

# A Framework for Evaluating Security in the Presence of Signal Injection Attacks

Ilias Giechaskiel, Youqian Zhang, and Kasper B. Rasmussen

University of Oxford, Oxford, UK

{`ilias.giechaskiel,youqian.zhang,kasper.rasmussen`}@cs.ox.ac.uk

**Abstract.** Sensors are embedded in security-critical applications from medical devices to nuclear power plants, but their outputs can be spoofed through signals transmitted by attackers at a distance. To address the lack of a unifying framework for evaluating the effect of such transmissions, we introduce a system and threat model for signal injection attacks. We further define the concepts of existential, selective, and universal security, which address attacker goals from mere disruptions of the sensor readings to precise waveform injections. We finally introduce an algorithm which allows circuit designers to concretely calculate the security level of real systems, and we apply our definitions and algorithm in practice using measurements of injections against a smartphone microphone. Overall, our work highlights the importance of evaluating the susceptibility of systems against signal injection attacks, and introduces both the terminology and the methodology to do so.

## 1 Introduction

In our daily routine we interact with dozens of sensors: from motion detection in home security systems and tire pressure monitors in cars, to accelerometers in smartphones and heart rate monitors in smartwatches. The integrity of these sensor outputs is crucial, as many security-critical decisions are taken in response to the sensor values. However, specially-crafted adversarial signals can be used to remotely induce waveforms into the outputs of sensors, thereby attacking pacemakers [6], temperature sensors [2], smartphone microphones [5], and car-braking mechanisms [16]. These attacks cause a system to report values which do not match the true sensor measurements, and trick it into performing dangerous actions such as raising false alarms, or even delivering defibrillation shocks.

The root cause of these vulnerabilities lies in the unintentional side-effects of the physical components of a system. For example, the wires connecting sensors to microcontrollers behave like low-power, low-gain antennas, and can thus pick up high-frequency electromagnetic radiations. Although these radiations are considered “noise” from an electrical point of view, hardware imperfections in the subsequent parts of the circuit can transform attacker injections into meaningful waveforms. Specifically, these radiations are digitized along with the true sensor outputs, which represent a physical property as an analog electrical quantity. This digitization process is conducted by Analog-to-Digital Converters (ADCs),

which, when used outside of their intended range, can cause high-frequency signals to be interpreted as meaningful low-frequency signals.

Despite the potential that signal injection attacks have to break security guarantees, there is no unifying framework for evaluating the effect of such adversarial transmissions. Our work fills this gap through the following contributions:

1. We propose a system model which abstracts away from engineering concerns associated with remote transmissions, such as antenna design (Section 2).
2. We define security against adversarial signal injection attacks. Our definitions address effects ranging from mere disruptions of the sensor readings, to precise waveform injections of attacker-chosen values (Section 3).
3. We introduce an algorithm to calculate the security level of a system under our definitions and demonstrate it in practice by injecting “OK Google” commands into a smartphone (Section 4).
4. We discuss how our model can be used to inform circuit design choices, and how to interpret defense mechanisms and other types of signal injection attacks in its context (Section 5).

Overall, our work highlights the importance of testing systems against signal injection attacks, and proposes a methodology to test the security of real devices.

## 2 System and Adversary Model

Remote signal injection attacks pose new challenges from a threat-modeling perspective, since the electrical properties of systems suggest that adversaries cannot arbitrarily and precisely change any sensor reading. To create a threat model and define security in its context, we need to first abstract away from specific circuit designs and engineering concerns related to remote transmissions. To do so, we separate the behavior of a system into two different transfer functions. The first function describes circuit-specific behavior, including how adversarial signals enter the circuit (e.g., through PCB wires acting as antennas), while the second one is ADC-specific, and dictates how the signals which have made it into the circuit are digitized. We describe this model in greater detail in Section 2.1, taking a necessary detour into electrical engineering to show why our proposal makes for a good system model. We then explain some sources of measurement errors even in the absence of an adversary in Section 2.2 and finish by detailing the capabilities and limitations of the adversary in Section 2.3. Both sub-sections are crucial in motivating the security definitions of Section 3.

### 2.1 Circuit Model

Analog-to-Digital Converters (ADCs) are central in the digitization process of converting signals from the analog to the digital realm, and our circuit block diagram (Figure 1) reflects that. In the absence of an adversary, the ADC digitizes the sensor signal  $s(t)$  as well as the environmental noise  $n(t)$ , and transfers the

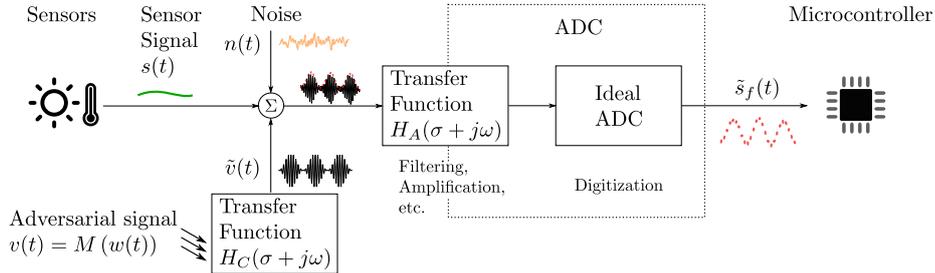


Fig. 1: System model: an adversarial signal  $v(t)$  enters the circuit and is transformed via  $H_C$ . It is digitized along with the sensor signal  $s(t)$  and the noise  $n(t)$  through an ADC-specific  $H_A$ . In successful attacks, the digitized signal will contain the demodulated version  $w(t)$  of the attacker signal  $v(t) = M(w(t))$ .

digital bits to a microcontroller. We model the ADC in two parts: an “ideal” ADC which simply digitizes the signal, and a transfer function  $H_A$ . This transfer function describes the internal behavior of the ADC, which includes effects such as filtering and amplification. The digitized version of the signal  $\tilde{s}_f(t)$  depends both on this transfer function, and the sampling frequency  $f$  of the ADC. An adversarial signal can enter the system (e.g., through the wires connecting the sensor to the ADC) and add to the sensor signal and the noise. This process can be described by a second, circuit-specific transfer function  $H_C$ , which transforms the adversarial signal  $v(t)$  into  $\tilde{v}(t)$ . Note that components such as external filters and amplifiers in the signal path between the point of injection and the ADC can be included in either  $H_A$  or  $H_C$ . We include them in  $H_A$  when they also affect the sensor signal  $s(t)$ , but in  $H_C$  when they are specific to the coupling effect.  $H_C$  and  $H_A$  are discussed in detail below.

**Circuit Transfer Function  $H_C$ .** To capture the response of the circuit to external signal injections, we introduce a transfer function  $H_C$ . This transfer function explains why the adversarial waveforms must be modulated, and why it is helpful to try and reduce the number of remote experiments to perform. For electromagnetic interference (EMI) attacks, the wires connecting the sensor to the ADC pick up signals by acting as (unintentional) low-power and low-gain antennas, which are resonant at specific frequencies related to the inverse of the wire length [7]. Non-resonant frequencies are attenuated more, so for a successful attack the adversary must transmit signals at frequencies with relatively low attenuation. For short wires, these frequencies are in the GHz range [7], so the low-frequency waveform  $w(t)$  that the adversary wants to inject into the output of the ADC  $\tilde{s}_f(t)$  may need to be modulated over a high-frequency carrier. We denote this modulated version of the signal by  $v(t)$ .

$H_C$  is also affected by passive and active components on the path to the ADC, and can also be influenced by inductive and capacitive coupling for small transmission distances, as it closely depends on the circuit components and their placement. Specifically, it is possible for 2 circuits with “the same components, circuit topology and placement area” to have different EMI behavior depending

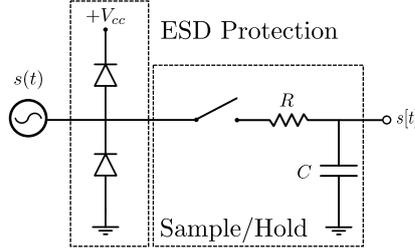


Fig. 2: The sample-and-hold mechanism of an ADC is an  $RC$  low-pass filter. Electrostatic Discharge (ESD) protection diodes can also introduce non-linearities.

on the component placement on the board [8]. Despite the fact that it is hard to mathematically model and predict the behavior of circuits in response to different signal transmissions,  $H_C$  can still be determined empirically using frequency sweeps. It presents a useful abstraction, allowing us to separate the behavior of the ADC (which need only be determined once, for instance by the manufacturer) from circuit layout and transmission details.

Note, finally, that  $H_C$  can also account for distance factors between the adversary and the circuit under test: due to the Friis transmission formula, as distance doubles, EMI transmission power needs to quadruple. This effect can be captured by increasing the attenuation of  $H_C$  by 6 dB, while defense mechanisms such as shielding can be addressed similarly. This approach allows us to side-step engineering issues of remote transmissions and reduce the number of parameters used in the security definitions we propose in Section 3.

**ADC Transfer Function  $H_A$ .** Every system with sensors contains one or more ADCs, which may even be integrated into the sensor chip itself. ADCs are not perfect, but contain components which may cause a mismatch between the “true” value at the ADC input and the digitized output. In this section, we describe how these components affect the digitization process.

Although there are many types of ADCs, every ADC contains three basic components: a “sample- or track-and-hold circuit where the sampling takes place, the digital-to-analog converter and a level-comparison mechanism” [10]. The sample-and-hold component acts as a low-pass filter, and makes it harder for an adversary to inject signals modulated at high frequencies. However, the level-comparison mechanism is essentially an amplifier with non-linearities which induces DC offsets, and allows low-frequency intermodulation products to pass through. These ADC-specific transformations, modeled through  $H_A$ , unintentionally demodulate high-frequency signals which are not attenuated by  $H_C$ .

**Sample-And-Hold Filter Characteristics.** A sample-and-hold (S/H) mechanism is an  $RC$  circuit connected to the analog input, with the resistor and the capacitor connected in series (Figure 2). The transfer function of the voltage across the capacitor is  $H_{S/H}(j\omega) = \frac{1}{1+j\omega RC}$ , and the magnitude of the gain is  $G_{S/H} = \frac{1}{\sqrt{1+(\omega RC)^2}}$ . As the angular frequency  $\omega = 2\pi f$  increases, the gain is reduced: the S/H mechanism acts as a low-pass filter. The  $-3$  dB cutoff frequency

is thus  $f_{cut} = \frac{1}{2\pi RC}$ , which is often higher than the ADC sampling rate (Appendix A). Hence, “aliasing” occurs when signals beyond the Nyquist frequency are digitized by the ADC: high-frequency signals become indistinguishable from low-frequency signals which the ADC can sample accurately.

**Amplifier Non-Linearities.** Every ADC contains amplifiers: a comparator, and possibly buffer and differential amplifiers. Many circuits also contain additional external amplifiers to make weak signals measurable. All these amplifiers have harmonic and intermodulation non-linear distortions [12], which an adversary can exploit. Harmonics are produced when an amplifier transforms an input  $v_{in}$  to an output  $v_{out} = \sum_{n=1}^{\infty} a_n v_{in}^n$ . In particular, if  $v_{in} = \hat{v} \cdot \sin(\omega t)$ , then:

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

This equation shows that “the frequency spectrum of the output contains a spectral component at the original (fundamental) frequency, [and] at multiples of the fundamental frequency (harmonic frequencies)” [12]. Moreover, the output includes a DC component, which depends only on the even-order non-linearities of the system. Besides harmonics, intermodulation products arise when the input signal is a sum of two sinusoids (for instance when the injected signal sums with the sensor signal):  $v_{in} = \hat{v}_1 \cdot \sin(\omega_1 t) + \hat{v}_2 \cdot \sin(\omega_2 t)$ . In that case, the output signal contains frequencies of the form  $n\omega_1 \pm m\omega_2$  for integers  $n, m \neq 0$ . These non-linearities demodulate attacker waveforms, even when they are modulated on high-frequency carriers.

**Diode Rectification.** Figure 2 shows that the input to ADC can contain reverse-biased diodes to ground and  $V_{cc}$  to protect the input from Electrostatic Discharge (ESD). When the input to the ADC is negative, or when it exceeds  $V_{cc}$ , the diodes clamp it, causing non-linear behavior. When the sensor signal  $s(t)$  is positive, this behavior is also asymmetric, causing a DC shift [12], which compounds with the amplifier non-linearities.

**Conclusion.** All ADCs contain the same basic building blocks, modeled through  $H_A$ . Although the sample-and-hold mechanism should attenuate high-frequency signals beyond the maximum sampling rate of the ADC, non-linearities due to ESD diodes and amplifiers in the ADC cause DC offsets and the demodulation of signals through harmonics and intermodulation products. Appendix A exemplifies these effects through experiments with different types of ADCs.

## 2.2 Sampling Errors in the Absence of an Adversary

The digitization process through ADCs entails errors due to quantization and environmental noise. Quantization errors exist due to the inherent loss of accuracy in the sampling process. An ADC can only represent values within a range, say between  $V_{min}$  and  $V_{max}$  volts, with a finite binary representation of  $N$  bits, called the *resolution* of the ADC. In other words, every value between  $V_{min}$  and  $V_{max}$  is mapped to one of the  $2^N$  values that can be represented using  $N$  bits.

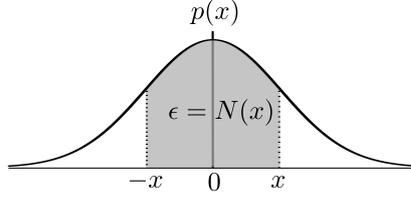


Fig. 3: Noise probability distribution  $p(x)$ . The shaded area represents the probability  $\epsilon = N(x) = \Pr[|n(t)| \leq x]$ .

As a result, there is a *quantization error* between the true sensor analog value  $s$  and the digitized value  $\tilde{s}$ . The maximum value of this error is

$$Q = \frac{V_{max} - V_{min}}{2^{N+1}} \geq |s - \tilde{s}| \quad (1)$$

The second source of error comes from environmental noise, which may affect measurements. We assume that this noise, denoted by  $n(t)$ , is independent of the signal being measured, and that it comes from a zero-mean distribution, i.e., that the noise is *white*. The security definitions we introduce in Section 3 require an estimate of the level of noise in the system, so we introduce some relevant notation here. We assume that  $n(t)$  follows a probability distribution function (PDF)  $p(x)$ , and define  $N(x)$  as the probability that the noise is between  $-x$  and  $x$ , as shown in Figure 3, i.e.,

$$N(x) = \Pr[|n(t)| \leq x] = \int_{-x}^x p(u) du$$

Note that typically the noise is assumed to come from a normal distribution, but this assumption is not necessary in our models and definitions.

We are also interested in the inverse of this function, where given a probability  $0 \leq \epsilon < 1$ , we want to find  $x \geq 0$  such that  $N(x) = \epsilon$ . For this  $x$ , the probability that the noise magnitude falls within  $[-x, x]$  is  $\epsilon$ , as also shown in Figure 3. Because for some distributions there might be multiple  $x$  for which  $N(x) = \epsilon$ , we use the smallest such value:

$$N^{-1}(\epsilon) = \inf\{x \geq 0 : N(x) = \epsilon\} \quad (2)$$

Since  $N(x)$  is an increasing function, so is  $N^{-1}(\epsilon)$ .

To account for repeated measurements, we introduce a short-hand for sampling errors, which we denote by  $E_s(t)$ . The sampling errors depend on the sensor input into the ADC  $s(t)$ , the sampling rate  $f$ , the discrete output of the ADC  $\tilde{s}_f(t)$  as well as the conversion delay  $\tau$ , representing the time the ADC takes for complete a conversion:

$$E_s(t) = \begin{cases} |\tilde{s}_f(t + \tau) - s(t)| & \text{if a conversion starts at } t \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

### 2.3 Adversary Model

Our threat model and definitions can capture a range of attacker goals, from attackers who merely want to disrupt sensor outputs, to those who wish to inject precise waveforms into a system. We define these notions precisely in Section 3, but here we describe the attacker capabilities based on our model of Figure 1. Specifically, in our model, the adversary can only alter the transmitted adversarial signal  $v(t)$ . He/she cannot directly influence the sensor signal  $s(t)$ , the (residual) noise  $n(t)$ , or the transfer functions  $H_A$  and  $H_C$ . The adversary knows  $H_A$ ,  $H_C$ , and the distribution of the noise  $n(t)$ , although the true sensor signal  $s(t)$  might be hidden from the adversary (see Section 3.2). The only constraint placed on the adversarial signal is that the attacker is only allowed to transmit signals  $v(t)$  whose peak voltage level is bounded by some constant  $V_{PK}^{Adv}$ , i.e.,  $|v(t)| \leq V_{PK}^{Adv}$  for all  $t$ . We call this adversary a  $V_{PK}^{Adv}$ -bound adversary, and all security definitions are against such bounded adversaries.

We choose to restrict voltage rather than restricting power or distance, as it makes for a more powerful adversarial model. Our model gives the adversary access to any physical equipment necessary (such as powerful amplifiers and highly-directional antennas), while reducing the number of parameters needed for our security definitions of Section 3. Distance/power effects can be compensated directly through altering  $V_{PK}^{Adv}$ , or indirectly by integrating them into  $H_C$ , as discussed in Section 2.1.

## 3 Security Definitions

Using the model of Figure 1, we can define security in the presence of signal injection attacks. The  $V_{PK}^{Adv}$ -bound adversary is allowed to transmit any waveform  $v(t)$ , provided that  $|v(t)| \leq V_{PK}^{Adv}$  for all  $t$ : the adversary is only constrained by power budget requirements. Whether or not the adversary succeeds in injecting the target waveform  $w(t)$  into the output of the system depends on the transfer functions  $H_C$  and  $H_A$ . For a given system described by  $H_A$  and  $H_C$ , there are three outcomes against an adversary whose only restriction is power:

1. The adversary can disturb the sensor readings, but cannot precisely control the measurement outputs, an attack we call *existential injection*. The lack of existential injections can be considered *universal security*.
2. The adversary can inject a target waveform  $w(t)$  into the ADC outputs with high fidelity, performing a *selective injection*. If the adversary is unable to succeed, the system is *selectively secure* against  $w(t)$ .
3. The adversary can *universally inject* any waveform  $w(t)$ . If there is any non-trivial waveform for which he/she fails, the system is *existentially secure*.

This section sets out to precisely define the above security notions by accounting for noise and quantization error (Equation (1)). Our definitions capture the intuition that systems are secure when there are no adversarial transmissions,

and are “monotonic” in power, i.e., systems are more vulnerable against adversaries with access to higher-powered transmitters. Our definitions are also monotonic in noise. In other words, in environments with low noise, even a small disturbance of the output is sufficient to break the security of a system. Section 3.1 evaluates whether an adversary can disturb the ADC output away from its correct value sufficiently. Section 3.2 then formalizes the notion of selective security against target waveforms  $w(t)$ . Finally, Section 3.3 introduces universal injections by defining what a non-trivial waveform is.

### 3.1 Existential Injection, Universal Security

The most primitive attack is a simple disruption of the sensor readings. There are two axes in which this notion can be evaluated: adversarial power and probability of success.<sup>1</sup> For a fixed probability of success, we want to determine the least amount of power for which an attack is successful. For a fixed power, we want to find the probability of a successful attack. Alternatively, if we fix both power and probability of success, we want to determine if a system is secure against disruptive signal injection attacks.

The definition for universal security is a formalization of the above intuition, calling a system secure when, even in the presence of injections (bounded by adversarial power), the true analog sensor value and the ADC digital output do not deviate by more than the quantization error and the noise, with sufficiently high probability. Mathematically:

**Definition 1 (Universal Security, Existential Injection).** For  $0 \leq \epsilon < 1$ , and  $V_{PK}^{Adv} \geq 0$ , we call a system **universally**  $(\epsilon, V_{PK}^{Adv})$ -secure if

$$\Pr \left[ E_s(t) \geq Q + N^{-1} \left( \frac{\epsilon + 1}{2} \right) \right] \leq \frac{\epsilon + 1}{2} \quad (4)$$

for every adversarial waveform  $v(t)$ , with  $|v(t)| \leq V_{PK}^{Adv}$  for all  $t$ .  $Q$  is the quantization error of the system,  $N^{-1}$  is the noise distribution inverse defined in Equation (2), and  $E_s$  is the sampling error as defined by Equation (3). The probability is taken over the duration of the attack, i.e., at each sampling point within the interval  $t_{start} \leq t \leq t_{end}$ . We call a successful attack an **existential injection**, and simply call a system **universally  $\epsilon$ -secure**, when  $V_{PK}^{Adv}$  is implied.

We first show that in the absence of injections, the system is universally  $\epsilon$ -secure for all  $0 \leq \epsilon < 1$ . Indeed, let  $x = N^{-1} \left( \frac{\epsilon + 1}{2} \right)$ , so that  $\Pr [ |n(t)| \leq x ] = \frac{\epsilon + 1}{2}$ . Then, in the absence of injections,

$$\begin{aligned} \Pr \left[ E_s(t) \geq Q + N^{-1} \left( \frac{\epsilon + 1}{2} \right) \right] &= \Pr [ |n(t)| \geq x ] = \\ &= 1 - \frac{\epsilon + 1}{2} = \frac{1 - \epsilon}{2} \leq \frac{\epsilon + 1}{2} \end{aligned}$$

<sup>1</sup> Success is probabilistic, as noise is a random variable.

which holds for all  $0 \leq \epsilon < 1$ , as desired. This proof is precisely the reason for requiring a noise level and probability of at least 50% in the definition: the proof no longer works if  $(1 + \epsilon)/2$  is replaced by just  $\epsilon$ . In other words, mere noise would be classified as an attack by the modified definition.

**Power.** We now show that a higher adversarial power budget can only make the system more vulnerable. Indeed, if a system is universally  $(\epsilon, V_1)$ -secure, then it is universally  $(\epsilon, V_2)$ -secure for  $V_2 \leq V_1$ . For this, it suffices to prove the contrapositive, i.e., that if a system is not universally  $(\epsilon, V_2)$ -secure, then it is not universally  $(\epsilon, V_1)$ -secure. For the proof, let  $v(t)$  be an adversarial waveform with  $|v(t)| \leq V_2$  such that Equation (4) does not hold, which exists by the assumption that the system is not universally  $(\epsilon, V_2)$ -secure. Then, by the transitive property,  $|v(t)| \leq V_1$ , making  $v(t)$  a valid counterexample for universal  $(\epsilon, V_1)$  security.

**Probability.** The third property we show is probability monotonicity, allowing us to define a ‘‘critical threshold’’ for  $\epsilon$ , above which a system is universally secure (for a fixed  $V_{PK}^{Adv}$ ), and below which a system is not universally secure. Indeed, for fixed  $V_{PK}^{Adv}$ , if a system is universally  $(\epsilon, V_{PK}^{Adv})$ -secure, then it is universally  $(\epsilon + \delta, V_{PK}^{Adv})$ -secure for  $0 \leq \delta < 1 - \epsilon$ , as

$$\begin{aligned} Pr \left[ E_s(t) \geq Q + N^{-1} \left( \frac{\epsilon + \delta + 1}{2} \right) \right] &\leq \\ Pr \left[ E_s(t) \geq Q + N^{-1} \left( \frac{\epsilon + 1}{2} \right) \right] &\leq \frac{\epsilon + 1}{2} \leq \frac{\epsilon + \delta + 1}{2} \end{aligned}$$

because  $N^{-1}$  is increasing. The contrapositive is, of course, also true: if a system is not universally secure for a given  $\epsilon$ , it is also not universally secure for  $\epsilon - \delta$  with  $0 \leq \delta \leq \epsilon$ .

**Thresholds.** For a given security level  $\epsilon$ , then, we can talk about the maximum (if any)  $V_{PK}^{Adv}$  such that a system is universally  $(\epsilon, V_{PK}^{Adv})$ -secure, or conversely the minimum (if any)  $V_{PK}^{Adv}$  such that a system is not universally  $(\epsilon, V_{PK}^{Adv})$ -secure. This is the **critical universal voltage level**  $V_c$  for the given  $\epsilon$ . Moreover, for any  $V_{PK}^{Adv}$ , there is a unique **critical universal security threshold**  $\epsilon_c$  such that the system is universally  $(\epsilon, V_{PK}^{Adv})$ -secure for  $\epsilon_c < \epsilon < 1$  and not universally  $(\epsilon, V_{PK}^{Adv})$ -secure for  $0 \leq \epsilon < \epsilon_c$ . By convention we take  $\epsilon_c = 0$  if the system is secure for all  $\epsilon$ , and  $\epsilon_c = 1$  if there is no  $\epsilon$  for which the system is secure. This critical threshold indicates the security level of a system: the lower  $\epsilon_c$  is, the better a system is protected against signal injection attacks.

### 3.2 Selective Injection and Security

The second definition captures the notion of security against specific target waveforms  $w(t)$ : we wish to find the probability that a  $V_{PK}^{Adv}$ -bounded adversary can make  $w(t)$  appear in the output of the ADC. Conversely, to define security in this context, we must make sure that the digitized signal  $\tilde{s}_f(t)$  differs from the waveform  $s(t) + w(t)$  with high probability, even if plenty of noise is allowed. There are two crucial points to notice about the waveform  $w(t)$ . First,  $w(t)$  is not

the raw signal  $v(t)$  the adversary is transmitting, as this signal undergoes two transformations via  $H_C$  and  $H_A$ . Instead,  $w(t)$  is the signal that the adversary wants the ADC to think that it is seeing, and is usually a demodulated version of  $v(t)$  (see Figure 1). Second,  $w(t)$  does not necessarily cancel out or overpower  $s(t)$ , because that would require predictive modeling of the sensor signal  $s(t)$ . However, if the adversary can predict  $s(t)$  (e.g., by monitoring the output of the ADC, or by using identical sensors), we can then ask about security against the waveform  $w'(t) = w(t) - s(t)$  instead. Given this intuition, we can define selective security as follows:

**Definition 2 (Selective Security, Selective Injection).** For  $0 \leq \epsilon < 1$ , and  $V_{PK}^{Adv} \geq 0$ , a system is called **selectively**  $(\epsilon, w(t), V_{PK}^{Adv})$ -**secure** if

$$\Pr \left[ E_{s+w}(t) \geq Q + N^{-1} \left( \frac{(1-\epsilon)+1}{2} \right) \right] > \frac{2-\epsilon}{2} \quad (5)$$

for every adversarial waveform  $v(t)$ , with  $|v(t)| \leq V_{PK}^{Adv}$  for all  $t$ , where the probability is taken over the duration of the attack.  $Q$  is the quantization error of the system,  $N^{-1}$  is the noise distribution inverse defined in Equation (2), and  $E_{s+w}(t) = |\tilde{s}_f(t+\tau) - s(t) - w(t)|$  during sampling periods, and 0 otherwise. We call a successful attack a **selective injection**, and simply call a system selectively  $\epsilon$ -secure, when  $V_{PK}^{Adv}$  and  $w(t)$  are clear from context.

This definition is monotonic in adversarial power, and the probability of success, allowing us to talk about “the” probability of success for a given waveform:

**Power.** The same argument as above shows that increasing  $V_{PK}^{Adv}$  can only make a secure system insecure, but not vice versa, i.e., that if a system is selectively  $(\epsilon, w(t), V_1)$ -secure, then it is selectively  $(\epsilon, w(t), V_2)$ -secure for  $V_2 \leq V_1$ . We can thus define the **critical selective voltage level**  $V_c^w$  for a given  $\epsilon$  and  $w(t)$ .

**Probability.** If a system is selectively  $\epsilon$ -secure (against a target waveform and power budget), then it is selectively  $(\epsilon + \delta)$ -secure for  $0 \leq \delta < 1 - \epsilon$ , because

$$\begin{aligned} P &= \Pr \left[ E_{s+w}(t) \geq Q + N^{-1} \left( \frac{1 - (\epsilon + \delta) + 1}{2} \right) \right] \\ &\geq \Pr \left[ E_{s+w}(t) \geq Q + N^{-1} \left( \frac{1 - \epsilon + 1}{2} \right) \right] \\ &> \frac{2 - \epsilon}{2} \geq \frac{2 - (\epsilon + \delta)}{2} \end{aligned}$$

If the system is not selectively  $\epsilon$ -secure, then it is not selectively  $(\epsilon - \delta)$ -secure.

Given the above, for a given waveform  $w(t)$  and fixed  $V_{PK}^{Adv}$ , we can define a waveform-specific **critical selective security threshold**  $\epsilon_c^w$  such that the system is vulnerable for all  $\epsilon^w$  with  $0 \leq \epsilon^w < \epsilon_c^w$  and secure for all  $\epsilon^w$  with  $\epsilon_c^w < \epsilon^w < 1$ . By convention we take  $\epsilon_c^w = 0$  if there is no  $\epsilon$  for which the system is vulnerable, and  $\epsilon_c^w = 1$  if there is no  $\epsilon$  for which the system is secure.

**Threshold Relationship.** The critical universal threshold of a system  $\epsilon_c$  is related to the critical selective threshold  $\epsilon_c^0$  against the zero waveform  $w(t) = 0$  through the equation  $\epsilon_c^0 = 1 - \epsilon_c$ . Indeed, if a system is not universally  $\epsilon$ -secure, then  $P = \Pr [E_s(t) \geq Q + N^{-1} (\frac{\epsilon+1}{2})] > \frac{\epsilon+1}{2}$ , so  $\frac{2-(1-\epsilon)}{2} = \frac{\epsilon+1}{2} < P = \Pr [E_{s+0}(t) \geq Q + N^{-1} (\frac{(1-(1-\epsilon))+1}{2})]$ , making the system is selectively  $(1 - \epsilon)$ -secure for the zero waveform. Conversely, if a system is selectively  $(1 - \epsilon)$ -secure for the zero waveform, then it is not universally  $\epsilon$ -secure. The fact that a low critical universal threshold results in a high critical selective threshold for the zero threshold is not surprising: it is easy for an adversary to inject a zero signal by simply not transmitting anything.

### 3.3 Universal Injection, Existential Security

The final notion of security is a weak one, which requires that the adversary cannot inject at least one “representable” waveform into the system, i.e., one which is within the ADC limits. We can express this more precisely as follows:

**Definition 3 (Representable Waveform).** *A waveform  $w(t)$  is called **representable** if it is within the ADC voltage levels, and has a maximum frequency component bounded by the Nyquist frequency of the ADC. Mathematically,  $V_{min} \leq w(t) \leq V_{max}$  and  $f_{max} \leq f_s/2$ .*

Using this, we can define security against at least one representable waveform:

**Definition 4 (Existential Security, Universal Injection).** *For  $0 \leq \epsilon < 1$ , and  $V_{PK}^{Adv} \geq 0$ , a system is called **existentially**  $(\epsilon, V_{PK}^{Adv})$ -**secure** if there exists a representable waveform  $w(t)$  for which the system is selectively  $(\epsilon, w(t), V_{PK}^{Adv})$ -secure. We call a system existentially  $\epsilon$ -secure when  $V_{PK}^{Adv}$  is clear. If there is no such  $w(t)$ , we say that the adversary can perform any **universal injection**.*

As above, power and probability are monotonic in the opposite direction.

**Power.** If a system is existentially  $(\epsilon, V_1)$ -secure, then it is  $(\epsilon, V_2)$ -secure for  $V_2 \leq V_1$ . By assumption, there is a representable  $w(t)$  such that the system is selectively  $(\epsilon, w(t), V_1)$ -secure. By the previous section, this system is  $(\epsilon, w(t), V_2)$ -secure, concluding the proof.

**Probability.** If a system is existentially  $(\epsilon_1, V)$ -secure, then it is  $(\epsilon_2, V)$ -secure for  $\epsilon_1 \leq \epsilon_2$ . By assumption, there is a representable  $w(t)$  such that the system is selectively  $(\epsilon_1, w(t), V)$ -secure. By the previous section, the system is also  $(\epsilon_2, w(t), V)$ -secure, as desired.

**Thresholds.** Extending the definitions of the previous sections, for fixed  $\epsilon$  we can define a **critical existential voltage level**  $V_c^{exist}$  below which a system is existentially  $\epsilon$ -secure, and above which the system is existentially  $\epsilon$ -insecure. Similarly, for a fixed adversarial voltage we can define the **critical existential security threshold**  $\epsilon_c^{exist}$ , above which the system is existentially secure, and below which the system is insecure.

Table 1: The adversary can easily disturb the smartphone output (existential injection), and inject human speech (universal injection). Selective injections of sines are less precise than exponentials of the same frequency.

Injection	Resulting Signal	Crit. Thres.
Existential	$w(t) \neq 0$	0.892
Selective	$w(t) = e^{\sin(2\pi f_m t)}$	0.747
Selective	$w(t) = \sin(2\pi f_m t)$	0.562
Universal	“OK Google” commands	$\leq 0.562$

In some cases, security designers may wish to adjust the definitions to restrict target waveforms (and hence existential security counterexamples) even further. For instance, we might wish to check whether an adversary can inject all waveforms which are sufficiently bounded away from 0, periodic waveforms, or waveforms of a specific frequency. The proofs for power and probability monotonicity still hold, allowing us to talk about universal security against  $\mathcal{S}$ -representable waveforms: waveforms which are representable and also in a set  $\mathcal{S}$ .

## 4 Security Evaluation of a Smartphone Microphone

In this section, we illustrate how our security definitions can be used to determine the security level of a commercial, off-the-shelf smartphone microphone. We first introduce an algorithm to calculate the critical selective security threshold  $\epsilon_c^w$  against a target waveform  $w(t)$  in Section 4.1. We then use the algorithm to calculate the critical thresholds of a smartphone in Section 4.2. Finally, we comment on universal security in Section 4.3, where we show that we are able to inject complex “OK Google” commands. We summarize our results in Table 1.

### 4.1 Algorithm for Selective Security Thresholds

In this section, we introduce an algorithm to calculate the critical selective security threshold  $\epsilon_c^w$  of a system against a target waveform  $w(t)$ , using a transmitted signal  $v(t)$ . The first step in calculating the security level is determining the noise distribution. To that end, we collect  $N$  measurements of the system output  $\tilde{s}_f(t)$  during the injection and pick one as the *reference* signal. We then pick  $1 \leq k \leq N - 2$  of them to calculate the noise (*estimation* signals), while the remaining are used to verify our calculations (*validation* signals).

Our algorithm first removes any DC offset and re-scales the measurements so that the root-mean-square (RMS) voltages of the signals are the same. The repeated measurements are then phase-aligned, and we calculate the distance between the reference signal and the estimation signals. The average of this distance should be very close to 0, as the signals are generated in the same way. However, the standard deviation  $\sigma$  is non-zero, so we can model noise as following a zero-mean normal distribution  $n(t) \sim N(0, \sigma^2)$ . We can then find the critical threshold between the reference signal and any target *ideal* waveform

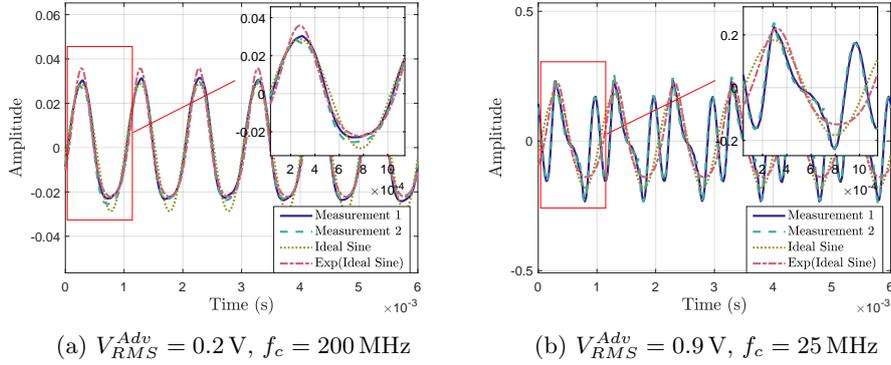


Fig. 4: Clean (4a) and Distorted (4b) waveforms injected into the smartphone, with ideal sine and exponential sine functions for comparison.

$w(t)$  as follows: we first detrend, scale, and align the ideal signal to the reference waveform, as with the estimation signals. Then, we calculate the errors (distance) between the ideal and the reference signal. Finally, we perform a binary search for different values of  $\epsilon$ , in order to find the largest  $\epsilon$  for which Equation (5) does not hold: this is the critical threshold. To calculate the inverse of the noise, we use the percentile point function  $ppf(\epsilon)$ , which is the inverse of the cumulative distribution function, and satisfies  $N^{-1}(\epsilon) = ppf((1 + \epsilon)/2)$ . Note that since the critical universal threshold  $\epsilon_c$  is related to the selective critical threshold of the zero waveform  $\epsilon_c^0$  through  $\epsilon_c = 1 - \epsilon_c^0$  (Section 3.2), the same algorithm can be used to calculate  $\epsilon_c$ .

## 4.2 Existential and Selective Injections into a Smartphone

We demonstrate how our algorithm can be used in a realistic setup using a Motorola XT1541 Moto G3 smartphone. We directly inject modulated  $f_m = 1 \text{ kHz}$  signals, collecting  $N = 10$  measurements of  $2^{15}$  sample points per run, and record the data at a frequency of  $f_s = 44.1 \text{ kHz}$ . We first modulate  $f_m$  over  $f_c = 200 \text{ MHz}$  using an output level of  $V_{RMS}^{Adv} = V_{PK}^{Adv} / \sqrt{2} = 0.2 \text{ V}$ . This injection is demodulated well by the smartphone and has a similarity (as indicated by the Pearson Correlation Coefficient) of over 0.98 compared to a pure 1 kHz tone. We call this example the “clean” waveform. The second injection, which we call the “distorted” waveform, uses  $f_c = 25 \text{ MHz}$ ,  $V_{RMS}^{Adv} = 0.9 \text{ V}$ , and has a similarity of less than 0.55 to the ideal tone. Example measurements of these signals and “ideal” signals (see below) are shown in Figure 4.

The algorithm first calculates the noise level using the reference signals. As expected, the error average is very close to 0 (usually less than  $10^{-6}$ ), while the standard deviation  $\sigma$  is noticeable at around 0.0015. Taking the reference signals as the target signal  $w(t)$ , the critical selective thresholds are close to 1. In other words, even if the injected waveforms do not correspond to “pure” signals, the

Table 2: Mean and std. deviation  $(\mu, \sigma)$  of critical selective thresholds  $\epsilon_c^w$  for different target signals  $w(t)$ . Injections using the clean waveform are always more successful than with the distorted waveform. Validation signals are injected with high fidelity, and are better modeled by an exponential rather than a pure sine.

Waveform	Validation	Ideal Sine	$e^{\text{Ideal Sine}}$	$w(t) \neq 0$
Clean	(0.98, 0.03)	(0.56, 0.04)	(0.75, 0.06)	(0.89, 0.01)
Distorted	(0.95, 0.09)	(0.31, 0.05)	(0.34, 0.05)	(0.71, 0.04)

adversary can inject them with high fidelity: the system is not selectively secure against them with high probability.

We also tried two signals as the signal  $w(t)$  that the adversary is trying to inject: a pure 1 kHz sine wave, and an exponential of the same sine wave. The averages and standard deviations for the calculated thresholds over all combinations of  $k$  and reference signals are shown in Table 2. As we would expect, the thresholds for the distorted waveform are much lower than the values for the clean waveform: the signal is distorted, so it is hard to inject an ideal signal. We also find that the exponential function is a better fit for the signal we are seeing, and can better explain the harmonics. Table 2 also includes the critical universal injection threshold based on the two waveform injections. This threshold is much higher for both waveforms, as injections disturb the ADC output sufficiently, even when the demodulated signal is not ideal.

### 4.3 Universal Injections on a Smartphone

In this section, we demonstrate that the smartphone is vulnerable to the injection of arbitrary commands, which cause the smartphone to behave as if the user initiated an action. We first inject a modulated recording of “OK Google, turn on the flashlight” into the microphone port, checking both whether the voice command service was activated in response to “OK Google”, and whether the desired action was executed. We repeat measurements 10 times, each time amplitude-modulating the command at a depth of  $m = 100\%$  with  $V_{RMS}^{Adv} = 0.6$  V on 26 carrier frequencies  $f_c$ : 25 MHz, 50 MHz, and 100 – 2400 MHz at a step of 100 MHz. The voice-activation feature (“OK Google”) worked with 100% success rate (10/10 repetitions) for all frequencies, while the full command was successfully executed for 23 of the 26 frequencies we tested (all frequencies except  $f_c \in \{1.3, 2.0, 2.4\text{GHz}\}$ ). Increasing the output level to  $V_{RMS}^{Adv} = 0.9$  V, increased success rate to 25/26 frequencies. Only  $f_c = 2.4$  GHz did not result in a full command injection, possibly because the Wi-Fi disconnected in the process.

We repeated the above injections, testing 5 further commands to (1) call a contact; (2) text a contact; (3) set a timer; (4) mute the volume; and (5) turn on airplane mode. The results remained identical, regardless of the actual command to be executed. As a result, all carrier frequencies which are not severely attenuated by  $H_C$  (e.g., when coupling to the user’s headphones) are vulnerable to injections of complex waveforms such as human speech.

## 5 Discussion

We now discuss how our work can inform design choices. To start, choosing the right ADC directly impacts the susceptibility to signal injection attacks. As shown in Appendix A, some ADCs distort the demodulated output, and are thus more resilient to clean sinusoidal injections. Moreover, other ADCs require fine-grained control over the carrier frequency of injection. As the adversarial signal is transformed through the circuit-specific transfer function  $H_C$ , the adversary may not have such control, resulting in a more secure system.

Having chosen the appropriate ADC based on cost, performance, security, or other considerations, a designer needs to assess the impact of  $H_C$ . Prior work has shown that even small layout or component changes affect the EMI behavior of a circuit [3,8,19]. Since the ADC behavior can be independently determined through direct power injections, fewer experiments with remote transmissions are required to evaluate the full circuit behavior, and how changes in the circuit’s topology influence the system’s security.

Our selective security definition and algorithm address how to determine the vulnerability of a system against specific waveforms. Universal security, on the other hand, allows us to directly compare the security of two systems for a fixed adversarial power budget through their critical universal security thresholds. Moreover, given a probability/threshold  $\epsilon$ , we can calculate the critical universal voltage level, which is the maximum output level for which a system is still universally  $\epsilon$ -secure.

Our smartphone case study showed that our framework can be used in practice with real systems, while our “OK Google” experiments demonstrated that less-than-perfect injections of adversarial waveforms can have the same effect as perfect injections. This is because there is a mismatch between the true noise-level of a system and the worst-case noise-level that the system expects. In other words, injections worked at all carrier frequencies, even when the demodulated output was noisy or distorted. This is a deliberate, permissive design decision, which allows the adversary to succeed with a range of different and noisy waveforms  $w(t)$ , despite small amplitudes and DC offsets.

Although not heavily discussed in this paper, our model and definitions are general enough to capture alternative signal injection techniques. For instance, electro-mechanical sensors have resonant frequencies which allow acoustic injection attacks [18,21].  $H_C$  can account for such imperfections in the sensors themselves, attenuating injection frequencies which are not close to the resonant frequencies. Our system model also makes it easy to evaluate countermeasures and defense mechanisms in its context. For example, shielding increases the attenuation factor of  $H_C$ , thereby increasing the power requirements for the adversary (Section 2.1). Alternatively, a low-pass filter (LPF) before the ADC and/or amplifier changes  $H_A$ , and attenuates the high-frequency components which would induce non-linearities. Note, however, that even moving the pre-amplifier, LPF and ADC into the same IC package does not fully eliminate the vulnerability to signal injection attacks (Section 6) as the channel between the analog sensor and the ADC cannot be fundamentally authenticated.

## 6 Related Work

Although no frameworks and security definitions have been previously proposed for signal injection attacks, there is extensive literature on attacking sensors/components using electromagnetic (EM), acoustic, and optical injections [4], including EM attacks on IMDs to deliver defibrillation shocks and inject audio signals into webcams and bluetooth headsets [6]. More recently, [14] attacked GPIO pins on microcontrollers, also proposing an analytical model of the magnetic field in EM induction attacks.

Shoukry et al. [16] presented an attack against Anti-Lock Braking Systems (ABS) by generating a magnetic field that replaces the real signal, by a malicious one. The same researchers later presented an authentication system called PyCRA for active sensors [17], which depends on non-spoofable physical and computational delay limits, although [15] suggests that PyCRA would require high computational overhead in practice.

[5] used EM transmissions to trigger voice commands in smartphones by emitting AM-modulated signals that get picked up by the user’s hands-free headset. Other researchers have injected voice commands through modulating signals on ultrasound frequencies [23], or by playing two tones at different ultrasound frequencies, and exploiting non-linearities in components [13].

Acoustic attacks which transmit sound waves at the resonant frequencies of gyroscopes can incapacitate [18] or precisely control [20] drones, with attackers who account for sampling rate drifts being able to control the outputs of accelerometers for longer periods of time [21]. Finally, optical attacks can be used to spoof medical infusion pump measurements [9], and cause autonomous cars and unmanned aerial vehicles (UAVs) to drift or fail [1,11,22].

## 7 Conclusion

Sensors guide many of our choices, and we often blindly trust their values. However, it is possible to spoof their outputs through electromagnetic or other signal injection attacks. To address the lack of a unifying framework describing the susceptibility of devices to such attacks, we defined a system and adversary model for signal injections. Our model is the first to abstract away from specific environments and circuit designs and presents a strong adversary who is only limited by transmission power. It also makes it easy to discuss and evaluate countermeasures in its context and covers different types of signal injection attacks.

Within our model, we defined existential, selective, and universal security, capturing effects ranging from mere disruptions of the ADC outputs to precise injections of all waveforms. We showed in practice that our definitions can be used to evaluate the security level of an off-the-shelf smartphone, and we introduced a novel algorithm to calculate “critical” thresholds, which express how close an injected signal is to the ideal signal. In response to the emerging signal injection threat, our work paves the way towards a future where security can be quantified and compared through our methodology and security definitions.

Table 3: The ADCs used in our experiments cover a range of different properties.

ADC	Manufacturer	Package	Type	Bits	Max $f_s$	$f_{cut}$
TLC549	Texas Instruments	DIP	SAR	8	40 kHz	2.7 MHz
ATmega328P	Atmel	Integrated	SAR	10	76.9 kHz	0.1-11.4 MHz
Artix7	Xilinx	Integrated	SAR	12	1 MHz	5.3 MHz
AD7276	Analog Devices	TSOT	SAR	12	3 MHz	66.3 MHz
AD7783	Analog Devices	TSSOP	$\Delta\Sigma$	24	19.79 Hz	[50,60 Hz]
AD7822	Analog Devices	DIP	Flash	8	2 MHz	128.4 MHz

## Appendix A ADC Response $H_A$ to Malicious Signals

As explained in Section 2, an adversary trying to inject signals remotely into a system typically needs to transmit modulated signals over high-frequency carriers. As  $H_C$  is unique to each circuit and needs to be re-calculated even for minor changes to its components and layout [3], the first step to determine the system vulnerability is to understand the behavior  $H_A$  of the ADC used.

To do so, we inject signals generated via a Rohde & Schwarz SMC100A/B103 signal generator directly into 6 ADCs and determine their demodulation characteristics. The ADCs come from 4 manufacturers in different packages and are controlled via different protocols. The maximum sampling rate  $f_s$  of the ADCs ranges from a few Hz to several MHz, while the resolution ranges from 8 to 24 bits. The ADC types include Delta-Sigma ( $\Delta\Sigma$ ), half-flash, and successive approximation (SAR). Table 3 shows these properties along with the  $-3$  dB cutoff frequency  $f_{cut}$ , calculated using the  $R, C$  parameters in the ADCs’ datasheets.

We inject sinusoidal signals of different frequencies  $f_m$ , which have been amplitude modulated (AM) on different carrier frequencies  $f_c$ . In other words, we consider the intended signal to be  $w(t) = \sin(2\pi f_m t)$ , the sensor signal to be absent ( $s(t) = 0$ ), and evaluate how “close”  $w(t)$  is to the ADC output  $\tilde{s}_f(t)$ . Due to space restrictions, we only describe typical results for each ADC.

**ATmega328P.** Figure 5 presents two example measurements of outputs of the ATmega328P, both in the time domain and in the frequency domain. The input to the ADC is a  $f_m = 1$  Hz signal modulated over different high-frequency carriers. As shown in the frequency domain (bottom of Figure 5), the fundamental frequency  $f_m$  dominates all other frequencies, so the attacker is able to inject a signal of the intended frequency into the output of the ADC. However, the output at both carrier frequencies has strong harmonics at  $2f_m, 3f_m, \dots$  Hz, which indicates that the resulting signal is not pure. Moreover, there is a residual high-frequency component, which is attenuated as the carrier frequency  $f_c$  increases. Finally, there is a frequency-dependent DC offset caused, in part, by the ESD diodes, while the peak-to-peak amplitude of the measured signal decreases as the carrier frequency increases. This is due to the low-pass filtering behavior of the sample-and-hold mechanism, which also explains why we are only able to demodulate signals for carrier frequencies until approximately 150 MHz.

**TLC549.** The TLC549 (Figure 6a) also demodulates the injected signal, but still contains harmonics and a small high-frequency component.

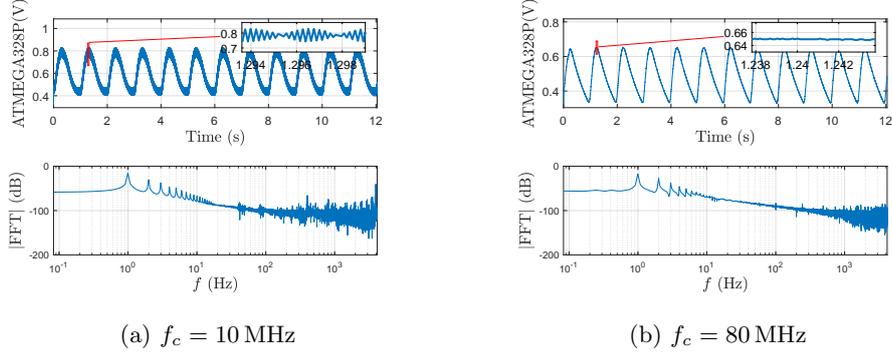


Fig. 5: Example ATmega328P output for power  $P = 0$  dBm, signal frequency  $f_m = 1$  Hz, and modulation depth  $m = 50\%$ . The signal exhibits the correct fundamental frequency, but also contains strong harmonics and a high-frequency component, which is attenuated as the carrier frequency  $f_c$  increases.

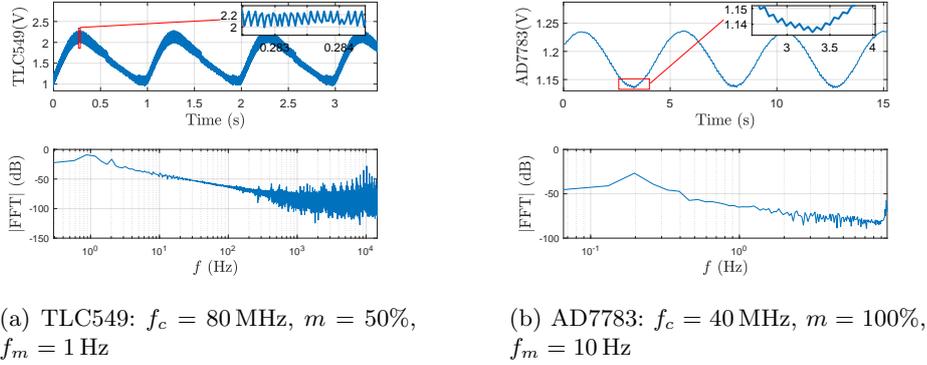


Fig. 6: Example TLC549 (6a) and AD7783 (6b) outputs for a transmission power of  $P = 5$  dBm. Both ADCs demodulate the injected signal, but present harmonics and some high-frequency components. The AD7783 signal is aliased.

**AD7783.** As the AD7783 (Figure 6b) only has a sampling frequency of  $f_s = 19.79$  Hz, aliasing occurs when the baseband signal exceeds the Nyquist frequency  $f_s/2$ . For example, when the baseband frequency is  $f_m = 10$  Hz, the fundamental frequency dominating the measurements is of frequency  $2f_m - f_s = 20 - 19.79 = 0.21$  Hz, with a high-frequency component of  $f_s - f_m = 9.79$  Hz.

**AD7822, AD7276, Artix7.** The three remaining ADCs contain strong high-frequency components which dominate the low-frequency signal. Their outputs appear to be AM-modulated, but at a carrier frequency which is below the ADC's Nyquist frequency. However, with manual tuning of the carrier frequency, it is possible to remove this high frequency component, causing the ADC to demodulate the input. This is shown for the Flash ADC AD7822 in Figure 7, where we change the carrier frequency  $f_c$  in steps of 100 Hz.

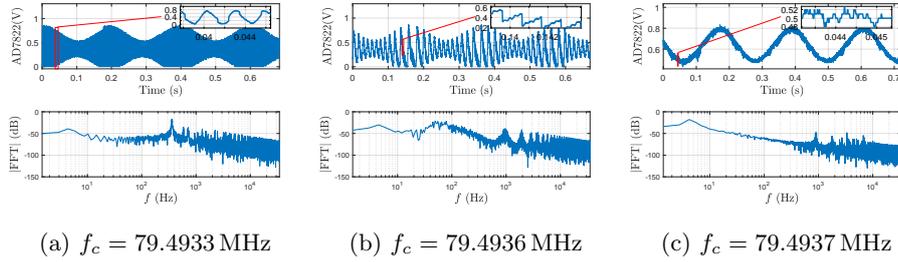


Fig. 7: Example AD7822 output for power  $P = -5$  dBm, signal frequency  $f_m = 5$  Hz, and depth  $m = 50\%$ . Signal demodulation requires a fine-tuned  $f_c$ .

**Conclusion.** The results of our experiments lead to the following observations:

1. **Generality** – All 6 ADCs tested are vulnerable to signal injections at multiple carrier frequencies, as they demodulate signals, matching the theoretical expectations of Section 2.1. As the ADCs are of all major types and with a range of different resolutions and sampling frequencies, the conclusions drawn should be valid for other ADC chips.
2. **Low-Pass Filter** – Although all ADCs exhibited low-pass filtering characteristics, the maximum vulnerable carrier frequency for a given power level was multiple times the cut-off frequency of the  $RC$  circuit at the input of the ADC. This extended the frequency range that an attacker could use for transmissions to attack the system.
3. **Power** – The adversary needs to select the power level of transmissions carefully: too much power in the input of the ADC can cause saturation and/or clipping of the measured signal. Too little power, on the other hand, results in output that looks like noise or a zero signal.
4. **Carrier Frequency** – Some ADCs were vulnerable at any carrier frequency that is not severely attenuated by the sample-and-hold mechanism. For others, high-frequency components dominated the intended baseband signal of frequency  $f_m$  in the ADC output for most frequencies. Even then, carefully-chosen carrier frequencies resulted in a demodulated ADC output.

## References

1. Davidson, D., Wu, H., Jellinek, R., Singh, V., Ristenpart, T.: Controlling UAVs with sensor input spoofing attacks. In: USENIX Workshop on Offensive Technologies (WOOT) (2016)
2. Fu, K., Xu, W.: Risks of trusting the physics of sensors. *Commun. ACM* **61**(2), 20–23 (2018)
3. Gago, J., Balcells, J., González, D., Lamich, M., Mon, J., Santolaria, A.: EMI susceptibility model of signal conditioning circuits based on operational amplifiers. *IEEE Transactions on Electromagnetic Compatibility* **49**(4), 849–859 (2007)
4. Giechaskiel, I., Rasmussen, K.B.: SoK: Taxonomy and challenges of out-of-band signal injection attacks and defenses. arXiv 1901.06935 (2019)
5. Kasmi, C., Lopes-Esteves, J.: IEMI threats for information security: Remote command injection on modern smartphones. *IEEE Transactions on Electromagnetic Compatibility* **57**(6), 1752–1755 (2015)

6. Kune, D.F., Backes, J.D., Clark, S.S., Kramer, D.B., Reynolds, M.R., Fu, K., Kim, Y., Xu, W.: Ghost talk: Mitigating EMI signal injection attacks against analog sensors. In: IEEE Symposium on Security and Privacy (S&P) (2013)
7. Leone, M., Singer, H.L.: On the coupling of an external electromagnetic field to a printed circuit board trace. *IEEE Transactions on Electromagnetic Compatibility* **41**(4), 418–424 (1999)
8. Lissner, A., Hoene, E., Stube, B., Guttowski, S.: Predicting the influence of placement of passive components on EMI behaviour. In: European Conference on Power Electronics and Applications (2007)
9. Park, Y.S., Son, Y., Shin, H., Kim, D., Kim, Y.: This ain't your dose: Sensor spoofing attack on medical infusion pump. In: USENIX Workshop on Offensive Technologies (WOOT) (2016)
10. Pelgrom, M.J.M.: *Analog-to-Digital Conversion*. Springer Publishing Company, Incorporated, 3rd edn. (2017)
11. Petit, J., Stottelaar, B., Feiri, M., Kargl, F.: Remote attacks on automated vehicles sensors: Experiments on camera and LiDAR. *Black Hat Europe* (2015)
12. Redouté, J.M., Steyaert, M.: *EMC of Analog Integrated Circuits*. Springer Publishing Company, Incorporated, 1st edn. (2009)
13. Roy, N., Hassanieh, H., Roy Choudhury, R.: BackDoor: Making microphones hear inaudible sounds. In: International Conference on Mobile Systems, Applications, and Services (MobiSys) (2017)
14. Selvaraj, J., Dayanikli, G.Y., Gaunkar, N.P., Ware, D., Gerdes, R.M., Mina, M.: Electromagnetic induction attacks against embedded systems. In: Asia Conference on Computer and Communications Security (ASIACCS) (2018)
15. Shin, H., Son, Y., Park, Y., Kwon, Y., Kim, Y.: Sampling race: Bypassing timing-based analog active sensor spoofing detection on analog-digital systems. In: USENIX Workshop on Offensive Technologies (WOOT) (2016)
16. Shoukry, Y., Martin, P.D., Tabuada, P., Srivastava, M.B.: Non-invasive spoofing attacks for anti-lock braking systems. In: Cryptographic Hardware and Embedded Systems (CHES) (2013)
17. Shoukry, Y., Martin, P.D., Yona, Y., Diggavi, S., Srivastava, M.B.: PyCRA: Physical challenge-response authentication for active sensors under spoofing attacks. In: Conference on Computer and Communications Security (CCS) (2015)
18. Son, Y., Shin, H., Kim, D., Park, Y., Noh, J., Choi, K., Choi, J., Kim, Y.: Rocking drones with intentional sound noise on gyroscopic sensors. In: USENIX Security Symposium (2015)
19. Sutu, Y.H., Whalen, J.J.: Statistics for demodulation RFI in operational amplifiers. In: International Symposium on Electromagnetic Compatibility (EMC) (1983)
20. Trippel, T., Weisse, O., Xu, W., Honeyman, P., Fu, K.: WALNUT: Waging doubt on the integrity of MEMS accelerometers with acoustic injection attacks. In: IEEE European Symposium on Security and Privacy (EuroS&P) (2017)
21. Tu, Y., Lin, Z., Lee, I., Hei, X.: Injected and delivered: Fabricating implicit control over actuation systems by spoofing inertial sensors. In: USENIX Security Symposium (2018)
22. Yan, C., Xu, W., Liu, J.: Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle. *DEFCON* (2016)
23. Zhang, G., Yan, C., Ji, X., Zhang, T., Zhang, T., Xu, W.: DolphinAttack: Inaudible voice commands. In: Conference on Computer and Communications Security (CCS) (2017)