WALNUT: Waging Doubt on the Integrity of MEMS Accelerometers with Acoustic Injection Attacks

Timothy Trippel, Ofir Weisse, Wenyuan Xu*, Peter Honeyman, Kevin Fu

Computer Science and Engineering, University of Michigan
*Computer Science and Engineering, University of South Carolina

`https://spqr.eecs.umich.edu/walnut/`

Abstract—**Cyber-physical systems depend on sensors to make automated decisions. Resonant acoustic injection attacks are already known to cause malfunctions by disabling MEMS-based gyroscopes. However, an open question remains on how to move beyond denial of service attacks to achieve full adversarial control of sensor outputs. Our work investigates how analog acoustic injection attacks can damage the digital *integrity* of a popular type of sensor: the capacitive MEMS accelerometer. Spoofing such sensors with intentional acoustic interference enables an out-of-spec pathway for attackers to deliver chosen digital values to microprocessors and embedded systems that blindly trust the unvalidated integrity of sensor outputs. Our contributions include (1) modeling the physics of malicious acoustic interference on MEMS accelerometers, (2) discovering the circuit-level security flaws that cause the vulnerabilities by measuring acoustic injection attacks on MEMS accelerometers as well as systems that employ on these sensors, and (3) two software-only defenses that mitigate many of the risks to the integrity of MEMS accelerometer outputs.**

We characterize two classes of acoustic injection attacks with increasing levels of adversarial control: *output biasing* and *output control*. We test these attacks against 20 models of capacitive MEMS accelerometers from 5 different manufacturers. Our experiments find that 75% are vulnerable to output biasing, and 65% are vulnerable to output control. To illustrate end-to-end implications, we show how to inject fake steps into a Fitbit with a $5 speaker. In our *self-stimulating attack*, we play a malicious music file from a smartphone's speaker to control the on-board MEMS accelerometer trusted by a local app to pilot a toy RC car. In addition to offering hardware design suggestions to eliminate the root causes of insecure amplification and filtering, we introduce two low-cost software defenses that mitigate output biasing attacks: *randomized sampling* and *180° out-of-phase sampling*. These software-only approaches mitigate attacks by exploiting the periodic and predictable nature of the malicious acoustic interference signal. Our results call into question the wisdom of allowing microprocessors and embedded systems to blindly trust that hardware abstractions alone will ensure the integrity of sensor outputs.

*Corresponding faculty author.

1. Introduction

With the proliferation of motion-driven applications and Microelectromechanical systems (MEMS) technologies, MEMS accelerometers have been widely used in cyber-physical systems, such as implantable medical devices, automobiles, avionics, and even critical industrial systems [1], [2], [3], [4], [5], [6]. These systems deploy layers of software that abstract away hardware details to collect and analyze data provided by sensors, and then autonomously react to sensor data in real time. The software assumes that the underlying hardware is behaving according to specification, and the common practice is to inherently trust the output from sensors. After years of effort towards encouraging better security practices in software, developers are becoming more diligent in hardening software to security vulnerabilities, but fewer methodologies exist in the sensor hardware domain.

It is already known that acoustic interference can cause denial of service (DoS) attacks against MEMS gyroscopes [7]. Building upon this previous knowledge, our paper questions current assumptions about the *integrity* of sensory data, and specifically explores the data integrity of *MEMS accelerometers* with a focus of answering the following questions: (1) How can an adversary achieve fine grained control over a sensor's output? (2) How well will system software cope with untrustworthy measurement of motion? (3) How could sensors be designed differently to eliminate the integrity issues? (4) What can be done to protect legacy sensors? Answering these questions is challenging yet critical to securing cyber-physical systems, and the learned insights can guide future design choices and methodologies to mitigate security risks introduced by deploying MEMS sensors in cyber-physical systems.

MEMS accelerometers have a sensing mass, connected to springs, that is displaced when the sensor is accelerated. Acoustic waves propagate through the air, and exhibit forces on physical objects in their path. If the acoustic frequency is tuned correctly, it can vibrate the accelerometer's sensing mass, altering the sensor's output in a predictable way. To systematically analyze the vulnerabilities of MEMS accelerometers, we model the impact of acoustic interference on the sensor's entire architecture, including both the sensing mass and signal conditioning components. We identify

IEEE computer society

two problematic components in the signal conditioning path of typical MEMS accelerometers (i.e., insecure low-pass filters and insecure amplifiers) that lead to two types of adulterated outputs: fluctuating measurements and constant measurements. These two components not only explain the root cause of DoS attacks [7] but also enable us to design two additional attack classes: sensor *output biasing* and *output control* that permit increasing levels of adversarial control over the output of MEMS accelerometers. Of the 20 models of accelerometers we tested, our experiments show that 75% are vulnerable to output biasing attacks (i.e., insecure low pass filters enable false fluctuating output measurements under acoustic interference), and 65% are vulnerable to output control attacks (i.e., insecure amplifiers enable false constant output measurements under acoustic interference). At the software system level, our experiments demonstrate the ease of injecting acoustic interference into an Android smartphone's accelerometer to take control of an app that drives an RC car. We also demonstrate a proof of concept end-to-end acoustic attack by injecting 3,000 steps per hour into a Fitbit. The results confirm our concerns that system software does not adequately validate the integrity of sensory data—blindly trusting the output of sensors by default.

Defending against malicious acoustic interference by applying acoustic dampening materials to sensors was previously investigated [7], [8]. Other defense mechanisms exist to thwart sensor-spoofing attacks in scenarios where the actuator and sensor operate in tandem [9]. Other common approaches to deal with signal interference include averaging or filtering. All of these techniques are either impractical (increases packaging size), not applicable (the sensor must operate with an actuator in a closed loop system), or insufficient (cannot filter out all interference) in defending against all proposed acoustic injection attacks. Therefore, we offer two types of defenses: (1) hardware solutions, whereby the acoustic injection attacks can be eliminated if the MEMS sensors are designed with security in mind, i.e., each component on the signal conditioning path is chosen with larger operation parameters, and (2) software solutions for retroactively protecting vulnerable MEMS accelerometers already deployed in various devices and systems. We evaluate our software defense mechanisms on vulnerable MEMS accelerometers showing that output biasing attacks can be mitigated.

Hardware vulnerabilities in MEMS accelerometers expose attack vectors that can compromise the integrity of the autonomous decision making path of cyber-physical systems. Our contributions address security concerns in MEMS accelerometers:

- We model the adversarial physics of acoustic injection attacks on MEMS accelerometers and identify hardware constraints on the signal conditioning path that lead to adulterated sensor outputs.
- We design two types of component-level attack classes: *output biasing* and *output control* that exploit independent hardware design flaws in MEMS
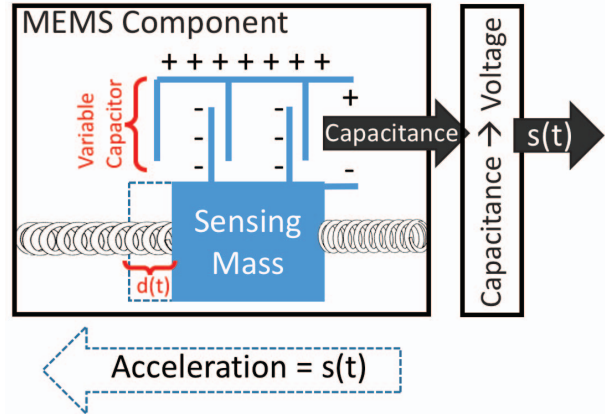

Figure 1. **Functional Diagram of Capacitive MEMS Accelerometer based on [10], [11].** When accelerated, the displacement of the mass creates an electrical signal due to a change in capacitance. The measured acceleration, $s(t)$, relates to the displacement of the mass, $d(t)$, according to Newton's second law of motion, $F = m \cdot a$, and Hooke's law, $F = -k_s \cdot d$.

accelerometers. We test our attacks on 20 different MEMS accelerometer models and show that 75% are vulnerable to output biasing attacks, and 65% are vulnerable to output control attacks. We also demonstrate a proof of concept, end-to-end, acoustic attack on two vulnerable systems pertaining to health, wellness, and the Internet of Things.

- We suggest hardware design practices that can increase the difficulty required to successfully mount an acoustic injection attack. We propose and measure the effectiveness of two low-cost software defenses that mitigate acoustic injection attacks: *randomized sampling* and *180° out-of-phase sampling*.

2. Background

In this work, we focus on a specific accelerometer: the capacitive MEMS accelerometer. Capacitive MEMS accelerometers are traditionally implemented using a variable capacitive structure [10], [11], as shown in Figure 1, and are manufactured using MEMS technology: a process by which micro-mechanical structures are machined into integrated circuit (IC) packages along with other electrical components. These sensors measure acceleration using the displacement of a mass connected to springs. This displacement is translated to a continuous voltage signal. In accordance with Newton's second law of motion, $F = m \cdot a$, and Hooke's law, $F = -k_s \cdot d$, the acceleration voltage signal is: $a = \frac{-k_s d}{m}$.

Additional processing is required for the electrical acceleration signals to interface with components external to the accelerometer, e.g. microprocessors. Figure 2 illustrates a typical design of the signal conditioning path in a MEMS accelerometer [10]. Prior to *digitization* via an Analog-to-Digital Converter (ADC, component D in Figure 2), analog signal are typically *amplified* (component C in Figure 2) and *low-pass filtered* (LPF, component D in Figure 2). Like any

4

Figure 2. **Typical architecture of the signal conditioning path in a MEMS accelerometer based on [10].** The change in capacitance measured by a sensing mass (Fig. 1) is converted to a voltage, amplified, filtered, and digitized. Without stage D aliasing can occur, enabling *output biasing* attacks. Signal clipping at C can introduce a DC component into the acceleration signal, enabling *output control* attacks.

Figure 3. **Acoustic Interference Disturbs Acceleration Measurements.** True acceleration and acoustic interference can both displace the mass, creating electrical acceleration signals. The measured acceleration, $\hat{s}(t)$, is a linear combination of true acceleration, $s(t)$, and acoustic acceleration, $s_a(t)$.

circuit component, the amplifier and ADC have limitations. Amplifiers have upper and lower bounds; when the input signal exceeds these bounds, *signal clipping* occurs, and abnormal acceleration readings are reported. Likewise, the ADC has requirements that must be met. According to the Nyquist sampling theorem, a minimum sampling rate is required to avoid misinterpreting an analog signal represented in digital form, also known as *signal aliasing*. Therefore, it is common practice to place an LPF prior to an ADC, to filter out high frequency signal components and enforce the Nyquist requirement.

Both analog and digital accelerometers are available on the market. Analog accelerometers output the analog signals from the amplifier directly, while digital sensors typically contain an LPF and ADC. We use an analog sensor to help us understand how acoustic waves interact with the sensing mass–spring structure.

3. Threat Model

We assume the goal of an attacker is to hijack the control of the systems that are driven by acceleration sensor data.

Attack Scope. Access controls to sensor data have tightened because of privacy concerns raised by previous research [12], [13], [14], [15], [16]. Thus, we assume that attackers can neither directly access the digitized sensor readings nor physically touch the sensors. Instead, we assume attackers exploit vulnerabilities by emitting nearby acoustics to affect the integrity of sensor data, i.e., analog signals on the signal conditioning path before being digitized.

Sensor Access. Although we assume attackers do not gain physical access to a specific targeted device containing a MEMS accelerometer, we do allow an adversary to gain access to a substantially identical device to study acoustic attack capabilities. In our attacks, we do not assume the more powerful adversary such as a lunch-time attack where an adversary has temporary physical access. However, we assume the attacker is able to reverse engineer a sample device to extract the exact model of MEMS accelerometer and profile the accelerometer's behavior under different acoustic frequencies and amplitudes. This leads to a key question to the success of the attacks: *to what extent will two instances of the same device behave in a similar way when they are subjected to the same acoustic signals?*
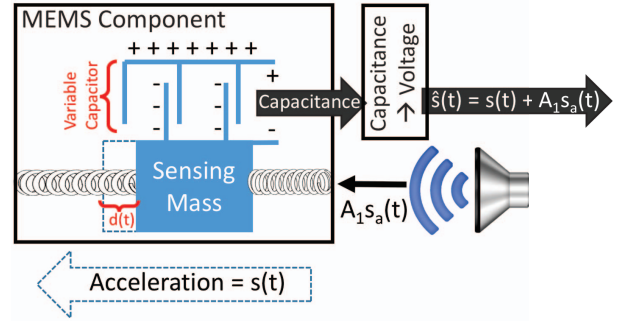
Speaker Access. We assume that the attacker is able to induce sound in the vicinity of the victim device, at frequencies in the human audible to ultrasonic range (2–30 kHz). This can be done by applying the sound externally, or by playing sounds from speaker in the vicinity of the target sensors. This might be done via means of remote software exploitation (e.g., remotely affecting the multimedia software in a phone or a car) or by a *drive-by ditty* where a user is tricked into playing malicious music either by email or a web page with autoplay audio enabled. The attacker is also able to synthesize any shape, i.e., varying amplitude and phase, of acoustic signal within the stated frequency range.

4. Attack Modeling

Acoustic attacks are possible because capacitive MEMS accelerometers use the displacement of a mass as a proxy for measuring acceleration. Figure 3 shows the MEMS component of a typical accelerometer. When the sensing mass is displaced, an electrical signal is generated, $\hat{s}(t)$. Primarily, the mass is displaced by forces resulting from true acceleration (i.e., physical motion). However, forces from acoustic pressure waves can also displace the mass. Because of this, we denote electrical acceleration signals generated by true acceleration: $s(t)$, and those generated by acoustic interference: $s_a(t)$. Using these representations, we model how acoustic interference impacts the electrical acceleration signals generated by MEMS accelerometers, and validate our model. Then, we describe the goal of acoustic injection attacks and provide an overview of the challenges of conducting these attacks.

4.1. Modeling Acoustic Effects on Accelerometers

We develop a model for how an electrical acceleration signal, generated by a capacitive MEMS accelerometer, is distorted by acoustic noise. We model the measured acceleration as a linear combination of the *true acceleration* and *acoustic acceleration*. Namely, for a true acceleration signal

5

Figure 4. **Experimental Setup for Evaluating our Model of Electrical Acceleration Signal Generation.** The accelerometer is simultaneously subjected to both a 70 Hz vibration (*true acceleration*) and to 2.9 kHz acoustic noise (to generate *acoustic acceleration*). Note: a modification of this setup, removing the vibrating platform, is used for experiments in Sections 5–8.


Figure 5. **Acceleration Signal Generation Model.** a) A 70 Hz sinusoidal mechanical vibration signal stimulates *true acceleration*. b) A sinusoidal, on–off modulated, acoustic interference signal stimulates *acoustic acceleration*. c) Using the experimental setup in Fig. 4, a MEMS accelerometer subjected to both mechanical vibrations (true acceleration) and acoustic interference (acoustic acceleration) outputs an acceleration signal that is a linear combination of the stimuli signals shown in (a) and (b). Note that all plots are in the time domain.

$s(t)$, and the acoustic acceleration signal $s_a(t)$, the measured acceleration signal $\hat{s}(t)$ is:

$$\hat{s}(t) = s(t) + A_1 \cdot s_a(t) \qquad (1)$$

A_1 is the attenuation of the acoustics in transit to the target device. For an acoustic frequency F_a, played at amplitude A_0, and phase ϕ, the acoustic acceleration generated is modeled as $s_a(t) = A_0 \cdot cos(2\pi F_a t + \phi)$. Therefore the measured acceleration is:

$$\hat{s}(t) = s(t) + A_1 A_0 \cdot cos(2\pi F_a t + \phi) \qquad (2)$$

Evaluating the Model. We evaluate the model, in Equation 2, with the experimental setup shown in Figure 4. An analog MEMS accelerometer, the ADXL337, was placed on top of a vibration platform vibrating at 70 Hz, simulating an example of *true acceleration* on the sensor. An off-the-shelf tweeter speaker [17] was suspended 10 cm above the sensor to decouple the sensor from mechanical vibrations emanating from the speaker. The output of the sensor was sampled by an Arduino microcontroller's ADC at a sampling rate of 7 kHz. The samples were logged by a computer connected to the Arduino. The experimental setup was placed inside an acoustic isolation chamber to avoid external noise. Outside the chamber, a commodity audio amplifier [18] amplified a 2.9 kHz acoustic signal that was supplied to the speaker. To allow visual distinction between the true acceleration and acoustically stimulated acceleration, the acoustic signal was on/off modulated at 0.5 Hz.

Results. Figure 5a depicts the 70 Hz sinusoidal physical vibration signal input to the vibrating platform. Figure 5b shows the sinusoidal, on–off modulated, acoustic interference signal input to the speaker. Figure 5c depicts the acceleration signal measured when the acoustic noise is played in conjunction with the 70 Hz vibration. The measured acceleration is a linear combination of the true acceleration and artificial acoustic acceleration, supporting our model.

4.2. Maximizing the Acoustic Disturbance

The goal of an attacker is to maximize the acoustic disturbance on MEMS accelerometers, or maximize the

attenuation coefficient, A_1, in our model. The attenuation coefficient, A_1, is a function of acoustic frequencies. Physics allows the attacker to achieve the maximum acoustic disturbance by exploiting a mechanical property of a vibrating mass–spring system — *resonance*. Vibrating these systems at their resonant frequencies achieves maximum displacement of the mass, i.e., $A_1 = 1$. *To substantially displace the sensing mass using acoustics, the acoustic frequency must match the mechanical resonant frequency of the sensor.* For the previous experiment, 2.9 kHz was the resonant frequency of the ADXL337.

5. Acoustic Attack Building Blocks

Based on our model, it seems plausible an attacker may use acoustics to spoof output measurements from MEMS accelerometers, and tamper with systems that utilize such sensors. However, there are several challenges:

- **Process Variation:** The attacker can obtain a different instance of the exact model of accelerometer to determine its resonant frequency. *How do resonant frequencies of MEMS accelerometers vary with process variation? Or is the resonant frequency characteristic of each model similar?*

- **Controlling the Artificial Acceleration**: As our model shows, acceleration signals resulting from acoustic interference are of the same frequency as the acoustic waves that created them. *How do the artificial acceleration signals get distorted or removed by downstream signal conditioning components? How can an attacker leverage the predictability of acoustic acceleration to achieve fine grained control over an accelerometer's output?*

- **Altering the Behavior of Software through Accelerometers**: *How can an attacker influence the behavior of software that takes input from an accelerometer?*
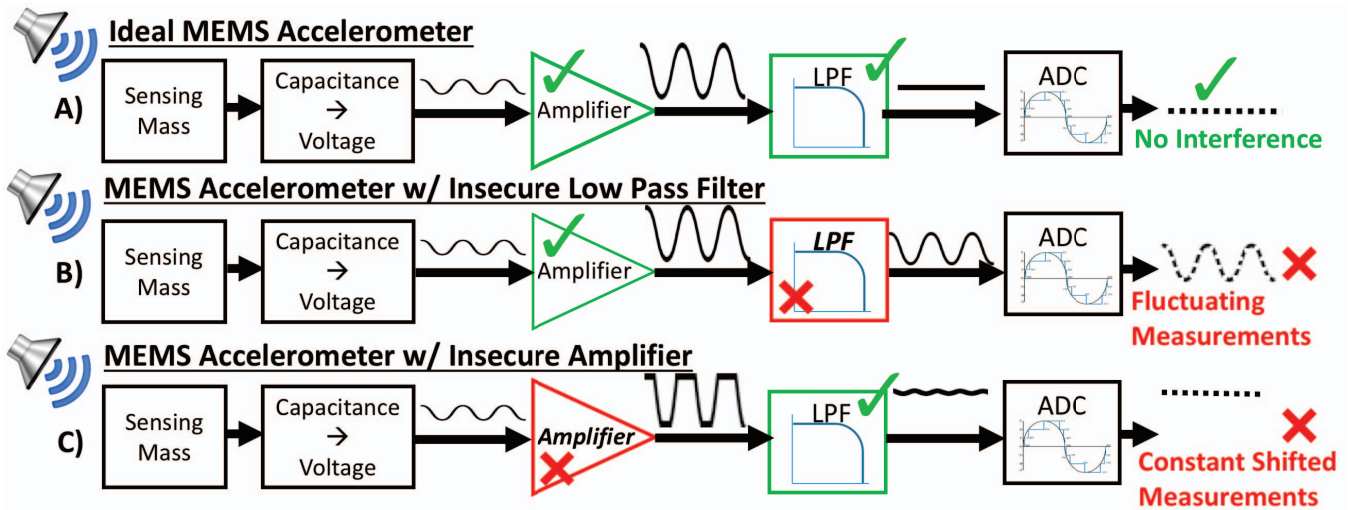
6

Figure 6. **Examples of Signal Distortion from Hardware Deficiencies**: A) An ideal MEMS accelerometer filters out all high frequency interference at its low pass filter (LPF), removing all acoustically generated acceleration signals. B) An accelerometer with an insecure LPF does not fully suppress acoustic acceleration signals and the digital output measurements are sinusoidally fluctuating. C) An insecure amplifier asymmetrically clips high amplitude acoustic acceleration signals, introducing a DC component into the amplified signal. The DC component is not removed by the LPF. However, the LPF removes the high frequency components present in the sharp corners of the clipped signal, resulting in a low-frequency, low-amplitude signal with non-zero DC offset. The digitized measurements are mostly constant and shifted.

Assuming a linear model of acceleration signal generation, this section predicts the impacts of downstream signal conditioning hardware on the digital representation of these signals. Our experiments show that because of security deficiencies in an accelerometer's signal conditioning hardware, digitized acoustic acceleration measurements may manifest themselves in two ways: *fluctuating acceleration* as if the chip is under high vibration and *constant shifted acceleration* as if the chip is on a launching rocket. These two types of falsified output will serve as the building blocks for the full-fledged attacks.

5.1. Signal Conditioning Hardware Deficiencies

The two critical hardware components typically included in a MEMS accelerometer's signal conditioning path are: an *amplifier* and a *low pass filter (LPF)*, components C and D in Figure 2 respectively.

In an ideal case — when the amplifier and LPF work perfectly — any injected acoustic acceleration signals are removed by the signal conditioning hardware before being digitized and do not pass through to end systems, as show in Figure 6a. However, in reality these components have physical limits. Specifically, each accelerometer has a limit regarding the maximum amplitude and frequency of acceleration it can measure. Exceeding these limits distorts their acceleration measurements.

Low Pass Filter. To prevent high frequency noise from contaminating ADC samples, designers typically include an analog LPF before the ADC (component D in Figure 2). An ideal analog LPF filters out all frequencies above a designated cutoff frequency, F_{cutoff}, while passing all frequencies below. To enforce the Nyquist requirement, LPFs are designed to only pass frequencies which are half that of the ADC's sampling rate, F_s, i.e., $F_{\text{cutoff}} = \frac{1}{2}F_s$. However, in practice, it is impossible to manufacture an LPF that passes all frequencies up to F_{cutoff} (e.g., *exactly* half the sampling frequency) and completely blocks all frequencies above F_{cutoff}. Instead, there is a range of frequencies around F_{cutoff} which are attenuated but not completely removed. Acoustic acceleration signals can be affected by the LPF in one of two ways:

1) **Insecure LPF**: The accelerometer's LPF is designed with a cut-off frequency that is either above, or too close to the resonant frequency of the sensor. The sinusoidal acoustic acceleration signal, whose frequency matches the accelerometer's resonant frequency, is not completely attenuated by the LPF. It slips through to the ADC where it is usually undersampled, as shown in Figure 6b.
2) **Secure LPF**: The acoustic acceleration signal's frequency is well above the cut-off frequency of the LPF and is completely attenuated.

Acoustic acceleration signals directly correspond to the acoustic frequency which generated them (Section 4.1). If the LPF is insecurely designed (1) the false output acceleration measurements will be sinusoidally *fluctuating*.

Amplifier. Ideally, the input range of the amplifier is large enough to handle any signal the sensing mass can produce. In reality, the amplifier is typically chosen to cope with the maximum specified acceleration. This exposes an attack surface. Resonant acoustic interference can displace the sensing mass enough to create a high amplitude acceleration signal that exceeds the dynamic range of the amplifier.

Thus, acoustic acceleration signals can be potentially be distorted. We classify two types of amplifiers:

1) **Insecure Amplifier:** Previous research has shown MEMS accelerometers to report false measurements when signal clipping occurs from exceeding the dynamic range of its amplifier [10], [19], [20]. The causality stems from the introduction of a DC component into the output signal of the saturated amplifier, as illustrated in Figure 6c. This DC component is not removed by the LPF, however, the sharp clipped edges, i.e., the high frequency components, are attenuated. Additionally, when the accelerometer's LPF is securely designed, i.e., the cutoff frequency is much lower than the resonant frequency, the non-clipped portion of the acoustic acceleration signal is also attenuated. Given the construction of the amplifiers [20], clipping can be asymmetrical, and what slips through to the ADC *resembles* a low-amplitude sinusoid with non-zero DC offset. The digital output measurements are mostly constant and non-zero, as reported by [10].

2) **Secure Amplifier:** When the unamplified acceleration signal is within the dynamic range of the amplifier, clipping does not occur. The acceleration signal remains undistorted.

In summary, under resonant acoustic interference the sensor may report three types of measurements: true measurements and two types of falsified measurements. The false sensor measurements are due to insecurities in hardware components, as shown in Figure 6:

A) **True Measurements**: The accelerometer's amplifier tolerates the high amplitude acceleration signals generated under resonant acoustic interference, i.e., no signal clipping occurs. The accelerometer's resonant frequency is much greater than the LPF's cutoff frequency. The LPF attenuates high frequency acoustic acceleration signals.

B) **Fluctuating False Measurements**: No signal clipping is observed at the amplifier. The LPF does not completely attenuate high frequency acoustic acceleration signals. Acoustic acceleration signals are under-sampled by the ADC.

C) **Constant Shifted False Measurements**: Signal clipping occurs at the amplifier introducing a non-zero DC component into the amplified signal. A securely designed LPF passes DC signals and blocks high frequency signals. A mostly constant, non-zero, signal is sampled by the ADC.

Recall that acoustic acceleration is only generated when the sound waves displace the sensing mass, i.e., when the acoustic frequency matches the resonant frequency of the sensing structure. Only then will fluctuating (2) and constant (3) false measurements be observed. Conversely, resonant frequencies can be identified when accelerometers exhibit these phenomena. We test 40 widely used MEMS accelerometers to experimentally demonstrate the above behaviors MEMS accelerometers exhibit when their acoustic resonant frequencies are played.

5.2. Finding Resonant Frequencies

A sensor at rest should measure constant acceleration of $0\,g$ along the X and Y axes and $1\,g$ along the Z axis, accounting for gravity. At a given frequency, if output measurements deviate from normal, i.e., they are *fluctuating* or *constantly shifted*, that frequency is considered a resonant frequency. By sweeping an acoustic frequency range and acquiring several acceleration measurements at each frequency, both scenarios can be observed. Fluctuating measurements are observable by calculating the standard deviations of multiple samples at each frequency. Constant shifted measurements are observable by calculating the means of multiple samples taken at each frequency.

We survey 40 widely used MEMS accelerometers: 2 instances each of 20 different models from 5 different manufacturers, including both analog and digital sensors, to determine their resonant frequencies. A frequency where the standard deviation or mean deviates from normal by at least $0.1g$, less than 5% of the typical noise margin, is classified a resonant frequency of that accelerometer model.

5.3. Experimental Setup

Our experimental setup is identical to the setup in Figure 4, absent the vibrating platform. All 40 MEMS accelerometers (both digital and analog) were attached to a table. The experiments were conducted in an acoustic isolation chamber to avoid external acoustic effects. Each sensor was oriented to experience $0\,g$ along the X and Y axes and $1\,g$ along the Z axis, due to gravity. Outside the chamber, a commodity audio amplifier [18] amplified single frequency acoustic signals generated by a function generator. The amplifier drove an off-the-shelf tweeter speaker [17] inside the acoustic chamber. The speaker was suspended 10 cm above the sensor to decouple the sensor from mechanical vibrations emanating from the speaker. All digital accelerometers were connected via a serial peripheral interface (SPI) or inter-integrated circuit (I2C) bus to an Arduino microcontroller running a driver program. Analog accelerometers were sampled using the Arduino's ADC. While at rest, each accelerometer was subjected to single tone acoustic frequencies from 2 kHz–30 kHz, at 50 Hz intervals. At each frequency interval, 256 acceleration readings were acquired along all possible axes at a sampling rate of at least 400 Hz. As a baseline, 256 acceleration readings were also acquired without sound. All acceleration samples were logged by a Python script running on a computer connected to the Arduino microcontroller.

To determine the resonant frequencies, the speaker was operated near its maximum amplitude, around 110 dB Sound Pressure Level (SPL). To ensure that the speaker produced all sounds at similar SPL, we validated the speaker's frequency response using a measurement microphone with a
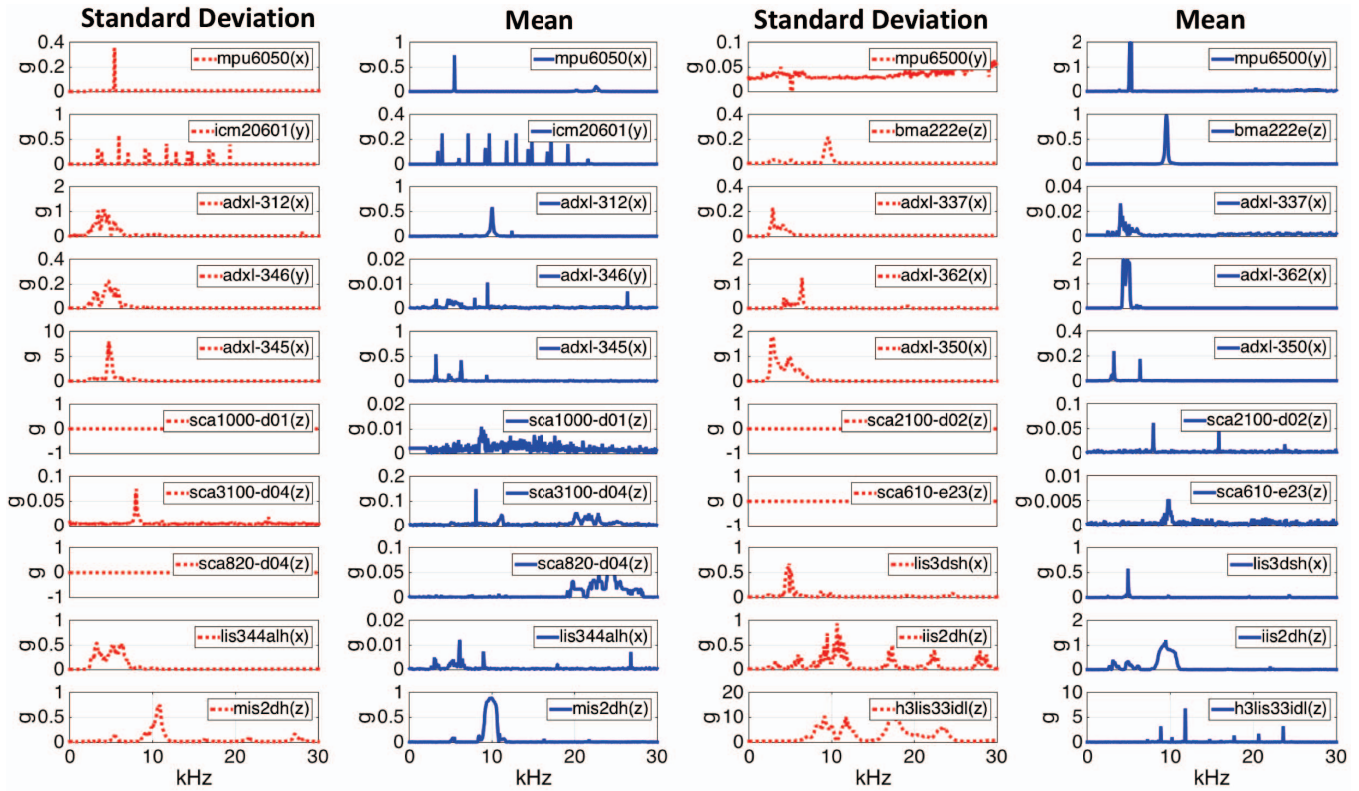
8

Figure 7. **Effects of Acoustic Interference at Different Frequencies on the Behavior of Accelerometer Output Measurements.** Peaks in either standard deviation (dotted red) or mean (solid blue) indicate the acoustic interference caused the accelerometer to generate false acceleration measurements. Acoustic frequencies where false measurements are observed are mechanical resonant frequencies of that accelerometer. Peaks in standard deviation indicate the accelerometer has an insecure LPF (Fig. 6b), resulting in sinusoidally *fluctuating* false measurements. Peaks in mean indicate the accelerometer has an insecure amplifier (Fig. 6c), resulting in *constant shifted* false measurements. The axis of acceleration displayed in each plot is in each respective legend. The mean plots are normalized to the value of acceleration when the accelerometer was at rest (0 g if along X or Y axis; 1 g if along Z axis).

frequency response of 4 Hz–100 kHz [21]. The speaker's frequency response was relatively flat (at 110 db SPL) across its entire range, from 1.8 kHz to 30 kHz.

5.4. Results

The means and standard deviations of the 256 raw data samples taken at each frequency interval are plotted in Figure 7. Of the 20 sensor models we tested, 15 exhibited standard deviation spikes of at least $0.1g$ and 13 experienced mean spikes of at least $0.1g$. We observe the following from these results:

1) Both instances of the same sensor model behaved identically. Therefore, the results of only a single instance of each sensor model is shown in Figure 7.
2) Resonant frequencies can fall in a range, not only a single frequency.
3) Several sensors have multiple resonant frequencies.
4) Several sensors have resonant frequencies which result in all combinations of *constant shifted measurements* (mean spike) and/or *fluctuating measurements* (standard deviation spike).
5) Most sensors that were not affected by acoustic interference are physically larger than sensors that

were affected. This indicates the MEMS feature size may affect its susceptibility to acoustic interference.

In summary, acoustic resonant frequencies stimulate MEMS accelerometers to output false measurements that are either *fluctuating* or *constantly shifted*.

6. Controlling Accelerometer Output

Although the ultimate goal of an adversary is to control a sensor-driven autonomous system, an intermediate goal is to demonstrate direct control of the digital time series data output by a sensor. Thus, we ask the following question: *Given a function that represents the desired sensor output signal, how does one design acoustic interference to mimic said function?* In this section, we show how to utilize the predictability of both types of false measurements (fluctuating or constant) to control the time series output of a sensor. Our key contribution is the identification of two distinct classes of acoustic injection attacks, *output biasing* and *output control* attacks based on controlling fluctuating or constant false measurements, respectively. Table 1 summarizes our results on the extent to which sensors are vulnerable to what attack.

TABLE 1. ACCELEROMETER RESONANT FREQUENCIES: UNDER RESONANT ACOUSTIC INTERFERENCE, AN OUTPUT BIASING ATTACK CLASS INDICATES A SENSOR'S FALSIFIED MEASUREMENTS FLUCTUATE (INSECURE LPF) WHILE AN OUTPUT CONTROL ATTACK CLASS INDICATES CONSTANT FALSIFIED MEASUREMENTS ARE OBSERVED (INSECURE AMPLIFIER). TWO INSTANCES OF EACH SENSOR WERE TESTED.

Model	Type	Typical Usage	Resonant Frequency (kHz)			Amplitude (g)∗	Attack Class‡		
			X	Y	Z		X	Y	Z
Bosch - BMA222E	Digital	Mobile devices, Fitness	5.1–5.35	–	9.4–9.7	1	B	–	BC
STM - MIS2DH	Digital	Pacemakers, Neurostims	–	–	8.7–10.7	1	–	–	BC
STM - IIS2DH	Digital	Anti-theft, Industrial	–	–	8.4–10.8, ...	1.2	–	–	BC
STM - LIS3DSH	Digital	Gaming, Fitness	4.4–5.2	4.4–5.6	9.8–10.2	1.6	BC	BC	BC
STM - LIS344ALH	Analog	Antitheft, Gaming	2.2–6.6	2.2–5.7	2.2–5.6	0.6	B	B	B
STM - H3LIS331DL	Digital	Shock detection	–	–	11–13, ...	5.2	–	–	BC
INVN - MPU6050	Digital	Mobile devices, Fitness	5.35	–	–	0.75	BC	–	–
INVN - MPU6500	Digital	Mobile devices, Fitness	5.1, 20.3	5.1–5.3	–	1.9	BC	C	–
INVN - ICM20601	Digital	Mobile devices, Fitness	3.8, ...	3.3, ...	3.6, ...	1.1	BC	BC	BC
ADI - ADXL312	Digital	Car Alarm, Hill Start Aid	3.2–5.4	2.95–4.75	9.5–10.1	1.3	B	B	BC
ADI - ADXL337	Analog	Fitness, HDDs	2.85–3.1	3.8–4.4	–	0.8	B	B	–
ADI - ADXL345	Digital	Defense, Aerospace	4.4-5.4	3.1–6.8	4.4–4.7	7.9	BC	BC	B
ADI - ADXL346	Digital	Medical, HDDs	4.3–5.1	6.1	4.95, ...	1.75	B	B	B
ADI - ADXL350	Digital	Mobile devices, Medical	2.5–6.3	2.5–4	2.5–6.8	1.8	B	B	B
ADI - ADXL362	Digital	Hearing Aids	4.2–6.5, ...	4.3–6.5, ...	4.5–6.5	1.4	BC	BC	BC
Murata - SCA610	Analog	Automotive	–	–	–	–	–	–	–
Murata - SCA820	Digital	Automotive	24.3	–	–	0.13	C	–	–
Murata - SCA1000	Digital	Automotive	–	–	–	–	–	–	–
Murata - SCA2100	Digital	Automotive	–	–	–	–	–	–	–
Murata - SCA3100	Digital	Automotive	7.95	–	8	0.15	C	–	C

∗ Amplitude is taken as the maximum false output measurement observed.   – Experiments found no resonance
‡ **B** = Output Biasing Attack; **C** = Output Control Attack (Red Highlight)   ... Additional ranges of resonance elided
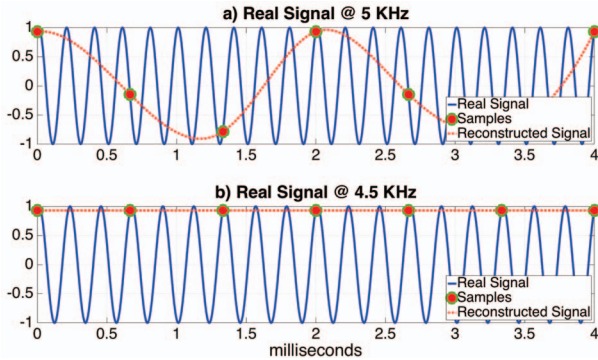STM = STMicroelectronics; ADI = Analog Devices; INVN = InvenSense

Figure 8. **Examples of Signal Aliasing.** A) Sampling a 5 kHz analog signal at 1.5 kHz results in a 500 Hz signal alias. B) A 1.5 kHz sampling rate applied to a 4.5 kHz (integer multiple) analog signal yields a constant DC (0 Hz) alias.

Figure 9. **Examples of Amplitude and Phase Modulation (AM and PM).** A) Amplitude modulation (AM) encodes the information signal in the envelope of the carrier. *AM acoustic interference can only spoof either all positive or negative acceleration.* B) Phase modulation (PM) encodes the information signal in the phase of the carrier. *Unlike AM, PM allows an attacker to utilize the full range of the carrier, and therefore spoof both positive and negative acceleration signals.*

6.1. Output Biasing Attack

The output biasing attack utilizes sampling deficiencies at the ADC and gives an adversary control over the accelerometer's output for several seconds. This attack pertains to accelerometers that experience *fluctuating* false measurements at their resonant frequencies due to insecure LPFs (Figure 6b). To perform an output biasing attack, an adversary must accomplish two goals:

1) **Stabilize** fluctuating false measurements into constant measurements by shifting the acoustic resonant frequency to induce a DC alias at the ADC.
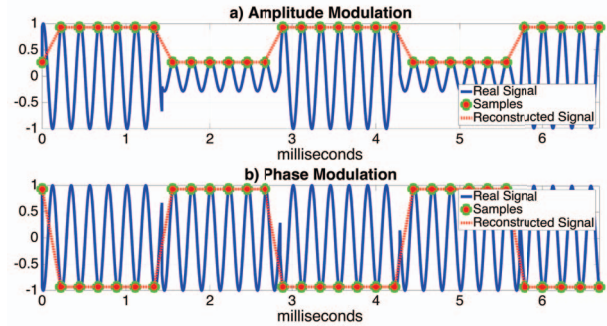2) **Reshape** the desired output signal by modulating it on top of the acoustic resonant frequency.

The first step can be accomplished through *signal aliasing*. The second step can be realized with *signal modulation*.

Signal Aliasing. Aliasing is the misinterpretation of an analog signal caused by digitizing it with an inadequate sampling rate. According to the Nyquist sampling theorem, an analog signal with maximum frequency component F_{max} must be sampled at a minimum rate of $2 \cdot F_{max}$ to avoid signal aliasing. Figure 8a illustrates aliasing with a 5 kHz sinusoid and a sampling rate of 1.5 kHz. Reconstructing this signal from the digital samples results in a 500 Hz aliased signal. When the frequency of the analog signal is an integer multiple of the sampling frequency, a constant DC (direct

10

Figure 10. **Spelling WALNUT: Output Biasing Attack on Sensors with Inaccurate ADCs**. We demonstrate the output biasing attack can control the X-axis acceleration signals of the A) ADXL350 and B) ADXL345 accelerometers for over a second, spoofing the sensor to spell out "WAL-NUT". This attack leverages a security-flaw in the low pass filters of specific accelerometers. Each accelerometer was positioned with the Z-axis aligned with gravity, so the X-axis output should have measured 0 g. Sensors with inaccurate ADCs cause the acoustically stimulated acceleration signal to inconsistently alias to varying almost-DC signals, hence the WALNUT signal is slightly distorted.

Figure 11. **Spelling WALNUT: Output Biasing Attack on Sensors with Accurate ADCs**. We demonstrate the output biasing attack can control the X-axis acceleration signals output from the A) ADXL337 and B) LIS344ALH accelerometers for over 5 seconds, spoofing the sensor to spell out "WALNUT". This attack leverages a security-flaw in the low pass filters of specific accelerometers. Each accelerometer was positioned with the Z-axis aligned with gravity, so the X-axis output should have measured 0 g. Given an insecure LPF, sensors with accurate ADCs are more vulnerable than those with inaccurate ADCs because an attacker can more easily guess the sampling phase when it is stable, hence the WALNUT signal is less distorted than in Fig. 10

current, 0 Hz) alias is encountered. Figure 8b illustrates this phenomenon with a 4.5 kHz sinusoid sampled at 1.5 kHz.

Signal Modulation. Signal modulation is used to transmit arbitrary information signals over another carrier signal. Here we focus on *amplitude* and *phase* modulation, which utilize constant frequency carrier signals. Assume a sinusoidal carrier signal $f_c(t) = A \cdot sin(2\pi t f + \phi)$, with t the time, f the frequency, and ϕ a constant phase offset:

1) **Amplitude Modulation** (AM) consists of varying the amplitude, A, of the carrier signal over time according to the amplitude of the information signal being transmitted. The amplitude, A, becomes a time-domain function, $A(t)$, resulting in the modulated signal: $S_{AM} = A(t) \cdot sin(2\pi t f + \phi)$. Figure 9a illustrates amplitude modulating a square wave on top of a sinusoidal carrier frequency, f_c.

2) **Phase Modulation** (PM) consists of varying the phase, ϕ, of the carrier signal over time according to the amplitude of the information signal being transmitted. The phase ϕ becomes a time-domain function, $\phi(t)$, resulting in the modulated signal $S_{PM}(t) = A \cdot sin(2\pi t f + \phi(t))$. Figure 9b illustrates phase modulating a square wave on top of a sinusoidal carrier frequency, f_c.

Biasing the Output. Here we explain the two steps of the output biasing attack: 1) stabilize fluctuating false measurements by producing a DC alias of the acoustic acceleration signal, and 2) modulate the desired accelerometer output signal over the acoustic resonant frequency. We demonstrate the output biasing attack by spoofing a MEMS accelerometer to output a signal spelling "WALNUT".

Step 1) Converting the fluctuating false measurements into constant false measurements is accomplished by inducing a DC alias of the acceleration signal at the ADC

(Figure 8b). A DC alias of a periodic analog signal is observed if the analog signal's frequency is an integer multiple of the sampling frequency, F_{samp}. An accelerometer's ADC sampling rate, F_{samp}, is fixed. The sampling times at discrete intervals k, can be denoted $t_k = k \cdot \frac{1}{F_{samp}}$. *Given the resonant frequencies of a MEMS accelerometers are often not a single frequency, but a range, an attacker can find a small frequency deviation f_ϵ such that the acoustic frequency $F_a = F_{res} + f_\epsilon$ is still within the resonance zone.* Selecting F_a in a way that it is an integer multiple of the sampling rate, F_{samp}, results in a DC alias, shifting the output of the sensor to a constant value. Therefore, if $F_a = F_{res} + f_\epsilon = N \cdot F_{samp}$ where $N \in \{1, 2, 3...\}$, the measured acceleration signal is then:

$$\begin{aligned}
\hat{s}(t_k) &= s(t_k) + A_1 \cdot s_a(t_k) \\
&= s(t_k) + A_1 A_0 \cdot cos(2\pi F_a t_k + \phi) \\
&= s(t_k) + A_1 A_0 \cdot cos(2\pi N k + \phi) \\
&= s(t_k) + A_1 A_0 \cdot cos(\phi)
\end{aligned} \quad (3)$$

For example, if the resonant frequency and sampling rate are $F_{res} = 3280$ Hz, $F_{samp} = 150$ Hz, one can select the deviation to be $f_\epsilon = 20$ Hz, such that $F_a = 3280 + 20 = 3300 = 22 \cdot F_{samp}$, to achieve a DC-aliased time series output.

Step 2) The attacker employs either amplitude or phase modulation techniques to further shape the output signal of the accelerometer. Regarding output biasing attacks, PM allows an attacker to use the full amplitude of the carrier frequency to modulate the desired signal, where AM utilizes only the upper or lower half of the carrier signal (Figure 9). An attacker must use PM to stimulate an acceleration signal that has both negative and positive components.

11

Limitations. Note that an attacker can control the acoustic interference phase ϕ in a relative, but not absolute manner. They can increase or decrease the phase, but always relative to the sampling phase, ϕ_{samp}, which they do not control or know. Hand tuning ϕ to be synchronized with ϕ_{samp} requires feedback from the accelerometer under attack. Figure 8b illustrates that the maximum bias amplitude is reached when samples are taken at the peaks of the acoustically stimulated acceleration signal. The less ϕ_{samp} drifts over time, the more stable the attack. With some sensors, it is possible to tweak F_a so that the DC-aliased output is maintained for up to 30 seconds.

Evaluation. We evaluated the output biasing attack on all sensors that yielded fluctuating output measurements at their resonant frequencies (standard deviation spikes in Figure 7). The same experimental setup shown in Figure 4 was used, absent the vibrating platform. The acoustic interference frequency was adjusted around the resonant frequency, specific to each sensor, until the fluctuating measurements stabilized. Using a function generator, a piecewise-linear signal spelling "WALNUT" was modulated over the acoustic resonant frequencies.

Results — Sensors with an Inaccurate ADC. Figure 10 illustrates the output biasing attack on two digital accelerometers with inaccurate ADCs, the ADXL350 and ADXL345. Spoofed acceleration signals, spelling "WALNUT", with peak-to-peak amplitudes of 10 g, were achieved for 1–2 seconds. These accelerometers, and all digital accelerometers tested, had inaccurate ADCs that did not take samples at precise time intervals, i.e. ϕ_{samp} fluctuates. This limits an attackers ability to achieve control over a sensor's output for more than 1–2 seconds. Note that PM was used to output the "WALNUT" signal on the ADXL350, while AM was used on the ADXL345. As a result, the spoofed acceleration ranges from -5 g to 5 g using PM on the ADXL350, while the ADXL345 only sees acceleration in the positive range, 0 g to 10 g. AM can either spoof all positive or all negative acceleration, since only the upper or lower envelope of the AM carrier signal is utilized.

Results — Sensors with an Accurate ADC. Figure 11 illustrates the output biasing attack on two analog accelerometers interfaced with accurate ADCs, the ADXL337 and LIS344ALH. Spoofed acceleration signals spelling "WALNUT", with peak-to-peak amplitudes of 1 g, were achieved on both sensors for tens of seconds. These analog accelerometers were interfaced with accurate ADCs that took samples at precise time intervals. This made it easier to maintain a consistent DC-aliased output signal for several tens of seconds. PM was used to attack both sensors, simply to yield the highest peak-to-peak amplitude possible. *Note how the spoofed acceleration signals on sensors with accurate ADCs compares to the spoofed signals on sensors with inaccurate ADCs (Figure 11 vs. 10).*

6.2. Output Control Attack

The output control attack gives an adversary indefinite full control of an accelerometer's output. This attack is
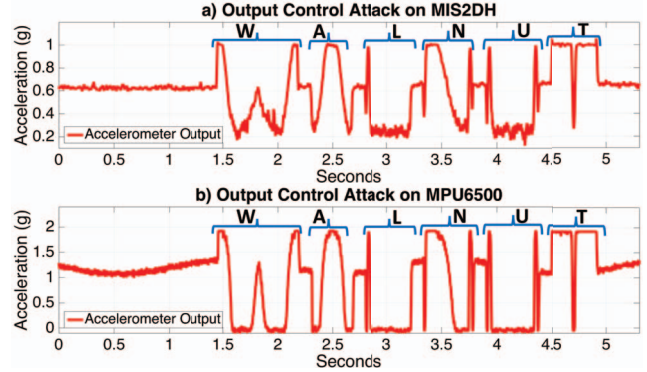

Figure 12. **Spelling WALNUT: Output Control Attack.** We demonstrate the output control attack achieving full indefinite control over the X-axis acceleration signals of the a) MIS2DSH and b) MPU6500 accelerometers, spoofing the sensor to spell out "WALNUT". Each accelerometer was positioned with the Z axis aligned with gravity, so the X axis output should have measured 0 g. This attack leverages a security-flaw in the amplifier of specific accelerometers. The attacker does not need to know anything about the sampling regime of the ADC, hence the WALNUT signal is the least distorted compared with Figs. 10 and 11.

applicable to accelerometers that exhibit constant shifted false measurements at their resonant frequencies due to insecure amplifiers (Figure 6c). No signal aliasing at the ADC is needed, since the false output measurements are already stable and constant. This allows an adversary to control the acceleration output indefinitely. To perform an output control attack, an adversary need accomplish one goal: reshape the desired sensor output signal by modulating it over the resonant frequency.

Controlling the Output. Achieving fine grain control over a sensor's output requires using amplitude modulation. Amplitude modulation modulates the amplitude of clipping at the amplifier, which is effectively demodulated by the LPF. Regardless of the ADC's sampling regime, an attacker has full control over the sensor's output. With PM, the amplitude of clipping does not change. Hence, AM yields a more effective attack.

Evaluation. We evaluated the output control attack on all sensors that demonstrated constant false output measurements (mean spikes in Figure 7). The same experimental setup shown in Figure 4 was used, absent the vibrating platform. A signal spelling "WALNUT" was amplitude modulated over each sensor's acoustic resonant frequency.

Results. Figure 12 illustrates the chosen output attack on two accelerometers tested, the MIS2DSH and MPU6500. Spoofed acceleration signals, spelling "WALNUT" with peak-to-peak amplitudes of up to 1 g were achieved on both sensors. Note how stable the acoustically stimulated output signal is compared with the signals spoofed by output biasing attacks in Figures 10 and 11.

7. Attacking Embedded Devices

The ultimate goal of an attacker is to leverage accelerometer hardware vulnerabilities to stealthily control software

12

Device Model	Resonant Frequency (kHz)			Amplitude (g)
	X	Y	Z	
MPU6500 Sensor Only	5.1, 20.3	5.1–5.3	None	1.9
Galaxy S5	5.25–5.55	5.35	None	2
Galaxy Note 3	5.3–5.4	None	None	0.4

Acoustic Attack on Smartphone Drives RC Car

Figure 13. **Smartphone Attacking its own Accelerometer to Control an RC Car.** An Android phone runs an application that controls an RC car based on the phones orientation, measured by its internal MEMS accelerometer. Simultaneously, a malicious audio file is playing over the phone's speaker, mounting an output control attack on the phone's accelerometer. The RC car is essentially piloted by the audio file.

running on embedded devices. Embedded software applications often assume trustworthy input from accelerometers to make automated or closed-loop decisions. We demonstrate two system-level attacks using acoustic injection: (1) controlling a smartphone application that drives an RC car by playing a malicious music file on the phone, and (2) controlling a Fitbit fitness tracker to earn financial rewards by playing tones from an external speaker. Unlike our previous experiments, there is no external speaker for the smartphone attack. Instead, our attack uses the built-in speaker in the smartphone to play a music file that hijacks control of the accelerometer's output. We refer to this special subclass of vulnerability as a *self-stimulation attack* when a vulnerable system overtly co-locates a transmitter near a sensor by design—making standoff distances effectively zero meters.

7.1. Packaging Effects on Resonant Frequencies

Attacking an accelerometer buried in an embedded device raises an important question: *Does the packaging change the acoustic resonant frequency at all?* Here we demonstrate that packaging an accelerometer inside an embedded device only slightly alters its resonant frequencies. We analyzed two different smartphones with the same MEMS accelerometer model (MPU6500): the Samsung Galaxy S5 and Galaxy Note 3. We evaluated the acoustic vulnerabilities of accelerometers inside the phones using the same experimental setup we used for evaluating sensors (Figure 4), minus the vibrating platform. Each phone reported real time acceleration data via an Android application (Wireless IMU) that transmitted the data over a UDP stream to a nearby computer, rather than through an Arduino microcontroller. Table 2 summarizes the results of our experiments, and compares our results with the results from attacking the sensor alone. Evidently, the acoustic resonant frequency of an accelerometer mostly stands apart from its packaging, though the amplitude of acoustic acceleration can be attenuated by packaging.

7.2. Smartphone Controlled RC Car

To demonstrate the self-stimulation attack on the smartphone we attempted to hijack control of a smartphone application that makes use of the phone's accelerometer to pilot a wireless RC car. Numerous inexpensive RC cars are controlled with smartphone applications. These applications allow users to tilt the phone in the direction

they want to steer the car. This functionality employs the phone's MEMS accelerometer. The accelerometer measures the phone's physical orientation in relation to gravity. The application translates this information into digital commands that are sent to the car via WiFi or Bluetooth. The goal is to use the phone's speaker to spoof acceleration measurements that would trigger the RC car application to send commands to the car — commanding the car to go forwards, backwards, and to stop. This notion of an application (playing music) contaminating the behavior of another application (steering an RC car) running simultaneously violates basic Android data and privilege separation principles. This attack demonstrates a unique write side channel.

Evaluation. The experimental setup is shown in Figure 13. An RC car, Samsung Galaxy S5 smartphone, and computer were all placed on the same local area network. The Samsung Galaxy S5 phone contains an MPU6500 accelerometer, a sensor that is vulnerable to the output control attack. The phone ran three Android applications from the Google Play store: 1) RC car controlling application (i-Spy Toys), 2) accelerometer monitoring application (Wireless IMU), and 3) an application that played audio files (WavePad Audio Editor). The car controlling application polled the orientation state of the accelerometer and sent digital commands to the car over a TCP connection. The accelerometer monitoring application sent UDP packets with accelerometer measurements to the computer in real time. The audio application played a malicious WAV file that had been pre-loaded on the phone.

The RC car application monitors and reacts to X-axis acceleration. When the user tilts the phone flat or upright, i.e. the X-axis acceleration is $0\,g$ or $1\,g$ respectively, the application sends forward or backwards commands to the car. When the phone is approximately at a $30°$ angle, the X-axis acceleration is $0.3\,g$ and the application sends stop commands to the car.

Results. The phone was placed in an upright position (X-axis aligned with gravity). The malicious WAV file contained an AM acoustic interference signal designed to drive the car forward and backward, shown in Figure 14a. The
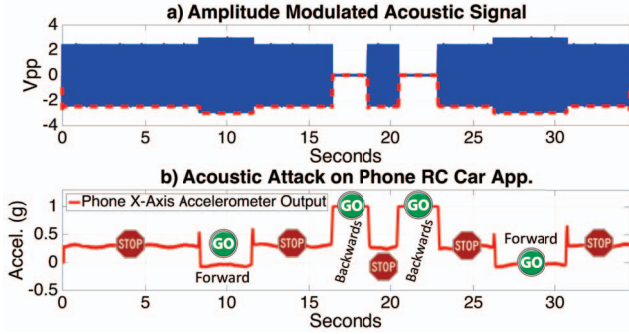
13

Figure 14. **Controlling an RC Car with an Output Control Attack on a Samsung Galaxy S5.** A smartphone controlled RC car reacts to commands its given over WiFi. The car behaves according to the phone's X-axis orientation towards gravity. The amplitude modulated acoustic signal in (A) is used to mount an output control attack that controls the phone's accelerometer output. The false acceleration measurements (B) trick the application to send forward/stop/backward commands to the RC car.

acoustic interference was played over the phone's speaker. Figure 14b shows the X-axis acceleration spoofed by the malicious audio file, and how the RC car reacted.

7.3. Free Fitbit Rewards

Several companies, including Walgreens and Higi, incentivize people to exercise by offering rewards programs that tether to their personal fitness tracking wristbands and monitor their daily physical activity. These fitness tracking wristbands use accelerometer driven pedometers [1] to count the number of steps the user takes over the course of a day. Rather than exploiting software vulnerabilities to spoof step counts [22], we demonstrate how one can spoof approximately 3,000 steps an hour on a Fitbit One [23] fitness tracker using acoustic interference and earn free rewards.

We opened a Higi.com account and tethered a Fitbit One device to the account. Using a similar setup as shown in Figure 4, absent the vibrating platform, acoustic interference at the resonant frequency of the Fitbit's accelerometer was played for approximately 40 minutes. No signal aliasing or modulation was needed as simply spoofing fluctuating false measurements was sufficient to register thousands of false steps. We were able to register 2,100 steps in that time and earn 21 rewards points on Higi.com without walking a single step. Due to ethical considerations, we have not claimed any of these rewards and have notified the respective manufactures about such flaws.

8. Defending Against Acoustic Attacks

Acoustic attacks exploit security vulnerabilities in the hardware components of MEMS accelerometers. Going forward, building secure sensors may eradicate this acoustic threat vector. However, vulnerable MEMS accelerometers are currently already deployed in many devices and systems. In this section we provide both *hardware design suggestions* and *software defense mechanisms* to increase the

Defense Mechanism	Output Biasing	Output Control
Secure LPF & Amplifier	✓	✓
Acoustic Dampening Materials	✓	✓
Randomized Sampling	✓	
180° Out-of-Phase Sampling	✓	

difficulty of mounting acoustic injection attacks on MEMS accelerometers. Table 3 summarizes the effectiveness of each suggestion and mechanism in thwarting each proposed attack. It is important to note that though some of the defenses we propose may not completely eradicate acoustic vulnerabilities, they will certainly increase the exploitation difficulty for the adversary.

8.1. Hardware Design Suggestions

Both kinds of acoustic injection attacks, output biasing and output control, exploit hardware deficiencies in the signal conditioning components. Specifically the LPF, amplifier, and mechanical sensing structures of MEMS accelerometers are negatively impacted by resonant acoustic interference (Figure 6). Designing these components to better tolerate acoustic interference would make MEMS accelerometers resilient to our attacks.

Secure Low Pass Filter. Output biasing attacks leverage signal aliasing at the ADC to control the accelerometer's output, a capability that should be suppressed by low pass filtering the analog acceleration signal prior to digitization. Low pass filters are designed to pass low frequency signals while blocking high frequency signals. They have three important frequency ranges: 1) *pass band*, 2) *transition band*, and 3) *stop band*. The pass band does not block any frequencies in its range. Frequencies in the transition band are increasingly attenuated, and frequencies in the stop band are completely blocked. The frequency that marks the transition point between the pass band and transition band is known as the cutoff frequency, F_{cutoff}.

A properly designed analog LPF should have a cut-off frequency of less than half of the ADC sampling rate, i.e., $F_{cutoff} = \frac{F_{samp}}{2}$, to prevent signal aliasing. The sampling rates of most accelerometers we analyzed were less than 1.5 kHz, implying the maximum frequency acceleration signal they could accurately measure was less than 750 Hz. Most accelerometers also exhibited resonant frequencies greater than 2.5 kHz. Three scenarios explain why the LPFs we encountered in the sensors we analyzed do not always filter out high frequency acoustic interference:

1) **No LPF Exists**: Designers did not include an LPF in the signal conditioning path at all (unlikely).
2) **Signal Clipping at the Amplifier**: The amplifier was not securely designed to account for high amplitude acoustic noise, causing signal clipping to be observable. Signal clipping introduces a DC component into the output signal which slips through the LPF.

14

3) **Resonant Frequency Lies in the LPF's Transition Band**: the resonant frequency of the accelerometer lies within the LPF's transition band. As a result, the LPF does not fully attenuate the acoustic interference.

The solution to scenario 1 (though this scenario is unlikely) is straightforward: add an LPF. The solution to scenario 2 is discussed in the following section. Lastly, scenario 3 is the most difficult to address. Designing an LPF that has a transition band that does not overlap the accelerometer's resonant frequency can be accomplished in three ways: 1) *lower the cutoff frequency*, 2) *narrow the transition band*, or 3) *design the mass-spring sensing structure to exhibit a higher resonant frequency*.

All three have different limitations. The first lowers the frequency limit of vibrations an accelerometer can measure. The second requires adding many extra components, eventually for little to no added benefit. Finally, the last is possible but requires stiffening the spring and losing sensitivity [24].

Secure Amplifier. Output control attacks leverage signal clipping at the amplifier to introduce a DC component into the acceleration signal which slips through any subsequent LPF. This is prevented in two ways:

1) **More Tolerant Amplifier**: Design an amplifier that can accept the large amplitude inputs that are generated under acoustic interference.
2) **Pre-filter Amplifier Inputs**: Filter acoustic resonant frequencies prior to the amplifier with another LPF or band-stop filter.

The first solution is potentially limited by size, power, and cost. The larger the amplifier circuitry, the more power and chip area it consumes. These increase sensor cost and decrease deployability. The second solution, which some designs do employ [10], is limited by the cost of adding more components, but may not increase power consumption.

Acoustic Dampening Materials. Attenuating acoustic waves before they penetrate sensor packaging can prevent acoustic acceleration signals from being generated at all. Surrounding accelerometer ICs with acoustic dampening materials, such as synthetic foam [7], [8], can shield it from acoustic noise. The limitation here is size: acoustic dampening foam takes up space, a scarce resource in most embedded systems.

8.2. Software Defense Mechanisms

Redesigning hardware to tolerate acoustic interference is not an option for accelerometers already deployed in the field. For a subset of these sensors we provide two different defense mechanisms that can be implemented in software and deployed as firmware updates: *randomized sampling* and *180° out-of-phase sampling*. These solutions are only capable of preventing output biasing attacks, where acoustic acceleration signals have not been distorted by amplifier clipping. They work by eliminating an attacker's ability to achieve a DC signal alias at the ADC. Each defense


Figure 15. **Example of Randomized and** $180°$ **Out-of-Phase Sampling:** A) Sampling at random times within the resonant frequency period prohibits an attacker's ability to control sensor outputs with DC aliasing. B) Taking 2 samples $180°$ out-of-phase, with respect to the resonant frequency, will yield samples symmetric around the true acceleration value. Averaging the samples, in both mechanisms, cancels out the disturbance.

mechanism takes advantage of the requirement that only acoustic resonant frequencies can displace the sensing mass, and that these frequencies are known at design time. For that reason, we consider only sensors that exhibit false fluctuating measurements under acoustic interference. Both solutions assume the device has control over the sampling regimes of its sensors, i.e., they employ analog sensors and software controlled ADCs (several microcontrollers allow software to trigger the ADC to take a sample [25]).

Randomized Sampling. Randomized sampling eliminates the predictability of an ADC's sampling regime. Instead of setting an ADC to sample at a fixed interval, randomized sampling adds a random amount of delay to the beginning of each sampling period. This prevents an attacker from tuning the resonant frequency to induce a DC alias, i.e., step 1 (stabilize) of the output biasing attack (Section 6.1). Randomized sampling intentionally amplifies the effect of having an inaccurate ADC. Computing a moving average over several samples then smooths the fluctuating measurements.

An adversary performing an output biasing attack stabilizes the fluctuating false acceleration measurements by tuning the acoustic frequency such that it is an integer multiple of the sampling frequency (Equation 3). Defeating this attack, we add random delay, t_{delay}, to the sampling time, t_k, s.t. t_{delay} is uniformly distributed in $[0, \frac{1}{F_{\text{res}}}]$. Recall that the acoustic frequency F_a is close to the resonant frequency: $F_a \approx F_{\text{res}}$. Therefore, setting the sampling times $t_k^* = t_k + t_{\text{delay}}$ results in a symmetrical distribution of $\hat{s}(t_k)$ over a full cycle of acoustically stimulated acceleration measurements, $cos(2\pi F_a t_k + \phi)$.

Figure 15a illustrates the concept of randomized sampling. The resulting distribution of $\hat{s}(t_k)$ is not uniformly distributed over $[s(t_k) - s_a(t_k), s(t_k) + s_a(t_k)]$, but rather it is *symmetric* around the value of true acceleration, $s(t_k)$. Hence, computing a moving average of several samples filters out periodic acoustic acceleration but not true acceleration. Randomized sampling does not destroy valid peri-
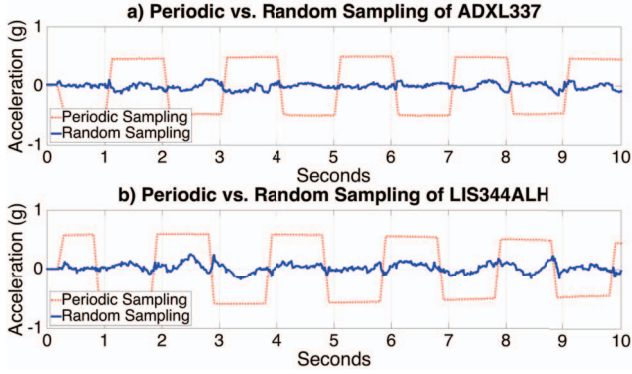
15

Figure 16. **Randomized Sampling:** Both periodic (red dotted line) and randomized sampling (solid blue line) are employed on ADCs interfaced to two analog accelerometers, the ADXL337 and LIS344ALH. Simultaneously, output biasing attacks were crafted on both sensors to induce artificial square wave output signals. The bogus square wave acceleration signals are attenuated by deploying randomized sampling.

Figure 17. **180° Out-of-Phase Sampling**: Both periodic (red dotted line) and 180° out-of-phase (solid blue line) sampling are employed on ADCs interfaced to two analog accelerometers, the ADXL337 and LIS344ALH. Simultaneously, output biasing attacks were crafted on both sensors to induce artificial square wave output signals. The bogus square wave acceleration signal is attenuated by deploying out-of-phase sampling.

odic acceleration signals, i.e. vibrations within $[0, \frac{F_{\text{samp}}}{2}]$, because in most cases, the maximum frequency of true acceleration is much smaller than the resonant frequency.

Some MEMS accelerometers exhibit multiple resonant frequencies. For these sensors, the random delay added to the sampling time, t_{delay}, should be uniformly distributed in $[0, \frac{1}{F_{\text{lcm}}}]$, where F_{lcm} is the least common multiple of all resonant frequencies exhibited by the device. No matter what resonant frequency the attacker uses, $\hat{s}(t_k)$ remains symmetrically distributed around the true acceleration value.

180° Out-of-Phase Sampling. One hundred eighty degree out-of-phase sampling attenuates acceleration signals with frequencies around a given sensor's resonant frequency. It acts as a simple band-stop filter in software. An ADC performing out-of-phase sampling takes two samples at a 180° phase delay with respect to the resonant frequency F_{res}. Namely, two samples are taken at times $t_k, t_k + t_{\text{delay}}$ where $t_{\text{delay}} = \frac{1}{2 \cdot F_{\text{res}}}$. The true acceleration measurement value is then computed by taking the average: $s_k = \frac{1}{2}(s(t_k) + s(t_{\text{delay}}))$. Figure 15b illustrates the out-of-phase sampling concept.

Following step 1 (stabilize) of the output biasing attack (Section 6.1), an adversary chooses an acoustic frequency approximately equal to the resonant frequency, $F_a \approx F_{\text{res}}$. Out-of-phase sampling is analogous to a notch filter around the resonant frequency range. Given an acoustic acceleration signal, $s_a(t_k)$:

$$
\begin{aligned}
s_a(t_k + t_{\text{delay}}) &= A_0 A_1 cos(2\pi F_a(t_k + t_{\text{delay}}) + \phi) \\
&= A_0 A_1 cos(2\pi F_a t_k + \pi + \phi) \quad (4) \\
&= -s_a(t_k)
\end{aligned}
$$

Stated otherwise, the value of two samples of acoustically stimulated acceleration taken 180° out-of-phase are opposites. Assuming the maximum frequency of the true acceleration signal, $s(t)$, is much smaller than the resonant frequency, then $s(t)$ will be the same across two out-of-phase samples while the acoustically stimulated accelera-
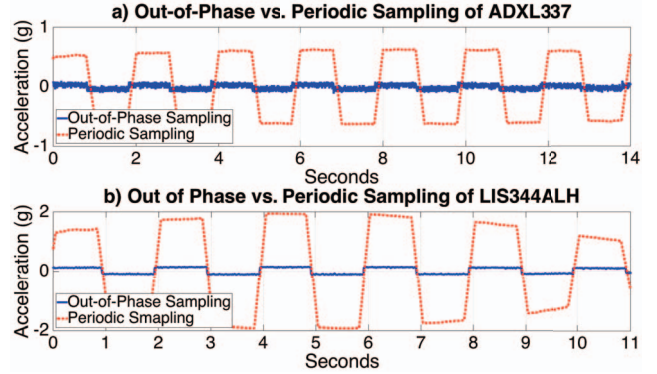
tion, $s_a(t)$, is not. Namely, $s(t_k) \approx s(t_k + t_{\text{delay}})$ and $s_a(t_k) = -s_a(t_k + t_{\text{delay}})$. Averaging the out-of-phase samples yields:

$$
\frac{1}{2}(\hat{s}(t_k) + \hat{s}(t_k + t_{\text{delay}})) \approx \frac{1}{2}(2s(t_k) + 0) = s(t_k) \quad (5)
$$

The measured acceleration signal after averaging is approximately the same as the true acceleration signal $s(t)$.

Implementation & Evaluation. Both sampling mechanisms assume software can control the sampling regimes of the sensors, i.e. an analog sensor sampled by software controlled ADCs. We demonstrate randomized sampling and out-of-phase sampling on two analog accelerometers, the ADXL337 and LIS344ALH, interfaced to ADCs embedded in the Arduino microcontroller. We use the same experimental setup described in Figure 4, without the vibrating platform. For randomized sampling, the ADC was programmed to add a random delay, t_{delay}, at the beginning of each sampling cycle according to the resonant frequency of the respective accelerometer. Conversely, for out-of-phase sampling the ADC was configured to take two samples at exactly $1/F_{\text{res}}$ seconds apart. Output biasing attacks were performed to create bogus square wave acceleration signals on both sensors. Figures 16 and 17 show the effectiveness of random and out-of-phase sampling, respectively, vs. normal periodic sampling at filtering out the maliciously spoofed square waves.

9. Related Work

Attacking and defending systems at the analog sensor and actuator layer can be classified as *analog cybersecurity*. Acoustic injection attacks represent one type of attack in the analog cybersecurity quiver. Analog cybersecurity owes its heritage to research in the side channel analysis and fault injection attack communities that grew from research on smartcard security in the 1990s [26]. Most recently, Shoukry et al. [9] develop a sensor authentication scheme, called

16

PyCRA, to thwart attacks on sensor–actuator systems. They categorize analog cybersecurity attacks into sensor spoofing and eavesdropping attacks, and offer ways to identify spoofed signals. From there they are able to subtract spoofed signals from the perceived signals to recover the original sensor measurements. As Shin et al. point out [27], PyCRA makes assumptions about an adversary's capabilities and can be defeated when these assumptions are violated. In contrast to PyCRA, we offer techniques to eliminate the attacker's ability to spoof sensor measurements to begin with.

Spoofing Analog Sensors. Earlier sensor spoofing attacks focus on the system level, manipulating a system's behavior by altering a sensor's environmental perception. In the unmanned aerial vehicle (UAV) space, Son et al. [7] demonstrate intentional acoustic interference on MEMS gyroscopes in drones, causing the them to crash. Davidson et al. [28] spoof optical flow sensors on UAVs with intense light to control their lateral motion. In the medical device domain, Park et al. [29] utilize intentional infrared interference to trigger medical infusion pumps to over deliver medicine to patients. Foo Kune et al. [30] show how carefully crafted electromagnetic interference (EMI) can be injected into signal digitization circuitry inside implantable medical devices to control the delivery of pacing and defibrillation shocks. In the automotive area, Shoukry et al. [31] demonstrate how to deliver false readings to anti-lock braking systems (ABS) via the magnetic wheel speed sensors using EMI. Lastly, Yan et al. spoof various Tesla autopilot subsystems with intentional ultrasonic and EMI interference to cause safety critical malfunctions. Rather than separately analyzing individual systems that utilize sensors for analog vulnerabilities, our work takes a fundamental approach: exploring the analog vulnerabilities of the sensors themselves for the purpose of defending the systems that employ them.

Intentional and Unintentional Interference. Engineering researchers undergo great efforts to design robust systems that are resilient to interference, including electromagnetic and acoustic. Boneh et al. demonstrate how computational faults induced by interference can break cryptographic protocols [26]. Consequently, understanding the potential threats of interference on systems, devices, and sensors is vital. Giri et al. classify intentional EMI threats into categories of frequency range, level of sophistication, and effects on targeted systems [32]. Delsing et al. explore the vulnerability of sensor networks to intentional EMI [33]. The effects of unintentional and intentional EMI on implantable medical devices have also been investigated [30], [34], [35]. Dean et al. and Castro et al. characterize the effects of high power acoustic noise on MEMS gyroscopes [36], [37], [38], and Soobramaney and Castro et al. develop mechanisms to mitigate acoustic interference on MEMS gyroscopes using acoustic dampening materials [8], [38]. Soobramaney also demonstrates a defense mechanism that utilizes a modified gyroscope to respond to only acoustic interference to cancel the interference signal from the true signal [8]. To the best of our knowledge, we are the first to demonstrate how intentional acoustic interference on MEMS accelerometers can be leveraged to *control* their output.

Information Leakage. Information leakage from physical properties, or *side-channels*, of computing systems are also relevant to analog cybersecurity. Recent studies show that gyroscopes and accelerometers can leak personal information [12], [13], [14], [15], [16]. Michalevsky et al. show that gyroscopes in smart-phones can be used as a microphone to eavesdrop on conversations [16]. Marquardt et al. demonstrate that smart-phone accelerometers leak enough information to infer keystrokes from a nearby keyboard [12]. Similarly, Owusu and Aviv show smart-phone accelerometer information leakage can be leveraged to infer user touch-screen gestures and key presses to leak passwords and PIN codes to unlock phones [14], [15]. Dey et al. found that process variation in accelerometers yields a unique fingerprint that can uniquely identify a device [13]. These efforts are a reminder that physical attacks on analog sensors render securing data integrity, authentication, and confidentiality between sensors and microprocessors challenging.

10. Conclusion

Because MEMS accelerometers use displacement as a proxy for measuring acceleration, malicious acoustic interference at resonant frequencies can damage the integrity of a sensor's digital outputs. Our work models the physics of acoustic injection attacks on MEMS accelerometers, validated by measuring the outputs of sensors subjected to our acoustic interference. Our experiments show that subtle hardware security flaws in amplification and filtering circuits of the signal conditioning path represent the fundamental root causes of the vulnerabilities. These pervasive security flaws lead to two unusual classes of sensor vulnerabilities: *output biasing* and *output control*. In our acoustic tests of 20 accelerometer models from 5 manufacturers, we found 75% are vulnerable to output biasing attacks and 65% vulnerable to output control attacks. We also demonstrate proof-of-concept end-to-end attacks with physical consequences. In terms of compromising data integrity, we show how to inject fake steps into a Fitbit fitness tracker to earn financial rewards. In terms of affecting control systems, we accomplish a *self-stimulating attack* whereby we play malicious music files from a smartphone's speaker to control an app that drives an RC car. To reduce the risks of attacks on the integrity of MEMS accelerometers, we recommend hardware design suggestions to increase the security of amplifiers and filters and mitigate acoustic attacks on the next generation of sensors. For sensors already deployed in the field, we offer two low-cost software defense mechanisms to prevent output biasing attacks: *randomized sampling* and *180° out-of-phase sampling*. Our software defense mechanisms can protect all accelerometers vulnerable to output biasing attacks, but not output control attacks.

Acknowledgments

References

[1] N. Zhao, "Full-featured pedometer design realized with 3-axis digital accelerometer," *Analog Dialogue*, 2010.

[2] A. Wright and T. Pratt, "Dell raises the bar in shock-resistant hard drives," Dell Inc., Tech. Rep., 2008, http://www.dell.com/downloads/global/vectors/2008_freefallprotection.pdf.

[3] Analog Devices, "The five motion senses: Using MEMS inertial sensing to transform applications," Analog Devices, Tech. Rep., 2009, http://www.analog.com/media/en/technical-documentation/white-papers/The_Five_Motion_Senses.pdf.

[4] P. Wang, G. V. Naccarelli, M. R. Rosen, N. M. Estes III, D. L. Hayes, and D. E. Haines, *New Arrhythmia Technologies*. Wiley, 2008.

[5] M. Kraft and N. M. White, *MEMS for automotive and aerospace applications*. Elsevier, 2013.

[6] A. Kuznetsov, Z. Abutidze, B. Portnov, V. Galkin, and A. Kalik, "Development of MEMS sensors for aircraft control systems," *Gyroscopy and Navigation*, 2011.

[7] Y. Son, H. Shin, D. Kim, Y. Park, J. Noh, K. Choi, J. Choi, and Y. Kim, "Rocking drones with intentional sound noise on gyroscopic sensors," in *Proc. USENIX Security Symposium*, 2015.

[8] P. Soobramaney, "Mitigation of the effects of high levels of high-frequency noise on mems gyroscopes," Ph.D. dissertation, Auburn University, 2013.

[9] Y. Shoukry, P. Martin, Y. Yona, S. Diggavi, and M. Srivastava, "Pycra: Physical challenge-response authentication for active sensors under spoofing attacks," in *Proc. ACM CCS*, 2015.

[10] M. L. Shaw, "Accelerometer overload considerations for automotive airbag applications," SAE Technical Paper, Tech. Rep., 2002.

[11] M. Andrejašič, "MEMS accelerometers," University of Ljubljana, Tech. Rep., 2008, http://mafija.fmf.uni-lj.si/seminar/files/2007_2008/MEMS_accelerometers-koncna.pdf.

[12] P. Marquardt, A. Verma, H. Carter, and P. Traynor, "(sp) iPhone: decoding vibrations from nearby keyboards using mobile phone accelerometers," in *Proc. ACM CCS*, 2011.

[13] S. Dey, N. Roy, W. Xu, R. R. Choudhury, and S. Nelakuditi, "AccelPrint: Imperfections of accelerometers make smartphones trackable," in *NDSS*, 2014.

[14] E. Owusu, J. Han, S. Das, A. Perrig, and J. Zhang, "ACCessory: password inference using accelerometers on smartphones," in *ACM HotMobile*, 2012.

[15] A. J. Aviv, B. Sapp, M. Blaze, and J. M. Smith, "Practicality of accelerometer side channels on smartphones," in *Proc. ACSAC*, 2012.

[16] Y. Michalevsky, D. Boneh, and G. Nakibly, "Gyrophone: Recognizing speech from gyroscope signals," in *Proc. USENIX Security Symposium*, 2014.

[17] Pyramid Car Audio, "Tw28 bullet horn tweeter," http://www.pyramidcaraudio.com/sku/TW28/300-Watt-Aluminum-Bullet-Horn-in-Enclosure-wSwivel-Housing.

[18] Yamaha Corporation, "R-S201 Receiver User Manual."

[19] J. J. Rychcik, J. E. Vandemeer, and M. L. Shaw, "Characterizing input saturation in low-g accelerometers," 2002, http://archives.sensorsmag.com/articles/0502/68/main.shtml#ref1.

[20] Analog Devices, "Avoiding op amp instability problems in single-supply applications," Analog Devices, Tech. Rep., 2001, http://www.analog.com/media/en/analog-dialogue/volume-35/number-1/articles/avoiding-op-amp-instability-problems.pdf.

[21] National Instruments Inc., "G.R.A.S. 46BE 1/4" CCP Free-field Standard Microphone Set," http://www.ni.com/pdf/manuals/G.R.A.S._46BE.pdf.

[22] M. Rahman, B. Carbunar, and M. Banik, "Fit and vulnerable: Attacks and defenses for a health monitoring device," *arXiv 1304.5672*, 2013.

[23] FitBit, "FitBit One," https://www.fitbit.com/one.

[24] J. Voldman, "Case study: A capacitive accelerometer," in *OpenCourseWare*, 2013.

[25] Texas Instruments, "Tiva TM4C129XNCZAD microcontroller data sheet (Rev. B)," June 2014, http://www.ti.com/lit/ds/symlink/tm4c129xnczad.pdf.

[26] D. Boneh, R. A. DeMillo, and R. J. Lipton, "On the importance of checking cryptographic protocols for faults," in *EUROCRYPT*, 1997.

[27] H. Shin, Y. Son, Y. Park, Y. Kwon, and Y. Kim, "Sampling race: Bypassing timing-based analog active sensor spoofing detection on analog-digital systems," in *Usenix WOOT*, 2016.

[28] D. Davidson, H. Wu, R. Jellinek, V. Singh, and T. Ristenpart, "Controlling UAVs with sensor input spoofing attacks," in *Usenix WOOT*, 2016.

[29] Y. Park, Y. Son, H. Shin, D. Kim, and Y. Kim, "This ain't your dose: Sensor spoofing attack on medical infusion pump," in *Usenix WOOT*, 2016.

[30] D. Foo Kune, J. Backes, S. S. Clark, D. Kramer, M. Reynolds, K. Fu, Y. Kim, and W. Xu, "Ghost talk: Mitigating EMI signal injection attacks against analog sensors," in *IEEE Symp. on Security & Privacy*, 2013.

[31] Y. Shoukry, P. Martin, P. Tabuada, and M. Srivastava, "Non-invasive spoofing attacks for anti-lock braking systems," in *International Workshop on Cryptographic Hardware and Embedded Systems*, 2013.

[32] D. Giri and F. Tesche, "Classification of intentional electromagnetic environments," *IEEE Trans. Electromagnetic Compatibility*, 2004.

[33] J. Delsing, J. Ekman, J. Johansson, S. Sundberg, M. Bäckström, and T. Nilsson, "Susceptibility of sensor networks to intentional electromagnetic interference," in *Intl. Zurich Symp. on EMC*, 2006.

[34] D. L. Hayes, P. J. Wang, D. W. Reynolds, N. M. Estes, J. L. Griffith, R. A. Steffens, G. L. Carlo, G. K. Findlay, and C. M. Johnson, "Interference with cardiac pacemakers by cellular telephones," *New England Journal of Medicine*, 1997.

[35] S. Lee, K. Fu, T. Kohno, B. Ransford, and W. H. Maisel, "Clinically significant magnetic interference of implanted cardiac devices by portable headphones," *Heart Rhythm*, 2009.

[36] R. N. Dean, G. T. Flowers, A. S. Hodel, G. Roth, S. Castro, R. Zhou, A. Moreira, A. Ahmed, R. Rifki, B. E. Grantham *et al.*, "On the degradation of MEMS gyroscope performance in the presence of high power acoustic noise," in *IEEE Intl. Symp. on Industrial Electronics*, 2007.

[37] R. N. Dean, S. T. Castro, G. T. Flowers, G. Roth, A. Ahmed, A. S. Hodel, B. E. Grantham, D. A. Bittle, and J. P. Brunsch, "A characterization of the performance of a MEMS gyroscope in acoustically harsh environments," *IEEE Trans. Industrial Electronics*, 2011.

[38] S. Castro, R. Dean, G. Roth, G. T. Flowers, and B. Grantham, "Influence of acoustic noise on the dynamic performance of MEMS gyroscopes," in *ASME International Mechanical Engineering Congress and Exposition*, 2007.

18