

Taxonomy and Challenges of Out-of-Band Signal Injection Attacks and Defenses

Ilias Giechaskiel, Kasper B. Rasmussen

in *IEEE Communications Surveys & Tutorials*, vol. 22, no. 1, pp. 645–670,
Firstquarter 2020.

©2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
DOI:10.1109/COMST.2019.2952858

Taxonomy and Challenges of Out-of-Band Signal Injection Attacks and Defenses

Ilias Giechaskiel and Kasper B. Rasmussen

Abstract—Recent research has shown that the integrity of sensor measurements can be violated through out-of-band signal injection attacks. These attacks target the conversion process from a physical quantity to an analog property—a process that fundamentally cannot be authenticated. Out-of-band signal injection attacks thus pose previously-unexplored security risks by exploiting hardware imperfections in the sensors themselves, or in their interfaces to microcontrollers. In response to the growing-yet-disjointed literature in the subject, this article presents the first survey of out-of-band signal injection attacks. It focuses on unifying their terminology and identifying commonalities in their causes and effects through a chronological, evolutionary, and thematic taxonomy of attacks. By highlighting cross-influences between different types of out-of-band signal injections, this paper underscores the need for a common language irrespective of the attack method. By placing attack and defense mechanisms in the wider context of their dual counterparts of side-channel leakage and electromagnetic interference, this study identifies common threads and gaps that can help guide and inform future research. Overall, the ever-increasing reliance on sensors embedded in everyday commodity devices necessitates that a stronger focus be placed on improving the security of such systems against out-of-band signal injection attacks.

Index Terms—Out-of-Band, Signal Injections, Hardware Imperfections, Mixed-Signal Systems, Survey, Attacks and Defenses

I. INTRODUCTION

MATHEMATICALLY secure algorithms can be broken in practice due to a mismatch between the high-level system model used for analysis and the real-world environment on which code runs. For example, data-dependent electromagnetic, optical, and acoustic emanations, as well as variations in power consumption can reveal the information processed by a device [1], with or without the help of intentional faults [2]–[4]. However, attacks exploiting hardware imperfections are not limited to side-channel leakage of confidential data: recent research has shown that it is possible to target the integrity of sensor measurements in a similar *out-of-band* fashion.

These out-of-band signal injection attacks can be performed using electromagnetic radiation exploiting circuits unintentionally acting as receiver antennas [5]–[7], as well as optical [8], [9] and acoustic [10]–[12] emissions targeting flaws in the conversion process from physical properties into electrical ones. The systems attacked have been equally diverse, and include medical devices [5], drones [10], [13], hard drives [12] and cameras [14], among others. However, despite the wide

range of attack methods and devices targeted, research in the field thus far has been disjointed.

The ad-hoc nature of this type of research might stem, in part, from the fact that out-of-band signal injection attacks have so far only been (openly) conducted in a lab environment. However, out-of-band attacks have still garnered the interest of technological and mainstream news publications outside of the academic community [15]–[20]. They have even prompted national agencies to issue Computer Emergency Readiness Team (CERT) advisories [21]. In some cases, the effects of out-of-band attacks can be fatal: for instance, Foo Kune et al. have demonstrated that low-power attacker signals can trick cardiac implantable electrical devices into causing pacing inhibition and defibrillation shocks [5]. Although the techniques used are not identical to those of mass-produced sonic repellents [22] or commercial [23] and military [24] jammers, out-of-band signal injection attacks share the same potential for weaponization. As a result, to bring attention to these potentially severe issues, and to help designers better protect future hardware devices, this study conducts the first comprehensive survey of out-of-band signal injection attacks.

A. Survey Scope

This article focuses on out-of-band signal injection attacks, which target the connections between sensors, actuators, and microcontrollers, or exploit imperfections in the hardware itself. Although the term is defined precisely in Section II, we note here a few key features of the attacks investigated in this survey. The first property of out-of-band signal injection attacks is that they aim to *change* values processed by a system, rather than infer them. This fact distinguishes them from the side-channel [1] and fault-injection attacks [2]–[4] mentioned in the introduction.

The second feature is that out-of-band attacks do not change the measured quantity itself. For instance, using electromagnetic signals to change the audio recorded by a microphone [5] is an example of an out-of-band signal injection attack, but heating a temperature sensor with an open flame is not.

The final characteristic highlights the physical aspect of out-of-band signal injection attacks. In other words, the attacks studied in this paper alter sensor measurements or actuator inputs at the hardware layer instead of the protocol layer. As a result, spoofing attacks of unauthenticated, digital communication interfaces are out-of-scope. For example, wireless transmissions can be used to spoof the pressure of car tires and trigger warning lights [25], alter the flow of insulin injections [26], or change pacemaker settings to deliver shock commands with implantable cardiac defibrillators [27]. However, because these interfaces can easily be protected with

Manuscript received April 29, 2019; revised August 30, 2019; accepted November 03, 2019. Date of publication November 12, 2019; date of current version November 12, 2019. Corresponding author: Ilias Giechaskiel.

The authors are with the Department of Computer Science, University of Oxford, Oxford, UK (e-mail: ilias.giechaskiel@cs.ox.ac.uk; kasper.rasmussen@cs.ox.ac.uk).

Digital Object Identifier 10.1109/COMST.2019.2952858

cryptography, they are not considered in this survey. For similar reasons, jamming, e.g., train signal controls [28], or over-powering legitimate GPS signals [29] are out-of-scope because they rely on intentional communication interfaces. As a result, they do not exploit hardware imperfections, instead relying on high-power, in-band electromagnetic transmissions. Finally, relay attacks on LiDARs, radars, and sonars [9], [14], [30]–[32] are also not out-of-band signal injection attacks. This is because, in order to spoof the distance between the attacker and the victim, they depend on winning a race between the adversarial signal and the true reflected pulses.

That said, these related research areas can offer invaluable insight into novel injection techniques and possible countermeasures. As a result, cross-disciplinary connections are made throughout this work, but with a clear focus on how they impact out-of-band signal injection attacks and defenses.

B. Contributions & Organization

Despite the growing literature in the area and extensive parallels between prior work in hardware security research and out-of-band signal injection attacks and defenses, no other article has made these connections explicit, or traced their evolution through time, theme, or approach. Our survey fills these gaps through the following contributions:

- 1) It unifies the diverse terminology used by different works (Section II) and summarizes the threat model (Section III) to create a common language through which to discuss attack and defense mechanisms.
- 2) It proposes the first chronological and thematic evolution of out-of-band signal injection attacks (Section VIII), which highlights cross-influences between electromagnetic (Section IV), conducted (Section V), acoustic (Section VI), and other (Section VII) attacks.
- 3) It creates a taxonomy of countermeasures introduced to prevent and detect out-of-band attacks (Section IX).
- 4) It places attacks and defenses in the wider context of side-channel leakage and electromagnetic interference attacks (Section X). Using these insights, this study identifies gaps in the experimental approach of published research, and proposes concrete steps to overcome these challenges in the future (Section XI).

II. CHOICE OF TERMINOLOGY

The terms used to describe the numerous acoustic, electromagnetic, and optical attacks on sensor-to-microcontroller and microcontroller-to-actuator interfaces have so far been inconsistent, with some works not even naming the attacks at all [33], [34]. This section sets out to identify and unify the nomenclature used as a first step towards providing a common language through which to compare the various works. As the threat model (Section III) and the causes of vulnerability (Section VIII) will reveal, the commonalities in the attack techniques highlight a need for an all-encompassing term irrespective of the method of injection. In other words, although some of the more restrictive terms are appropriate for describing specific attacks, we find that doing so can hide potential insights that arise from considering different types of

attacks jointly. The term we have chosen for this unification is *out-of-band signal injection attacks*:

Definition 1 (Out-of-Band Signal Injection Attacks). *Out-of-band signal injection attacks are adversarial manipulations of interfaces not intended for communication involving sensors/actuators that cause a mismatch between the true physical property being measured/acted upon and its digitized version.*

To motivate our definition, the term *injection* was chosen because it captures the fact that values reported by a system are altered; it is not channel-specific; and it has already been adopted by different works [5], [11], [13], [35]–[44]. The *out-of-band* qualifier is necessary to distinguish the attacks studied in this survey from signal injection attacks on sensors using pulse reflections such as LiDARs [9], [14], [31], signal injection attacks on the physical layer of communication protocols [45], and false data injection attacks [28], [46]. As explained in Section I-A, these attacks are out-of-scope, as they do not depend on hardware vulnerabilities, but instead use external communication interfaces.

By contrast, our term captures attacks which target interfaces using signals outside of their intended frequency of operation. It includes ultraviolet or infrared light against cameras which should only be recording the visible part of the spectrum, and ultrasonic injections against microphones meant to be recording only audible sounds: these attacks transmit signals that are literally outside the operational band. It also includes electromagnetic signals against systems without any (intentional) antennas, and acoustic attacks against gyroscopes and accelerometers: these are also out-of-band, since they inject signals through channels other than the ones used by the sensor to measure the physical property. Our use of the out-of-band modifier is therefore consistent with the definition for out-of-band covert communication [47]. It has also recently been used by Tu et al. [43] to describe acoustic attacks on inertial sensors, further motivating its choice in this survey.

It should be noted that earlier work [31], [48] has proposed a subdivision of signal injections attacks into *regular-channel attacks*, which target the sensor structure itself by “using the same type of physical quantity sensed”, *transmission-channel attacks*, which target the connection between the sensor output and the measurement setup, and *side-channel attacks*, where the sensors themselves are targeted, but “by physical stimuli other than those they are supposed to sense”. We do not adopt this categorization, as it generally follows the medium of injection (optical, electromagnetic, and acoustic respectively). Moreover, regular-channel attacks are usually in-band, while side-channel attacks have an overloaded meaning.

We similarly find (*intentional*) *interference* [5], [10], [12], [36], [38], [40], [42], [49] to be unsuitable as a term because: (a) it does not make it clear that the attackers can in some cases inject waveforms of their choosing; and (b) Intentional Electromagnetic Interference (IEMI) has an established meaning in Electromagnetic Compatibility (EMC) literature [50], [51]. As Section X indicates, IEMI attacks often use high-power, destructive transmissions, and therefore have a different aim than out-of-band signal injection attacks.

The term (*sensor*) *spoofing* [8], [9], [14], [43] was also

TABLE I: Terminology used by different works to describe out-of-band signal injection attacks.

Terminology	Example References
Injection	[5], [13], [35], [38], [41]–[44]
Intentional Interference	[5], [10], [12], [36], [38], [40], [42], [49]
Non-Linearity	[11], [54], [55], [68]
Spoofing	[8], [9], [14], [43]
Other (See Text)	[7], [53], [67]

avoided for similar reasons: it has an overloaded meaning in authentication contexts and with in-band signal injection attacks [6], [52]. Moreover, it does not capture the physical aspect of injections, and does not accurately describe coarse-grained attacks which lead to saturation of a sensor.

Other terms used have been specific to the particular channel which is being exploited, including *induction attacks* [7], *acoustic resonance* [53], and *(acoustic) non-linearity* [11], [54], [55]. Such terms were not selected because they are channel-specific, and focus on the mechanism of the attack, rather than the effect. Similarly, methodology-inspired terms which have been avoided include Radio Frequency Injection (RFI), Direct Power Injection (DPI), and other terminology that arises in immunity or susceptibility literature against (non-adversarial) electromagnetic interference (EMI) [56]–[66].

Finally, the term *transduction attacks*, proposed by Fu and Xu [67] to mean attacks which “exploit a vulnerability in the physics of a sensor to manipulate its output or induce intentional errors” has not yet received mainstream recognition. It also does not necessarily make it clear that the attack may target the interface between the sensor and the rest of the system, instead of just the sensor itself. The various terms which have been used to describe out-of-band signal injection attacks are shown in Table I, along with example references.

Since the majority of attacks in the literature are on sensors rather than actuators, we will refer to both of them collectively as sensors for brevity, and will distinguish between the two only when it is necessary to do so, i.e., when there is a divergence in the attack methodology.

III. SYSTEM AND ADVERSARY MODEL

Systems depend on sensors and actuators to interface with their external environment. They therefore require a conversion of a physical property (e.g., temperature or speed) to or from an electrical quantity (such as voltage or resistance). This electrical measurement is typically analog in nature, and is digitized by an Analog-to-Digital Converter (ADC) before it is processed. Although modern cryptography has mostly solved the problem of secure communication between digital interfaces, there is no way to authenticate the measurement itself, or the analog component of the connection between the sensor or actuator and a microcontroller. This lack of authentication, coupled with hardware imperfections, can be exploited for out-of-band signal injection attacks.

Conceptually, the sensor, the ADC, and the microcontroller perform logically distinct functions, but all three can be fully encapsulated into the same Integrated Circuit (IC) chip. Although this chip presents a digital interface to third parties,

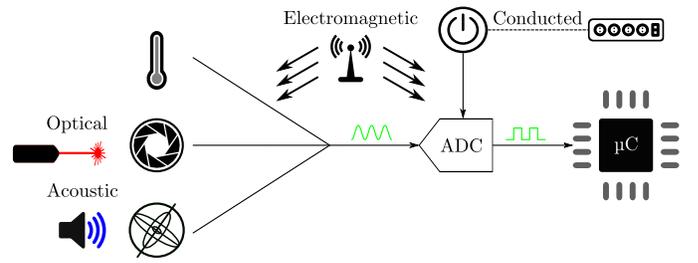


Fig. 1: System model for out-of-band signal injection attacks. Remote and conducted adversarial electromagnetic emanations, optical emissions, and acoustic waves can attack the sensors themselves, or the interfaces connecting sensors to microcontrollers through Analog-to-Digital Converters (ADCs).

which can be protected by cryptographic protocols, the sensor itself can still be vulnerable to out-of-band signal injection attacks. For example, acoustic attacks targeting the resonant frequencies of gyroscopes and accelerometers have proven to be effective even against digital ICs [10], [13], [43].

As explained in Section II, attackers are not allowed to manipulate the property being measured itself (e.g., the radiation being measured by a Geiger counter): in the language of Shoukry et al. [69], the property measured itself is *trusted*, although the measurement itself is not. Nonetheless, attackers are allowed to transmit signals outside of the limits being sensed, which are still interpreted as valid measurements. For example, an attacker can produce ultrasound waves which are picked up by a microphone recording human speech [11], [54], [55], or shine infrared (IR) light into a camera capturing the visible part of the spectrum [14]. However, attacking a microphone with audible sound, or a camera with visible light is not allowed under this threat model, since the attacker is manipulating the property being sensed in-band. A secondary goal for some attacks is therefore undetectability or *concealment* [69]. Attacks also need to be *non-invasive* [69], and preclude direct physical access to the system under attack. It should be noted that different attack techniques have different distance requirements, with electromagnetic attacks theoretically having a longer range compared to optical and acoustic attacks. We defer the discussion of distance and related considerations to Section XI.

Figure 1 summarizes the channels that have successfully been exploited in the literature thus far. Some of the out-of-band signal injection attacks use electromagnetic (EM) waves to penetrate the wires connecting sensors and microcontrollers (Section IV), or the power circuit of the device (Section V). The same effects can sometimes also be achieved through shared power lines (conducted attacks). Finally, other attacks may target the sensors themselves through sound (Section VI), and alternative means such as infrared light (Section VII).

The effectiveness of different attacks, however, has mostly been evaluated in a qualitative fashion so far (e.g., whether the system was tricked into performing an action or not). Recent research, however, has attempted to mathematically define security against out-of-band signal injection attacks, and therefore quantify their success [44]. Specifically, Giechaskiel

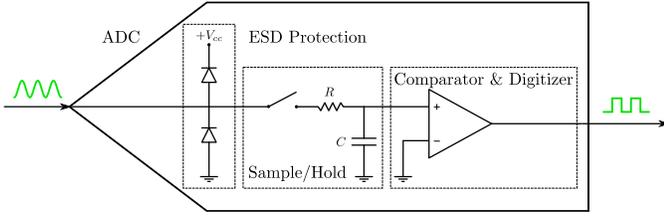


Fig. 2: Analog-to-Digital Converter (ADC) model: non-linearities due to Electrostatic Discharge (ESD) protection diodes and amplifiers (e.g., comparators) counteract low-pass filtering effects of the sample-and-hold mechanism and can unintentionally demodulate high-frequency input signals.

et al. [44] introduced probabilistic definitions which attempt to capture the fidelity with which an adversary can make a target waveform appear at the input of the microcontroller of Figure 1. These definitions address both coarse-grained attacks which merely disrupt sensor measurements and fine-grained attacks with precise waveform injections. Although the definitions abstract away from specific hardware considerations, they still model the discrepancy between an adversary’s target waveforms and the actual signal transmitted, which is necessary due to the behavior of the underlying circuits.

The specific circuit properties that allow adversarial signals to be injected into a device vary with the injection method and module targeted. For example, EM attacks typically depend on unintentional antennas in Printed Circuit Board (PCB) traces [5], while acoustic attacks exploit resonance in gyroscopes [10] and accelerometers [13]. As each attack exploits unique properties of the target device, the details of specific hardware imperfections in sensors and other modules are only expanded upon in subsequent sections. This section instead discusses common features of different attacks at a high level. One of these commonalities is that the vulnerability of systems to out-of-band signal injection attacks depends on both: (a) how adversarial signals are received by the devices under attack; and (b) how these signals are digitized.

Giechaskiel et al. [44] recently proposed a general circuit model which uses two transfer functions to separate these two aspects of vulnerability. The first transfer function describes circuit-specific transformations that an adversarial signal undergoes. For example, for EM attacks, these transformations include the (unintentional) low-power, low-gain antenna-like behavior of PCB traces connecting sensors to ADCs [5], [44]. The second transfer function, on the other hand, is ADC-specific, and summarizes the artifacts of the digitization process. These two transfer functions dictate that for a successful injection, attacker signals typically need to be transmitted over high-frequency carriers, and be demodulated into low-frequency, meaningful waveforms. According to the work of Giechaskiel et al., components within ADCs are the culprits for these demodulation effects [44]. The main constituents of an ADC that contribute to its demodulation characteristics are therefore summarized in Figure 2.

According to Pelgrom, an ADC uses three basic components to convert analog signals into digital ones: a “sample- or track-and-hold circuit where the sampling takes place, the digital-

to-analog converter and a level-comparison mechanism” [70]. Level-comparison amplifiers contribute to the demodulation properties of ADCs [44] due to non-linear distortions, including *harmonics* and *intermodulation products* [71].

Harmonics are responsible for producing “spectral components at multiples of the fundamental [input] frequency” [71]. As an example, for a sinusoidal of angular frequency $\omega = 2\pi f$, harmonics transform the input $v_{in} = \hat{v} \cdot \sin(\omega t)$ into:

$$v_{out} = \left(\frac{a_2 \hat{v}^2}{2} + \frac{3a_4 \hat{v}^4}{8} + \dots \right) + \left(a_1 \hat{v} + \frac{3a_3 \hat{v}^3}{4} + \dots \right) \sin(\omega t) - \left(\frac{a_2 \hat{v}^2}{2} + \frac{a_4 \hat{v}^4}{2} + \dots \right) \cos(2\omega t) + \dots \quad (1)$$

As Equation (1) shows, the output also contains a Direct Current (DC) component, which depends solely on the “even-order nonlinear behavior” [71] of the system.

DC shifts can also be the result of reverse-biased diodes at the input of an ADC. These diodes protect the circuit from Electrostatic Discharge (ESD) by clamping negative voltages to ground, and inputs exceeding the maximum allowed voltage to V_{cc} . This behavior of ESD diodes can cause a (non-linear) DC shift [71], which attackers can also exploit [7].

Intermodulation distortions, on the other hand, arise when the input signal contains signals of two different frequencies. For example, a sum of two sinusoids $v_{in} = \hat{v}_1 \cdot \sin(\omega_1 t) + \hat{v}_2 \cdot \sin(\omega_2 t)$ may represent an adversarial signal laid on top of the legitimate sensor signal [44]. Non-linearities and trigonometric identities would then dictate that the output signal contains frequencies of the form $n\omega_1 \pm m\omega_2$ for integers n, m . As Giechaskiel et al. note, both types of “non-linearities demodulate attacker waveforms, even when they are modulated on high-frequency carriers” [44]. This fact makes them crucial for out-of-band signal injection attacks.

The final component of ADCs which is relevant to the adversarial injections studied in this article is the sample-and-hold circuitry shown in Figure 2. In its simplest form, the sample-and-hold mechanism consists of a resistor and a capacitor (*RC circuit*) connected to the input of the ADC. The transfer function of the voltage across the capacitor is therefore $H_{S/H}(j\omega) = \frac{1}{1+j\omega RC}$. This dictates that as the angular frequency ω increases, the gain $G_{S/H} = \frac{1}{\sqrt{1+(\omega RC)^2}}$ is reduced [44]. To put it differently, the sample-and-hold mechanism acts as a low-pass filter, counteracting the *aliasing* effect, which occurs when input signals are faster than half the sampling rate of the ADC (*Nyquist frequency*). One would expect this filtering behavior to reduce the vulnerability of systems to out-of-band signal injection attacks. However, the filter’s cutoff frequency in practice “is often much higher than the sampling rate of the ADC” [44], necessitating additional anti-aliasing filters before the ADC input [70].

Overall, imperfections in the sensors themselves, the ADCs, or the connections between them can result in high-frequency signals being interpreted as meaningful low-frequency ones. As modulation over high-frequency signals is often necessary to enter the targeted circuit [5], [44], the demodulation properties of ADCs allow remote attackers without physical access to inject signals into a system in an out-of-band fashion.

The subsequent sections discuss these imperfections in greater detail, with a focus on the specific method of injection.

IV. ELECTROMAGNETIC TRANSMISSIONS

The antenna-like behavior of wires and traces is extensively studied in the fields of Electromagnetic Compatibility (EMC) and Electromagnetic Immunity. This is done to ensure interoperability between the various household electrical appliances, by guaranteeing that devices neither cause nor are susceptible to undue interference [72]. Such research has shown an inverse relationship between the length of microstrip PCB traces and the frequencies at which the traces are resonant [73]. However, to better predict the response of PCB traces to external EM fields, many additional parameters are important in practice. For instance, “field incidence, polarization angles, and the magnitude and phase of the impedances loading the microstrip terminations” [74] are all useful in modeling trace behavior.

Although going into the details of such antenna models is outside the scope of this article, the effects of unintentional antennas have been central in the security community. Until now, unintentional *transmitting* antennas have been key for side-channel analysis: data-dependent emissions can reveal the information processed by a device to a remote attacker [75]. However, wires conversely acting as unintentional *receiving* antennas have only become a focal point of research more recently: out-of-band signal injection attacks have demonstrated that the antenna-like behavior of wires between sensors and microcontrollers can result in adversarial EM signals being interpreted as legitimate measurements.

As the systems targeted are not intended for communication, the phenomenon is known as *back-door coupling* [5], [76]–[79]. In back-door coupling, the “radiation couples through imperfections (apertures) in an electromagnetic shield, giving rise to a diffuse and complex field pattern within the shielded structure” [76]. Consequently, predicting the susceptibility of systems against back-door coupling is a hard task “without detailed testing, although properties averaged over frequency bands can be predicted” [80]. In other words, although the resonant behavior of simple geometric structures (e.g., lines and rectangles) has been extensively studied [81], extensive experiments are necessary to identify the extent to which intermodulation products appear [79]. Such products of diodes and other non-linear components can act as potential envelope detectors causing systems to take the wrong safety-critical actions [5]. As a result, these effects are a concern for more than just compliance with EMC regulations.

Implantable medical devices (IMDs) are an example of a safety-critical system where external electromagnetic interference (EMI) can cause physical harm to people. As a result, there is extensive research on the EMI behavior of various IMDs [82]–[94]. Some of these works have even pinpointed the properties of “non-linear circuit elements” in pacemakers as the culprits for demodulating RF signals produced by cell phones [86], [95]. However, the consequences of intentional out-of-band electromagnetic signal injection attacks on IMDs were only first identified in 2009 by Rasmussen et al. in the context of a distance-bounding protocol [33].

The proposed protocol used ultrasound transmissions to place guarantees on the distance between two communicating parties, but it was determined that an EM signal could “induce a current in the audio receiver circuit just as if the IMD received a sound signal” [33]. This would break protocol properties which depend on the speed of sound constant: adversarial transmissions propagating at the speed of light allow an adversary to operate from a longer distance. This attack is perhaps the first out-of-band electromagnetic signal injection, since it utilized EM emanations to attack an ultrasound-based protocol: unintentional antennas were found on the path “from the reception circuit to the piezo element”, which was effectively “working as a microphone” [33].

Although the attack by Rasmussen et al. was more of a side-note to an otherwise-secure protocol [33], Foo Kune et al.’s seminal 2013 “Ghost Talk” paper [5] made such adversarial injections the focal point of research. It showed that EM emissions could affect Electrocardiogram (ECG) measurements and cause IMDs to deliver fatal defibrillation shocks [5]. Foo Kune et al. succeeded in injecting arbitrary analog measurements, making a marked improvement in the literature compared to coarse replay and jamming attacks on IMDs [26], [27], [96], [97]. Their approach also significantly differed from high-power Intentional Electromagnetic Interference (IEMI) leading to the transient upset or destruction of commercial equipment [50], [76]–[78], [98]–[103].

The attack on IMDs by Foo Kune et al. [5] used low-power (≤ 10 W), low-frequency (kHz range) EM signals which coupled to the leads of ECGs and Cardiac Implantable Electrical Devices (CIEDs). In the open air, the distance achieved was up to 1.67m, but when submerging the devices in saline (to approximate the composition of the human body), successful attacks were limited to less than 10cm. The signals emitted targeted the baseband (i.e., the frequency of operation) of the IMDs directly, and thus did not make use of the non-linearities identified above and in Section III. However, this attack should still be considered out-of-band, as the IMD leads were meant to require physical contact for measurements, and should not be reacting to remote electromagnetic transmissions. In other words, this attack is more akin to coupling to a multimeter’s probes to cause wrong voltage readings remotely, rather than causing interference to a WiFi device by transmitting at 2.4GHz: although both require external stimuli, the mode of injection is different from the intended mode of operation.

Foo Kune et al. also conducted out-of-band attacks against webcams and bluetooth headsets that were up to 1m away from an 80mW source [5]. The authors transmitted modulated signals over high-frequency carriers (in the hundreds of MHz) that make use of unintentional antennas on the path between the microphone and the amplifier. Non-linearities then demodulated the input signals and produced intelligible audio output. This output overpowered legitimate conversations via the headset and fooled music identification services as well as automated dial-in services (by emulating key presses via modulated Dual Tone Multiple Frequency (DTMF) signals) [5].

Kasmi and Lopes Esteves similarly targeted smartphone microphones, but with a goal of triggering voice commands (e.g., “OK Google”, “Hey Siri”) by emitting Amplitude-

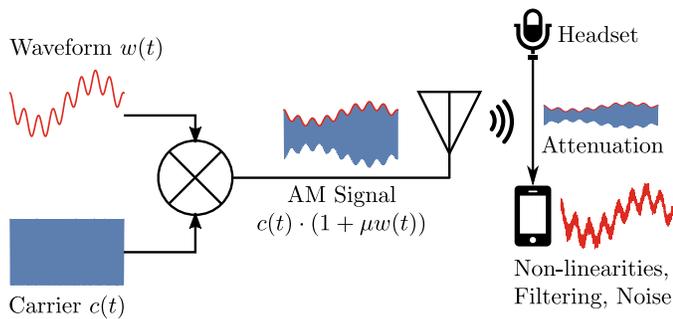


Fig. 3: Basic operating principle of an electromagnetic out-of-band signal injection attack against a microphone. Amplitude-Modulated (AM) signals are transmitted using an antenna, and are picked up (and attenuated) by headphone cables or PCB traces. Non-linearities in amplifiers coupled with low-pass filters remove the carrier wave and down-convert the target signal. The demodulated signal can be used to break ultrasound protocols [33], fool music services [5], or inject voice commands [38], despite the additional noise introduced.

Modulated (AM) signals [38]. These signals get picked up by the user’s hands-free headset, are then demodulated due to non-linearities, and finally get executed by the software voice-processing service. It is interesting to note that by default, “a long hardware button press is required for launching the service” [38]. However, a Frequency-Modulated (FM) signal at the same frequency can also emulate this headphone button press [38], allowing the attack to be fully carried out remotely.

The attack by Kasmi and Lopes Esteves used *front-door coupling*, as the “radiation couples to equipment intended to communicate or interact with the external environment” [76]. This is because headphones can be used as FM antennas and can thus not be effectively shielded. It should be noted that the field strength required was in the order of $25 - 30 \text{ V m}^{-1}$, which is close to the limit for human safety, and an order of magnitude higher than the required immunity level (3 V m^{-1}) [38]. This illustrates that high powers might still be required for reasonable attack distances: in a subsequent work, the authors noted that their attack requires a power of 40W for a distance of 2m, and 200W for a distance of 4m [40].

Figure 3 shows an example of an electromagnetic out-of-band signal injection attack, which summarizes the attacks against microphone sensors. It shows that the desired attacker waveform $w(t)$ needs to be modulated over a high-frequency carrier $c(t)$, so that the signal can be picked up with relatively low attenuation by the victim device’s wires. For example, amplitude modulation with a modulation depth $0 < \mu \leq 100\%$ can be used to couple to wired headphones. Through non-linearities in the phone’s internal amplifier, as well as low-pass filtering effects, the target waveform is preserved at the output of the digital signal processing (DSP) chip. Although this process introduces noise, the demodulated signal can still be distinguished by online music services [5], imitate voice-initiated commands, which are then executed by the phone [38], or break protocol guarantees [33]. Although the above works targeted microphones, out-of-band electromag-

netic signal injections are not sensor- or device-specific. For example, subsequent sections discuss proof-of-concept EM-based injections against temperature sensors [104], [105]. However, most attacks have primarily used amplitude modulation, leading to a question that has yet not been addressed:

Open Question 1: What is the optimal modulation scheme for out-of-band electromagnetic signal injection attacks? In other words, can Frequency Modulation (FM), Phase Modulation (ϕ M), or other schemes be used instead of Amplitude Modulation (AM)?

Another largely-unexplored research area is that of magnetic emissions. Specifically, magnetic attacks had largely been ignored in the literature, until Shoukry et al. demonstrated that it is possible to confuse Anti-Lock Braking Systems (ABS) [6]. This is done by exposing the magnetic-based wheel speed sensors to an in-band attacker-generated magnetic field at close proximity [6]. Doing so can alter speed measurements, potentially veering cars off the road [6].

Although the work by Shoukry et al. is in-band and limited in distance, it inspired subsequent work in out-of-band attacks: Selvaraj et al. recently conducted the first attacks on actuators (rather than sensors) through EM transmissions [7]. Specifically, an unmodulated sawtooth waveform was chosen to cause a “sharp decrease, for a very small amount of time” at the target servo [7]. Because the servo is controlled using Pulse Width Modulation (PWM), a waveform of the same frequency (50Hz) therefore results in a one-way (clockwise) rotation.

This attack has a few limitations: changing the attacking frequency to 60Hz causes the servo to “change positions randomly” [7]. Moreover, relatively high powers are required: a 10V (peak-to-peak) waveform is insufficient, so a 50W amplifier and a 1-to-6 step-up transformer are necessary. Moreover, one of the servo wires is wrapped around the toroid transferring the EM signal. Although “the same effect was observed when a length of the wire was placed within a solenoid”, “producing an effect at a distance requires the proper selection of a field directivity element” [7]:

Open Question 2: How can one precisely control an actuator in both directions, and at a distance?

Selvaraj et al. additionally proposed an analytical model of electromagnetic induction attacks for sensors and actuators, with a focus on the magnetic rather than the electric field [7]. To support their model, they further conducted experiments against General Purpose Input/Output (GPIO) pins of microcontrollers in analog and digital modes. They showed that 1.82W transmissions of unmodulated signals at frequencies between 0–1000MHz can result in a DC offset, even when the microcontroller is at a distance of up to 1m from the source. This indicates that an adversary can successfully inject signals over a wide range of frequencies, without having precisely determined the resonance behavior of the system.

Selvaraj et al. [7] were only concerned with the average power received and not time-dependent signals. This is in contrast to work on the demodulating effects of amplifiers

(Section V), which depends on inter-modulation products and harmonics. Instead, ESD diodes were identified as the culprits for the resulting DC offset, due to clipping non-linearities. However, it is not clear whether the same methodology can induce attacker-desired, time-varying waveforms through modulated (in amplitude or otherwise) transmissions: in the language of Giechaskiel et al. [44], Selvaraj et al. performed an *existential injection* which disturbs the ADC readings, but not a *selective injection* of attacker-chosen waveforms, unlike the earlier work of Foo Kune et al. [5].

As a final point of note, researchers have identified that coupling into the wiring interconnects within ICs is possible [74]. However, this disturbance “can be neglected up to several gigahertz”, since ICs are “usually smaller than a few centimeters” [74]. This leads to another research question:

Open Question 3: Is it possible to conduct out-of-band signal injection attacks into digital ICs which integrate sensors, ADCs, and microcontrollers?

Although this question has been answered in the affirmative for acoustic attacks (Section VI), it remains open for electromagnetic ones.

V. CONDUCTED SIGNALS

A different class of out-of-band signal injection attacks requires an indirect physical connection between the attacker and the victim, such as a shared power line. Unlike their radiated counterparts of Section IV, these *conducted* attacks do not require signals to be picked up by unintentional receiving antennas in the path between sensors and microcontrollers. Instead, signals are propagated along conductors primarily on the powering circuit, which can transfer through crosstalk or coupling to paths containing non-linearities. This propagation of electrical disturbances through structures and cables is studied in Transmission-Line Theory [106]–[108], and through the Baum-Liu-Tesche (BLT) equation [109]–[111].

Much like electromagnetic attacks, out-of-band conducted signal injection attacks have also been primarily experimental in nature. Their methodology often follows that of susceptibility literature, which predicts a device’s response to high-frequency radio signals. Systems tested include microcontrollers, ADCs, and other embedded devices which contain I/O and power pins. The goal of such research is to quantify immunity to radiated and conducted EM disturbances, and is typically concerned with the average power received by the embedded system, similar to the work by Selvaraj et al. [7] summarized in Section IV.

To avoid legal and practical considerations related to electromagnetic transmissions, the experimental approach followed is known as Direct Power Injection (DPI), and consists of injecting harmonic disturbances from a few kHz to a couple of GHz and measuring the relationship between forward power and frequency. Multiple works have shown that as the frequency of the input increases, immunity to DPI also

increases [58], [61]–[63], [65], [112].¹ In other words, higher frequencies generally require higher forward power injections for the same level of susceptibility. This was also true of the (remote) injections by Selvaraj et al. [7], discussed above.

A similar methodology can be applied to evaluate the demodulation characteristics of amplifiers and transistors [56], [57], [60], [66], and therefore better predict the fidelity with which attackers can inject target waveforms, both in the conducted and in the radiated settings. This was recently done by Giechaskiel et al. [44] for six ADCs, with a view on how to exploit the demodulating effect for out-of-band signal injection attacks. It was shown that ADCs of three different types from four manufactures, and with different resolutions and sampling frequencies can all demodulate AM waveforms [44]. Generally, it was determined that the fundamental frequency persists along with its harmonics and some high-frequency components, even for carriers which are multiple times the ADCs’ sampling and cutoff frequencies [44].

It is worth noting, however, that the different ADCs do not behave identically. For instance, some ADCs require fine-tuning of the carrier frequency, with 100Hz making a difference as to whether the injected signal is fully demodulated or not [44]. On the other hand, some ADCs are vulnerable across the spectrum, i.e., for all frequencies which do not get severely attenuated to filtering effects [44]. Moreover, as discussed in an extended version of the same paper [116], some ADCs “result in more sawtooth-like output”, and are therefore “more resilient to clean sinusoidal injections”. Finally, the same extended version demonstrates that attacks can also be performed remotely, without following DPI methodology: a 10dBm (10mW) transmission can be demodulated by a receiver amplifier at small distances (5 cm).

In their “Trick or Heat” work, Tu et al. [104] also conducted DPI experiments on operation amplifiers, but with a view on how to exploit rectification effects for out-of-band signal injection attacks on temperature sensors. They, too, determined that as the frequency increases, the magnitude of the AC voltage decreases, while “EMI signals at specific frequencies induce a significant DC offset” [104]. Moreover, for a given frequency of injection, power and the induced DC offset are “locally proportional”, though the rate of change “gradually decreases as the power of injected EMI signals grows” [104]. However, power and DC offset are not always positively correlated, even for remote transmissions: for some frequencies, the induced DC offset is negative [104].

Having characterized the behavior of individual amplifiers, Tu et al. turned their attention to different types of thermal sensors, including Negative Temperature Coefficient (NTC) thermistors, shielded and unshielded K-type thermocouples, and Resistance Temperature Detectors (RTDs). With a 35dBm (3.2W) electromagnetic source, Tu et al. succeeded in changing the reported temperature of various devices by at least 0.5°C [104]. The systems attacked included, among others, newborn incubators, soldering irons, and 3D printers, which were placed at distances of up to 6m.

¹Immunity behavior is different for EMI-induced offsets through the ground plane for amplifiers [113] and precision voltage references [114]. See the survey by Ramdani et al. [115] for more information on RF immunity models.

In some experiments, a thick wall was present between the transmitting device and an infant incubator under attack. It was shown that, even in this setup, an adversary can increase the measured skin temperature by 3.4°C or decrease it by 4.5°C [104], again demonstrating the potentially fatal consequences of out-of-band attacks. Most of the attacks by Tu et al. used unmodulated transmissions, and therefore only looked at the relationship between frequency and DC offset, or the relationship between power and DC offset. However, when investigated jointly, amplitude modulation was capable of causing selective injections [44]. In other words, Tu et al. [104] spelled “HI” in the output of the temperature sensor by appropriately modulating the amplitude of the transmission.

In a different strand of research, Lopes Esteves and Kasmi demonstrated how to inject voice commands (“OK Google”) into a smartphone through conducted means [40]. Specifically, the attack exploited the fact that on the device’s circuit board, the phone’s USB charging port is physically close to the audio frontend, where demodulation (envelope detection) can take place due to non-linearities [40]. As a result, back-door coupling occurs, either due to “a re-radiation of the interference from the USB circuitry bypassing the physical isolation by parasitic coupling (crosstalk) or the possible sharing of the V_{cc} and GND networks on the PCB” [40].

Open Question 4: What properties of the power circuit and related layout considerations make systems vulnerable to conducted out-of-band signal injections?

The methodology used was inspired by experiments on the propagation of conducted disturbances and on EM injections into power cables. Specifically, amplitude-modulated signals were injected at various locations of the power network, i.e., at different plug points on the same strip and on extension cords. The phone was left charging either on a computer USB port, or through a wall adapter. Experiments were repeated both with a magnetic injection probe (directly coupling to cables), and a “custom coupler made with capacitors, resistors and a high-frequency transformer” [40]. In all cases, it was determined that the smartphone can demodulate (and execute) commands carried on the 200–250 MHz range at distances up to 10m, even with only a 0.5W source. Such conducted attacks therefore significantly lower the power requirements and increase the injection distance compared to the same authors’ remote EM attack on smartphones [38].

True Random Number Generators (TRNGs) which are based on Ring Oscillators (ROs) are also vulnerable to conducted signal injection attacks. ROs are composed of an odd number of logical NOT gates chained together in a ring formation, where the output of the last gate is used as the input to the first gate. The value of the loop between any two stages of the RO oscillates between true and false. The frequency of oscillation is influenced by the delay of the logic gates and the delay between the RO’s stages, which are in turn influenced by small variations in the manufacturing process, as well as voltage, and temperature (PVT) [117].

As a result, by XORing several ring oscillators together, one can exploit the randomness of the phase jitter to create

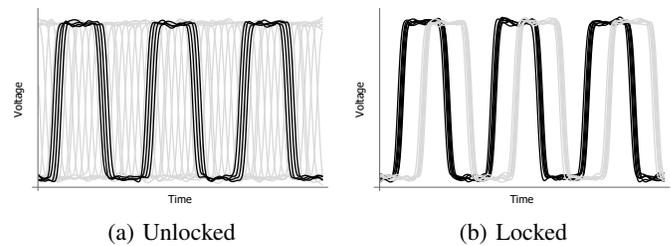


Fig. 4: Example waveforms for two Ring Oscillators (ROs) with frequency locking (a) absent or (b) present.

a TRNG [118]. However, due to the frequency dependence on voltage, a suitable signal can lead to frequency locking of the oscillators [119], [120], removing the differences in the randomness of the jitter. Markettos and Moore first conducted the attack in practice in 2009 by directly injecting 24 MHz signals into the power supply of two ring oscillators composed of discrete logic chips [35]. Moreover, they succeeded in biasing TRNGs even in secure microcontrollers and smartcards: a sinusoidal wave of 1 V peak-to-peak (2.5 mW) at 24.04 MHz was enough to cause a 5 V EMV Chip and Pin smartcard to fail statistical tests of randomness [35].

Time-varying signals are not always necessary: a constant (DC) power supply voltage can also lead to locking of ring oscillators in an under-volted Field-Programmable Gate Array (FPGA). This is because there is a “dependence of the frequency of one oscillator on the current peaks caused by rising and falling edges of the second oscillator” [121]. Figure 4 illustrates what happens when two ring oscillators frequencies lock: during normal operation (Figure 4a), the ring oscillator values “slide past each other, minimising the likelihood of two rings transitioning together” [35]. However, when a frequency- or voltage-based attack causes the ROs to lock (Figure 4b), their relationship becomes predictable, biasing the TRNG.

It should be noted that although under-/over-power attacks are usually considered fault attacks, in this case the ring oscillators are still functioning properly, but the entropy of the TRNG is reduced due to less jitter present [39], [122]. An interesting new class of such remote under-voltage attacks on TRNGs has recently surfaced. Because ring oscillators have the potential to increase the delay of FPGA elements by causing voltage drops, they can also cause timing violations, thereby reducing the randomness of TRNGs [123]. Such an attack does not require equipment for physical injections. Instead, the adversary only needs co-located (but logically and physically isolated) circuits on the same FPGA as the target TRNG. This setup reflects multi-tenant cloud designs [123], and presents new challenges for the protection of shared FPGAs against software-only attacks without physical access.

Although the above attacks generally alter the power supply directly, the same outcome can be achieved through EM emanations targeting the wires connecting the various stages of the ring oscillators [36], [37], [49]. This requires micro-probes at very close proximity to the ring oscillators (in the order of $100\mu\text{m}$ from the FPGA packaging), so as to localize the effects of the injection [36], [49]. However, TRNGs are also vulnerable against EM injections into power supply cables:

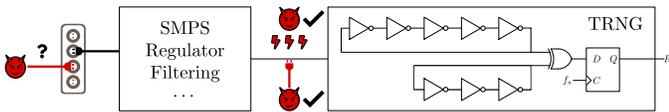


Fig. 5: True Random Number Generators (TRNGs) based on Ring Oscillators (ROs) are vulnerable to frequency locking: electromagnetic and conducted signals into power supply cables can bias the randomness outputs. Are attacks through a shared mains power supply network also possible?

Osuka et al. demonstrated that an injection probe wrapped around the DC power supply cable of a TRNG can also bias the TRNG [42]. Although the design only used two ring oscillators composed of discrete logic chips,² injections were successful even when the probe was placed at a distance of 40cm from the ROs, with a power of only 25.2dBm (0.33 W).

It is worth highlighting that although the conducted voice command injection attack by Lopes Esteves and Kasmi was performed over shared power lines [40], all existing attacks on TRNGs bypass AC-to-DC rectification and voltage regulation. This leads to the following question for future research:

Open Question 5: Is it possible to bias True Random Number Generators through conducted out-of-band signal injection attacks on the primary side of power supplies (mains voltage), as shown in Figure 5?

VI. ACOUSTIC EMANATIONS

Research into out-of-band acoustic signal injection attacks has primarily focused on: (a) attacking electro-mechanical devices by causing vibrations at their resonant frequencies; and (b) exploiting microphone non-linearities for inaudible voice commands. In the former category, Micro-Electro-Mechanical Systems (MEMS) gyroscopes and accelerometers have been a popular target for acoustic resonance attacks.

MEMS gyroscopes operate through “vibrating mechanical elements to sense rotation” [124]. In other words, MEMS gyroscopes contain oscillating structures which, when rotated, appear to have a measurable force (called the Coriolis force) exerted on them [125]. These mechanical resonators “generate and maintain a constant linear or angular momentum”, so that “when the gyroscope is subjected to an angular rotation, a sinusoidal Coriolis force at the frequency of drive-mode oscillation is induced in the sense direction” [125]. This force is exerted in a different direction from the moving direction, and, depending on the type of the gyroscope, the angular rotation can be estimated through changes in capacitance, piezoresistive effects, etc. [124], [125]. Out-of-band acoustic signal injections transmit sounds at the resonant frequencies of MEMS sensors, causing them to report incorrect values.

Early research into the properties of MEMS gyroscopes had shown that high-power acoustic noise at or near the resonant

frequency can degrade the performance of the sensor [126]–[128]. However, the security effects of intentional sound transmissions were not explored until the 2015 “Rocking Drones” paper by Son et al. [10]. Initially, the effect was a simple denial-of-service (DoS) attack on drones. It was caused by the transmission of single-tone sound waves at the resonant frequency of drones’ gyroscopes, so there was no control over their movements. The distance was also short, at 10cm using a speaker producing a Sound Pressure Level (SPL) of up to 113dB at the target frequencies.

Proof-of-concept control was then first demonstrated in a Black Hat conference presentation against gyroscopes in virtual reality (VR) headsets and self-balancing vehicles [34]. Tu et al.’s “Injected and Delivered” paper later became the first academic work to control gyroscopes in a much more fine-grained fashion [43]. This research allowed control for long periods of time (up to the minute range), and at long distances (up to 7.8m with a maximum SPL of 135dB) [43].

Tu et al. noticed that single-tone frequencies (like the ones used by Son et al. [10]) result in an oscillating discrete (digitized) output, which destabilizes equipment. In other words, a simple transmission at the resonant frequency is a type of DoS attack because the angular velocity (as measured by the gyroscope) fluctuates between positive and negative values (Figure 6a). However, it is possible to remove these negative components by decreasing the transmission amplitude during the corresponding measurements. This is called a *Side-Swing* attack, which “proportionally [amplifies] the induced output in the target direction and attenuate[s] the output in the opposite direction” [43] (Figure 6b).

Instead of attenuating signals during half of the transmission period, one can also control “the induced output by manipulating the phase of the digital signal with repetitive phase pacing” in a *Switching* attack [43] (Figure 6c). As this is accomplished in practice by changing the tonal frequency instead of attenuating the amplitude, a Switching attack contributes twice as much to the overall change in direction as a Side-Swing attack. This is shown by looking at the accumulating heading angle in Figure 7. It should be noted that by accounting for drifts in the sampling rate of the ADC (which are amplified during adversarial injections [43]), both attacks can control the gyroscopic output for longer periods of time.

In response to the rising interest in acoustic vulnerabilities, Khazaaleh et al. [129] created a mathematical model to explain the resonance response of gyroscopes. They showed that “the misalignment between the sensing and driving axes of the gyroscope is the main culprit behind the vulnerability of the gyroscope to ultrasonic attacks” [129]. More precisely, because “the sensing direction is not exactly orthogonal to the driving direction, some of the energy gets coupled to the sensing direction” [129]. This causes a false reading, which is typically corrected “by employing a demodulator in the readout circuit” [129]. When the transmission frequency is slightly different from the sensing frequency, the gyroscope generates “measurable output”, whose frequency equals “the difference between the driving frequency and the frequency of the acoustic signal” [129]. As shown experimentally, this model also explains why it is better to transmit near the reso-

² Much like the original work by Marketos and Moore [35]. However, Bayon et al. [36], [49] targeted a more realistic TRNG composed of 50 ROs.

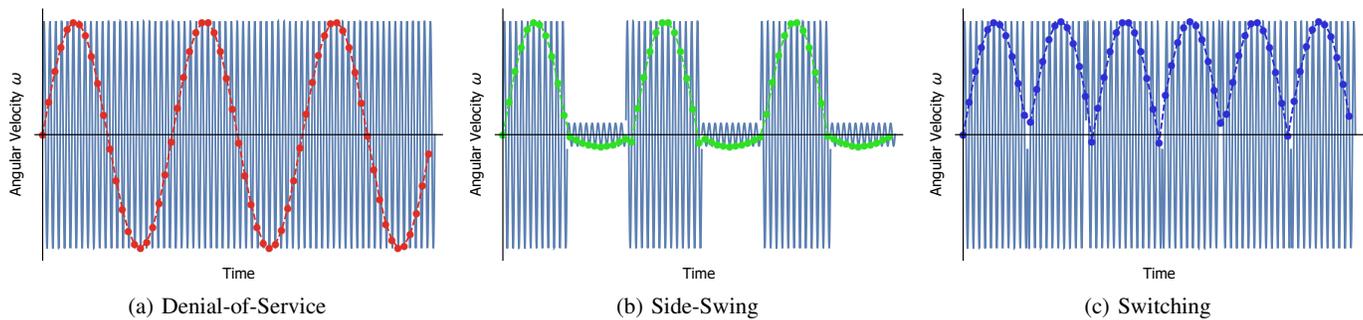


Fig. 6: Different acoustic injection approaches against gyroscopes. For (a) a Denial-of-Service (DoS) attack, a single-tone transmission at the gyroscope’s resonant frequency suffices. This results in oscillating digital measurements of angular velocity, and can destabilize equipment [10], [43]. To remove the negative measurement components, one can either (b) decrease the transmission amplitude in a *Side-Swing* attack [43], or (c) change the frequency of transmission for a *Switching* attack [43]. The cumulative effects of these approaches are shown in Figure 7.

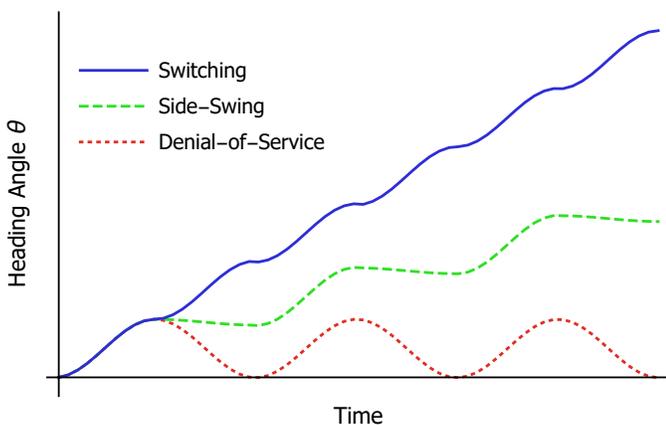


Fig. 7: Effects of the three attacks of Figure 6. In a Denial-of-Service (DoS) attack, the accumulating heading angle fluctuates, while it continues increasing for both attacks proposed by Tu et al. [43]. In Switching attacks, the angle increases at twice the rate of Side-Swing attacks.

nant frequency rather than exactly at it [129]. It also suggests that low-pass filters or differential measurements through additional proof masses are ineffective countermeasures against out-of-band acoustic signal injection attacks [129].

Although Tu et al. succeeded in controlling gyroscopes in phones, scooters, stabilizers, screwdrivers, and VR headsets among others, only DoS attacks were successful against accelerometers [43]. MEMS accelerometers consist of spring-mass systems, so acceleration results in a deflection of the seismic mass. This deflection “is detected by means of capacitive elements, the capacitances of which change with deflection, or by piezoresistive elements that detect strain induced by the motion of the seismic mass through a change in resistor values” [130]. Acoustic vibrations at the resonant frequencies of the spring-mass systems can also displace the suspended mass, making them vulnerable to out-of-band attacks. According to Trippel et al., insecure amplifiers and low-pass filters (LPFs) prior to the accelerometer ADCs can demodulate both Amplitude-Modulated (AM) and Phase-Modulated (ϕ M) attacker injections [13]. These insecurities are the results of

clipping non-linearities and permissive filtering respectively, and allow for both biasing and control attacks.

Trippel et al. used accelerometers to spell words, naming their work “WALNUT” for the output of the spoofed sensor measurements. They were also able to control off-the-shelf devices, such as remote-controlled (RC) cars, and Fitbit fitness tracking wristbands [13]. Although spoofing step counts might seem innocuous, companies often offer financial rewards for health-related activity [13], so cheating devices (which do not yet exploit out-of-band effects) are already being sold [131]. Most attacks by Trippel et al. [13] were performed at distances of 10cm, with a speaker producing an SPL of 110dB. The duration of control over the output of the MEMS sensors was often limited to 1–2s (and up to 30s) due to sampling rate drifts. Moreover, it was shown that the three axes do not behave identically to acoustic injections. For example, there are some MEMS devices for which only the x -axis responds to acoustic transmissions, while others are vulnerable in all three axes, but at different resonant frequencies.

Although Trippel et al. [13] attacked a single sensor in one direction at a time, Nashimoto et al.’s “Sensor CON-Fusion” investigated whether *sensor fusion* using a Kalman Filter can improve the robustness of measurements [41]. It was shown that “while sensor fusion introduces a certain degree of attack resilience, it remains susceptible” to combined acoustic and electromagnetic injections [41]. Specifically, Nashimoto succeeded in simultaneously controlling the roll, pitch, and yaw (the three angular axes in aircraft nomenclature) by fusing the outputs of an accelerometer, a gyroscope, and magnetometer. However, in non-simulated environments, “there is an error in the roll angle”, and “the resulting inclination does not last long” [41]. Although fusion is further explored in the context of defense mechanisms (Section IX), the above discussion leads to the following research question:

Open Question 6: Is it possible to use acoustic injections to precisely control MEMS gyroscope and accelerometer measurements in all three directions simultaneously and/or for longer periods of time?

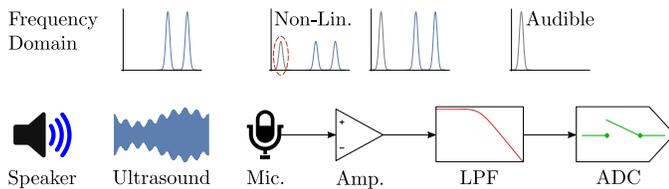


Fig. 8: High-level overview of ultrasonic attacks against microphones [11], [55], [133], [134]. A speaker transmits inaudible tones, but intermodulation products are produced due to non-linearities in the microphone and amplifier. A Low-Pass Filter (LPF) removes ultrasound frequencies, so the Analog-to-Digital Converter (ADC) records only audible by-products.

In a different strand of research, ten years after a video demonstrating that shouting in a data center causes unusually high disk I/O latency [132], Shahrads et al. showed that acoustic transmissions can cause vibrations in Hard-Disk Drives (HDDs) [53]. These vibrations result in read and write errors at distances up to 70cm with a sound level of 102.6dBA [53]. As a result, they can make systems unresponsive, even leading to Blue Screen of Death (BSOD) errors [53].

Although Shahrads et al. primarily focused on the effect of the angle of transmission [53], research conducted in parallel more precisely pinpointed the root cause of the issue using Finite Element Analysis [12]. Specifically, it was shown that (audible) acoustic waves “can displace a read/write head or disk platter outside of operational bounds, inducing throughput loss”, even though the displacement is of only a few nanometers [12]. In addition, in their “Blue Note” paper, Bolton et al. also used ultrasonic transmissions to attack the shock sensors which are meant to protect HDDs during sudden drops by “parking the read/write head” [12]: modern HDDs contain “piezo shock sensors or MEMS capacitive accelerometers” to “detect sudden disturbances” [12], and can also be attacked through ultrasound transmissions at their resonant frequencies. Through these malicious acoustic attacks, SPL levels of up to 130dB cause HDDs at distances of 10cm to become unresponsive, thus disabling laptops and video recorders [12].

Not all acoustic attacks transmit at resonant frequencies. By contrast, other research targets microphone non-linearities to cause inaudible sound to be recorded. Microphones operate by converting the mechanical deformation of a membrane (caused by the air pressure of a sound wave) into a capacitive change, which produces an Alternating Current (AC) signal [11]. This process also has non-linearities which produce second-order components [11], including harmonics and intermodulation products, as discussed in Section III.

Early work on acoustic attacks primarily focused on adversarial control of machine learning in speech recognition systems [135], [136]. Such research did not take advantage of non-linearities, and was in-band, as the transmissions were audible (although indecipherable by humans). However, later investigations revolved around ensuring that the transmitted frequencies are beyond the human-audible range (20kHz). This was first accomplished by Zhang et al. in their “DolphinAttack” paper, which exploited non-linearities in microphone sensors [11], as shown in Figure 8.

The work by Zhang et al. showed how to transform (modulated) ultrasound transmissions into valid commands which were executed by speech recognition systems such as Apple Siri, Google Now, Microsoft Cortana, and Amazon Alexa [11]. The same authors later expanded on their attack by testing different setups, and increasing the attack distance from 1.7m to 19.8m [134]. This was done by replacing the 125dB source with an ultrasonic transmitter array and amplifier outputting 1.5W to increase sound pressure. Above this transmission power, the attack becomes audible due to non-linearities in the transmission medium (air) and the source speakers [11].

Earlier, Roy et al. had also noted that non-linearities in speakers make it harder for an adversary to increase the attack distance: “increasing the transmit power at the speaker triggers non-linearities at the speaker’s own diaphragm and amplifier, resulting in an audible [output]” [55]. Instead, multiple speakers in the form of an ultrasonic speaker array can be used to attack voice recognition systems including Amazon Alexa and Google Now at a distance of up to 7.6m using a 6W source [55]. The attack works by partitioning the audio spectrum across the various speakers in a way that “reduces the audible leakage from any given speaker” while minimizing the total leakage power [55]. This prevents any of the non-linearities (and the transmitted signal itself) from being audible. It should be noted that if multiple non-cooperating ultrasonic sources are emitting simultaneously, intermodulation distortions can create audible byproducts [68], allowing for the detection of potential attacks.

Parallel to the 2017 DolphinAttack paper [11], similar research was in progress at Princeton [133]: Song and Mittal also succeeded in injecting inaudible voice commands to an Amazon Echo and an Android phone. Although they accomplished relatively long distances (3.54m with an input power of 23.7W), their work remains in poster format. Moreover, prior to their inaudible voice commands work [55], Roy et al. also used inaudible ultrasound transmissions to record audible sounds. Specifically, in their “BackDoor” work, they proposed a high-bandwidth covert channel that operates at up to 1.5m using a 2W source [54]. Although covert channels are not discussed in this survey, the work by Roy et al. is included due to the methodology which naturally led to their later work.

More concretely, instead of using amplitude modulation over a single frequency like Zhang et al. [11], Roy et al. simultaneously played two ultrasound tones whose shadows create audible sounds (only sensed by microphones) due to non-linearities [54]. They showed that amplitude modulation could not be used due to non-linearities in the ultrasound transmitters themselves, which would result in audible signals. Instead, further pre-computation was required to remove the “ringing effect”, where “the transmitted sound becomes slightly audible even with FM modulation” [54]. These results point towards the next open question [134]:

Open Question 7: How can non-linear acoustics in the transmission medium (air and speakers) be avoided to further extend the range of inaudible attacks?

VII. OPTICAL AND THERMAL MANIPULATIONS

Although electromagnetic, conducted, and acoustic attacks form the majority of out-of-band signal injection attacks, there has been some research on optical attacks, as well as temperature attacks which bias RO-based TRNGs. In the former category, attacks exploit permissive filtering and poor shielding in interfaces which only expect ambient environmental conditions. Most papers so far have targeted sensors in an in-band fashion: out-of-scope research includes attacks on sonars, radars, and LiDARs [9], [14], [31], [32], as well as visible-light attacks on cameras in unmanned aerial vehicles (UAVs) [52] and cars [9], [14], [32].

Researchers have also hypothesized that excessive light injections would blind car cameras and confuse auto-controls in automated vehicles [137]. Indeed, limited success against cameras has been achieved using Ultraviolet (UV) and Infrared (IR) lasers up to 2m away [14]. However, attacks were only possible in dark environments, and the results were not reproducible with invisible lasers against other makes and models of CMOS cameras [9].

In a different strand of research, Park et al. showed that some medical infusion pumps are not well-protected against adversarial optical injections [8]. Specifically, in order to measure the flow rate of the medicine being administered, pumps are fitted with drip measurement sensors. These sensors consist of an IR emitter and receiver facing each other. When a drop passes through the sensor, the IR receiver temporarily senses less light due to diffusion, allowing for the rate to be measured. However, because the sensor is not well-enclosed, an adversary can shine an IR laser into the sensor, causing these drops to be undetected, thereby saturating the sensor. By then un-blinding the sensor, the attackers can also trick the firmware into detecting fake drops, and bypass alarms.

Adversaries can therefore selectively both over- and under-infuse a patient for an extended period of time, and for a variety of normal flow volumes. However, this can only be done with coarse-grained control over the real flow rate. Most of the experiments by Park et al. were conducted at a distance of 10cm. Success was nevertheless reported up to 12m away using a 30mW IR laser pointer [8]. These results show that optical attacks can reach meaningful distances, but are, of course, limited by line-of-sight considerations.

It should be noted that the attack by Park et al. should be considered to be out-of-band, as the pump was not meant to receive external stimuli, in contrast to, for example, LiDARs, which are supposed to interact with external objects. In other words, a LiDAR depends on its surroundings to reflect its transmitted pulses and therefore infer the distance to the interfering objects. On the other hand, the drip sensor and (part of) the intravenous (IV) tube could be enclosed and shielded from the environment. This naturally leads to the defense mechanisms proposed by Park et al. [8], which are discussed in Section IX, and which beg the following question:

Open Question 8: Are all out-of-band optical attacks a matter of poor filtering and shielding?

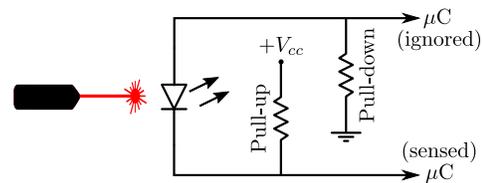


Fig. 9: Unconventional circuit layout to convert a Light Emitting Diode (LED) into a photodiode [138]. Pull-up and -down resistors can be internal to microcontroller input pins.

Under somewhat unrealistic assumptions, the answer to the above question might be “no”. In their typical mode of operation, Light Emitting Diodes (LEDs) convert current to light. However, they can also function in the reverse by producing current when illuminated [139], [140]. In preliminary experiments, Loughry recently showed that this behavior can be exploited for an optical covert channel [138]. The setup is certainly unconventional: both LED pins are connected to microcontroller GPIO pins, which are configured as inputs with (internal) pull-up and pull-down resistors, as shown in Figure 9. According to Loughry, seven out of ten LEDs tested responded to laser and light of different wavelengths, with some LEDs producing measurable current at both ends (anode and cathode) [138]. Whether out-of-band signal injections or fault attacks can exploit this effect is yet to be seen [138].

The final class of attacks exploits the dependence of ring oscillators (ROs) on temperature to reduce the entropy of the TRNG. It is only mentioned here for completeness, as distance requirements would dictate physical access to the device under attack. Early work in the area showed that statistical randomness tests of RO-based TRNGs would fail for certain FPGA temperatures [141]. More detailed experimentation conducted a few years later using different heat-transfer methods (resistor heater, Peltier cooler, and liquid nitrogen) then showed that “the hotter the temperature, the larger the bias” [142]. Martín et al. [39] also investigated the effect of temperature across multiple TRNG designs based on Self-Timed Rings (STRs), which do not exhibit the frequency locking effects discussed earlier [143]. It was shown that the effects of temperature increases on the randomness of STR-based TRNGs were not significant, due to a combination of a decrease in frequency and an increase in jitter due to thermal noise [39].

Since devices were operating within their specifications, these biasing effects could be considered out-of-band attacks which operate at limited distances. This is in contrast to, for example, Martín et al.’s work which investigates the entropy of TRNGs in response to ionizing radiation [144]. However, remote conducted attacks on TRNGs are possible using local voltage drops [123]. Whether it is possible to reproduce these effects using heating circuits [145] remains an open question:

Open Question 9: Can software-only thermal effects bias RO-based TRNGs in multi-tenant FPGAs?

Overall, the efficacy of out-of-band optical thermal injections has been limited so far, due to the limited nature of vulnerable interfaces and proximity considerations respectively.

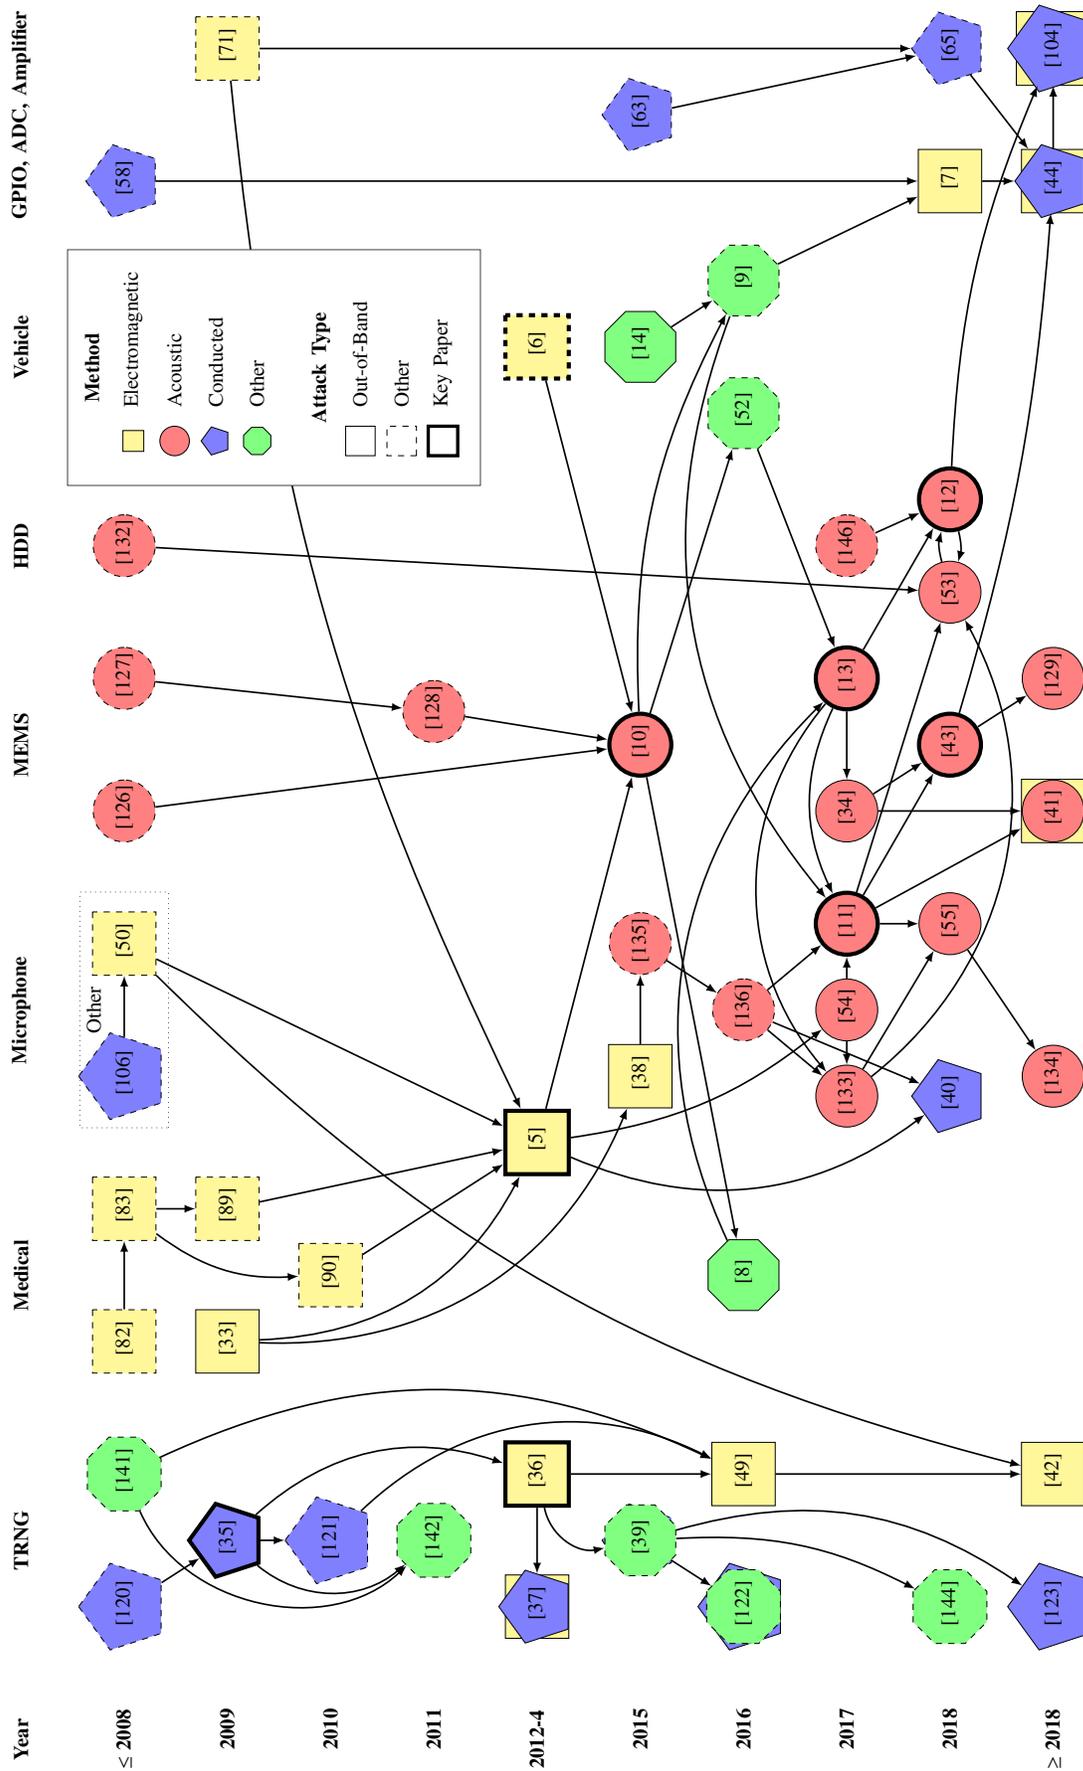


Fig. 10: Evolutionary and thematic taxonomy of out-of-band signal injection attacks, categorized by topic and methodology. Cross-influences between out-of-band signal injections and some key related works are also shown. Highly influential papers are discussed more extensively in Section VIII.

VIII. TAXONOMY OF ATTACKS

This section presents a taxonomy of out-of-band signal injection attacks, tracing their evolution through time and topic, and identifying commonalities in their methodology and source of vulnerability. We first highlight thematic and evolutionary cross-influences between the various works studied in this survey (Figure 10), and then categorize the key hardware imperfections that make them possible (Table II). We start with a citation graph of out-of-band attacks and related work in Figure 10, where an edge $X \rightarrow Y$ indicates that X is cited by Y . To reduce clutter, if a paper X is cited by both Y and Z , and Y is cited by Z , then no arrow from X to Y is drawn.

Figure 10 reveals that cross-influences are not limited to attacks. Instead, there is a general trend of earlier research observing the effects of non-adversarial interference, with later work actively exploiting the same phenomenon for signal injection attacks. For example, multiple works had identified the effects of electromagnetic interference on medical devices [82], [83], [89], [90] (with more papers discussed in Section IV), but Foo Kune et al. [5] were the first to recognize the effect as a security concern rather than a safety and reliability one. Similarly, Dean et al. commented on the effect of acoustic noise on gyroscopes [127], [128], but Son et al. [10] later used the same effect to destabilize drones.

Observation 1: Out-of-band signal injection attacks identify the effect of noise on systems and find novel ways to amplify it through hardware imperfections.

The graph further shows that out-of-band signal injection attack topics and methodology has both been inspired by and has itself inspired research exploiting more traditional avenues of attack. For example, investigations into acoustic signal injection attacks [11], [54], [55] have been influenced by research exploiting machine learning algorithms which respond to commands which are audible but indecipherable by humans [135], [136]. In addition, out-of-band acoustic attacks on MEMS sensors [10], [13] have both inspired and been influenced by in-band optical attacks on unmanned aerial vehicles (UAVs) [52]. Moreover, after the effects of acoustic noise on the security of gyroscopes were first identified [10], subsequent work improved the level of control over the output [43], and attacked additional types of MEMS sensors [13], and electro-mechanical devices such as HDDs [12].

Observation 2: After an exploitable hardware imperfection has been identified, determining its root cause opens up new avenues of attacks across different domains, and with alternative methodologies.

Figure 10 further highlights a few key works which in the opinion of the authors have played a central role in the development of the field. The first set of such papers [5], [10], [11], [35], [36] was chosen due to the high number of citations they have received overall (≥ 95) and from other out-of-band attack research (≥ 5).³ Specifically, the works by Markettos

and Moore [35] and by Bayon et al. [36] were chosen because they successfully biased TRNGs in the conducted and EM settings respectively, going beyond earlier theoretical work on oscillator locking [119], [120]. Their work led to the development of a new branch of attacks, which has so far developed rather independently of other out-of-band signal injections, as shown in Figure 10.

Foo Kune et al.’s work on adversarial electromagnetic interference [5] also features prominently in Figure 10, having been cited by almost all out-of-band signal injection attacks that were published after it (with the exception of TRNG research). Foo Kune et al. were the first to successfully exploit non-linearities and unintentional antennas in remote electromagnetic injection attacks, which were non-adversarial in prior work (e.g., [82], [83]) or only mentioned in passing [33].

With over 160 citations since 2017, the research by Zhang et al. [11] has been very influential in the realm of out-of-band acoustic attacks against microphones. Unlike earlier work on covert communication [54], and indecipherable-yet-audible commands [135], [136], Zhang et al.’s “DolphinAttack” exploited microphone non-linearities for inaudible injections.

Research on acoustic attacks is perhaps more mature against MEMS sensors, in great part due to early work by Son et al., who first showed how to disrupt gyroscopes [10]. Moreover, Trippel et al.’s research has also significantly furthered the state-of-the-art in acoustic injections by controlling the output of accelerometers for short periods of time [13]. As a result, the work of Tu et al., which showed how to extend the duration of control [43], is in the second set of works highlighted in Figure 10. This set contains recent studies whose novelty and potential has not yet received mainstream attention.⁴ For example, the techniques proposed by Tu et al. to overcome ADC sampling rate drifts should be applicable to other methods of injection, and against different types of targets. We place the work by Bolton et al. [12] in the same category, as it managed to bridge research on HDD attacks (e.g., [53]) with attacks on MEMS sensors, and contained significant insights into why resonance attacks against hard drives work.

The final work highlighted in Figure 10 is the in-band attack of Shoukry et al. against an Anti-lock Braking System (ABS) [6]. It has been included not just for its high citation count (> 135 , of which 10 are out-of-band attacks), but because it is the first EM paper to focus on the magnetic field rather than the electric field. As a result, it serves as inspiration to recent out-of-band magnetic attacks [7], [41], which we hope will be explored more in the future (Section XI).

We also summarize the various out-of-band signal injection attacks along with factors which contribute to them in Table II. The table further notes the maximum power used and distance achieved for an attack, including the level of control over the resulting signal. Effects range from theoretical attacks which are only partially realized to practical attacks which disrupt, bias, or completely control the output. As Table II indicates, information on the attack setup was often hard to find, sometimes completely missing, and often had to be identified by

³Citation counts are current as of 7 Nov. 2019 according to Google Scholar.

⁴Trippel et al.’s 2017 work [13] lies between the two categories, already having over 75 citations, 12 of which are from other attack papers.

TABLE II: Sources of vulnerability, methods, and effects for out-of-band signal injection attacks along with maximum power and distance. For each potential source of vulnerability, \checkmark signifies that the paper claims the vulnerability contributes to the attack, while $-$ that it does not. The attack methods used are 🔊 acoustic, ⚡ conducted, 📶 electromagnetic, and 🔦 optical. The effects achieved include 🚫 disruption, 🎯 bias, and 🎯 control of the output, while ◯ denotes a theoretical or partially-realized attack. Resonance (\dagger) includes frequency locking for ring oscillators, and unintentional wire antennas for EM attacks.

Authors	Year	Ref.	Target	Power	Distance	Method	Effect	Resonance (\dagger)	Non-Linearity	Improper Filtering	Poor Shielding	Insecure Algorithm
Rasmussen et al.	2009	[33]	Medical	(unspecified)	(unspecified)	📶	◯	\checkmark	$-$	$-$	\checkmark	$-$
Petit et al.	2015	[14]	Camera	(unspecified)	2.00m	🔦	🎯	$-$	$-$	\checkmark	$-$	$-$
Cao et al.	2016	[122]	TRNG	(under-volt)	(direct)	⚡	🎯	\checkmark	$-$	$-$	$-$	$-$
Wang et al.	2017	[34]	MEMS	(unspecified)	(unspecified)	🔊	🎯	\checkmark	\checkmark	\checkmark	\checkmark	$-$
Nashimoto et al.	2018	[41]	MEMS	(unspecified)	(unspecified)	🔊📶	🎯	\checkmark	$-$	\checkmark	\checkmark	\checkmark
Mahmoud and Stoilović	2019	[123]	TRNG	(140K ROs)	(internal)	⚡	🎯	\checkmark	$-$	$-$	\checkmark	$-$
Markettos and Moore	2009	[35]	TRNG	0.002 W	(direct)	⚡	🎯	\checkmark	$-$	\checkmark	$-$	\checkmark
Bayon et al.	2012	[36]	TRNG	0.003 W	100 μm	📶	🎯	\checkmark	$-$	$-$	$-$	$-$
Buchovecká et al.	2013	[37]	TRNG	≤ 0.563 W	“near”	⚡📶	🎯	\checkmark	$-$	$-$	$-$	$-$
Foo Kune et al.	2013	[5]	Med. & Mic.	10.000 W	1.67m	📶	🎯	\checkmark	\checkmark	\checkmark	\checkmark	$-$
Kasmi and Lopes Esteves	2015	[38]	Microphone	200.000 W	4.00m	📶	🎯	\checkmark	$-$	$-$	\checkmark	\checkmark
Bayon et al.	2016	[49]	TRNG	0.003 W	100 μm	📶	🎯	\checkmark	$-$	$-$	$-$	$-$
Park et al.	2016	[8]	Medical	0.030 W	12.00m	🔦	🎯	$-$	$-$	$-$	\checkmark	\checkmark
Roy et al.	2017	[54]	Microphone	2.000 W	1.50m	🔊	◯	$-$	\checkmark	$-$	$-$	$-$
Song and Mittal	2017	[133]	Microphone	23.700 W	3.54m	🔊	🎯	$-$	\checkmark	$-$	$-$	$-$
Lopes Esteves and Kasmi	2018	[40]	Microphone	0.500 W	10.00m	⚡	🎯	\checkmark	\checkmark	\checkmark	$-$	\checkmark
Osuka et al.	2018	[42]	TRNG	0.331 W	0.40m	⚡📶	🎯	\checkmark	$-$	$-$	$-$	$-$
Roy et al.	2018	[55]	Microphone	6.000 W	7.62m	🔊	🎯	$-$	\checkmark	$-$	$-$	\checkmark
Selvaraj et al.	2018	[7]	GPIO	1.820 W	1.00m	📶	🎯	\checkmark	\checkmark	\checkmark	\checkmark	$-$
Giechaskiel et al.	2019	[44]	ADC	0.010 W	0.05 m	⚡📶	◯	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
Tu et al.	2019	[104]	Temp. & Amp.	3.162 W	6.00m	⚡📶	🎯	\checkmark	\checkmark	\checkmark	\checkmark	$-$
Yan et al.	2019	[134]	Microphone	1.500 W	19.80m	🔊	🎯	$-$	\checkmark	$-$	$-$	\checkmark
Son et al.	2015	[10]	MEMS	113 dB	0.10m	🔊	🎯	\checkmark	$-$	$-$	\checkmark	$-$
Trippel et al.	2017	[13]	MEMS	110 dB	0.10m	🔊	🎯	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
Zhang et al.	2017	[11]	Microphone	125 dB	1.75m	🔊	🎯	$-$	\checkmark	$-$	$-$	\checkmark
Bolton et al.	2018	[12]	HDD	130 dB	0.10m	🔊	🎯	\checkmark	$-$	$-$	\checkmark	\checkmark
Shahrad et al.	2018	[53]	HDD	103 dB	0.70m	🔊	🎯	\checkmark	$-$	$-$	\checkmark	$-$
Tu et al.	2018	[43]	MEMS	135 dB	7.80m	🔊	🎯	\checkmark	$-$	\checkmark	\checkmark	\checkmark
Khazaaleh et al.	2019	[129]	MEMS	94 dB	0.11m	🔊	🎯	\checkmark	$-$	$-$	\checkmark	\checkmark

looking up the datasheets of the signal generators, antennas, and amplifiers used. This lack of experimental details is further discussed in the context of future research in Section XI.

We identify five key aspects of vulnerability that attacks exploit: resonance; non-linearity; improper filtering; poor shielding; and insecure algorithms. All attacks which do not target microphones and optical sensors depend on resonance of some sort: this can be acoustic resonance of mechanical structures, electromagnetic resonant frequencies of unintentional antennas, or the existence of locking frequencies for ring oscillators. Other attacks depend on non-linearities of amplifiers, microphones, and speakers to demodulate high-frequency signals. This is because resonant frequencies are often much higher than those of desired injection waveforms.

In addition, many works identify improper filtering, particularly prior to ADCs and amplifiers, as well as poor shielding as factors for out-of-band attacks. Finally, in some cases, insecure

sampling and processing algorithms exacerbate the problem by making it easier for an adversary to trick the system under attack into performing a dangerous action. These sources of vulnerability form the basis for many of the proposed countermeasures, which we discuss in detail in Section IX.

IX. ANALYSIS OF COUNTERMEASURES

Although the literature on out-of-band attacks is quite broad, research on defenses has been more sparse. Section IX-A first summarizes the state-of-the-art in countermeasures specific to out-of-band injections. Section IX-B then expands our discussion by introducing general protective and preventive approaches which remain applicable in this context.

A. Out-of-Band Defense Mechanisms

The works that have investigated countermeasures against out-of-band signal injection attacks have noted that a combi-

TABLE III: Summary of evaluated (\checkmark), proposed (\times), and absent ($-$) countermeasures against 🔊 acoustic, ⚡ conducted, 📶 electromagnetic, and 🌟 optical out-of-band signal injection attacks.

Authors	Year	Ref.	Target	Method	Robust Hardware	Better Sampling	Sensor Fusion	Improved Filtering	More Shielding	Anomaly Detection
Marketos and Moore	2009	[35]	TRNG	⚡📶	\times	-	-	\times	\times	-
Rasmussen et al.	2009	[33]	Medical	📶	-	-	-	-	\times	-
Bayon et al.	2012	[36]	TRNG	📶	-	-	-	-	-	-
Buchovecka et al.	2013	[37]	TRNG	⚡📶	-	-	-	-	-	-
Foo Kune et al.	2013	[5]	Medical	📶	\checkmark	-	-	\checkmark	\checkmark	\checkmark
Kasmi and Lopes Esteves	2015	[38]	Microphone	📶	\times	-	-	-	\times	\times
Petit et al.	2015	[14]	Camera	🌟	-	-	\times	\times	-	-
Shoukry et al.	2015	[69]	Active Sensors	📶	-	\checkmark	-	-	-	-
Son et al.	2015	[10]	MEMS	🔊	\times	-	-	-	\checkmark	-
Bayon et al.	2016	[49]	TRNG	📶	-	-	-	-	-	-
Cao et al.	2016	[122]	TRNG	⚡	-	-	-	-	-	-
Park et al.	2016	[8]	Medical	🌟	-	-	-	-	\times	\times
Shin et al.	2016	[48]	Medical	🌟	-	-	\times	-	\times	-
Roy et al.	2017	[54]	Microphone	🔊	-	-	-	-	-	-
Song and Mittal	2017	[133]	Microphone	🔊	-	-	-	-	-	-
Trippel et al.	2017	[13]	MEMS	🔊	\times	\checkmark	-	\times	\times	-
Wang et al.	2017	[34]	MEMS	🔊	\times	-	\times	-	\times	\times
Zhang et al.	2017	[11]	Microphone	🔊	\checkmark	-	-	-	-	\checkmark
Bolton et al.	2018	[12]	HDD	🔊	\times	-	\times	-	\checkmark	\checkmark
Lopes Esteves and Kasmi	2018	[40]	Microphone	⚡	-	-	-	\times	-	\times
Nashimoto et al.	2018	[41]	MEMS	🔊📶	-	-	\checkmark	-	-	-
Osuka et al.	2018	[42]	TRNG	⚡📶	-	-	-	\times	\times	-
Roy et al.	2018	[55]	Microphone	🔊	-	-	-	-	-	\checkmark
Selvaraj et al.	2018	[7]	GPIO	📶	-	-	-	\times	\times	-
Shahrad et al.	2018	[53]	HDD	🔊	\times	-	-	-	\times	\times
Tu et al.	2018	[43]	MEMS	🔊	-	\times	\times	\times	\times	-
Giechaskiel et al.	2019	[44]	ADC	⚡📶	\times	\times	-	\times	\times	\times
Khazaaleh et al.	2019	[129]	MEMS	🔊	\times	-	-	-	\times	\times
Mahmoud and Stoilović	2019	[123]	TRNG	⚡	-	-	-	-	-	\times
Muniraj and Farhood	2019	[147]	Servo	📶	-	\checkmark	-	-	-	\checkmark
Tu et al.	2019	[104]	Temp. & Amp.	⚡📶	-	-	\times	\times	\times	\checkmark
Yan et al.	2019	[134]	Microphone	🔊	\checkmark	-	-	-	-	\checkmark
Tharayil et al.	2019	[148]	MEMS	🔊📶	-	-	\checkmark	-	-	\checkmark
Zhang and Rasmussen	2020	[105]	Generic Sensors	📶	-	\checkmark	-	\times	\times	-

nation of prevention and detection techniques both in software and in hardware are necessary to improve security. We have divided the proposed defense mechanisms into six categories: more resilient hardware; improved sampling algorithms; sensor fusion and duplication; better filtering; additional shielding; and anomaly detection of measurements and the environment. We discuss each category in detail below, and summarize the various proposals per paper in Table III. As the table indicates, much of the discussion has been theoretical, with few works evaluating countermeasures in practice.

Observation 3: The effectiveness of proposed countermeasures remains mostly theoretical, as practical implementations are often limited in scope, with superficial discussion of monetary and computational costs.

Robust Hardware: In response to resonance and non-linearity vulnerabilities, various works have proposed preventive improvements in the hardware itself to make it more robust and less susceptible to attacks. One of these improvements against

electromagnetic attacks reduces asymmetries in differential inputs to a system. By doing so, attacker transmissions are received almost identically by the two unintentional receiving antennas, and are severely attenuated. For example, Marketos and Moore recommend reducing the asymmetries in ring oscillators through “carefully balanced transistors”, or the use of differential ones, which are “less affected by supply and substrate noise” [35]. Similarly, Foo Kune et al. found that using differential rather than single-ended comparators attenuated signals by up to 30 dB [5]. Although signals could still be injected, the power requirements to do so increased significantly, thereby raising the bar for attackers.

Another approach is to change the sensors themselves, rather than attempt to improve the physical layout of a circuit. For example, both Shahrad et al. [53] and Bolton et al. [12] note that replacing Hard-Disk Drives (HDDs) with Solid-State Drives (SSDs) thwarts acoustic resonance attacks due to a lack of moving parts. In a similar vein, Zhang et al. noted that the iPhone 6 Plus resisted their inaudible voice commands, since it is “designed to suppress any acoustic signals whose

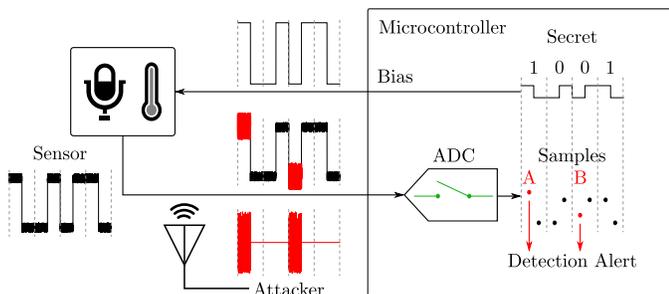


Fig. 11: High-level overview of the defense mechanism by Zhang and Rasmussen [105]. Oversampling by a factor of $2n$ and selectively turning sensors on and off allows detection of out-of-band electromagnetic attacks: without knowing the secret sequence (1001 here for $n=4$), the adversary will cause inconsistent (A) or unexpected (B) non-zero samples.

frequencies are in the ultrasound range” [11].

Finally, better frontends with fewer non-linearities are less sensitive to EMI noise [38], [44] and sonic injections [11], [13], [34]. They can therefore make it harder for adversaries to inject their desired signals into the system. Such general designs are discussed in greater detail in Section IX-B.

Better Sampling: Many papers have proposed improvements in the sampling technique to make it harder for an adversary to predict how a high-frequency signal will be converted to a low-frequency one. In 2015, Shoukry et al. proposed an alternative method of sampling active sensors called “PyCRA” (for Physical Challenge-Response Authentication) to detect signal injection attacks [69]. Active sensors “perform some action to evoke and measure a physical response from some measurable entity”, and include, for instance, magnetic encoders measuring angular velocity. Shoukry et al.’s proposal revolves around physical challenges to prove the absence of adversarial transmissions. Specifically, when the actuator is off (*silenced*), there should be no measured quantity unless an attack is taking place. By only shutting down the actuator for a small period of time, PyCRA can detect attackers without compromising the quality of sensor measurements and actuation results. Adversaries cannot stop transmissions in time due to physical and computational delay limits, allowing PyCRA to identify them [69]. It should be noted that Shin et al. have suggested that PyCRA would require high computational overhead in practice, both in terms of the minimum sampling rate needed to hit those physical limits, and for the detectors themselves [48]. Moreover, PyCRA requires active sensors, so it primarily protects against in-band attacks.

However, a similar proposal by Zhang and Rasmussen recently showed how to protect both powered and non-powered passive sensors [105]. The key idea is to use a secret bitstream to selectively turn off the sensor and observe whether the measured signal has been altered by electromagnetic injections, as shown in Figure 11. More concretely, for each sensor measurement, $2n$ ADC samples are taken, corresponding to an n -bit secret sequence. Each secret bit is Manchester-encoded, so that a 0-bit is represented as the pair (0,1), corresponding to turning the sensors off for one sample, and then turning

them on (“biasing”) for another sample. A 1-bit is similarly encoded as the pair (1,0), first turning on the sensor, and then turning it off during the second sample.

When the sensor is turned off, all samples should be close to zero, within some noise- and device-dependent tolerance. Moreover, for fast-enough sampling frequencies and slow-enough sensor signals, the n samples when the sensor is on should be close to each other. As a result, to inject a single measurement successfully, an attacker needs to correctly predict the n secret bits, which only happens with probability 2^{-n} for a randomly chosen bit sequence. By using a switch, non-powered passive sensors can also be adapted to use this approach. Moreover, spikes in the frequency domain allow Zhang and Rasmussen to detect attacker transmissions even for non-constant sensor signals [105]. Overall, by oversampling for each sensor measurement, noise can be distinguished from adversarial signals with probabilistic guarantees.

An alternative approach to prevent attackers from injecting their desired waveforms into a system is to add randomness to the sampling process, especially for ADCs which are only vulnerable for limited carrier frequencies [44]. The effect is essentially one of “having an inaccurate ADC” [13], allowing a moving average to filter out injected periodic signals. This is similar to sampling with a “dynamic sample rate”, defeating the side-swing and switching attacks of Tu et al. [43], which were explained in Section VI. Out-of-phase sampling has also been proposed as a band-stop filter to reject frequencies near an accelerometer’s resonant frequency, thereby removing attacker-injected DC offsets [13].

It should be noted that protecting against signal injection attacks into actuators has not been studied as extensively in literature. However, Muniraj and Farhood recently proposed a detection method based on a watermarking scheme that slightly alters the actuation parameters [147], similar to the proposal of Zhang and Rasmussen [105]. Under attack, the measured response of the system does not match the effect of the watermark, allowing detection. The same paper also suggested pseudo-randomly changing the pulse frequency of the Pulse Width Modulation (PWM) signals, making an actuator attack harder to accomplish. Overall, these methods only alter the shape of the waveform that an adversary can inject rather than the root cause of the vulnerability itself. They are therefore not sufficient countermeasures by themselves.

Sensor Fusion: A few works have suggested that using sensors of different types (*fusion*) or multiple sensors of the same type (*duplication*) with different vulnerable frequency ranges will make injections harder [12], [14], [34], [43], [48]. This is because an adversary would need to mount multiple simultaneous attacks, which potentially interfere destructively. However, in “Sensor CON-Fusion”, Nashimoto et al. [41] showed that a fusion algorithm based on Kalman filters could be circumvented (Section VI). As a result, better techniques are needed to protect against adversarial injections, instead of simply faulty readings [149], [150]. Tharayil et al. proposed an improved such fusion algorithm, which takes into account the mathematical relations between the underlying physical quantities. This allows them to link measurements by a gyroscope and a magnetometer in a way which can detect adversarial

injections without any hardware modifications [148].

Other researchers have proposed that additional sensors be used to measure and counteract attacker signals. These additional components should not be identical to the vulnerable sensors: as Khazaaleh et al. show, many MEMS gyroscopes integrate a second, “identical proof mass to perform differential measurements” and “eliminate unwanted vibrations” [129]. However, they still remain vulnerable to ultrasonic attacks.

Instead, “an additional gyroscope [...] that responds only to the resonant frequency” may be able to remove the resonance effect from the main gyroscope [10]. Similarly, microphones might be able to detect (and potentially cancel) resonant frequencies to protect MEMS gyroscopes [34] and HDDs [12]. This approach would be hard in practice, at least for hard drives: the area to be protected would need to cover “the read/write head [completely] as it moves across the disk”, and the sound wave to be generated would be potentially large, raising many issues about its implementation [12].

Filtering: Most papers studied highlight the need for better filters to reduce the vulnerable frequency range against conducted [35], [40], [44], [104], electromagnetic [5], [7], [42], [44], [104], [105], acoustic [13], [43], and optical attacks [14]. However, only Foo Kune et al. [5] have performed systematic experiments studying the effect of filtering on the efficacy of out-of-band signal injection attacks. To start with, Foo Kune et al. noted that adding a low-pass filter in their experiments against bluetooth headsets allowed audio signals to pass, but attenuated the injected electromagnetic signal by 40dB. Moreover, they proposed an adaptive filtering mechanism which uses the measured signal and the ambient EMI level to cancel the attacker-injected waveform. Using a Finite Impulse Response (FIR) filter, the algorithm estimates this waveform, and allows quick recovery of the original signal, after an onset period at the beginning of the attack. It should be noted, however, that filters might not be effective against MEMS sensors: Khazaaleh et al. noted that “false readings could not be attenuated by adding a 10Hz low-pass filter”, despite the resonant frequency being in the kHz range [129].

Shielding: Better separation from the environment also improves protection against out-of-band signal injection attacks. For this reason, it has been recommended by most authors investigating attacks on sensors other than microphones. This shielding may come in the form of physical isolation [8], [10], [43], [53], [129], better acoustic dampening materials [12], [13], [34], [43], or radio frequency shielding [5], [7], [33], [35], [38], [42], [44], [48], [104], [105]. For instance, Foo Kune et al. demonstrated a 40dB attenuation of the injected signal, even when the shielding had “large imperfections” [5]. These openings (e.g., for wires to pass through) result in “major degradations in the shielding” [151]. Indeed, Selvaraj et al. noted that “while a light sensor can function in a mesh-based Faraday cage, magnetic shielding would prevent light from reaching the sensor” [7]. In addition, Bolton et al. showed that dampening foam “significantly reduced an HDD’s susceptibility to write blocking”, but “did not attenuate lower frequency signals” [12]. Moreover, the foam led to an increased temperature of 10°C, which can also result in disk failure. As a result, “it is often necessary to use a combination

of shielding and other protective measures” [151].

Anomaly Detection: Instead of trying to prevent signal injection attacks, some works have proposed better software-level processing of sensor signals, primarily for anomaly detection, with or without additional hardware. One such approach is to estimate the ambient level of electromagnetic [5], [38], [104], optical [8] and acoustic [11], [12], [34], [53], [55], [134], [152] emissions. For example, Park et al. noted that saturation attacks can be detected simply “by checking whether the light intensity exceeds the preset maximum level” [8]. Foo Kune et al. [5] further investigated the use of additional (intentional) antennas or reference conductors to measure the levels of EMI radiation. This estimate can then be used by their adaptive filtering algorithm [5], which was discussed above.

In a similar vein, Tu et al. recommended the addition of a superheterodyne AM receiver to create a tunable EM detector [104]. This detector was shown to be useful in estimating and compensating errors in the measurements. Other detection mechanisms can operate with existing hardware: for example, Khazaaleh et al. noted that “sensing fingers”, which are already used to measure displacement in the y-axis, can detect large displacements caused by resonance [129].

The question of how systems should behave once an attack has been detected has large been side-stepped by many works. However, Bolton et al. introduced an algorithm to augment the hard drive feedback controller and compensate for intentional acoustic interference [12]. The addition of this attenuation controller reduces the position errors of the read and write heads to within the accepted tolerance levels, and allows the HDD to operate in the presence of an attack.

Another way of detecting attacks is to use machine learning classifiers [148]. However, such classifiers can be prone to false positives, and will miss precise waveform injections. As a result, it is often necessary to look for artifacts that would not be present during the normal operation of a sensor, such as harmonics and low or high frequency components [11], [44], [55], [134], [152]. This might not always be as straightforward as simply detecting energy at low or high frequencies that are only present due to non-linearities: for sophisticated attackers, defense mechanisms need to exploit the properties of the legitimate signal itself. For example, Roy et al. showed that “voice signals exhibit well-understood patterns of fundamental frequencies”, which are not present in attacks and environmental noise. As a result, they can be used to detect acoustic commands generated by ultrasound signals [55]. Similarly, the *absolute refractory period* is hard for an attacker to spoof precisely via EM injections [5]. This period represents the time span after a contraction during which the cardiac tissue will not contract again. As shown by Foo Kune et al. [5], it can be used to distinguish between a real and adversarial signals.

Finally, more restrictive processing of sensor data can also help mitigate signal injection attacks. For example, safe defaults when the sensor output is deemed as untrustworthy [5], [8] can reduce the effects of successful attacks on health- and safety-critical actions taken by systems. Similarly, less permissive choices in the design of voice interfaces can prevent non-targeted attacks from succeeding. As an example, adding voice authentication and custom keywords can prevent

command injections into smartphones [38], [40], [44].

Observation 4: Until more resilient components replace vulnerable ones, defense-in-depth is necessary to protect against signal injection attacks. This can be accomplished through better filtering and shielding to prevent attacks, and through better sampling, fusion, and anomaly detection algorithms to identify them.

B. Other Defensive Approaches

As out-of-band signal injection attacks are closely connected with different areas of research (Section X), there is extensive overlap in the proposed countermeasures. For example, before Foo Kune et al. [5] proposed an adversarial EMI detector, Wan et al. [153], [154] introduced a similar design to “increase the immunity of a microcontroller-based system in a complex electromagnetic environment”. Moreover, to protect against LiDAR attacks, Shin et al. proposed sensor fusion and redundancy, fail-safe defaults, better shielding (by reducing receiving angle), and randomized pinging directions and waveforms [31]. Similarly, Davidson et al. proposed sensor fusion and an improved optical flow algorithm to protect against optical in-band sensor spoofing [52]. Moreover, Blue et al. detected (audible) command injections by identifying a frequency band which is produced by electronic speakers, but is absent in human speech [152]. In fact, detecting unique features of the sensed property is a common defense mechanism for general sensor manipulation attacks, such as those against Smart Grid power plants [155], [156], or unmanned aircraft systems [147]. However, such approaches require a theoretical system model, and assume an adversary who cannot inject data obeying this model.

Observation 5: Defense mechanisms for in-band attacks, excessive environmental noise, and faulty sensors are often directly applicable to out-of-band signal injection attacks and vice versa.

In a different strand of research, Redouté and Richelli have proposed some guidelines for improving immunity against EM interference attacks [157], [158]. These recommendations could be applied in the context of general out-of-band attacks:

- 1) *Filter induced signals before the non-linear device.* This suggestion is not limited to amplifiers, but can be used in other setups, including power transistors [159]. It has been shown to result in an up to a 12.5× reduction in EMI-induced offsets [157], [160], [161], but may require bulky passive components, adding noise to the circuit.
- 2) *Linearize the stage generating the DC shift,* for example, by using amplifiers with a wider common mode input range, resulting in better linear behavior [162].
- 3) *Prevent the accumulation of DC shift,* for instance by addressing the slow rate asymmetry and parasitic capacitances [163]–[165].
- 4) *Compensate and remove the induced offset,* for example, using cross-connected differential pairs [164].

As discussing all possible EMI-resistant amplifier designs is out-of-scope, the interested reader should refer to various comparative works [165]–[167] as a starting point. Similarly, one should refer to advances in gyroscopic technologies [124], [168], [169] which do not use MEMS constructions, or reduce sensitivity to random vibrations: as Khazaaleh et al. noted [129], removing the “misalignment between the sensing and driving axes” will make systems more secure against out-of-band acoustic attacks.

Observation 6: More accurate and sensitive hardware that is robust to environmental influences is a natural defense mechanism against out-of-band attacks.

X. ADDITIONAL RELATED WORK

As this study contains the first survey of out-of-band signal injection attacks, this section shows the close connections with side-channel leakage and electromagnetic interference. For example, using insights into the resonant frequencies of gyroscopes, Farshteindiker et al. showed that unprivileged websites could act as covert channel receivers, even at very low sampling frequencies of 20Hz [170]. Block et al. improved the design by not requiring external equipment for the attack, instead relying on the smartphone’s speaker and accelerometer [171]. Matyunin et al. then used the same effect for cross-device tracking using ultrasonic transmissions at or near the resonant frequencies of gyroscopes [172]. Moreover, Michalevsky et al. showed that MEMS gyroscope measurements are sensitive to acoustic signals in their vicinity [173]. As a result, they can be used to distinguish between different speakers, and, in part, the content of the speech [173] due to conducted vibrations of the loudspeakers used [174], [175].

In other words, the same source of vulnerability which can be used to destabilize [10] and control [13], [34], [43] gyroscopes and accelerometers can be used for covert channel communication [170], [171], tracking [172], and speaker identification [173], [174]. Similarly, instead of using microphone non-linearities for command injections [11], [55], [134], Shen et al. [176] and Chen et al. [177] leveraged them to protect users’ privacy by jamming nearby recording devices.

The countermeasures proposed in the works above mirror those of Section IX-A, and include anti-aliasing filters, shielding, and sensor fusion. Moreover, suggestions to increase noise in side- and regular-channel emissions parallel out-of-band defense mechanisms based on reducing the sampling accuracy. For example, decreasing “the fidelity of the input audio” can prevent against inaudible voice injection attacks [136]. Similarly, fonts which minimize emissions at high frequencies [178]–[180] exploit the human eye sensitivity to “low spatial frequencies” [178]. As EM emanations of video display units (“TEMPEST”) [181]–[183] mostly convey “the high-frequency part of the video signal” [178], images are transformed in a way that is almost transparent to human viewers, but prevents the reconstruction from side-channel listeners. Researchers have likewise shown that adding certain patterns to video frames [184] or the flashing of LED lights [185]

can reduce the fidelity of reconstructed images from camera recordings, while not influencing regular viewers as much. Finally, in some respects, anomaly detection resembles statistical and machine learning approaches to detect covert channels, and can therefore draw inspiration from seemingly unrelated disciplines, such as timing and storage network covert channels [186].

Observation 7: The sources of vulnerabilities for out-of-band signal injection attacks are often the same as those for hardware-based covert- and side-channel attacks. This allows the same techniques to be reused for attacks and defenses across disciplines.

Other research has indicated that devices which are typically used as actuators can actually effectively function as sensors. LEDs can function as photodiodes [138], while speakers [187] and HDDs [188] can both be converted into microphones. Although all three attacks have so-far required the assistance of malware, further research is required to identify the implications for out-of-band signal injection attacks.

Observation 8: Reuse of off-the-shelf equipment in unconventional setups expands the surface for signal injection attacks exploiting hardware imperfections.

The lines between electromagnetic interference and out-of-band signal injection attacks are also blurred. This is in part because the self-classification of attacks depends primarily on the research community with which an author is aligned, rather than the end result of the injection. For example, the voice injection command attacks of Kasmi and Lopes Esteves [38], [40] are categorized by their authors as Intentional Electromagnetic Interference (IEMI) attacks, despite the relatively low power used, and the lack of upsets or destruction of equipment. Similarly, Osuka et al. considered their work to be in the IEMI realm [42], even though they biased the randomness of a TRNG. This fact also partially explains why research on out-of-band attacks against TRNGs has largely ignored attacks against other targets and vice versa.

In general, this mismatch of expectations often results in unexplored avenues of research, as can be seen, for example, in the IEMI attacks on UAVs of Lopes Esteves et al. [189]: although there is a strong inverse correlation between the battery temperature reading and the strength of the electric field, the authors do not further investigate how to precisely control the sensor output. Moreover, as was explained in Section IX-B, research into electromagnetic interference can provide insights into how to build more resilient hardware, even when the hardware is only tested against “unintentional parasitic signals and does not take a malicious behavior of an attacker into account” [40].

Observation 9: The proposed terminology based on the outcome rather than the method of injection can help systematize attack and defense approaches, and reveal previously unexplored connections.

XI. FUTURE DIRECTIONS

Despite the amount of research conducted on out-of-band signal injection attacks, there is no common methodology to evaluate how susceptible systems are to them. This is in contrast to related disciplines, such as side-channel analysis [75], direct power injection and near-field scan immunity [59], fault injection attacks [4], and IEMI attacks [190]. Indeed, although many papers sweep through frequencies to find the resonant ones [5], [13], some do not adopt this terminology [10], [12], [53], and do not specify how wide the frequency steps should be. This proves to be particularly problematic, as some attack windows “are as narrow as a few Hertz” [53].

Recently, Tu et al. [43] provided a more detailed methodology for acoustic injection attacks, which starts with a *profiling* stage. During this phase, single-tone sounds are transmitted, and are swept at an interval of 10Hz. The devices targeted remain stationary during the profiling stage. Further increments of 1Hz or smaller can be used near the resonant frequencies to estimate the sampling frequency of the ADC, and account for its drift. The next stage involves *synchronizing* to a frequency which is close to a multiple of the ADC sampling rate. This step is followed by *manipulating* the attack parameters, and *adjusting* them in response to drifts. Although this approach provides some common ground for evaluation, several questions remain unanswered, especially when assessing countermeasures to claim that a system is secure. These questions include: what the frequency range itself should be; what the step should be for wide ranges; what modulation method to use and with what parameters; and whether there are other factors that would need to be examined during experimentation. For instance, for electromagnetic attacks, the incident angle of the EM field and the distance of attacks can have a profound impact on their success, especially as they relate to generalizing from the near- to the far-field.

Observation 10: A precise experimental procedure which specifies sweep, modulation, and other parameters is needed for out-of-band signal injection attacks.

The question of the maximum feasible attack distance has mostly been of theoretical interest, with practical attacks often limited to a few centimeters. Even though EM attacks should in theory have a longer range than acoustic and optical attacks, the converse appears to be true in the experiments conducted by the works studied in this survey (Table II). There is also a worrying trend of assuming that more power and more expensive equipment easily translates to a long-range attack. For example, Tu et al. [43] claim that with more speakers, gyroscopes can be attacked from an 8× longer distance, but as Roy et al. [55] showed, doing so is not a trivial engineering concern, if the inaudibility of injections is to be maintained.

Similarly, although Foo Kune et al. [5] claim that a 20dB gain directional antenna and a 1W source can attack equipment at distances of up to 50m, these estimates seem optimistic: according to Lopes Esteves and Kasmi [40], a 200W source is required for a distance of 4m for remote command injections [38]. What is more, high-powered EM sources have

the potential to cause faults in other equipment and be harmful to human life. As a result, determining how to inject precise signals from a distance is particularly challenging. These problems become even more pronounced for magnetic attacks on actuators, which have been limited so far.

Observation 11: Dedicated facilities and test equipment for long-range experimentation are needed.

Note that as many of the systems targeted are safety- and mission-critical, we can expect that, in the future, some devices may be required to undergo a certification process. Indeed, a CERT alert warning of MEMS susceptibility to ultrasonic resonance [21] highlights that out-of-band signal injection vulnerabilities are a concern for governments and corporations alike. Regulations will thus pave the way for an expanding industry around facilities and test equipment for EMC immunity against adversarial injections.

As the currently published work often leaves experimental details under-specified, reproducibility becomes a significant challenge: for instance, as discussed in Section VIII, details on the power used were often not readily available, but required searching through datasheets. Moreover, the duration of attacks was also often not specified. Trippel et al. [13] reported that some of their attacks against accelerometers only work for a couple of seconds before the attack fails. Sampling rate drifts thus necessitated manual tuning, or more sophisticated attacking techniques, such as those proposed by Tu et al. [43]. However, without details of the function generator specifications, it would be hard to know whether some of the issues are caused by poor clock accuracy of the generator. This problem is bound to become even more pronounced when using Software-Defined Radio (SDR) and other low-end commodity hardware for attack weaponization.

Minor variations in the construction of devices can also have significant effects on the sensors' behavior, and will potentially impact the reproducibility of attacks. For instance, Dey et al. [191] showed that otherwise identical accelerometers can be tracked due to slightly different performance characteristics. Giechaskiel et al. [44] recently introduced security definitions to address the lack of directly comparable metrics describing the outcome of injection attacks. However, the overall absence of experimental details, coupled with monetary costs and legal requirements associated with using the electromagnetic spectrum, make security research into out-of-band signal injection attacks a challenging space for new researchers to enter.

Observation 12: Reproducibility through common metrics which allow for direct comparison of the effects of injection and standardized experimental setups are necessary to advance the state-of-the-art.

Besides the defense mechanisms of Section IX, to protect future devices for attack, new security-sensitive products must take a fundamentally different approach to trusting the outputs of sensors. In the words of Fu and Xu, there is a need to "shift from component-centric security to system-centric tolerance of untrustworthy components", perhaps taking note of advances

in fault-tolerant literature [67]. Fu and Xu also recommend that sensor outputs be "continuously checkable by software for adversarial influence", such as through internal debugging information that is hidden from accessible APIs [67]. They further highlight the need for interdisciplinary teams and education [67]. Indeed, until new hardware is deployed, many cross-disciplinary solutions will be required to prevent, detect, and mitigate attacks. As out-of-band signal injection attacks become more powerful, collaboration will be necessary to address the multifaceted research influences of the field.

Observation 13: Interdisciplinary research quantifying the effectiveness of countermeasures is needed to inform future hardware and software design choices.

XII. CONCLUSION

Our ever-increasing reliance on sensors and actuators highlights the need for a comprehensive look into electromagnetic, conducted, acoustic, and optical out-of-band signal injection attacks. These attacks cause a mismatch between a physical property being measured by a sensor or acted upon by an actuator and its digitized version. Out-of-band signal injection attacks can be used to control or disrupt drones, hard drives, and medical devices, among others, with potentially fatal consequences on human life. In light of the importance of such attacks, this paper took the first step towards unifying the diverse and expanding research through a taxonomy of attacks, defenses, and terminology. Our work revealed interdisciplinary influences between seemingly disparate topics, and also made several observations that can inform future research in the area. Overall, better experimental and reporting procedures are necessary for direct comparisons of the effects of attack and defense mechanisms.

REFERENCES

- [1] R. Spreitzer, V. Moonsamy, T. Korak, and S. Mangard, "Systematic classification of side-channel attacks: A case study for mobile devices," *IEEE Communications Surveys & Tutorials (COMST)*, vol. 20, no. 1, pp. 465–488, Feb 2018.
- [2] A. Barenghi, L. Breveglieri, I. Koren, and D. Naccache, "Fault injection attacks on cryptographic devices: Theory, practice and countermeasures," *Proceedings of the IEEE (JPROC)*, vol. 100, no. 11, pp. 3056–3076, Nov 2012.
- [3] D. Karaklajić, J.-M. Schmidt, and I. M. R. Verbauwhede, "Hardware designer's guide to fault attacks," *IEEE Transactions on Very Large Scale Integration Systems (TVLSI)*, vol. 21, no. 12, pp. 2295–2306, Dec 2013.
- [4] B. Yuce, P. Schaumont, and M. Wittenman, "Fault attacks on secure embedded software: Threats, design, and evaluation," *Journal of Hardware and Systems Security (HASS)*, vol. 2, no. 2, pp. 111–130, Jun 2018.
- [5] 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 *IEEE Symposium on Security and Privacy (S&P)*, 2013.
- [6] Y. Shoukry, P. Martin, P. Tabuada, and M. Srivastava, "Non-invasive spoofing attacks for anti-lock braking systems," in *International Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 2013.
- [7] J. Selvaraj, G. Y. Dayamikli, N. P. Gaunkar, D. Ware, R. M. Gerdes, and M. Mina, "Electromagnetic induction attacks against embedded systems," in *ACM ASIA Conference on Computer and Communications Security (ASIACCS)*, 2018.

- [8] Y. Park, Y. Son, H. Shin, D. Kim, and Y. Kim, "This ain't your dose: Sensor spoofing attack on medical infusion pump," in *USENIX Workshop on Offensive Technologies (WOOT)*, 2016.
- [9] C. Yan, W. Xu, and J. Liu, "Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle," DEF CON, 2016.
- [10] Y. Son, H. Shin, D. Kim, Y. Park, J. Noh, K. Choi, J. Choi, and Y. Kim, "Rocking drones with intentional sound noise on gyroscopic sensors," in *USENIX Security Symposium*, 2015.
- [11] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu, "DolphinAttack: Inaudible voice commands," in *ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [12] 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 *IEEE Symposium on Security and Privacy (S&P)*, 2018.
- [13] 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 *IEEE European Symposium on Security and Privacy (EuroS&P)*, 2017.
- [14] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote attacks on automated vehicles sensors: Experiments on camera and LiDAR," Black Hat Europe, 2015.
- [15] Forbes, "Want to ruin someone's Oculus Rift fun? Fire this sonic gun at their head," <https://www.forbes.com/sites/thomasbrewster/2017/07/11/alibaba-researchers-attack-facebook-vr-with-soundwaves>, Created: 2017-07-11. Accessed: 2019-11-07.
- [16] The Register, "Boffins Rickroll smartphone by tickling its accelerometer," <https://www.theregister.co.uk/2017/03/15/boffins%5Frickroll%5Fsmartphone%5Fby%5Ftickling%5Fits%5Faccelerometer>, Created: 2017-03-15. Accessed: 2019-11-07.
- [17] Ars Technica, "Sounds bad: Researchers demonstrate 'sonic gun' threat against smart devices," <https://arstechnica.com/gadgets/2017/07/sounds-bad-researchers-demonstrate-sonic-gun-threat-against-smart-devices>, Created: 2017-07-28. Accessed: 2019-11-07.
- [18] The New York Times, "It's possible to hack a phone with sound waves, researchers show," <https://www.nytimes.com/2017/03/14/technology/phone-hacking-sound-waves.html>, Created: 2017-03-14. Accessed: 2019-11-07.
- [19] Fox News, "Sonic weapon knocks drones right out of the sky," <https://www.foxnews.com/tech/sonic-weapon-knocks-drones-right-out-of-the-sky>, Created: 2017-11-26. Accessed: 2019-11-07.
- [20] The Inquirer, "Sonic attacks can bork hard disks and crash Windows and Linux," <https://www.theinquirer.net/inquirer/news/3033287/sonic-and-ultrasonic-attacks-can-crash-hard-disks-and-windows-and-linux>, Created: 2018-05-31. Accessed: 2019-11-07.
- [21] Cybersecurity and Infrastructure Security Agency, "ICS-ALERT-17-073-01A: MEMS accelerometer hardware design flaws (update A)," <https://www.us-cert.gov/ics/alerts/ICS-ALERT-17-073-01A>, Created: 2017-04-11. Accessed: 2019-11-07.
- [22] National Public Radio, "Can you hear it? Sonic devices play high-pitched noises to repel teens," <https://www.npr.org/2019/07/10/739908153/can-you-hear-it-sonic-devices-play-high-pitched-noises-to-repel-teens>, Created: 2019-07-10. Accessed: 2019-11-07.
- [23] British Broadcasting Corporation, "Dewsbury driver who used speed camera jammer jailed," <https://www.bbc.co.uk/news/uk-england-leeds-47202419>, Created: 2019-02-11. Accessed: 2019-11-07.
- [24] The Wall Street Journal, "U.S. downed Iranian drone with new technology," <https://www.wsj.com/articles/u-s-downed-iranian-drone-with-new-technology-11563579400>, Created: 2019-07-19. Accessed: 2019-11-07.
- [25] I. Rouf, R. Miller, H. Mustafa, 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 *USENIX Security Symposium*, 2010.
- [26] C. Li, A. Raghunathan, and N. K. Jha, "Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system," in *IEEE International Conference on e-Health Networking, Applications and Services (Healthcom)*, 2011.
- [27] D. Halperin, T. S. Heydt-Benjamin, B. Ransford, S. S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. H. Maisel, "Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses," in *IEEE Symposium on Security and Privacy (S&P)*, 2008.
- [28] S. Lakshminarayana, J. S. Karachiwala, S.-Y. Chang, G. Revadigar, S. L. S. Kumar, D. K. Y. Yau, and Y.-C. Hu, "Signal jamming attacks against communication-based train control: Attack impact and countermeasure," in *ACM Conference on Security & Privacy in Wireless and Mobile Networks (WiSec)*, 2018.
- [29] K. C. Zeng, S. Liu, Y. Shu, D. Wang, H. Li, Y. Dou, G. Wang, and Y. Yang, "All your GPS are belong to us: Towards stealthy manipulation of road navigation systems," in *USENIX Security Symposium*, 2018.
- [30] R. Chauhan, R. M. Gerdes, and K. Heaslip, "Demonstration of a false-data injection attack against an FMCW radar," in *Embedded Security in Cars (ESCAR)*, 2014.
- [31] H. Shin, D. Kim, Y. Kwon, and Y. Kim, "Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications," in *International Conference on Cryptographic Hardware and Embedded Systems (CHES)*, 2017.
- [32] 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 (IoT-J)*, vol. 5, no. 6, pp. 5015–5029, Dec 2018.
- [33] K. B. Rasmussen, C. Castelluccia, T. S. Heydt-Benjamin, and S. Capkun, "Proximity-based access control for implantable medical devices," in *ACM Conference on Computer and Communications Security (CCS)*, 2009.
- [34] Z. Wang, K. Wang, B. Yang, S. Li, and A. Pan, "Sonic gun to smart devices: Your devices lose control under ultrasound/sound," Black Hat USA, 2017.
- [35] A. T. Marketos and S. W. Moore, "The frequency injection attack on ring-oscillator-based true random number generators," in *International Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 2009.
- [36] P. Bayon, L. Bossuet, A. Aubert, V. Fischer, F. Poucheret, B. Robisson, and P. Maurine, "Contactless electromagnetic active attack on ring oscillator based true random number generator," in *International Workshop on Constructive Side-Channel Analysis and Secure Design (COSADE)*, 2012.
- [37] S. Buchovecká and J. Hlaváč, "Frequency injection attack on a random number generator," in *IEEE International Symposium on Design and Diagnostics of Electronic Circuits & Systems (DDECS)*, 2013.
- [38] C. Kasmi and J. Lopes Esteves, "IEMI threats for information security: Remote command injection on modern smartphones," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 57, no. 6, pp. 1752–1755, Dec 2015.
- [39] H. Martín, T. Korak, E. San Millán, and M. Hutter, "Fault attacks on STRNGs: Impact of glitches, temperature, and underpowering on randomness," *IEEE Transactions on Information Forensics and Security (TIFS)*, vol. 10, no. 2, pp. 266–277, Feb 2015.
- [40] J. Lopes 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. System Design & Assessment Note 48, Apr 2018, <http://ece-research.unm.edu/summa/notes/SDAN/SDAN0048.pdf>.
- [41] S. Nashimoto, D. Suzuki, T. Sugawara, and K. Sakiyama, "Sensor CON-Fusion: Defeating Kalman filter in signal injection attack," in *ACM ASIA Conference on Computer and Communications Security (ASIACCS)*, 2018.
- [42] S. Osuka, D. Fujimoto, Y.-I. Hayashi, N. Homma, A. Beckers, J. Balasch, B. Gierlich, and I. M. R. Verbauwhede, "EM information security threats against RO-based TRNGs: The frequency injection attack based on IEMI and EM information leakage," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 61, no. 4, pp. 1122–1128, Aug 2019.
- [43] Y. Tu, Z. Lin, I. Lee, and X. Hei, "Injected and delivered: Fabricating implicit control over actuation systems by spoofing inertial sensors," in *USENIX Security Symposium*, 2018.
- [44] I. Giechaskiel, Y. Zhang, and K. B. Rasmussen, "A framework for evaluating security in the presence of signal injection attacks," in *European Symposium on Research in Computer Security (ESORICS)*, 2019.
- [45] R. Jin and K. Zeng, "Physical layer key agreement under signal injection attacks," in *IEEE Conference on Communications and Network Security (CNS)*, 2015.
- [46] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in *ACM Conference on Computer and Communications Security (CCS)*, 2009.
- [47] B. Carrara and C. Adams, "Out-of-band covert channels—a survey," *ACM Computing Surveys (CSUR)*, vol. 49, no. 2, pp. 1–36, Jun 2016.
- [48] 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 *USENIX Workshop on Offensive Technologies (WOOT)*, 2016.

- [49] P. Bayon, L. Bossuet, A. Aubert, and V. Fischer, "Fault model of electromagnetic attacks targeting ring oscillator-based true random number generators," *Journal of Cryptographic Engineering (JCEN)*, vol. 6, no. 1, pp. 61–74, Apr 2016.
- [50] W. A. Radasky, C. E. Baum, and M. W. Wik, "Introduction to the special issue on high-power electromagnetics (HPEM) and intentional electromagnetic interference (IEMI)," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 46, no. 3, pp. 314–321, Aug 2004.
- [51] E. Savage and W. Radasky, "Overview of the threat of IEMI (intentional electromagnetic interference)," in *IEEE International Symposium on Electromagnetic Compatibility (EMC)*, 2012.
- [52] D. Davidson, H. Wu, R. Jellinek, T. Ristenpart, and V. Singh, "Controlling UAVs with sensor input spoofing attacks," in *USENIX Workshop on Offensive Technologies (WOOT)*, 2016.
- [53] M. Shahrad, A. Mosenia, L. Song, M. Chiang, D. Wentzloff, and P. Mittal, "Acoustic denial of service attacks on hard disk drives," in *Workshop on Attacks and Solutions in Hardware Security (ASHES)*, 2018.
- [54] N. Roy, H. Hassanieh, and R. R. Choudhury, "BackDoor: Making microphones hear inaudible sounds," in *International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2017.
- [55] N. Roy, S. Shen, H. Hassanieh, and R. R. Choudhury, "Inaudible voice commands: The long-range attack and defense," in *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2018.
- [56] Y.-H. Sutu and J. J. Whalen, "Statistics for demodulation RFI in operational amplifiers," in *IEEE International Symposium on Electromagnetic Compatibility (EMC)*, 1983.
- [57] H. Ghadamabadi, J. J. Whalen, R. Coslick, C. Hung, T. Johnson, W. Sitzman, and J. Stevens, "Comparison of demodulation RFI in inverting operational amplifier circuits of the same gain with different input and feedback resistor values," in *IEEE International Symposium on Electromagnetic Compatibility (EMC)*, 1990.
- [58] A. Boyer, S. Ben Dhia, and E. Sicard, "Modelling of a direct power injection aggression on a 16-bit microcontroller input buffer," in *International Workshop on Electromagnetic Compatibility of Integrated Circuit (EMC Compo)*, 2007.
- [59] —, "Modelling of a mixed signal processor susceptibility to near-field aggression," in *IEEE International Symposium on Electromagnetic Compatibility (EMC)*, 2007.
- [60] J. Gago, J. Balcells, D. González, M. Lamich, J. Mon, and A. Santolaria, "EMI susceptibility model of signal conditioning circuits based on operational amplifiers," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 49, no. 4, pp. 849–859, Nov 2007.
- [61] X.-L. Gao, C.-Y. Tian, L.-Y. Lao, Y.-H. Chen, and Y.-Y. Chen, "Improved direct power injection model of 16-bit microcontroller for electromagnetic immunity prediction," *Journal of Central South University of Technology*, vol. 18, no. 6, pp. 2031–2035, Dec 2011.
- [62] A. Ayed, T. Dubois, J.-L. Levant, and G. Duchamp, "Failure mechanism study and immunity modeling of an embedded analog-to-digital converter based on immunity measurements," *Microelectronics Reliability*, vol. 55, no. 9-10, pp. 2067–2071, Aug/Sep 2015.
- [63] —, "Immunity measurement and modeling of an ADC embedded in a microcontroller using RFIIP technique," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 57, no. 5, pp. 955–962, Oct 2015.
- [64] A. Boyer, B. Vrignon, and M. Cavarroc, "Modeling magnetic near-field injection at silicon die level," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 58, no. 1, pp. 257–269, Feb 2016.
- [65] S. Kennedy, M. R. Yuce, and J.-M. Redouté, "Susceptibility of flash ADCs to electromagnetic interference," *Microelectronics Reliability*, vol. 81, pp. 218–225, Feb 2018.
- [66] C. Pouant, F. Torrès, A. Reineix, P. Hoffmann, J. Raoult, and L. Chusseau, "Modeling and analysis of large-signal RFI effects in MOS resistors," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 61, no. 1, pp. 111–120, Feb 2019.
- [67] K. Fu and W. Xu, "Risks of trusting the physics of sensors," *Communications of the ACM (CACM)*, vol. 61, no. 2, pp. 20–23, Feb 2018.
- [68] C. Yan, K. Fu, and W. Xu, "On Cuba, diplomats, ultrasound, and intermodulation distortion," *Computers in Biology and Medicine*, vol. 104, pp. 250–266, Jan 2019.
- [69] Y. Shoukry, P. Martin, Y. Yona, S. Diggavi, and M. Srivastava, "PycRA: Physical challenge-response authentication for active sensors under spoofing attacks," in *ACM Conference on Computer and Communications Security (CCS)*, 2015.
- [70] M. J. M. Pelgrom, *Analog-to-Digital Conversion*, 3rd ed. Springer, 2017.
- [71] J.-M. Redouté and M. Steyaert, *EMC of Analog Integrated Circuits*, 1st ed. Springer, 2010.
- [72] M. I. Montrose, *EMC and the Printed Circuit Board: Design, Theory, and Layout Made Simple*, 1st ed. Wiley, 1999.
- [73] M. Leone and H. L. Singer, "On the coupling of an external electromagnetic field to a printed circuit board trace," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 41, no. 4, pp. 418–424, Nov 1999.
- [74] J. L. Lagos and F. Fiori, "Worst-case induced disturbances in digital and analog interchip interconnects by an external electromagnetic plane wave—Part I: Modeling and algorithm," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 53, no. 1, pp. 178–184, Feb 2011.
- [75] P. Kocher, J. Jaffe, B. Jun, and P. Rohatgi, "Introduction to differential power analysis," *Journal of Cryptographic Engineering (JCEN)*, vol. 1, no. 1, pp. 5–27, Apr 2011.
- [76] M. G. Bäckström and K. G. Lovstrand, "Susceptibility of electronic systems to high-power microwaves: Summary of test experience," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 46, no. 3, pp. 396–403, Aug 2004.
- [77] D. V. Giri and F. M. Tesche, "Classification of intentional electromagnetic environments (IEME)," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 46, no. 3, pp. 322–328, Aug 2004.
- [78] D. Månsson, R. Thottappillil, M. G. Bäckström, and O. Lundén, "Vulnerability of European rail traffic management system to radiated intentional EMI," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 50, no. 1, pp. 101–109, Feb 2008.
- [79] T. Wolfgramm, A. Manicke, and H. G. Krauthäuser, "Field coupling to nonlinear circuits in resonating structures," in *IEEE International Symposium on Electromagnetic Compatibility (EMC)*, 2015.
- [80] J. Benford, J. A. Swegle, and E. Schamiloglu, *High Power Microwaves*, 3rd ed. CRC Press, 2016.
- [81] D. Pozar, *Microwave Engineering*, 4th ed. Wiley, 2011.
- [82] W. Inrich, L. Batz, R. Müller, and R. Tobisch, "Electromagnetic interference of pacemakers by mobile phones," *Pacing and Clinical Electrophysiology (PACE)*, vol. 19, no. 10, pp. 1431–1446, Jun 2006.
- [83] D. L. Hayes, P. J. Wang, D. W. Reynolds, N. A. M. Estes, J. L. Griffith, R. A. Steffens, G. L. Carlo, G. K. Findlay, and C. M. Johnson, "Interference with cardiac pacemakers by cellular telephones," *New England Journal of Medicine (NEJM)*, vol. 336, no. 21, pp. 1473–1479, May 1997.
- [84] S. L. Pinski and R. G. Trohman, "Interference in implanted cardiac devices, part I," *Pacing and Clinical Electrophysiology (PACE)*, vol. 25, no. 9, pp. 1367–1381, Sep 2002.
- [85] —, "Interference in implanted cardiac devices, part II," *Pacing and Clinical Electrophysiology (PACE)*, vol. 25, no. 10, pp. 1496–1509, Oct 2002.
- [86] V. Barbaro, P. Bartolini, G. Calcagnini, F. Censi, B. Beard, P. S. Ruggera, and D. M. Witters, "On the mechanisms of interference between mobile phones and pacemakers: Parasitic demodulation of GSM signal by the sensing amplifier," *Physics in Medicine & Biology*, vol. 48, no. 11, pp. 1661–1671, Jun 2003.
- [87] A. Cheng, S. Nazarian, D. D. Spragg, K. Bilchick, H. Tandri, L. Mark, H. Halperin, H. Calkins, R. D. Berger, and C. A. Henrikson, "Effects of surgical and endoscopic electrocautery on modern-day permanent pacemaker and implantable cardioverter-defibrillator systems," *Pacing and Clinical Electrophysiology*, vol. 31, no. 3, pp. 344–350, Mar 2008.
- [88] L. Cohan, F. M. Kusumoto, and N. F. Goldschlager, "Environmental effects on cardiac pacing systems," in *Cardiac Pacing for the Clinician*, F. M. Kusumoto and N. F. Goldschlager, Eds. Springer, 2008, pp. 595–618.
- [89] S. Lee, K. Fu, T. Kohno, B. Ransford, and W. H. Maisel, "Clinically significant magnetic interference of implanted cardiac devices by portable headphones," *Heart Rhythm*, vol. 6, no. 10, pp. 1432–1436, Oct 2009.
- [90] S. J. Seidman, R. Brockman, B. M. Lewis, J. Guag, M. J. Shein, W. J. Clement, J. Kippola, D. Digby, C. Barber, and D. Huntwork, "In vitro tests reveal sample radiofrequency identification readers inducing clinically significant electromagnetic interference to implantable pacemakers and implantable cardioverter-defibrillators," *Heart Rhythm*, vol. 7, no. 1, pp. 99–107, Jan 2010.
- [91] J. Misiri, F. Kusumoto, and N. Goldschlager, "Electromagnetic interference and implanted cardiac devices: The nonmedical environment (part I)," *Clinical Cardiology*, vol. 35, no. 5, pp. 276–280, May 2012.
- [92] —, "Electromagnetic interference and implanted cardiac devices: The medical environment (part II)," *Clinical Cardiology*, vol. 35, no. 6, pp. 321–328, Jun 2012.

- [93] T. Buczkowski, D. Janusek, H. Zavala-Fernandez, M. Skrok, M. Kania, and A. Liebert, "Influence of mobile phones on the quality of ECG signal acquired by medical devices," *Measurement Science Review*, vol. 13, no. 5, pp. 231–236, Nov 2013.
- [94] S. Driessen, A. Napp, K. Schmiedchen, T. Kraus, and D. Stunder, "Electromagnetic interference in cardiac electronic implants caused by novel electrical appliances emitting electromagnetic fields in the intermediate frequency range: A systematic review," *EP Europace*, vol. 21, no. 2, pp. 219–229, Feb 2019.
- [95] F. Censi, G. Calcagnini, M. Triventi, E. Mattei, and P. Bartolini, "Interference between mobile phones and pacemakers: A look inside," *Annali dell'Istituto Superiore di Sanità*, vol. 43, no. 3, pp. 254–259, Sep 2007.
- [96] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, "They can hear your heartbeats: Non-invasive security for implantable medical devices," in *ACM Special Interest Group on Data Communication Conference (SIGCOMM)*, 2011.
- [97] M. Rushanan, A. D. Rubin, D. Foo Kune, and C. M. Swanson, "SoK: Security and privacy in implantable medical devices and body area network," in *IEEE Symposium on Security and Privacy (S&P)*, 2014.
- [98] R. Hoad, N. J. Carter, D. Herke, and S. P. Watkins, "Trends in EM susceptibility of IT equipment," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 46, no. 3, pp. 390–395, Aug 2004.
- [99] J. Delsing, J. Ekman, J. Johansson, S. Sundberg, M. G. Bäckström, and T. Nilsson, "Susceptibility of sensor networks to intentional electromagnetic interference," in *International Zürich Symposium on Electromagnetic Compatibility (EMCZUR)*, 2006.
- [100] D. Månsson, R. Thottappillil, T. Nilsson, O. Lundén, and M. G. Bäckström, "Susceptibility of civilian GPS receivers to electromagnetic radiation," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 50, no. 2, pp. 434–437, May 2008.
- [101] F. Sabath, "Classification of electromagnetic effects at system level," in *International Symposium and Exhibition on Electromagnetic Compatibility (EMC Europe)*, 2008.
- [102] F. Brauer, F. Sabath, and J. L. ter Haseborg, "Susceptibility of IT network systems to interferences by HPEM," in *IEEE International Symposium on Electromagnetic Compatibility (EMC)*, 2009.
- [103] L. Palíšek and L. Suchý, "High Power Microwave effects on computer networks," in *International Symposium and Exhibition on Electromagnetic Compatibility (EMC Europe)*, 2011.
- [104] Y. Tu, S. Rampazzi, B. Hao, A. Rodriguez, K. Fu, and X. Hei, "Trick or heat? Manipulating critical temperature-based control systems using rectification attacks," in *ACM Conference on Computer and Communications Security (CCS)*, 2019.
- [105] Y. Zhang and K. B. Rasmussen, "Detection of electromagnetic interference attacks on sensor systems," in *IEEE Symposium on Security and Privacy (S&P)*, 2020.
- [106] Y. V. Parfenov, L. N. Zdoukhov, W. A. Radasky, and M. Ianoz, "Conducted IEMI threats for commercial buildings," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 46, no. 3, pp. 404–411, Aug 2004.
- [107] D. Månsson, T. Nilsson, R. Thottappillil, and M. G. Bäckström, "Propagation of UWB transients in low-voltage installation power cables," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 49, no. 3, pp. 585–592, Aug 2007.
- [108] D. Månsson, R. Thottappillil, and M. G. Bäckström, "Propagation of UWB transients in low-voltage power installation networks," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 50, no. 3, pp. 619–629, Aug 2008.
- [109] F. M. Tesche, J. M. Keen, and C. M. Butler, "Example of the use of the BLT equation for EM field propagation and coupling calculations," *URSI Radio Science Bulletin*, vol. 2005, no. 312, pp. 32–47, Mar 2005.
- [110] F. M. Tesche, "Development and use of the BLT equation in the time domain as applied to a coaxial cable," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 49, no. 1, pp. 3–11, Feb 2007.
- [111] S.-J. Guo, L.-S. Wu, M. Tang, and J.-F. Mao, "Analysis of illuminated bent microstrip line based on Baum-Liu-Tesche (BLT) equation," in *IEEE Electrical Design of Advanced Packaging and Systems Symposium (EDAPS)*, 2015.
- [112] F. Lafon, F. De Daran, M. Ramdani, R. Perdriau, and M. Drissi, "Immunity modeling of integrated circuits: An industrial case," *IEICE Transactions on Communications*, vol. E93.B, no. 7, pp. 1723–1730, Jul 2010.
- [113] A. Richelli, G. Delaini, M. Grassi, and J.-M. Redouté, "Susceptibility of operational amplifiers to conducted EMI injected through the ground plane into their output terminal," *IEEE Transactions on Reliability*, vol. 65, no. 3, pp. 1369–1379, Sep 2016.
- [114] A. Richelli, L. Colalongo, L. Toninelli, I. Rusu, and J.-M. Redouté, "Measurements of EMI susceptibility of precision voltage references," in *Workshop on Electromagnetic Compatibility of Integrated Circuits (EMC Compo)*, 2017.
- [115] M. Ramdani, E. Sicard, A. Boyer, S. Ben Dhia, J. J. Whalen, T. H. Hubing, M. Coenen, and O. Wada, "The electromagnetic compatibility of integrated circuits—past, present, and future," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 51, no. 1, pp. 78–100, Feb 2009.
- [116] I. Giechaskiel, Y. Zhang, and K. B. Rasmussen, "A framework for evaluating security in the presence of signal injection attacks," arXiv 1901.03675, 2019, <http://arxiv.org/abs/1901.03675>.
- [117] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE Journal of Solid-State Circuits (JSSC)*, vol. 34, no. 6, pp. 790–804, Jun 1999.
- [118] B. Sunar, W. J. Martin, and D. R. Stinson, "A provably secure true random number generator with built-in tolerance to active attacks," *IEEE Transactions on Computers (TC)*, vol. 56, no. 1, pp. 109–119, Jan 2007.
- [119] R. Adler, "A study of locking phenomena in oscillators," *Proceedings of the IRE (JRPROC)*, vol. 34, no. 6, pp. 351–357, Jun 1946.
- [120] B. Mesgarzadeh and A. Alvandpour, "A study of injection locking in ring oscillators," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2005.
- [121] N. Bochard, F. Bernard, V. Fischer, and B. Valtchanov, "True-randomness and pseudo-randomness in ring oscillator-based true random number generators," *International Journal of Reconfigurable Computing*, vol. 2010, pp. 879 281:1–879 281:13, Dec 2010.
- [122] Y. Cao, V. Rožić, B. Yang, J. Balasch, and I. M. R. Verbaauwhede, "Exploring active manipulation attacks on the TERO random number generator," in *IEEE International Midwest Symposium on Circuits and Systems (MWSCAS)*, 2016.
- [123] D. Mahmoud and M. Stojilović, "Timing violation induced faults in multi-tenant FPGAs," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2019.
- [124] M. N. Armenise, C. Ciminelli, F. Dell'Olivo, and V. M. N. Passaro, *Advances in Gyroscope Technologies*, 1st ed. Springer, 2011.
- [125] C. Acar and A. Shkel, *MEMS Vibratory Gyroscopes: Structural Approaches to Improve Robustness*, 2nd ed. Springer, 2009.
- [126] S. T. Castro, R. N. Dean, G. Roth, G. T. Flowers, and B. E. Grantham, "Influence of acoustic noise on the dynamic performance of MEMS gyroscopes," in *International Mechanical Engineering Congress and Exposition (IMECE)*, 2007.
- [127] R. N. Dean, G. T. Flowers, A. S. Hodel, G. Roth, S. T. Castro, R. Zhou, A. Moreira, A. Ahmed, R. Rifki, B. E. Grantham, D. A. Bittle, and J. P. Brunsch, "On the degradation of MEMS gyroscope performance in the presence of high power acoustic noise," in *IEEE International Symposium on Industrial Electronics (ISIE)*, 2007.
- [128] R. N. Dean, S. T. Castro, G. T. Flowers, G. Roth, A. Ahmed, A. S. Hodel, B. E. Grantham, D. A. Bittle, and J. P. Brunsch, "A characterization of the performance of a MEMS gyroscope in acoustically harsh environments," *IEEE Transactions on Industrial Electronics (TIE)*, vol. 58, no. 7, pp. 2591–2596, Jul 2011.
- [129] S. Khazaaleh, G. Korres, M. Eid, M. Rasras, and M. F. Daqaq, "Vulnerability of MEMS gyroscopes to targeted acoustic attacks," *IEEE Access*, vol. 7, pp. 89 534–89 543, Jul 2019.
- [130] F. Laermer, "Mechanical microsensors," in *MEMS: A Practical Guide to Design, Analysis, and Applications*, J. G. Korvink and O. Paul, Eds. Springer, 2006, pp. 523–566.
- [131] Sixth Tone, "The gadget that boosts your step count while you nap," <https://www.sixthtone.com/news/1002530/the-gadget-that-boosts-your-step-count-while-you-nap->, Created: 2018-06-28. Accessed: 2019-11-07.
- [132] B. Gregg, "Unusual disk latency," <http://www.brendangregg.com/blog/2008-12-31/unusual-disk-latency.html>, 2008, Created: 2008-12-31. Accessed: 2019-11-07.
- [133] L. Song and P. Mittal, "POSTER: Inaudible voice commands," in *ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [134] 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)*, Mar 2019.
- [135] T. Vaidya, Y. Zhang, M. Sherr, and C. Shields, "Cocaine noodles: Exploiting the gap between human and machine speech recognition," in *USENIX Workshop on Offensive Technologies (WOOT)*, 2015.

- [136] N. Carlini, P. Mishra, T. Vaidya, Y. Zhang, M. Sherr, C. Shields, D. Wagner, and W. Zhou, "Hidden voice commands," in *USENIX Security Symposium*, 2016.
- [137] J. Petit and S. E. Shladover, "Potential cyberattacks on automated vehicles," *IEEE Transactions on Intelligent Transportation Systems (TITS)*, vol. 16, no. 2, pp. 546–556, Apr 2015.
- [138] J. Loughry, ("Oops! Had the silly thing in reverse")—Optical injection attacks in through LED status indicators," in *International Symposium and Exhibition on Electromagnetic Compatibility (EMC Europe)*, 2019.
- [139] F. M. Mims, III, "Bidirectional optoisolator puts two LEDs nose to nose," *Electronics*, vol. 52, no. 10, p. 127, May 1979.
- [140] —, "Sun photometer with light-emitting diodes as spectrally selective detectors," *Applied Optics*, vol. 31, no. 33, pp. 6965–6967, Nov 1992.
- [141] M. Šimka and P. Komenského, "Active non-invasive attack on true random number generator," in *PhD Student Conference and Scientific and Technical Competition of Students of FEI TU Košice, Košice, Slovakia*, 2006.
- [142] M. Soucarros, C. Canovas-Dumas, J. Clédière, P. Elbaz-Vincent, and D. Réal, "Influence of the temperature on true random number generators," in *IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, 2011.
- [143] A. Cherkaoui, V. Fischer, A. Aubert, and L. Fesquet, "A self-timed ring based true random number generator," in *IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)*, 2013.
- [144] H. Martín, P. Martín-Holgado, P. Peris-Lopez, Y. Morilla, and L. Entrena, "On the entropy of oscillator-based true random number generators under ionizing radiation," *Entropy*, vol. 20, no. 7, pp. 1–11, Jul 2018.
- [145] A. Agne, H. Hangmann, M. Happe, M. Platzner, and C. Plessl, "Seven recipes for setting your FPGA on fire – A cookbook on heat generators," *Microprocessors and Microsystems*, vol. 38, no. 8, pp. 911–919, Nov 2014.
- [146] T. Dutta and A. R. Barnard, "Performance of hard disk drives in high noise environments," *Noise Control Engineering Journal*, vol. 65, no. 5, pp. 386–395, Sep 2017.
- [147] D. Muniraj and M. Farhood, "Detection and mitigation of actuator attacks on small unmanned aircraft systems," *Control Engineering Practice*, vol. 83, pp. 188–202, Feb 2019.
- [148] K. S. Tharayil, B. Farshtindiker, S. Eyal, N. Hasidim, R. Hershkovitz, S. Houry, I. Yoffe (Iofedov), M. Oren, and Y. Oren, "Sensor defense in-software (SDI): Practical software based detection of spoofing attacks on position sensors," arXiv 1905.04691, 2019, <http://arxiv.org/abs/1905.04691>.
- [149] J. Park, R. Ivanov, J. Weimer, M. Pajic, and I. Lee, "Sensor attack detection in the presence of transient faults," in *ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*, 2013.
- [150] R. Ivanov, M. Pajic, and I. Lee, "Attack-resilient sensor fusion for safety-critical cyber-physical systems," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 15, no. 1, pp. 21:1–21:24, Feb 2016.
- [151] F. M. Tesche, M. V. Ianoz, and T. Karlsson, *EMC Analysis Methods and Computational Models*, 1st ed. Wiley, 1996.
- [152] L. Blue, L. Vargas, and P. Traynor, "Hello, is it me you're looking for? Differentiating between human and electronic speakers for voice interface security," in *ACM Conference on Security & Privacy in Wireless and Mobile Networks (WiSec)*, 2018.
- [153] F. Wan, F. Duval, H. Cao, X. Savatier, A. Louis, and B. Mazari, "Increase of immunity of microcontroller to conducted continuous-wave interference by detection method," *Electronics Letters*, vol. 46, no. 16, pp. 1113–1114, Aug 2010.
- [154] F. Wan, F. Duval, X. Savatier, A. Louis, and B. Mazari, "Electromagnetic interference detection method to increase the immunity of a microcontroller-based system in a complex electromagnetic environment," *IET Science, Measurement & Technology*, vol. 6, no. 4, pp. 254–260, Jul 2012.
- [155] S. Sridhar and G. Manimaran, "Model-based attack detection and mitigation for automatic generation control," *IEEE Transactions on Smart Grid (TSG)*, vol. 5, no. 2, pp. 580–591, Mar 2014.
- [156] J. Jiang and Y. Qian, "Defense mechanisms against data injection attacks in smart grid networks," *IEEE Communications Magazine*, vol. 55, no. 10, pp. 76–82, Oct 2017.
- [157] J.-M. Redouté and A. Richelli, "A fundamental approach to EMI resistant folded cascode operational amplifier design," in *International Symposium and Exhibition on Electromagnetic Compatibility (EMC Europe)*, 2013.
- [158] —, "A methodological approach to EMI resistant analog integrated circuit design," *IEEE Electromagnetic Compatibility Magazine*, vol. 4, no. 2, pp. 92–100, Aug 2015.
- [159] C. Bona and F. Fiori, "A new filtering technique that makes power transistors immune to EMI," *IEEE Transactions on Power Electronics*, vol. 26, no. 10, pp. 2946–2955, Oct 2011.
- [160] C. Walravens, S. Van Winchel, J.-M. Redouté, and M. Steyaert, "Efficient reduction of electromagnetic interference effects in operational amplifiers," *Electronics Letters*, vol. 43, no. 2, pp. 84–85, Jan 2007.
- [161] F. Michel and M. Steyaert, "Differential input topologies with immunity to electromagnetic interference," in *European Solid-State Circuits Conference (ESSCIRC)*, 2011.
- [162] S. Sbaraini, A. Richelli, and Z. M. Kovács-Vajna, "EMI susceptibility in bulk-driven Miller opamp," *Electronics Letters*, vol. 46, no. 16, pp. 1111–1113, Aug 2010.
- [163] A. Richelli, L. Colalongo, M. Quarantelli, and Z. M. Kovács-Vajna, "Robust design of low EMI susceptibility CMOS OpAmp," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 46, no. 2, pp. 291–298, May 2004.
- [164] F. Fiori, "Design of an operational amplifier input stage immune to EMI," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 49, no. 4, pp. 834–839, Nov 2007.
- [165] F. Michel and M. Steyaert, "Comparison of high impedance input topologies with low EMI susceptibility," *Analog Integrated Circuits and Signal Processing*, vol. 65, no. 2, pp. 299–309, Nov 2010.
- [166] M. J. van der Horst, W. A. Serdijn, and A. C. Linnenbank, *EMI-Resilient Amplifier Circuits*, 1st ed. Springer, 2014.
- [167] A. Richelli, "EMI susceptibility issue in analog front-end for sensor applications," *Journal of Sensors*, vol. 2016, pp. 1 082 454:1–1 082 454:9, 2016.
- [168] D. E. Serrano, M. F. Zaman, A. Rahafrooz, P. Hruday, R. Lipka, D. Younkin, S. Nagpal, I. Jafri, and F. Ayazi, "Substrate-decoupled, bulk-acoustic wave gyroscopes: Design and evaluation of next-generation environmentally robust devices," *Microsystems & Nanoengineering*, vol. 2, pp. 16015:1–16015:10, Dec 2016.
- [169] D. E. Serrano, R. Lipka, D. Younkin, P. Hruday, J. Tovera, A. Rahafrooz, M. F. Zaman, S. Nagpal, I. Jafri, and F. Ayazi, "Environmentally-robust high-performance tri-axial bulk acoustic wave gyroscopes," in *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2016.
- [170] B. Farshtindiker, N. Hasidim, A. Grosz, and Y. Oren, "How to phone home with someone else's phone: Information exfiltration using intentional sound noise on gyroscopic sensors," in *USENIX Workshop on Offensive Technologies (WOOT)*, 2016.
- [171] K. Block, S. Narain, and G. Noubir, "An autonomic and permissionless Android covert channel," in *ACM Conference on Security & Privacy in Wireless and Mobile Networks (WiSec)*, 2017.
- [172] N. Matyunin, J. Szefer, and S. Katzenbeisser, "Zero-permission acoustic cross-device tracking," in *IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, 2018.
- [173] Y. Michalevsky, D. Boneh, and G. Nakibly, "Gyrophone: Recognizing speech from gyroscope signals," in *USENIX Security Symposium*, 2014.
- [174] S. A. Anand and N. Saxena, "Speechless: Analyzing the threat to speech privacy from smartphone motion sensors," in *IEEE Symposium on Security and Privacy (S&P)*, 2018.
- [175] S. A. Anand, C. Wang, J. Liu, N. Saxena, and Y. Chen, "Spearphone: A speech privacy exploit via accelerometer-sensed reverberations from smartphone loudspeakers," arXiv 1907.05972, 2019, <http://arxiv.org/abs/1907.05972>.
- [176] 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, May 2019.
- [177] Y. Chen, H. Li, S. Nagels, Z. Li, P. Lopes, B. Y. Zhao, and H. Zheng, "Understanding the effectiveness of ultrasonic microphone jammer," arXiv 1904.08490, 2019, <http://arxiv.org/abs/1904.08490>.
- [178] M. G. Kuhn and R. J. Anderson, "Soft Tempest: Hidden data transmission using electromagnetic emanations," in *International Workshop on Information Hiding (IH)*, 1998.
- [179] H. Tanaka, O. Takizawa, and A. Yamamura, "Evaluation and improvement of the Tempest fonts," in *International Workshop on Information Security Applications (WISA)*, 2005.
- [180] H. Tanaka, "Information leakage via electromagnetic emanations and evaluation of Tempest countermeasures," in *International Conference on Information Systems Security (ICISS)*, 2007.
- [181] W. Van Eck, "Electromagnetic radiation from video display units: An eavesdropping risk?" *Computers & Security*, vol. 4, no. 4, pp. 269–286, Dec 1985.

- [182] M. G. Kuhn, "Electromagnetic eavesdropping risks of flat-panel displays," in *International Workshop on Privacy Enhancing Technologies (PET)*, 2004.
- [183] —, "Compromising emanations of LCD TV sets," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 55, no. 3, pp. 564–570, Jun 2013.
- [184] L. Zhang, C. Bo, J. Hou, X.-Y. Li, Y. Wang, K. Liu, and Y. Liu, "Kaleido: You can watch it but cannot record it," in *International Conference on Mobile Computing and Networking (MobiCom)*, 2015.
- [185] S. Zhu, C. Zhang, and X. Zhang, "Automating visual privacy protection using a smart LED," in *International Conference on Mobile Computing and Networking (MobiCom)*, 2017.
- [186] S. Wendzel, S. Zander, B. Fechner, and C. Herdin, "Pattern-based survey and categorization of network covert channel techniques," *ACM Computing Surveys (CSUR)*, vol. 47, no. 3, pp. 1–27, Apr 2015.
- [187] M. Guri, Y. Solewicz, A. Daidakulov, and Y. Elovici, "SPEAKE(a)R: Turn speakers to microphones for fun and profit," in *USENIX Workshop on Offensive Technologies (WOOT)*, 2017.
- [188] A. Kwong, W. Xu, and K. Fu, "Hard drive of hearing: Disks that eavesdrop with a synthesized microphone," in *IEEE Symposium on Security and Privacy (S&P)*, 2019.
- [189] J. Lopes Esteves, E. Cottais, and C. Kasmi, "Unlocking the access to the effects induced by IEMI on a civilian UAV," in *International Symposium and Exhibition on Electromagnetic Compatibility (EMC Europe)*, 2018.
- [190] D. Månsson, R. Thottappillil, and M. G. Bäckström, "Methodology for classifying facilities with respect to intentional EMI," *IEEE Transactions on Electromagnetic Compatibility (TEM C)*, vol. 51, no. 1, pp. 46–52, Feb 2009.
- [191] S. Dey, N. Roy, W. Xu, R. R. Choudhury, and S. Nelakuditi, "Accelerometer: Imperfections of accelerometers make smartphones trackable," in *Network and Distributed System Security Symposium (NDSS)*, 2014.



Ilias Giechaskiel recently submitted his PhD thesis at the Department of Computer Science at the University of Oxford. He was a Clarendon and Cyber Security Scholar at Kellogg College, and was also funded by the EPSRC and the Oxford CDT in Cyber Security. His dissertation, "Leaky Hardware: Modeling and Exploiting Imperfections in Embedded Devices" was supervised by Prof. Kasper B. Rasmussen. Ilias holds a Master's in Advanced Computer Science (with distinction) from the University of Cambridge, and a Bachelor's in Mathematics (summa cum laude) from Princeton University. During his PhD, Ilias was also a Visiting Assistant in Research at Yale University, where he worked under Prof. Jakub Szefer on FPGA covert channels. His interests in hardware security extend beyond academia: Ilias has participated in numerous security Capture-the-Flag (CTF) competitions, and has interned at Bloomberg, Microsoft, Dropbox, Microsoft Research, and Jump Trading.



Kasper B. Rasmussen is an Associate Professor in the Computer Science Department at the University of Oxford. He joined the department in 2013 and was awarded a University Research Fellowship from the Royal Society of London in 2015. Prior to being at Oxford, Kasper spent two years as a post-doc at University of California, Irvine. Kasper completed his Ph.D. with Prof. Srdjan Capkun in the Department of Computer Science at ETH Zurich, where he worked on security issues related to secure time synchronization and localization, with a particular focus on distance bounding. His thesis won the "ETH Medal" for outstanding dissertation from the Swiss Federal Institute of Technology, and he was additionally awarded the Swiss National Science Foundation (SNSF) Fellowship for prospective researchers.