Detection of Electromagnetic Signal Injection Attacks on Actuator Systems

Youqian Zhang
youqian.zhang@cs.ox.ac.uk
University of Oxford
Oxford, United Kingdom

Kasper Rasmussen
kasper.rasmussen@cs.ox.ac.uk
University of Oxford
Oxford, United Kingdom

ABSTRACT

An actuator is a device that converts electricity into another form of energy, typically physical movement. They are absolutely essential for any system that needs to impact or modify the physical world, and are used in millions of systems of all sizes, all over the world, from cars and spacecraft to factory control systems and critical infrastructure. An actuator is a "dumb device" that is entirely controlled by the surrounding electronics, e.g., a microcontroller, and thus cannot authenticate its control signals or do any other form of processing. The problem we look at in this paper is how the wires that connect an actuator to its control electronics can act like antennas, picking up electromagnetic signals from the environment. This makes it possible for a remote attacker to wirelessly inject signals (energy) into these wires to bypass the controller and directly control the actuator.

To detect such attacks, we propose a novel detection method that allows the microcontroller to monitor the control signal and detect attacks as a deviation from the intended value. We have managed to do this without requiring the microcontroller to sample the signal at a high rate or run any signal processing. That makes our defense mechanism practical and easy to integrate into existing systems. Our method is general and applies to any type of actuator (provided a few basic assumptions are met), and can deal with adversaries with arbitrarily high transmission power. We implement our detection method on two different practical systems to show its generality, effectiveness, and robustness.

CCS CONCEPTS

• **Security and privacy → Embedded systems security**.

KEYWORDS

electromagnetic signal injection; cyber-physical system security; actuator

ACM Reference Format:
Youqian Zhang and Kasper Rasmussen. 2022. Detection of Electromagnetic Signal Injection Attacks on Actuator Systems. In *25th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2022), October 26–28, 2022, Limassol, Cyprus.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3545948.3545949

RAID 2022, October 26–28, 2022, Limassol, Cyprus

ACM ISBN 978-1-4503-9704-9/22/10…$15.00
https://doi.org/10.1145/3545948.3545949

1 INTRODUCTION

Actuator systems are embedded in our daily lives to such an extent that it is hard to find an example of an electronic system that does not have actuators in some form. Anything from household devices like smart locks or robotic vacuum cleaners, to transportation to production to defense. Actuators are everywhere.

These devices interact with the physical world by converting an electric signal into some other form of energy, typically movement (e.g., motors), but a heater or a light source can also be considered an actuator. It is well known that the wires used to feed electrical signals to such devices can act as antennas, unintentionally picking up electromagnetic interference (EMI) [18, 25, 27] from the environment, or indeed from an attacker. This inherent vulnerability allows an attacker to *wirelessly* inject attacking signals into the wires, disturbing or changing the original signals. Please refer to Section 2 for more details about such signal injection attacks.

Since an actuator is simply an energy transducer, it cannot authenticate its input signals and will respond to whatever it receives, in the worst case resulting in the adversary being able to fully control the state of the actuator. It is easy to see how such an attack can be used, e.g., to rotate the motor in the smart lock to unlock a door; or force closed a fuel injection valve in a car to stop the car's engine. When the target system is complex and important, these attacks can be incredibly powerful and dangerous. For example, imagine the potential harm if an adversary could control critical industrial applications (e.g., robotic arms) or medical devices (e.g., pacemakers), or say, move the control surfaces of an airplane without pilot input.

Even though such attacks are complicated to perform in practice, and as a result are still rare, we need to find effective detection and mitigation strategies to deal with them before they become common. In the last few years, there has been work on detecting such attacks on sensors [12, 24, 36, 42, 51]. In this paper, we focus on detecting attacks on actuators, which is quite a bit harder. The reason is that when a sensor is attacked, the receiving device is a microcontroller that has the ability to run filters and algorithms, or use redundant measurements for added security. For actuators, that is not as easy. When actuators are attacked the receiving device is the actuator itself, and since actuators are "dumb devices" (it might just be a coil of wire, like in a motor), they do not have the ability to ignore malicious signals, even if such signals deviate from some usual pattern.

In this paper, we provide a novel detection method that uses common and inexpensive electrical components, making it possible

to apply our method at scale. The basic idea is to compare the signal to be protected with a reference, in order to identify when any external interference is present. However, this is not as easy as it sounds. First of all, an adversarial signal will affect any reference signals as well, and there are challenges with the sampling rate, bandwidth limits, and signal processing efforts that can make a trivial scheme unusable in practice. Our detection method solves all those problems and we are able to provide strong detection guarantees for (almost) any actuator system.

Our contributions of this work are summarized as follows:

- We create a universal and flexible system model that fits most (if not all) actuator systems. It allows us to capture the specific needs of any system by tuning parameters of the model (Section 3).
- We propose a general and lightweight detection method that uses differential amplifiers to detect electromagnetic signal injection attacks, and we show that it can provide the actuator system with a strong security guarantee (Section 4 and Section 5).
- We implement the proposed detection method on a speaker system and a motor control system, and we demonstrate the generality, feasibility, and robustness of our detection method (Section 6).

In the remaining parts of this work, we present a background on the electromagnetic signal injection attacks in Section 2. Furthermore, additional important issues are discussed in Section 7, and related work is summarized in Section 8. Finally, a conclusion of this work is drawn in Section 9.

2 BACKGROUND

Before introducing our detection method, we provide a brief background on electromagnetic signal injection attacks. We first explain how electromagnetic waves are injected into a victim system, and next, we explain how the injected signals influence the actuator, as well as presenting successful attacks in previous studies.

2.1 Electromagnetic Signal Injection

Electromagnetic fields can affect a metal conductor by inducing voltage changes into it, and this has been thoroughly studied in the area of "Electromagnetism". Besides antennas for wireless communications, the metal conductor also exists in devices in the form of wires (or traces) that connect electronic components. These wires can also act as antennas to capture environmental electromagnetic waves. A typical example is that a headphone is required while playing FM radio on a mobile phone, where the headphone cable acts as the antenna. Many researchers exploited such "antenna-like" behaviors of the wires to inject malicious signals into the circuits [8, 9, 17, 22, 25, 29, 31, 33, 37–40, 44, 46–48, 51].

Many factors affect the injection process, but the attack power and the attack frequency are the basic ones that an attacker tunes, as they determine the effectiveness and efficiency of the injection [50]. To cause effective impacts onto the circuits, the injected voltage needs to be strong enough. Since the injected voltage is proportional to the attack power [15], the more powerful the attacking signal is, the higher the injected voltage will be, and it is more likely the attack is effective. In addition, in order to maximize the injected voltage, the attack frequency must be the resonant frequency of the wire; at other frequencies, it will cost more attack power to achieve the same amount of injected voltage [25]. By properly tuning the attack power and the attack frequency, the attacker can inject arbitrary signals into the wires.

2.2 Circuit Response to Injected Signal

After the injection, a successful attack depends on how the circuits respond to injected signals. On the one hand, the injected signal can be within the frequency band in which the circuits are designed to operate, namely the operational band (in-band). Since the malicious voltage changes are within the operational band, the circuits respond to them directly. This will subsequently influence a signal that drives the actuator, further impacting the actuator responses. On the other hand, the injected signal can also be out of the operational band (out-of-band). However, in order to affect the circuits in an effective and predictable way, it is essential to cause voltage changes within the operational band. A well-studied method of transferring the out-of-band changes to the operational band is exploiting the nonlinear properties of electronic components in the victim circuits: the attacker first exerts an in-band malicious signal onto an out-of-band radio-frequency (RF) carrier to form the attacking signal; next, after the signal injection, the malicious signal is extracted from the attacking signal due to nonlinearities of electronic components such as amplifiers [22, 25, 44], electro-static discharge (ESD) circuits [9, 38], and analog-to-digital converters (ADCs) [16, 17]; as a result, the in-band malicious signal appears in the operational band, further affecting the actuator.

Here are some examples of manipulating actuators. Selvaraj et al. [38] demonstrated how to inject fine-tuned attacking signals into the target wire to precisely manipulate servo angles. They exploited the nonlinearities of the electro-static discharge (ESD) circuits to toggle the voltage level of the signal so that they can precisely control the signal pulse width that determines the servo angle. Dayanikli et al. [9] demonstrated similar attacks: they could remotely manipulate the signal that controls switches in an AC-to-DC power converter, which regulates power delivery to electric vehicles. The attack can forcibly toggle the on/off state of the switch. An irreparable result of the attack is causing a short circuit to burn the converter. In this work, we will demonstrate that an arbitrary audio signal can be injected into a speaker system, in which the nonlinearities of the audio amplifier is exploited (see details in Section 6.2); furthermore, we will show that a motor control system can be precisely controlled, in which the imperfection of transistors [4, 5, 13] are exploited (see details in Section 6.3). All of these show that it is not difficult to use electromagnetic signal injection attacks to manipulate the actuator behaviors precisely.

3 SYSTEM MODEL AND ADVERSARY MODEL

In this section, we introduce a general and flexible system model that fits most actuator systems. This allows us to capture the needs of any specific system by tuning the model parameters. We also present a comprehensive attacker model that, together with the system model, forms a flexible tool to describe signal injection attacks and defense mechanisms on actuator systems.

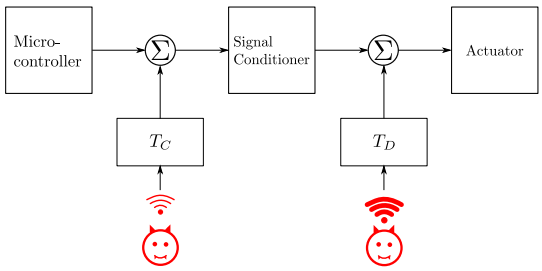

Figure 1: The actuator system consists of a microcontroller, a signal conditioner, and an actuator. The transfer functions T_C and T_D explain the control signal wire and the drive signal wire capturing the attacking signals, respectively.

3.1 System Model

We call a system that controls an actuator an "actuator system". In an actuator system, a microcontroller is the device used to regulate, command, and manage the behaviors of the actuator. Between the microcontroller and the actuator, there are circuits transforming the microcontroller output signal into a suitable signal to drive the actuator. For instance, such circuits may be for signal amplifications or waveform conversions. To capture all the characteristics of the circuits, we define a new device, called the signal conditioner. How this device works will differ from circuit to circuit, but we treat it as a black box. Therefore, our system model consists of three devices: a microcontroller, a signal conditioner, and an actuator; a block diagram of the system model is presented in Figure 1.

In the system model, wires are used to connect these devices: as shown in Figure 1, one wire is used to transmit the microcontroller output signal, which we call the *control signal*, to the signal conditioner; the other transmits the signal conditioner output signal, which we call the *drive signal*, to the actuator. Note that, in practice, the control signal wire often carries comparatively little power compared to the drive signal wire, as the voltage and the current of the microcontroller outputs are constrained to a few volts and milliamperes, respectively. Whereas the drive signal wire can carry high-power signals because some actuators consume significant power while working.

The operational frequency of the control signal and drive signal can vary significantly from system to system. Still, it is generally possible to define a normal operating range to which signals are confined in normal operation. This is important because while low/high pass filters can filter out adversarial injections at extreme frequencies, it is more difficult to filter out attacks in the operational range without affecting the valid control/drive signal. Our solution assumes that such an operational range can be defined and we call the upper limit of this range f_{max}. Note that we make no assumptions about what the value of this limit is, only that it exists. In Section 5, we discuss ways of extending this range way beyond the design limits of the electrical components.

3.2 Adversary Model

The attacker's goal is to affect the actuator by electromagnetic interference, i.e., inject an attacking signal into the system. The attacker can inject attacking signals into the actuator system remotely but cannot physically access or modify the actuator system. We grant the attacker full knowledge of the actuator system; specifically, the attacker can predict the waveform and timing (phase) of signals in the actuator system. The attacker can also craft any (physically possible) signal she wants.

In practice, signal injection can be rather complicated, especially from far away, but we deliberately grant the adversary extremely strong power to make sure our detection method works in every case. We manage this complexity using a transfer function that encapsulates any changes to the attacking signal caused by the injection process, e.g., frequency selectivity, attack distance, attenuation, spreading and convolution, etc., as shown in Figure 1. We do not limit the power available for the adversary. However, we do assume that there exists a lower limit, below which any injected signal no longer has a meaningful effect on the target system. We call this lower limit P_{min}. This power limit is set by the system designer to make sure that any injected signal above this limit is detected. It can be set arbitrarily low, but in order to successfully attack the system, the attacker must inject a signal with power higher than P_{min}.

The reason why we grant the attacker ideally strong abilities is that if such an attacker cannot avoid being detected by our proposed detection method, it is impossible for any other attackers who are no better than this ideally strong attacker to bypass the detection method.

3.3 Two Injection Points

In a particular physical system, there could exist multiple injection places through which attacking signals enter the system. Therefore, many electronic components will also be affected by the injections. However, only when these injections lead to effects on signals that directly determine responses of the system will the system be successfully manipulated by the attacker, and this has also been considered and shown in previous studies [11, 44]. Therefore, regardless of where the signal injection happens in the system, even if currents are induced in many places at once, it is possible to find an input signal that, when applied to one of the two wires in our model, produces the same effect. This means that without loss of generality, we can model any signal injection as if the attacking signal was injected into the control or drive signal wire through an appropriate transfer function. In practice, signal injection does in fact almost exclusively happen via these wires, because these are the most efficient "antennas" in the system, and thus where most of the energy is transferred.

There is an important difference between these two injection points. As mentioned previously, the power of the control signal is comparatively weak, so the adversary can more easily overshadow any valid signal on the control signal wire, and it will generally take less power to make changes that affect the actuator through this injection point. We define such an injection as a *control signal injection*. The second injection point is the drive signal wire. This wire will generally carry signals with higher power and more specialized waveforms. For some actuators, e.g., brushless electric motors, the drive signals are not only high-powered but also somewhat complex, and the timing between the different phases of the signal is very important to the operation of the motor. This means

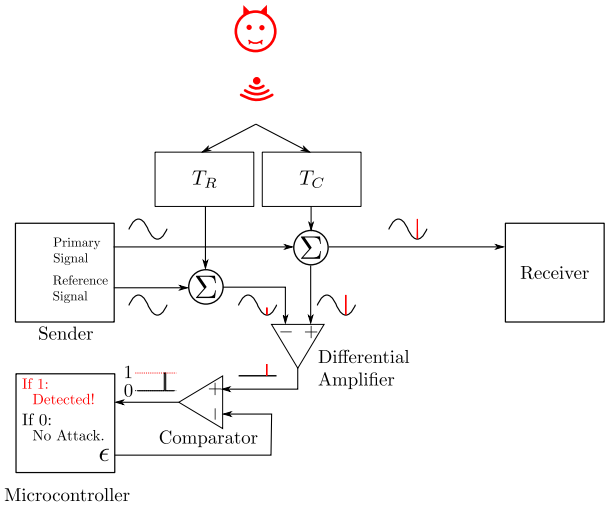

Figure 2: A differential amplifier compares the primary signal that is transmitted from the sender to the receiver with the reference signal. A comparator circuit further compares the differential amplifier output with a threshold ϵ, and the microcontroller determines whether attacks happen according to the binary results of the comparator.

that injection into this wire is more difficult and requires much more power from the adversary. We define such an injection as a *drive signal injection*.

Regarding the control signal injection, as indicated in Figure 1, only if the signal conditioner reacts to malicious changes in the control signal will the actuator be impacted. Thus, there is no possibility of performing a faster control signal injection where the circuits would not react. However, the drive signal injection impacts the actuator directly, meaning that such an injection can manipulate the actuator without any reaction from the circuits.

Despite the differences in attacker capabilities between the two injection points, our detection system, described in the following sections, works for both the control signal wire and the drive signal wire.

4 ATTACK DETECTION

As mentioned previously, our detection approach works for both the control signal wire and the drive signal wire. Rather than choose one of the two injection points for our description, we instead treat each wire as having a "sender" and a "receiver". The sender device is thus either the microcontroller or the signal conditioner, with the corresponding receiver being the signal conditioner or the actuator, respectively. We discuss any minor differences between the injection points in Section 4.4.

In order to detect a signal injected into the wire, we make the sender generate two identical signals, one primary signal (sent to the receiver) and one reference signal (used for detection). This idea is illustrated in Figure 2. A differential amplifier is then used to amplify any differences between the primary and reference signals. In the absence of an attack, these two signals should be identical and thus produce no output from the amplifier. However, if the primary wire is affected by an external signal, the difference will be amplified and can be detected using a comparator and a microcontroller. A very important requirement is that the reference wire is sufficiently different from the primary wire to make sure that the adversary cannot modify both in the same way. This can be easily accomplished by simply making the wires different lengths [3] in order to make them sensitive to different frequencies, but more significant difference can be achieved, e.g., with additional Radio Frequency (RF) shielding materials on the reference wire.

When no attack signal is present, the two input signals of the differential amplifier (the primary and the reference) are the same, and the differential amplifier output is zero. In reality, there will be a non-zero amount of noise, but the output is essentially zero. When an attack happens, the primary wire and the reference wire both pick up the attacking signal. But, because the two wires cannot be modified in the same way, captured by the two different transfer functions T_C and T_R shown in Figure 2, the two inputs to the differential amplifier will be different. This results in a non-zero signal on the output of the differential amplifier, and allows the microcontroller to detect the attack. It is essential to emphasize that simultaneously radiating two attacking signals, each carefully matched with the characteristics of a wire, will not cause the same injected signals into the two wires or avoid the detection. This is because the transfer functions essentially guarantees different frequencies of the injected signals, meaning the voltages in the wires change at different rates and eventually result in a voltage difference.

Please note that the differential amplifier is used in a novel way that is different from how it is commonly used in analog electronics, in which the differential amplifier is used to reduce equal interference (common-mode interference) onto its two inputs [34]. However, in our detection method, the two inputs are deliberately crafted such that the differential amplifier captures the attack interference, rather than mitigate it.

4.1 Modeling Differential Amplifier Output

The differential amplifier amplifies the difference of its input signals. We model this as the difference between the primary and the reference $\delta(t)$, plus additive white Gaussian noise $n(t)$, amplified by a constant gain G. To simplify the notation, we omit t hereafter. The output of the differential amplifier becomes:

$$o = G(\delta + n)$$

Given an attacking signal s, the signal that is injected in the primary wire is $T_C(s)$, and the signal that is injected in the reference wire is $T_R(s)$. In order to obtain a simple relationship between these two injected signals, we make the simplifying assumption that T_R can be expressed as being K times weaker than T_C. Therefore we can write:

$$\delta = T_C(s) - T_R(s) = T_C(s) - \frac{1}{K}T_C(s) = \frac{K-1}{K}T_C(s)$$

Thus, the output of the differential amplifier becomes

$$o = G\left(\frac{K-1}{K}T_C(s) + n\right) \tag{1}$$

Finally we take advantage of the fact that the power that is absorbed by the receiving antenna (the primary wire) is proportional

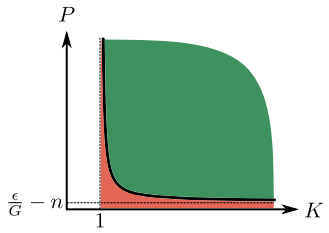

Figure 3: The minimum detectable attack power is expressed as a function of K. The detection method can detect attacks on and above the curve (green), but it cannot detect attacks below the curve (red). By decreasing (increasing) ϵ or increasing (decreasing) G, the horizontal asymptote can be moved down (up).

to the attack power P [15], to arrive at the final model for our detection system:

$$o = G\left(\frac{K-1}{K}P + n\right) \tag{2}$$

This equation gives us the output of the differential amplifier as a function of the main parameters of our detection system, namely the noise n, the gain of the differential amplifier G, the "difference" of the primary and reference wires K, and the power of the adversarial signal P.

4.2 Detection Rule and Choice of Parameters

According to Equation (2), when no attack signal is present, i.e., the attacker's power is 0, we have $o = Gn$. To make sure that small amounts of noise do not cause false positives, we define a threshold ϵ that the output of the amplifier must exceed in order to be detected as an attack. The actual detection is done by a comparator whose output is high when $o \geq \epsilon$ and low otherwise. This allows the output of the comparator to be fed into an interrupt pin of the microcontroller, as shown in Figure 2, and ensures that even attack signals with a very short duration can be detected efficiently without requiring the microcontroller to sample at a high rate. Moreover, regarding the detection latency, such a design can detect an attack immediately when the detection circuit captures the waveform difference.

Since the attack is detected if:

$$G\left(\frac{K-1}{K}P + n\right) \geq \epsilon$$

we can rearrange this to see that the detection method can detect any attack with power that fulfills the following requirement:

$$P \geq \left(\frac{\epsilon}{G} - n\right) \cdot \frac{K}{K-1} \tag{3}$$

From Inequality (3) we see that the minimum detectable power can be made arbitrarily small with appropriate choices of K, G, and ϵ. In the following, we describe the procedures for choosing these values.

The choice of K is relatively simple: bigger is better. A large K means that the difference between the two transfer functions T_C and T_R, which govern how an attacking signal affects the primary and references wires, is as big as possible. To get a sense of how the choice of K affects the detection performance, we plot the attacker's power P as a function of K in Figure 3. The detection method detects attacks on or above the curve (green region), while the attacks below the curve (red region) are not detected. We see that K does not have to be very high for the detection method to be effective, but it does have to be above 1, i.e., the primary and reference wires do have to differ.

The amplification of the differential amplifier G is dictated by the choice of amplifier used. Different amplifiers have different maximum gains, and typical values range from 100 to 300. Generally, G should be chosen as high as possible, although in noisy environments, it may be beneficial to reduce the amplification to reduce the sensitivity to noise.

As for choosing the detection threshold ϵ, it needs to be chosen such that environmental noise does not cause a detection event. Therefore, ϵ is chosen just high enough to make sure that false positives from noise are kept to a minimum; we show an example of this in our implementation in Section 6. Moreover, because noise environments are often complicated and change significantly over time, we emphasize that ϵ does not have to be constant. It can for example be adaptively adjusted to accommodate lower levels of ambient electromagnetic noise during the night. We further discuss this in Section 7.2.

4.3 Security Analysis

Recall from the adversary model that the goal of the attacker is to affect the actuator. In order to achieve this goal, the attacker must inject a signal with power of at least P_{min}. We prove that such attacks are always detected by our detection method as follows.

Substituting P_{min} into Inequality (3) we see that if the following inequality holds the attack is detected:

$$P_{min} \geq \left(\frac{\epsilon}{G} - n\right) \cdot \frac{K}{K-1}$$

To show that it is always possible to find values of K, ϵ, and G to make this inequality true, for any value of n we first make the observation that K can be made arbitrarily high independent of noise. Since $K/(K-1)$ approaches 1 for high enough values of K, we can pick a high value and reduce the above inequality to

$$G(P_{min} + n) \geq \epsilon$$

In addition, as mentioned in the previous subsection, it is a functional requirement that the detection threshold must not be triggered by the noise alone, i.e., the following must hold:

$$Gn < \epsilon$$

Both of the two inequalities above must be true in order to have a functional detection system. That gives the following constraint:

$$G(P_{min} + n) \geq \epsilon > Gn$$
$$P_{min} + n > n$$
$$P_{min} > 0$$

Thus for any value of $P_{min} > 0$, it is possible to find values of K, G, and ϵ that allows the detection system to detect any adversarial signal with power above P_{min} and at the same time do not trigger false positives from noise.

4.4 Differences Between Injection Points

Recall that one injection point is the control signal wire, and the other is the drive signal wire. The first difference is between the differential amplifiers at these two injection points. Recall that the control signal has a low voltage level, and as such, it is sufficient for the differential amplifier to have an input voltage range of several volts. However, the drive signal's voltage can go up to hundreds of volts, e.g., 380 V industrial motors. Thus, a differential amplifier with a large enough voltage input range is needed such that the tapped signal will not cause any damage to the differential amplifier. It is not hard to find such a differential amplifier in the market. Note that since the differential amplifier has a much higher impedance than the actuator, the tapping only draws a tiny portion of the control/drive signal, causing negligible impacts on the signal conditioner/actuator.

Another difference at these two injection points is that the drive signal can be much more complex than the control signal, and thus, it may be more complicated while deploying our approach for the drive signal. In the previously mentioned example of a brushless electric motor, the microcontroller produces one signal for controlling, while the signal conditioner needs to convert this solitary control signal into three different signals to drive the motor. In general, it is essential to deploy our approach to each signal to guarantee the security, which means one for the control signal and three for the drive signals. However, in many cases where the physical properties of the multiple wires are the same or very similar and they are put close to each other, it is tricky that protecting one wire is sufficiently enough, and doing so can significantly reduce the complexity of deploying our approach. This is because the attacker cannot selectively choose a wire to affect in these cases, and in other words, all of these identical or similar wires will be impacted by the attack. In the example of the brushless DC motor, its three drive signal wires are almost identical, and they are put very close to each other. Therefore, protecting any one of the wires with our approach is equivalent to protecting all three wires.

4.5 Attacks on Detection Circuit

Our defense mechanism has added circuitry to the system that could itself be the target of an injection attack. In this section, we demonstrate that this circuitry cannot be exploited by the attacker to achieve the injection.

First, we note that there is no path from our detection circuit to the actuator, so the only malicious action we have to consider is whether an adversary could inject a signal that would be hidden from detection because of interference in the detection circuit itself.

To analyze this, we define a new transfer function T for the main wire in the detection circuit. The resulting signal that is injected into the detection wire is then $T(s)$ when the adversary sends s. Please note that s also explains multiple attacking signals that are radiated by the adversary simultaneously, because the superimposition of these attacking signals makes them into one attacking signal. Note that there may also be multiple injection points as discussed in Section 3.3, but they can be modeled to the main wire, as o directly determines whether an attack happens or not. Therefore, the injected signal is superimposed onto o, described in Equation 1, making the modified differential amplifier output o':

$$\begin{aligned} o' &= T(s) + G\left(\frac{K-1}{K}T_C(s) + n\right) \\ &= \frac{G(K-1)}{K}\left(\frac{K}{G(K-1)}T(s) + T_C(s)\right) + Gn \end{aligned}$$

If the attacker wants to avoid detection o' must be zero (technically just less than ϵ, but basically zero). That means that the value in the parentheses must be zero, which in turn requires the following equation holds:

$$\frac{K}{G(K-1)}T(s) = -T_C(s) \tag{4}$$

The negative sign in Equation 4 implies that $T(s)$ and $T_C(s)$ must be 180 degrees out of phase, and this requires that the physical distance between the two corresponding wires is exactly half of the wavelength of the attacking signal s [3]. This is already a good argument for why an attacker cannot inject a signal that affects the actuator, and simultaneously cancel it out in the detection system, since the frequency would have to be in the 10-100s of GHz to get a half wavelength short enough. Such a high-frequency signal is way above what affects most actuators.

However just to make the point extra clear, let's assume that the attacker could in fact send a signal with a high enough frequency to make this work. Even then, the constant $\frac{K}{G(K-1)}$ in Equation 4 means that the signal injected into the detection system, $T(s)$, must be 100s of times stronger than $T_C(s)$, the signal injected into the actuator control signal itself. However, given that the two wires are so close to each other, and that the smaller of the two needs more power injected into it, it is impossible to achieve such a $T(s)$ in practice. As a result Equation 4 can never hold in practice.

For those two reasons (phase difference and relative power), no adversarial signal s can ever prevent its own detection due to interference with the detection circuitry itself.

5 EXTENDED MAXIMUM DETECTABLE FREQUENCY

Our detection method relies on a differential amplifier to help detect injected signals. Like all electronic components, a differential amplifier is designed to work within a particular frequency range. When choosing parameters for the detection system, a suitable differential amplifier should be used, which covers the entire range where adversarial signals are likely to be able to affect the actuator. However, on rare occasions, e.g., for very high frequency applications or if cost is a significant concern, it might be difficult to get a differential amplifier that fully covers the desired frequency range. For such cases, we have come up with a method to extend the maximum detectable frequency f_{max} beyond the normal upper bound of the differential amplifier.

Many previous studies [14, 26, 35, 49] have shown that a differential amplifier will still respond beyond its normal operational band, although the response is entirely different from the normal amplification within its design parameters. As the frequency increases beyond f_{max}, the peak amplitude of the differential amplifier output starts to decline, as the gain plummets to almost zero [23, 28]. This is shown in Figure 4 where the dashed curve depicts the change of the peak amplitude.

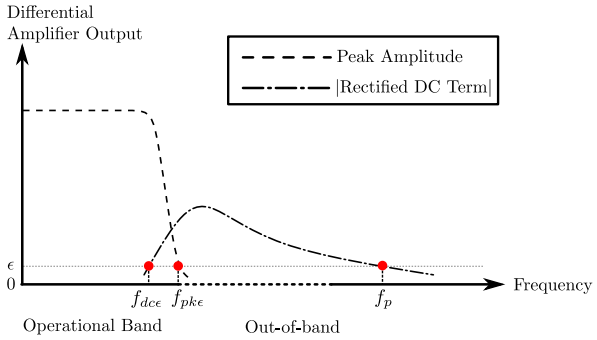

Figure 4: With constant attack power, the peak amplitude (dashed line) and the rectified DC term (dash-dotted line) of the differential amplifier output signal change along with the frequency. The maximum detectable frequency is extended to f_p.

Although the peak amplitude decreases to nothing, the output will gain a DC offset with respect to the normal ground state, shown in Figure 4 as the Rectified DC Term. This happens as the differential amplifier rectifies the high frequency signals [10, 13, 49]. The phenomenon is also known as radio-frequency (RF) rectification, and it is attributed to the nonlinear voltage-current characteristic of transistors that make up the differential amplifier [10]. Further increasing the frequency will eventually decrease the rectified DC term, which will ultimately become negligible when the frequency is high enough [14, 26, 35]. While this effect does eventually disappear, it allows us to extend the detection by hundreds or thousands of times higher than the upper bound of the operational band.

It is important to note that this phenomenon is not limited to a specific differential amplifier, but is true for many different designs, which has been experimentally verified in the literature [20, 49].

For our detection system to provide firm guarantees, it is essential to ensure no gap in the protected frequency band. Therefore we have to ensure that the DC offset rises enough to be detected before the normal peak amplitude of the differential amplifier goes to zero. In Figure 4, the frequency at which the magnitude of the rectified DC term exceeds ϵ is denoted as $f_{dc\epsilon}$, and the frequency at which the peak amplitude falls below ϵ is $f_{pk\epsilon}$. We show in Section 6 that we can easily achieve $f_{dc\epsilon} < f_{pk\epsilon}$ in practice.

6 IMPLEMENTATION

We implement our detection method on two practical and distinct actuator systems: a speaker system (in Section 6.2) and a motor control system (in Section 6.3). The objective of the implementation is to validate the feasibility of our detection method in practice. One of the reasons why we choose these two systems is that they are widely deployed in many critical applications: the speaker system can be found in applications in which sound information needs to be broadcast, such as mobile phones and car satellite navigation; the motor control system can be found in those that need to drive some mechanical structures, such as smart locks and insulin pumps. Implementing the detection method on these two systems also verifies its capabilities of handling different actuator systems

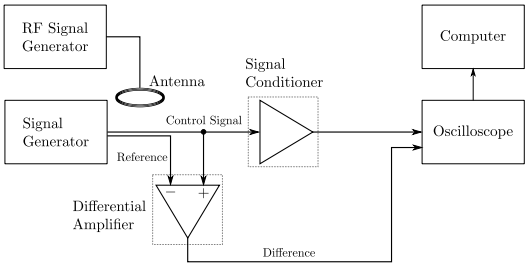

Figure 5: A setup of the actuator system. Devices in the dotted squares differ from system to system, and others are the same.

regardless of types of signals: sinusoidal signals (analog) are used in the speaker system, while pulses (digital) in the motor control system.

We first introduce how to build our own actuator systems on which we can quickly implement our detection method. Then, we show how to detect various attacking signals in each actuator system. We only demonstrate the control signal injection, as the drive signal injection is power-consuming and difficult to achieve with our equipment (please see detailed discussion in Section 7.4). Finally, a summary of the implementation of these two actuator systems is given in Section 6.4.

6.1 Setup

Based on the system model, we build a setup that can be easily configured into a speaker system or a motor control system, as shown in Figure 5. We use a signal generator to produce the control signal and the reference signal. The signal generator is functionally equivalent to the microcontroller. The benefit of using the signal generator is having easier control of signals regarding their frequencies, amplitudes, synchronization, etc.

The control signal is fed into a signal conditioner. The signal conditioner is different in these two systems: an audio power amplifier LM386 is used in the speaker system, and a brushed DC motor driver chip DRV8833 is used in the motor control system.

Regarding the actuator (either a loudspeaker or a motor), since its responses are deterministic and its input signal (i.e., the drive signal) sufficiently reflects the responses, we simply omit the actuator in the setup but use an oscilloscope to monitor and record the drive signal. An advantage of doing so is that different actuator systems can be quickly tested without extra work of using different methods to sense and process the actuator responses (e.g., a microphone to measure sound played by the speaker, or a hall-effect sensor to measure the speed of the motor). Moreover, a computer is used to process the data that is recorded by the oscilloscope. Note that the oscilloscope and the computer are used for the purpose of demonstrating the feasibility of the detection method, and are not used in the proposed applications.

Based on such an actuator system, we deploy our detection method to it. The control signal and the reference signal are fed into a differential amplifier, as shown in Figure 5. In the speaker system, we choose an AD623 with a gain of around 150 as the

differential amplifier because it is specifically designed to amplify small differences between its two inputs. As for the motor control system, a unity-gain differential amplifier AD629 is selected as the differential amplifier, as it can handle high-voltage inputs. The output of the differential amplifier is monitored and recorded by the oscilloscope, and the recorded data are sent to the computer for attack detection.

To achieve a large K, i.e., difference between the transfer functions of the control signal wire and the reference wire, we form a loop on the control signal wire to make it easier to pick up the attacking signal, and choose a short cable as the reference wire. Thus, the control signal wire is much more sensitive to the attacking signal than the reference wire. Note that it does not matter which wire is more sensitive because our detection method only requires the transfer functions to be different. Moreover, to guarantee that the control signal and the reference signal arrive at the differential amplifier at the same time, the tapping point is carefully chosen to ensure that the paths that feed these two signals into the differential amplifier have the same length.

Our setup is extremely flexible and allows us to easily experiment with different actuator types without having to build dedicated systems for each one. Despite being a lab setup we believe that our results accurately reflect the response of real commercial products.

6.2 Speaker System

In a speaker system, an audio signal is amplified and then broadcast. The objective of the attack is maliciously manipulating the waveform of the audio signal, and in the extreme case, can lead to the speaker system broadcasting any messages the attacker wishes.

6.2.1 Determining Threshold. The differential amplifier output is measured when no attack happens. The measurements show that the differential amplifier output signal amplitude is always below 2.4 mV. Since this value already includes all noise sources in our experimental environment, it is chosen as the threshold. The benefit of choosing this value as the threshold is that, on the one hand, it significantly reduces the possibility of which the noise accidentally triggers the detection; the false-positive rate remains at 0% as calculated from the measurements. On the other hand, this threshold is small enough to guarantee that the weakest attack that effectively impacts the actuator system is successfully detected, and this will be shown and explained in the experimental results as follows.

6.2.2 Direct Power Injection Attacks. The normal operational band of an audio amplifier is below the megahertz level, and low-frequency attacking signals are needed for in-band attacks. Due to the practical difficulty of injecting low-frequency attacking signals into the circuit wirelessly, we first demonstrate that the detection method can handle the in-band attacks using direct power injection (DPI) [16]. Note that in the following sections (Section 6.2.3 and Section 6.3.2), the attacking signals are injected wirelessly.

In order to show that any malicious frequency can be injected into the audio signal, the attack frequency is swept from 1 Hz to 10 MHz, and the peak-to-peak voltage of the attacking signal is from 10 mV to 100 mV. The reason why the highest attack frequency is set to 10 MHz, which is beyond the operational band of the audio amplifier, is to verify that no gap (as described in Section 5) exists in

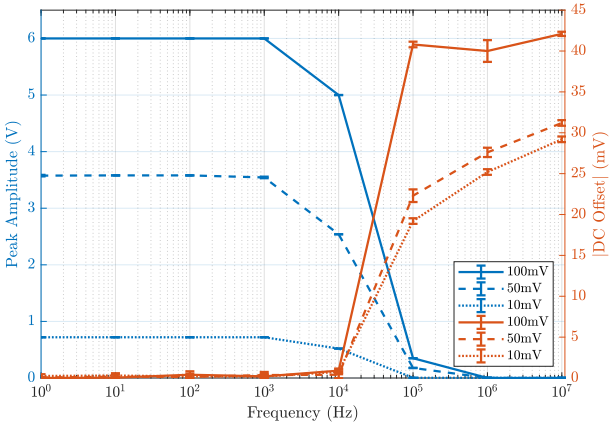

Figure 6: The peak amplitude (left y-axis) of the differential amplifier output drops to zero when the frequency of the attacking signal is far beyond the operational band of the audio amplifier; the DC offset (right y-axis) rises while increasing the frequency of the attacking signal. The peak-to-peak voltage of the attacking signal is from 10 mV **to** 100 mV

the frequency band. The reason why 10 mV is chosen as the weakest peak-to-peak amplitude of the attacking signal, is that the malicious change caused by an attack at this voltage is already around 49 dB weaker than the audio signal. Weaker attacking signals have little to no impact on the speaker system.

To demonstrate the impact of the attack on the differential amplifier in detail, we show both the peak amplitude and the DC offset in Figure 6. Each point in the figure represents the averaged peak amplitude or the averaged DC offset with a standard deviation. The first observation of the experimental results is related to the attack power: the peak amplitude and the DC offset increase (decrease) while the attack power increases (decreases). Concerning the attack frequency, when it is lower than 1 kHz, the peak amplitude is significantly larger than the threshold, which reveals the existence of the attacking signal. When the frequency is between 1 kHz and 1 MHz, the peak amplitude plummets, but it is still above the threshold; meanwhile, the DC offset rises above the threshold. When the frequency of the attacking signal reaches 1 MHz and beyond, the DC offset is well above the threshold, indicating the existence of the attack.

The experimental results validate the capabilities of the differential amplifier to detect attacks in the entire frequency range from DC to 10 MHz. We perform this experiment 240 times and all (240 out of 240) attacking signals are detected, making the true-positive rate is 100%. This shows that even for practical systems, the detection method provides strong protection against both in-band and out-of-band attacks.

6.2.3 Wireless Attacks. To test high-frequency attacks in a more realistic setting, we modulate a high frequency carrier with an audio signal and inject it wirelessly into the control and reference wires. An RF signal generator is used to produce the attacking signals, and they are radiated by a coil antenna, as shown in Figure 5. The antenna is placed around 2 cm above the control signal wire for

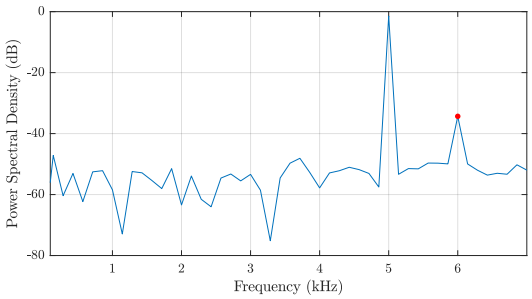

Figure 7: A 6 kHz malicious signal is successfully injected into the 5 kHz audio signal. The 6 kHz spike is highlighted by a red point in the frequency domain of the audio amplifier output. The power ratio between the 6 kHz frequency component and the 5 kHz is around −30 dB.

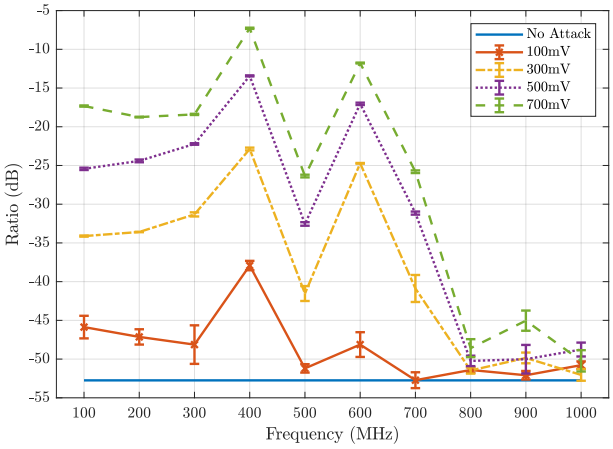

Figure 8: The power ratio between the malicious signal and the legitimate signal gradually decreases while increasing the frequency of the attacking signal. The peak-to-peak voltage of the attacking signal is from 100 mV to 700 mV.

the best possible energy transfer. That way we can use less power to achieve the wireless attack in our experiments. If an attacker is further away from the victim system, she needs more powerful attacking signals to achieve the attack.

To present a concrete attack, we choose to inject a 6 kHz malicious frequency into a 5 kHz audio signal. In Figure 7, an attack result is shown: in the frequency domain of the audio amplifier output, besides the legitimate 5 kHz frequency component, a malicious spike can be observed at 6 kHz. In order to quantify the impact of the attack, the power ratio between the malicious frequency component and the legitimate frequency component is measured, which can be expressed as the following equation:

$$impact = 10 \times \log_{10}\left(\frac{P_{malicious}}{P_{legitimate}}\right)$$

where P represents the power. The bigger the ratio is, the stronger the injected signal is, and the larger the impact of the attack is. When no attacking signal is presented, our measurements show that the *impact* remains at around −52.7 dB.

Different attacking signals are generated to test the performance of the detection method: the peak-to-peak voltage of the attacking signal is changed from 100 mV to 700 mV, and the carrier frequency of the RF signal is changed from 100 MHz to 1000 MHz. The impact of the attacks are numerically represented in Figure 8. When the attack frequency reaches 800 MHz, the *impact* is close to −52.7 dB, which means that attacks beyond this frequency will have little practical significance. We did conduct experiments beyond 1000 MHz, but the impact of the attacking signal beyond 1000 MHz are smaller, and hence we only focus on the frequency range within 1000 MHz.

Regarding the attack detection, the peak amplitudes of all measurements of the differential amplifier output are below the threshold. This is because the frequency of the attacking signal is already far beyond the operational band of the differential amplifier, as explained in Section 5. However, as shown in Figure 9, the DC offset of the differential amplifier output is well above the threshold throughout the range for all attacker signals other than 100 mV, indicating the existence of the attack. We see that the DC offset increases when the attack power is increased, so for attacking signals with peak-to-peak voltages of 300 mV, 500 mV, and 700 mV, the DC offsets are always above the threshold (solid blue line) regardless of the frequency. When the attacking signal is 100 mV, a few attacks fall below the threshold when the carrier frequencies reach 800 MHz and 900 MHz. Referring back to the impact of these two attacking signals in Figure 8, the ratios indicate that the impacts are so tiny that they are unlikely to have any significance for a practical system. Since our detection method successfully detects 389 out of 400 attacking signals, the true-positive rate is 97.25%.

In Figure 9, the curves of DC offsets vary up and down along the attack frequency. This is because the attacking signal is injected wirelessly instead of through DPI. The transfer function of the wire accounts for the ups and downs of the curves: the attacking signal is efficiently injected into the wire at specific frequencies where local maximum values of the DC offset reaches, but less efficient at other frequencies.

The experiment results show that the frequency range covered by the differential amplifier is easily large enough to protect the frequency band that the speaker system is vulnerable to. Our detection method shows the feasibility of detecting the attacking signals with frequencies from DC to far beyond the speaker system's operational band. Moreover, given the wireless injections, our detection method demonstrates its capabilities of handling real attack scenarios. We present concrete attacking signals that can precisely manipulate the audio frequencies, but it does not mean that our detection method can only handle these specific attacking signals. Any attacks that cause voltage changes of the differential amplifier output signal beyond the pre-determined threshold can be spotted immediately.

6.3 Motor Control System

In the motor control system, a pulse signal is used to control the rotating speed of the motor. The duty cycle of the pulse signal describes the amount of time that the signal is at the high-voltage level as a percentage of the total time of a cycle. The larger the duty cycle is, the faster the motor's rotation speed is. As mentioned in the setup, a motor driver is used as a signal conditioner to amplify the

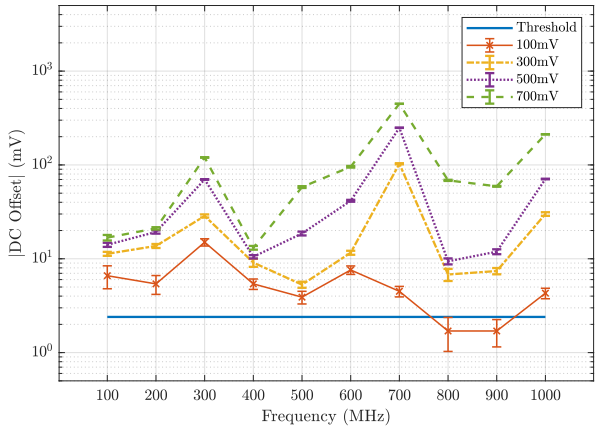

Figure 9: The DC offset of the differential amplifier output varies while changing the voltage level and the frequency of the attacking signal. The peak-to-peak voltage of the attacking signal is from 100 mV to 700 mV.

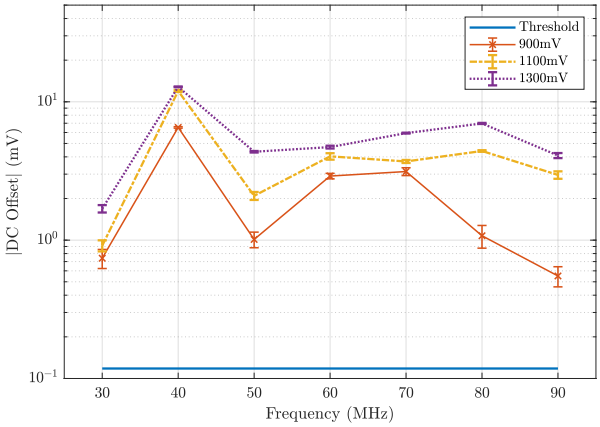

Figure 11: When an attack happens, the DC offset of the differential amplifier output is always above the threshold, implying detecting the attack.

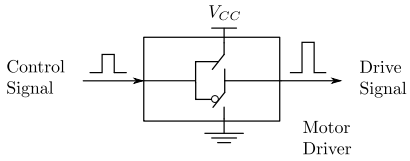

Figure 10: A motor driver is used to amplify a control signal to drive a motor.

control signal into a powerful drive signal to energize the motor. The motor driver is made of transistors, and for simplicity, as shown in Figure 10, they can be regarded as two switches that are connected in series and are controlled by the pulse signal. Since these two switches work in opposite ways, the output signal toggles between V_{CC} and the ground in the same pattern as the input signal. The attacker's objective is to manipulate the duty cycle and impact the functionality of the motor.

6.3.1 Determining Threshold. When no attack presents, the differential amplifier output signal is recorded, and the threshold is 0.17 mV. This threshold value is chosen as it makes the false-positive rate to its minimum (0%) in our experimental environment; also, this threshold is sufficiently large to spot the weakest attacks, as shown as follows.

6.3.2 Detection of Attacks. Since the differential amplifier is specifically designed to handle the input difference in its operational band, it is not difficult to detect the in-band attacks. We do not repeat the in-band attacks here but focus on the out-of-band attacks. Note that the out-of-band attacks are realized wirelessly.

In the experiments, the frequency of the attacking signal ranges from 30 MHz to 90 MHz, and the peak-to-peak voltage ranges from 900 mV to 1300 mV. The reason why the frequency of the attacking signal is below 90 MHz is that beyond this frequency, the motor driver never responds to the attacking signal, even though the peak-to-peak voltage of the attacking signal reaches its upper limit in the signal generator. The reason why the peak-to-peak voltage of the attacking signal is above 900 mV is that, below this voltage level, the attacking signal is too weak to affect the motor driver. In our experiment, the RF signals can cause the motor driver to output a low voltage level when a high voltage level should be outputted, thus reducing the duty cycle of the pulses. We can precisely control when to start and stop radiating the attacking signals, and hence, the duty cycle of the control signal can be precisely manipulated, further controlling the motor speed. Note that using other types of attacking signals can also increase the duty cycle [38]; however, the purpose of the experiment focuses on attack detection, and we do not further show and discuss how to control the motor speed.

Regarding the attack detection, both the peak amplitude and the DC offset of the differential amplifier output signal are checked. Under these out-of-band attacks, the peak amplitude is always below the threshold. However, as shown in Figure 11, the DC offset is always above the threshold, indicating an attack. All (210 out of 210) DC offsets are above the threshold, indicating that all attacking signals are detected. Therefore, the true-positive rate is 100%.

6.4 Summary of Implementation

The implementation of our detection method on the speaker system and the motor control system show the generality of our detection method regardless of the type of signal. The deployments also demonstrate the simplicity of implementing the detection method in practice. The high true-positive rates and low false-positive rates in the speaker system and the motor control system show the robustness of the detection method on different actuator systems.

7 DISCUSSION

In this section, we discuss different detection strategies and how an adaptive threshold can handle varying environmental noise. Moreover, we discuss the difficulties of canceling out an injected signal in circuits, as well as the difficulties of the drive signal injection.

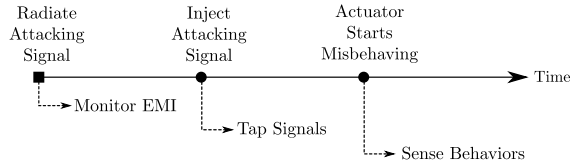

Figure 12: The timeline shows that an attack starts from radiating an attacking signal to actuator misbehaving. At different moments, the attack can be detected by different ways.

7.1 Different Detection Strategies

An electromagnetic signal injection attack has several distinct phases, each of which gives rise to different detection strategies. In the attack timeline shown in Figure 12, three moments are highlighted: the first moment is when the attacking signal is radiated; the second is when the attacking signal is injected into a wire in the target system; the third is when the actuator starts deviating from its intended activity. Each of the three moments marks the start of a new phase of the attack. The three phases should not be thought of as sequential since an attack of any meaningful duration will be in all three phases at once, but should rather be thought of as three opportunities to detect the attack.

In the first phase, when attacking signals are being radiated, the electromagnetic radiation can be detected in the environment using an antenna. Thus, the detection strategy is monitoring the environmental electromagnetic power level: if the power level is above a pre-defined threshold, or maybe outside a known noise profile, the attack is detected. This strategy has the potential to detect an attack early; however, it requires a monitoring device that can reliably detect adversarial interference at the frequency that would harm the target system, which is not always easy to achieve. Examples of this detection strategy are some of the anomaly detectors that will be introduced in Section 8.2.1.

In the second phase, when the attacking signals are successfully injected into the actuator system, the signals in the actuator system wires are changed. Since the attacking signals are not supposed to exist in the actuator system, these changes are a reliable indicator of an attack, if they can be measured. Our detection method uses the second form of detection, i.e., we detect signals that are successfully injected into the actuator system.

In the final phase, the actions of the actuator will deviate from what the system expects, assuming the attack is powerful enough to result in a measurable change. If the system can detect this behavior change, this can be used to detect the attack. This might be an attractive detection strategy since no effort is wasted on attacks that do not have a measurable effect on the target actuator. However, detecting such attacks typically requires extra sensors. By the nature of the detection method, it will only detect attacks after they have already affected the system. An example of this detection method is Muniraj and Farhood's work [30] that will be introduced in Section 8.2.3.

7.2 Adaptive Threshold

In our implementation of the detection method, we find a proper threshold by experiment and keep it constant while testing the performance of our detection method. The advantage of using a constant threshold is that once a proper threshold is found and determined, it is efficacious forever, and the designer never has to adjust it again. However, in some cases, the environmental noise varies significantly and complicatedly over time; for example, such noise may originate from the radiation of other complex circuits. To provide the actuator system with more flexibility, the designer can program the actuator system to adjust its threshold adaptively when necessary. Imagine a simple case: during the daytime, the noise is intense because of human activities (e.g., wireless communications, transportations), but at midnight when people sleep, the noise becomes relatively weak. During the daytime, the threshold can be slightly increased to allow more noise, and as such, it can avoid the noise frequently triggering the detection. At midnight, to restore the detection method to be more sensitive to attacks, the designer can program the actuator system to lower down the threshold.

Other environmental changes may also impact the detection circuits, and they can also be handled by the adaptive threshold. For example, in some harsh environments where the environmental temperature varies significantly, a temperature change will affect the resistance of metal conductors, further the voltages. Since the resistance of the two wires may not change consistently, the voltage difference between them increases when the temperature varies. To reduce the impact caused by the temperature, basically, materials with a small temperature coefficient should be chosen; for instance, regarding copper with a temperature coefficient of around 0.004, a change of 100 degree Celsius only lead to a change of 0.4 ohms in resistance. Furthermore, to cover the extra voltage difference that is caused by the temperature variation, the threshold can be adaptively set to a higher level in such an environment.

No matter how the designer adjusts the threshold adaptively, it is still essential to guarantee that the detection method meets the requirements as mentioned in previous sections: first, no noise triggers the detection accidentally; second, no attack that effectively impacts the actuator is missed.

7.3 Difficulty of Canceling Attacking Signals

An idea of mitigating the influence caused by attacks is generating an "anti-attack" signal to cancel out the attacking signal. The anti-attack signal and the attacking signal have the same frequency and amplitude, but they are 180 degrees out of phase. When the anti-attack signal and the attacking signal meet, they destruct each other. This idea is similar to the sound noise cancellation technology that is used in headphones. However, it is hard to realize such a cancellation regarding the electromagnetic interference. In the air, an electromagnetic signal propagates around the light speed; in the circuit, the speed halves. In addition, it takes time for the actuator system to capture the attacking signal and then generate the anti-attack signal for the cancellation. This means that the anti-attack signal always lags behind the attacking signal. It is difficult to synchronize the anti-attack signal with the attacking signal unless the microcontroller can predict the attacking signal.

7.4 Difficulty of Drive Signal Injection

As mentioned previously, compared with the control signal injections, a drive signal injection may require much more power if the actuator is power-consuming. We estimate the power of a drive signal injection as follows. According to datasheets of an off-the-shelf motor, it needs a drive signal that is around 4.5 W; as for a microcontroller, such as an Arduino Uno microcontroller, it can output a control signal that is only 0.1 W. For simplicity, we suppose that the attenuation on attacking signals is the same in those two injections. Then, the attacker needs to radiate at least $\frac{4.5\,\mathrm{W}}{0.1\,\mathrm{W}} = 45$ times more power to realize the drive signal injection than the control signal injection. This result implies that it is much more difficult and costly to conduct the drive signal injection than the control signal injection in practice.

Another evidence to show that the drive signal injection is hard to achieve is to regard the injection as wireless power transmission [41]. In wireless power transmission techniques, scientists specifically designed both antennas of the transmitter and the receiver to achieve the power transmission. Given the wire that works as a low-gain antenna in the actuator system, delivering enough power into the drive signal wire can be much more challenging.

8 RELATED WORK

Many countermeasures against electromagnetic signal injection attacks have been proposed and developed; however, it needs to be noted that protecting sensors has been much more extensively studied than actuators. The countermeasures can be categorized into two types: one is attenuation that aims to reduce attack impacts, and the other is detecting the existence of attacks.

8.1 Attenuation

Wrapping components with proper RF shielding materials is a common method to attenuate attacking signals [17, 22, 25, 29, 31, 33, 38, 40, 44, 46, 51]. However, the shielding materials provide finite attenuation [43], and a powerful attacker may still breach the protection by increasing her attack power. Although adding thicker shielding materials can increase the attenuation level, it will still challenge the weight and the size of the devices, especially for applications such as implantable medical devices and aviation. In addition to shielding materials, regarding traces in a printed circuit board (PCB), researchers suggested that via-fenced striplines can also eliminate attacking signals by a finite amount [8, 9].

Filtering is another prevalent solution to mitigate attacking signals. Low-pass filters can significantly attenuate out-of-band attacking signals [17, 25, 31, 38, 44, 51]. However, in-band attacking signals can still pass through the low-pass filters. Researchers also pointed out that the parasitics in surface mount components can convert the low-pass filter into a band-stop filter, which allows out-of-band attacking signals to pass [19]. Besides, Kune et al. [25] proposed to deploy an adaptive filtering mechanism [32] that makes use of knowledge about ambient electromagnetic emissions to attenuate the interference in sensor measurements. Crovetti and Musolino [6] also proposed a digital way to suppress the EMI-induced errors in the sensor measurements. However, it is challenging to have such digital methods for the actuator because it has no computational capabilities. Furthermore, Kune et al. [25] also recommended using differential rather than single-end comparator to attenuate the attacking signals in a finite frequency band, thereby raising the bar for attackers.

Note that our method can be used alongside these attenuation methods to provide additional protection.

8.2 Detection

8.2.1 Anomaly Detection. One detection method is to add a specific channel to monitor whether abnormal electromagnetic signals or activities appear. Note that although some of the following approaches are initially designed for sensors, similar ideas possibly also work in actuator systems. Researchers developed standalone detection systems that capture electromagnetic waves by dedicated antennas and then use intricate circuits to process the captured signals for detection [1, 2, 7]. Kune et al. [25] investigated using extra antennas or conductors to capture and measure attacking signals for detection, and the measurements can be then used by their adaptive filtering mechanism as mentioned previously. In a similar vein, Tu et al. [45] proposed adding a dummy sensor for detection and correction. In another work, Tu et al. [44] proposed leveraging the superheterodyne technique to create an anomaly detector to check whether sensor measurements carry malicious frequency components. Note that these approaches count on the knowledge about the waveforms of the attacking signals, which are usually high-frequency (e.g., MHz or GHz). Thus, they require electronic components (e.g., high-speed ADCs) that can properly handle high-frequency signals, as well as extra computing resources to process the captured signals for detection purposes, implying significant implementation overheads regarding both hardware and software.

As a comparison, our approach counts on the signal strength difference between the primary and the reference signals to detect attacks, rather than waveforms. This makes our approach gain advantages over the other approaches: first, a simple detection circuit made of differential amplifiers is used to catch the difference, and such a detection circuit has fewer hardware overheads; second, an interrupt pin of the microcontroller is configured to handle the output of the detection circuit to determine whether an attack happens, and it needs fewer computing resources. Besides, our approach does not require any RF interface to capture the attacking signals, thus avoiding the troubles of crafting the dedicated RF interfaces, as well as preventing extra attack power from entering the victim devices and causing other unwanted influence.

8.2.2 Detection Methods for Sensor Systems. Especially for sensor systems, Zhang and Rasmussen [51] proposed a generalized detection method that selectively turns off the sensor in a secret way to observe whether attacks alter the sensor measurements. Shoukry et al. [42] proposed similar detection methods, but they were designed for specific types of sensors, as well as requiring significant computational overheads. Succeeding studies [24, 36] further adapted these detection methods to more practical applications. Fang et al. [12] proposed adding unique noise (fingerprints) to sensor measurements and using machine learning techniques to detect the attacks.

In addition, in specific devices such as cardiac implantable electrical devices (CIED) [25] and smartphones [46], researchers utilized

users' reactions or behaviors while using these devices to identify the existence of attacks on the sensors. Several works mentioned that multiple built-in sensors of a device can react to variations of the electromagnetic environment, and the characteristics can be exploited to detect abnormal electromagnetic activities [21, 22]. Such a detection approach is also known as sensor fusion, which has been widely studied to detect signal injections that use other types of attacking signals such as ultrasonics and lasers [16, 50].

These detection methods work well for the sensors because the computational capabilities of the receiver (microcontroller) make the authentication possible. However, it is not easy to apply similar ideas to the actuator systems because the receiver (actuator) lacks computational capabilities to authenticate its input signals.

8.2.3 Detection Methods for Actuator Systems. Reliable sensor measurements can be used to indicate whether actuators are under attack. In unmanned aircraft systems, Muniraj and Farhood [30] proposed to artificially cause minor disturbances to the actuators at a random time and use sensors to capture the disturbances; unexpected disturbances imply attacks. However, this method trades off the stability of the whole system against its security. The same authors proposed another detection method that casts the actuator attack detection problem as an unknown input estimation problem and uses a two-stage extended Kalman filter to estimate actuator attacks from sensor measurements, requiring additional computational power. In addition to the two detection methods, the authors also proposed a method that adds randomness to control signals to improve the resilience of the actuator against malicious attacks.

Our approach outstrips these detection methods in terms of these three aspects. First, they require a complex model of the specific actuator system, which makes it difficult to be applied to other applications, whereas our approach is generalized for different actuator systems. Second, they need extra computing resources to run the detection algorithms, but we can use the interrupt mechanism of the microcontroller for detection, which is more efficient. Third, their detection methods always spot attacks after the actuator misbehaves; however, our approach detects the attacks earlier, thus possibly allowing the actuator systems to take proper measures to stop/mitigate the attacks.

9 CONCLUSION

In this paper, we have proposed a novel detection method that can detect electromagnetic signal injection attacks on actuator systems. This class of systems previously had to rely on physical security measures and signal decay, and had no meaningful security guarantees against a determined adversary. Our detection system fills this critical gap and provides strong detection guarantees to any actuator system. The core idea of our detection method is straightforwad: any diffenrence caused by external attacks between two identical signals (the primary signal and the reference signal) indicates the attacks. Our detection method provides provable guarantees against attacks, and can be tuned to any attack power and any amount of environmental noise. We have shown that our detection method provides the actuator system with a strong security guarantee, and an attacker who attempts to effectively manipulate the actuator system will always be detected by our detection method. Despite this, our detection method requires only a few cheap off-the-shelf electronic components and does not add any significant weight to the system it protects. This is important in many contexts, such as aviation and implantable medical devices. Moreover, the implementation of the detection method on a speaker system and a motor control system proves its generality for different actuator systems, as well as the effectiveness and the robustness in a practical setting.

REFERENCES

[1] Christian Adami, Christian Braun, Peter Clemens, M Joester, S Ruge, M Suhrke, HU Schmidt, and HJ Taenzer. 2014. HPM Detector System with Frequency Identification. In *2014 International Symposium on Electromagnetic Compatibility (EMC Europe)*. IEEE, 140–145.

[2] Christian Adami, Christian Braun, Peter Clemens, Michael Suhrke, HU Schmidt, and Achim Taenzer. 2011. HPM Detection System for Mobile and Stationary Use. In *EMC Europe 2011 York*. IEEE, 1–6.

[3] Constantine A Balanis. 2016. *Antenna Theory: Analysis and Design*. John Wiley & Sons, 145–176.

[4] Calogero Bona and Franco Fiori. 2009. EMIs-inducted Failures in MOS Power Transistors. In *2009 International Conference on Electromagnetics in Advanced Applications*. 564–567. https://doi.org/10.1109/ICEAA.2009.5297367

[5] Calogero Bona and Franco Fiori. 2010. A New Filtering Technique that Makes Power Transistors Immune to EMI. *IEEE Transactions on Power Electronics* 26, 10 (2010), 2946–2955.

[6] Paolo Crovetti and Francesco Musolino. 2021. Digital Suppression of EMI-Induced Errors in a Baseband Acquisition Front-End including Off-the-Shelf, EMI-Sensitive Operational Amplifiers. *Electronics* 10, 17 (2021), 2096.

[7] JF Dawson, ID Flintoft, P Kortoci, Linda Dawson, AC Marvin, MP Robinson, Mirjana Stojilovic, Marcos Rubinstein, Benjamin Menssen, Heyno Garbe, et al. 2014. A Cost-efficient System for Detecting An Intentional Electromagnetic Interference (IEMI) attack. In *2014 International Symposium on Electromagnetic Compatibility*. IEEE, 1252–1256.

[8] Gokcen Y Dayanikli. 2021. *Electromagnetic Interference Attacks on Cyber-Physical Systems: Theory, Demonstration, and Defense*. Ph. D. Dissertation. Virginia Tech.

[9] Gokcen Y Dayanikli, Rees Hatch, Ryan M Gerdes, Hongjie Wang, and Regan Zane. 2020. Electromagnetic Sensor and Actuator Attacks on Power Converters for Electric Vehicles. In *IEEE Workshop on the Internet of Safe Things*. IEEE.

[10] Analog Devices. 2009. RFI Rectification Concepts. https://www.analog.com/media/en/training-seminars/tutorials/MT-096.pdf.

[11] J Lopes Esteves and Chaouki Kasmi. 2018. Remote and silent voice command injection on a smartphone through conducted IEMI: Threats of smart IEMI for information security. *Wireless Security Lab, French Network and Information Security Agency (ANSSI), Tech. Rep* (2018).

[12] Kai Fang, Tingting Wang, Xiaochen Yuan, Chunyu Miao, Yuanyuan Pan, and Jianqing Li. 2022. Detection of Weak Electromagnetic Interference Attacks Based on Fingerprint in IIoT Systems. *Future Generation Computer Systems* 126 (2022), 295–304.

[13] Franco Fiori. 2014. Susceptibility of Smart Power ICs to Radio Frequency Interference. *IEEE Transactions on Power Electronics* 29, 6 (2014), 2787–2797.

[14] Marle L Forcier and Robert E Richardson. 1979. Microwave-rectification RFI Response in Field-effect Transistors. *IEEE Transactions on Electromagnetic Compatibility* (1979), 312–315.

[15] Harald T Friis. 1946. A Note on A Simple Transmission Formula. *Proceedings of the IRE* 34, 5 (1946), 254–256.

[16] Ilias Giechaskiel and Kasper Rasmussen. 2019. Taxonomy and Challenges of Out-of-band Signal Injection Attacks and Defenses. *IEEE Communications Surveys & Tutorials* 22, 1 (2019), 645–670.

[17] Ilias Giechaskiel, Youqian Zhang, and Kasper Rasmussen. 2019. A Framework for Evaluating Security in the Presence of Signal Injection Attacks. In *European Symposium on Research in Computer Security*. Springer, 512–532.

[18] Jasper P Goedbloed. 1992. *Electromagnetic Compatibility*. Prentice-Hall.

[19] Ryan Hurley. 2007. *Design Considerations for ESD/EMI Filters: II Low Pass Filters for Audio Filter Applications*. ON Semiconductor.

[20] Texas Instruments. 2013. AN-1698 A Specification for EMI Hardened Operational Amplifiers. https://www.ti.com/lit/an/snoa497b/snoa497b.pdf.

[21] Chaouki Kasmi and Jose Lopes-Esteves. 2015. Automated Analysis of the Effects Induced by Radio-frequency Pulses on Embedded Systems for EMC Functional Safety. In *2015 1st URSI Atlantic Radio Science Conference (URSI AT-RASC)*. IEEE, 1–1.

[22] Chaouki Kasmi and Jose Lopes-Esteves. 2015. IEMI Threats for Information Security: Remote Command Injection on Modern Smartphones. *IEEE Transactions on Electromagnetic Compatibility* 57, 6 (2015), 1752–1755.

[23] Charles Kitchin and Lew Counts. 2006. *A Designer's Guide to Instrumentation Amplifiers*. Analog Devices Norwood, MA.

[24] Sebastian Köhler, Richard Baker, and Ivan Martinovic. 2022. Signal Injection Attacks against CCD Image Sensors. In *ACM ASIA Conference on Computer and*

Communications Security. Association for Computer Machinery.

[25] Denis Foo Kune, John Backes, Shane S Clark, Daniel Kramer, Matthew Reynolds, Kevin Fu, Yongdae Kim, and Wenyuan Xu. 2013. Ghost Talk: Mitigating EMI Signal Injection Attacks against Analog Sensors. In *IEEE Symposium on Security and Privacy*. IEEE, 145–159.

[26] Curtis E Larson and James M Roe. 1979. A Modified Ebers-Moll Transistor Model for RF-interference Analysis. *IEEE Transactions on Electromagnetic Compatibility* (1979), 283–290.

[27] Marco Leone and Hermann L Singer. 1999. On the Coupling of An External Electromagnetic Field to A Printed Circuit Board Trace. *IEEE Transactions on Electromagnetic Compatibility* 41, 4 (1999), 418–424.

[28] Ron Mancini. 2003. *Op Amps for Everyone: Design Reference*. Newnes, 189–191.

[29] A Theodore Markettos and Simon W Moore. 2009. The Frequency Injection Attack on Ring-oscillator-based True Random Number Generators. In *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 317–331.

[30] Devaprakash Muniraj and Mazen Farhood. 2019. Detection and Mitigation of Actuator Attacks on Small Unmanned Aircraft Systems. *Control Engineering Practice* 83 (2019), 188–202.

[31] Saki Osuka, Daisuke Fujimoto, Yu-ichi Hayashi, Naofumi Homma, Arthur Beckers, Josep Balasch, Benedikt Gierlichs, and Ingrid Verbauwhede. 2018. EM Information Security Threats against RO-based TRNGs: The Frequency Injection Attack Based on IEMI and EM Information Leakage. *IEEE Transactions on Electromagnetic Compatibility* 61, 4 (2018), 1122–1128.

[32] John G Proakis. 2001. *Digital signal processing: principles algorithms and applications*. Pearson Education India, 500–519.

[33] Kasper Bonne Rasmussen, Claude Castelluccia, Thomas S Heydt-Benjamin, and Srdjan Capkun. 2009. Proximity-based Access Control for Implantable Medical Devices. In *Proceedings of the 16th ACM conference on Computer and communications security*. 410–419.

[34] Behzad Razavi. 2005. *Design of analog CMOS integrated circuits*. McGraw-Hill Education. 100 – 126 pages.

[35] Robert E Richardson. 1979. Modeling of Low-level Rectification RFI in Bipolar Circuitry. *IEEE Transactions on electromagnetic Compatibility* (1979), 307–311.

[36] Henri Ruotsalainen, Albert Treytl, and Thilo Sauter. 2021. Watermarking Based Sensor Attack Detection in Home Automation Systems. In *2021 26th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*. IEEE, 1–8.

[37] Jayaprakash Selvaraj. 2018. *Intentional Electromagnetic Interference Attack on Sensors and Actuators*. Ph. D. Dissertation. Iowa State University.

[38] Jayaprakash Selvaraj, Gökçen Y Dayanıklı, Neelam Prabhu Gaunkar, David Ware, Ryan M Gerdes, Mani Mina, et al. 2018. Electromagnetic Induction Attacks Against Embedded Systems. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*. ACM, 499–510.

[39] Haoqi Shan, Boyi Zhang, Zihao Zhan, Dean Sullivan, Shuo Wang, and Yier Jin. 2022. Invisible Finger: Practical Electromagnetic Interference Attack on Touchscreen-based Electronic Devices. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 1548–1548. https://doi.org/10.1109/SP46214.2022.00119

[40] Hocheol Shin, Yunmok Son, Youngseok Park, Yujin Kwon, and Yongdae Kim. 2016. Sampling Race: Bypassing Timing-based Analog Active Sensor Spoofing Detection on Analog-digital Systems. In *10th {USENIX} Workshop on Offensive Technologies ({WOOT} 16)*.

[41] Naoki Shinohara. 2014. *Wireless Power Transfer via Radiowaves*. Wiley Online Library.

[42] Yasser Shoukry, Paul Martin, Yair Yona, Suhas Diggavi, and Mani Srivastava. 2015. PyCRA: Physical Challenge-response Authentication for Active Sensors under Spoofing Attacks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 1004–1015.

[43] Frederick M Tesche, Michel Ianoz, and Torbjörn Karlsson. 1996. *EMC Analysis Methods and Computational Models*. John Wiley & Sons.

[44] Yazhou Tu, Sara Rampazzi, Bin Hao, Angel Rodriguez, Kevin Fu, and Xiali Hei. 2019. Trick or heat? Manipulating Critical Temperature-based Control Systems Using Rectification Attacks. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 2301–2315.

[45] Yazhou Tu, Vijay Srinivas Tida, Zhongqi Pan, and Xiali Hei. 2021. Transduction Shield: A Low-Complexity Method to Detect and Correct the Effects of EMI Injection Attacks on Sensors. In *ACM ASIA Conference on Computer and Communications Security*.

[46] Kai Wang, Richard Mitev, Chen Yan, Xiaoyu Ji, Ahmad-Reza Sadeghi, and Wenyuan Xu. 2022. GhostTouch: Targeted Attacks on Touchscreens without Physical Touch. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA. https://www.usenix.org/conference/usenixsecurity22/presentation/wang-kai

[47] Yuanda Wang, Hanqing Guo, and Qiben Yan. 2022. GhostTalk: Interactive Attack on Smartphone Voice System Through Power Line. In *Network and Distributed Systems Security (NDSS) Symposium*.

[48] David A Ware. 2017. *Effects of Intentional Electromagnetic Interference on Analog to Digital Converter Measurements of Sensor Outputs and General Purpose Input Output Pins*. Ph. D. Dissertation. Utah State University.

[49] Chunyu Wu, Guanghua Li, David J Pommerenke, Victor Khilkevich, and Gary Hess. 2018. Characterization of the RFI Rectification Behavior of Instrumentation Amplifiers. In *2018 IEEE Symposium on Electromagnetic Compatibility, Signal Integrity and Power Integrity (EMC, SI & PI)*. IEEE, 156–160.

[50] Chen Yan, Hocheol Shin, Connor Bolton, Wenyuan Xu, Yongdae Kim, and Kevin Fu. 2020. SoK: A Minimalist Approach to Formalizing Analog Sensor Security. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 233–248.

[51] Youqian Zhang and Kasper Rasmussen. 2020. Detection of Electromagnetic Interference Attacks on Sensor Systems. In *IEEE Symposium on Security and Privacy*. 203–216.