Communications Security (CCS)*. ACM, 2016, pp. 1528–1540.
- [39] X. Yuan, Y. Chen, Y. Zhao, Y. Long, X. Liu, K. Chen, S. Zhang, H. Huang, X. Wang, and C. A. Gunter, "Commandersong: A systematic approach for practical adversarial voice recognition," in *Proceedings of the 27th USENIX Security Symposium*, 2018, pp. 49–64.
- [40] K. Eykholt, I. Evtimov, E. Fernandes, B. Li, A. Rahmati, C. Xiao, A. Prakash, T. Kohno, and D. Song, "Robust physical-world attacks on deep learning visual classification," in *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 1625–1634.
- [41] K. Fu and W. Xu, "Risks of trusting the physics of sensors," *Communications of the ACM*, vol. 61, no. 2, pp. 20–23, 2018.
- [42] Y. Son, H. Shin, D. Kim, Y.-S. 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, pp. 881–896.
- [43] Y. Tu, Z. Lin, I. Lee, and X. Hei, "Injected and delivered: Fabricating implicit control over actuation systems by spoofing inertial sensors," in *Proceedings of the 27th USENIX Security Symposium*. USENIX Association, 2018, pp. 1545–1562.
- [44] Z. Wang, K. Wang, B. Yang, S. Li, and A. Pan, "Sonic gun to smart devices," *Black Hat USA*, 2017.
- [45] 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 19th International Conference on Cryptographic Hardware and Embedded Systems (CHES)*. Springer, 2017, pp. 445–467.
- [46] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote attacks on automated vehicles sensors: Experiments on camera and lidar," *Black Hat Europe*, 2015.
- [47] 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 34th IEEE Symposium on Security and Privacy (SP)*. IEEE, 2013, pp. 145–159.
- [48] C. Kasmi and J. L. Esteves, "IEMI threats for information security: Remote command injection on modern smartphones," *IEEE Transactions on Electromagnetic Compatibility*, vol. 57, no. 6, pp. 1752–1755, 2015.
- [49] I. Giechaskiel and K. B. Rasmussen, "Sok: Taxonomy and challenges of out-of-band signal injection attacks and defenses," *arXiv preprint arXiv:1901.06935*, 2019.
- [50] A. K. Sikder, G. Petracca, H. Aksu, T. Jaeger, and A. S. Uluagac, "A survey on sensor-based threats to internet-of-things (IoT) devices and applications," *arXiv preprint arXiv:1802.02041*, 2018.
- [51] I. Giechaskiel, Y. Zhang, and K. B. Rasmussen, "A framework for evaluating security in the presence of signal injection attacks," *arXiv preprint arXiv:1901.03675*, 2019.
- [52] Wikipedia, "List of sensors," https://en.wikipedia.org/wiki/List_of_sensors, 2019.
- [53] J. S. Wilson, *Sensor technology handbook*. Elsevier, 2004.
- [54] T. Grandke and W. H. Ko, *Sensors a Comprehensive Survey, Vol. 1: Fundamentals and General Aspects*. VCH, New York, 1989.
- [55] J. Eargle, *The Microphone Book: From mono to stereo to surround-a guide to microphone design and application*. CRC Press, 2012.
- [56] J. J. Engelsma, S. S. Arora, A. K. Jain, and N. G. Paulter, "Universal 3D wearable fingerprint targets: advancing fingerprint reader evaluations," *IEEE Transactions on Information Forensics and Security (TIFS)*, vol. 13, no. 6, pp. 1564–1578, 2018.
- [57] E.-C. Koch, "Review on pyrotechnic aerial infrared decoys," *Propellants, Explosives, Pyrotechnics*, vol. 26, no. 1, pp. 3–11, 2001.
- [58] K. Sharma and M. Paradakar, "The melamine adulteration scandal," *Food Security*, vol. 2, no. 1, pp. 97–107, Mar 2010.
- [59] 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 2nd IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2017, pp. 3–18.
- [60] N. S. Nise, *Control Systems Engineering*. John Wiley & Sons, 2007.
- [61] J. Fraden, *Handbook of modern sensors: physics, designs, and applications*. Springer Science & Business Media, 2004.
- [62] K. Ono, H. Cho, and T. Matsuo, "Transfer functions of acoustic emission sensors," *Journal of Acoustic Emission*, vol. 26, pp. 72–91, 2008.
- [63] B. Cochrun and A. Grabel, "A method for the determination of the transfer function of electronic circuits," *IEEE Transactions on Circuit Theory*, vol. 20, no. 1, pp. 16–20, 1973.
- [64] K. Ogata and Y. Yang, *Modern Control Engineering*. Prentice-Hall, 2002.
- [65] E. A. Parr, *Logic designer's handbook: circuits and systems*. Elsevier, 2013.
- [66] I. Sinclair and J. Dunton, *Electronic and Electrical Servicing, Second Edition: Consumer and Commercial Electronics*, 2nd ed. Newton, MA, USA: Newnes, 2007.
- [67] V. Tuzlukov, *Signal processing noise*. CRC Press, 2002.
- [68] W. Xu, C. Yan, W. Jia, X. Ji, and J. Liu, "Analyzing and enhancing the security of ultrasonic sensors for autonomous vehicles," *IEEE Internet of Things Journal*, vol. 5, no. 6, pp. 5015–5029, Dec 2018.
- [69] M. J. Wilson, S. R. Ford, and P. L. Rinaldo, *The ARRL Handbook for Radio Communications 2007*. Amer Radio Relay League, 2006.
- [70] C. Yan, W. Xu, and J. Liu, "Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle," *DEF CON*, 2016.
- [71] E. B. Joffe and K.-S. Lock, *Grounds for grounding: A circuit to system handbook*. John Wiley & Sons, 2011.
- [72] K. B. Rasmussen, C. Castelluccia, T. S. Heydt-Benjamin, and S. Capkun, "Proximity-based access control for implantable medical devices," in *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS)*. ACM, 2009, pp. 410–419.
- [73] Y.-C. Hu, A. Perrig, and D. B. Johnson, "Wormhole attacks in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 2, pp. 370–380, 2006.
- [74] J. L. Esteves and C. Kasmi, "Remote and silent voice command injection on a smartphone through conducted IEMI: Threats of smart IEMI for information security," Wireless Security Lab, French Network and Information Security Agency (ANSSI), Tech. Rep., 2018.
- [75] J. Selvaraj, G. Y. Dayanikli, N. P. Gaunkar, D. Ware, R. M. Gerdes, and M. Mina, "Electromagnetic induction attacks against embedded systems," in *Proceedings of the 2018 on Asia Conference on Computer and Communications Security (AsiaCCS)*. ACM, 2018, pp. 499–510.
- [76] J. Selvaraj, "Intentional electromagnetic interference attack on sensors and actuators," Ph.D. dissertation, Iowa State University, 2018.
- [77] D. A. Ware, "Effects of intentional electromagnetic interference on analog to digital converter measurements of sensor outputs and general purpose input output pins," Master's thesis, Utah State University, 2017.
- [78] NDT Resource Center, "Attenuation of sound waves," <https://www.nde-ed.org/EducationResources/CommunityCollege/Ultrasonics/Physics/attenuation.htm>.
- [79] H. Canada, "Guidelines for the safe use of ultrasound: Part ii—industrial and commercial applications safety code 24," 1991.
- [80] C. Yan, G. Zhang, X. Ji, T. Zhang, T. Zhang, and W. Xu, "The feasibility of injecting inaudible voice commands to voice assistants," *IEEE Transactions on Dependable and Secure Computing (TDSC)*, 2019.
- [81] C. Yan, K. Fu, and W. Xu, "On cuba, diplomats, ultrasound, and intermodulation distortion," *Computers in Biology and Medicine*, vol. 104, pp. 250–266, 2019.
- [82] L. Bjørnø, "Introduction to nonlinear acoustics," *Physics Procedia*, vol. 3, no. 1, pp. 5–16, 2010.
- [83] M. F. Hamilton, D. T. Blackstock *et al.*, *Nonlinear acoustics*. Academic press San Diego, 1998, vol. 1.
- [84] S. Castro, R. Dean, G. Roth, G. T. Flowers, and B. Grantham, "Influence of acoustic noise on the dynamic performance of MEMS gyroscopes," in *Proceedings of the ASME 2007 International Mechanical Engineering Congress and Exposition*. American Society of Mechanical Engineers, 2007, pp. 1825–1831.

- [85] D. Giri and F. Tesche, "Classification of intentional electromagnetic environments (IEME)," *IEEE Transactions on Electromagnetic Compatibility*, vol. 46, no. 3, pp. 322–328, 2004.
- [86] C. Bolton, S. Rampazzi, C. Li, A. Kwong, W. Xu, and K. Fu, "Blue note: How intentional acoustic interference damages availability and integrity in hard disk drives and operating systems," in *Proceedings of the 39th IEEE Symposium on Security and Privacy (SP)*, May 2018, pp. 1048–1062.
- [87] Y. Park, Y. Son, H. Shin, D. Kim, and Y. Kim, "This ain't your dose: Sensor spoofing attack on medical infusion pump," in *Proceedings of the 10th USENIX Workshop on Offensive Technologies (WOOT)*, 2016.
- [88] G. Smillie, "2 - analogue modulation principles," in *Analogue and Digital Communication Techniques*, G. Smillie, Ed. Oxford: Butterworth-Heinemann, 1999, pp. 15 – 45.
- [89] P. Horowitz and W. Hill, *The art of electronics*. Cambridge Univ. Press, 1989.
- [90] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu, "DolphinAttack: Inaudible voice commands," in *Proceedings of the 2017 ACM Conference on Computer and Communications Security (CCS)*. ACM, 2017.
- [91] N. Roy, S. Shen, H. Hassanieh, and R. R. Choudhury, "Inaudible voice commands: The long-range attack and defense," in *Proceedings of the 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. USENIX Association, 2018, pp. 547–560.
- [92] L. Song and P. Mittal, "Inaudible voice commands," *arXiv preprint arXiv:1708.07238*, 2017.
- [93] Texas Instruments, "System-level ESD protection guide," Texas Instruments, Tech. Rep., 2018.
- [94] C. Rostamzadeh, F. Canavero, F. Kashefi, and M. Darbandi, "Effectiveness of multilayer ceramic capacitors for electrostatic discharge protection," in *Compliance Magazine*, 2012.
- [95] R. Chauhan, "A platform for false data injection in frequency modulated continuous wave radar," Master's thesis, Utah State University, 2014.
- [96] Y. Shoukry, P. Martin, P. Tabuada, and M. Srivastava, "Non-invasive spoofing attacks for anti-lock braking systems," in *Proceedings of the 15th International Workshop on Cryptographic Hardware and Embedded Systems (CHES)*. Springer, 2013, pp. 55–72.
- [97] Y. S. Sakr, "Security and privacy in cyber-physical systems: Physical attacks and countermeasures," Ph.D. dissertation, UCLA, 2015.
- [98] D. Davidson, H. Wu, R. Jellinek, T. Ristenpart, and V. Singh, "Controlling UAVs with sensor input spoofing attacks," in *Proceedings of the 10th USENIX Workshop on Offensive Technologies (WOOT)*, 2016.
- [99] S. Nashimoto, D. Suzuki, T. Sugawara, and K. Sakiyama, "Sensor con-fusion: Defeating kalman filter in signal injection attack," in *Proceedings of the 2018 on Asia Conference on Computer and Communications Security (AsiaCCS)*. ACM, 2018, pp. 511–524.
- [100] N. Roy, H. Hassanieh, and R. Roy Choudhury, "Backdoor: Making microphones hear inaudible sounds," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, 2017, pp. 2–14.
- [101] H. Shen, W. Zhang, H. Fang, Z. Ma, and N. Yu, "Jamsys: Coverage optimization of a microphone jamming system based on ultrasounds," *IEEE Access*, vol. 7, pp. 67 483–67 496, 2019.
- [102] Y. Chen, H. Li, S. Nagels, Z. Li, P. Lopes, B. Y. Zhao, and H. Zheng, "Understanding the effectiveness of ultrasonic microphone jammer," *arXiv preprint arXiv:1904.08490*, 2019.
- [103] S. Maruyama, S. Wakabayashi, and T. Mori, "Tap 'n ghost: A compilation of novel attack techniques against smartphone touchscreens," in *Proceedings of the 40th IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 628–645.
- [104] C. E. Shannon, "Communication in the presence of noise," *Proceedings of the Institute of Radio Engineers*, vol. 86, no. 2, pp. 447–457, 1998.
- [105] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall Press, 2009.
- [106] 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 2015 ACM Conference on Computer and Communications Security (CCS)*, 2015, pp. 1004–1015.
- [107] D. Muniraj and M. Farhood, "Detection and mitigation of actuator attacks on small unmanned aircraft systems," *Control Engineering Practice*, vol. 83, pp. 188–202, 2019.
- [108] H. Shin, Y. Son, Y. Park, Y. Kwon, and Y. Kim, "Sampling race: Bypassing timing-based analog active sensor spoofing detection on analog-digital systems," in *Proceedings of the 10th USENIX Workshop on Offensive Technologies (WOOT)*, 2016.
- [109] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [110] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proceedings of the 7th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann Publishers Inc., 1981.
- [111] P. E. Pace, *Detecting and classifying low probability of intercept radar*. Artech House, 2009.
- [112] Y. Tu, S. Rampazzi, B. Hao, A. Rodriguez, K. Fu, and X. Hei, "Trick or heat? attack on amplification circuits to abuse critical temperature control systems," *arXiv preprint arXiv:1904.07110*, 2019.