



Injected and Delivered: Fabricating Implicit Control over Actuation Systems by Spoofing Inertial Sensors

Yazhou Tu, University of Louisiana at Lafayette; Zhiqiang Lin, Ohio State University; Insup Lee, University of Pennsylvania; Xiali Hei, University of Louisiana at Lafayette

<https://www.usenix.org/conference/usenixsecurity18/presentation/tu>

**This paper is included in the Proceedings of the
27th USENIX Security Symposium.**

August 15–17, 2018 • Baltimore, MD, USA

978-1-939133-04-5

**Open access to the Proceedings of the
27th USENIX Security Symposium
is sponsored by USENIX.**

Injected and Delivered: Fabricating Implicit Control over Actuation Systems by Spoofing Inertial Sensors

Yazhou Tu* Zhiqiang Lin† Insup Lee‡ Xiali Hei*
*University of Louisiana at Lafayette
†The Ohio State University
‡University of Pennsylvania

Abstract

Inertial sensors provide crucial feedback for control systems to determine motional status and make timely, automated decisions. Prior efforts tried to control the output of inertial sensors with acoustic signals. However, their approaches did not consider sample rate drifts in analog-to-digital converters as well as many other realistic factors. As a result, few attacks demonstrated effective control over inertial sensors embedded in real systems.

This work studies the out-of-band signal injection methods to deliver adversarial control to embedded MEMS inertial sensors and evaluates consequent vulnerabilities exposed in control systems relying on them. Acoustic signals injected into inertial sensors are out-of-band analog signals. Consequently, slight sample rate drifts could be amplified and cause deviations in the frequency of digital signals. Such deviations result in fluctuating sensor output; nevertheless, we characterize two methods to control the output: *digital amplitude adjusting* and *phase pacing*. Based on our analysis, we devise non-invasive attacks to manipulate the sensor output as well as the derived inertial information to deceive control systems. We test 25 devices equipped with MEMS inertial sensors and find that 17 of them could be implicitly controlled by our attacks. Furthermore, we investigate the generalizability of our methods and show the possibility to manipulate the digital output through signals with relatively low frequencies in the sensing channel.

1 Introduction

Sensing and actuation systems are entrusted with increasing intelligence to perceive the environment and react to it. Inertial sensors consisting of gyroscopes and accelerometers measure angular velocities and linear accelerations, which directly depict movements and orientations of a device. Therefore, systems equipped with inertial sensors are able to determine motional status and

make actuation decisions in a timely, automated manner. While inertial sensing allows a control system to actuate in response to environmental changes promptly, errors of inertial measurements could result in instantaneous actuations as well.

Micro-electro-mechanical systems (MEMS) gyroscopes are known to be susceptible to resonant acoustic interferences [41, 44, 45, 75]. Son et al. showed that a drone could be caused to crash by disturbing the gyroscope with intentional resonant sound [64]. Furthermore, Trippel et al. investigated the data integrity issue of MEMS accelerometers under acoustic attacks [68]. While they gained adversarial control over exposed accelerometers, few attacks demonstrated effective control over embedded sensors. Thus, it remains unrevealed that to what extent attackers could exploit embedded inertial sensors and possibly control the systems relying on them.

To achieve adversarial control over inertial sensors embedded in real systems, we need to consider several realistic factors: (a) *Attack setting*. Biasing attacks in [68] were conducted on exposed sensors connected to an Arduino board, making the sampling process and real-time sensor data accessible to attackers. In contrast, our work studies non-invasive attacks, implying that attackers cannot physically alter the system and can only infer necessary information about the sensor from observable phenomena. (b) *Sample rate*. The exact sample rate of embedded sensors could be difficult to access, and we find that slight drifts in the sample rate may cause troubles to attackers. (c) *Actuating direction*. While Trippel et al. [68] manipulated a smartphone controlled RC car by inducing sensor outputs in only one direction, most systems rely on inertial measurements in both directions for control purposes. In this work, we develop generalizable methods that could manipulate inertial measurements of embedded sensors and trigger actuations of different kinds of control systems in both directions.

Acoustic signals injected at resonant frequencies of inertial sensors are usually out-of-band signals, which

will be sampled by the analog-to-digital converter (ADC) with an insufficient sample rate. We characterize this kind of attacks as *out-of-band signal injections*, presenting several important features: (1) *Amplification of sample rate drifts*. We find that tiny drifts in the sample rate of an ADC could be amplified and cause more significant deviations in the frequency of the digital signal. Consequently, it could be difficult to induce and maintain a DC (Direct Current, 0 Hz) sensor output as in prior work [68]. The resulting digital signal serves as noises due to its oscillating nature; nevertheless, we perceive following properties to control it. (2) *Adjustable digital amplitudes*. Distortions caused by undersampling allow amplitudes of different digital samples within one cycle of oscillation adjustable. (3) *Phase pacing*. We find that a phase offset could be induced in the digital signal by switching the frequency of out-of-band analog signals.

Based on our analysis, we develop non-invasive attacks to manipulate the output of embedded inertial sensors as well as the derived information to deceive different kinds of control systems. We evaluate our attacks on 25 devices equipped with various models of inertial sensors from different vendors. Our experimental results show that 23 devices could be affected by acoustic signals and 17 of them are susceptible to implicit control. Our attack demonstrations include maliciously actuating the motor of self-balancing human transporters, manipulating a user’s view in virtual reality (VR) systems, spoofing a navigation system (Google Maps), etc. We have uploaded the demos of our proof-of-concept attacks at <https://www.youtube.com/channel/UCGMX3ZbE1V7BZYIX7RtF5tg>.

In summary, we list our contributions as follows:

- We devise two sets of novel spoofing attacks (*Side-Swing* and *Switching* attacks) against embedded MEMS inertial sensors to manipulate sensor outputs and the derived inertial information. The attacks are non-invasive and could deliver implicit control to different kinds of real systems relying on inertial sensors.
- We evaluate our attacks on 25 devices and find that 23 of them can be affected by acoustic signals, presenting different control levels. Our proof-of-concept attacks demonstrate adversarial control over self-balancing, aiming and stabilizing, motion tracking and controlling, navigation systems, etc.
- We propose the out-of-band signal injection model and methods to manipulate the oscillating digitized signal when an analog signal is sampled with an insufficient sample rate. We investigate the generalizability of our methods with a case study showing that attackers could manipulate the oscillating digitized signal by sending signals with relatively low frequencies through a universal sensing channel.




Figure 1: An illustration of acoustic injections on inertial sensors embedded in control systems. Injections of analog signals occur in the transducer. The signal will be digitized by the ADC before reaching the control system.

2 Inertial Sensors in Control Systems

MEMS inertial sensors use mechanical structures to detect inertial stimuli and generate electrical signals to depict it. MEMS accelerometers detect linear accelerations with a mass-spring structure. While MEMS gyroscopes use a similar structure to sense Coriolis accelerations a_{Cor} , an extra vibrating structure is used to drive the sensing mass with a velocity v , which is orthogonal to the sensing direction. The angular velocity ω causing the Coriolis acceleration can be derived by: $a_{Cor} = -2\omega \times v$.

Acoustic Injection. Although MEMS technology has significantly reduced the size, cost and power consumption of inertial sensors, the miniaturized mechanical structure could suffer from resonant acoustic interferences. Acoustic signals at frequencies close to the natural frequency of the mechanical structure could force the sensing mass into resonance. Displacements of the sensing mass are usually measured by capacitive electrodes and would induce electrical signals. The signal will then be digitized by the ADC and could possibly influence the control system, as shown in Figure 1.

Under resonance, the sensing mass is forced into vibrations at the same frequency as the external sinusoidal driving force (sound pressure waves). Therefore, the mass-spring structure of inertial sensors could serve as a receiving system for resonant acoustic signals and allow attackers to inject analog signals at specific frequencies. However, the ability of attackers towards adversarial control is still restricted in two aspects: (1) Attackers cannot inject arbitrary forms of analog signals. Since the injected analog signal is caused by mechanical resonance of the sensing mass, it would be a sinusoidal signal and always presents an oscillating pattern. (2) The digital signal cannot be controlled directly. Attackers could only induce specific digital signals by controlling the analog signal. This process is difficult to control especially in an embedded environment with limited information.

Control System. MEMS inertial sensors provide crucial feedback for control systems to make autonomous deci-

sions. Applications of MEMS gyros and accelerometers are very broad. Examples of these applications include human transporters, kinetic devices, robots, pointing systems for antennas, navigation of autonomous (robotic) vehicles, platform stabilization of heavy machinery, yaw rate control of wind-power plants, industrial automation units, and guidance of low-end tactical applications [55, 36, 58, 67]. Because of their ubiquitousness and criticality in control systems, it is important to examine MEMS inertial sensors' reliability and evaluate the resilience of control systems under sensor spoofing attacks.

This work evaluates non-invasive spoofing attacks against embedded MEMS inertial sensors on a wide range of control systems in consumer applications. The systems we investigate can be broadly divided into two categories: (1) *Closed-loop control systems*. The system continuously compares its current status with a goal status and tries to diminish the difference between them through actuations. (2) *Open-loop control systems*. The system simply follows inertial sensing information to make actuation decisions. Different instances of these systems will be evaluated in Section 6.

3 Threat Model

The objective of attackers is to spoof embedded inertial sensors and deliver adversarial control to the system. To achieve this, attackers need to induce specific digital signals to trigger actuations in the control system.

Non-invasiveness. The spoofing attack against embedded inertial sensors is non-invasive and can be implemented without physical contact to the target device. Attackers cannot physically alter the hardware, neither can they directly access or modify the sampling process as well as the sensor output. However, we assume that attackers can analyze the behavior of an identical device under acoustic effects before a real attack.

Audibility. The resonant frequencies of MEMS accelerometers are usually within the range of human hearing. However, the resonant frequencies of MEMS gyros are often in the ultrasound band (above 20 kHz). Therefore, acoustic signals used to attack gyros are inaudible. While resonant frequencies of gyros in several devices we test are between 19 to 20 kHz, they are still above the audible range of most adults [66].

Sound Source. Attackers can use consumer-grade speakers or transducers, directivity horns, and amplifiers to generate sound waves. The signal source can be a function generator, an Arduino board, or mini signal generator boards [22, 24]. We assume that the possible attack distance is several meters; attackers have sufficient resources, i.e., techniques or fund, to optimize the power, efficiency, directivity and emitting area of the

sound source. More capable attackers could use professional acoustic devices or highly customized acoustic amplification techniques to further improve the range as well as the effect of the attack.

4 Modeling and Analysis

In acoustic attacks, malicious analog signals injected into the transducer will be processed and digitized before reaching the control unit. Therefore, the effect of attacks depends on the attacker's ability to influence the digitized signal. In this section, we analyze the digitization process of out-of-band analog signals and propose general methods to control the oscillating digitized signal.

4.1 Digitization of Out-of-band Signals

Since the sensing mass under resonance is oscillating at the same frequency as sound waves, the resulting analog signal can be described by,

$$V(t) = A \cdot \sin(2\pi Ft + \phi_0) \quad (1)$$

where F is the frequency of resonant sound waves and the amplitude $A = A_0 k_a k_s$. A_0 is the amplitude of sound waves. The coefficients k_a and k_s represent the attenuation of acoustic energy during transmission and the sensitivity of the mechanical sensing structure respectively. This analog signal will then be sampled by the ADC. Assuming F_S is the sampling rate, and $t_0 = 0, t_1 = \frac{1}{F_S}, \dots, t_i = \frac{i}{F_S}, \dots$, are sampling times, the digitized signal will be,

$$V[i] = A \cdot \sin(2\pi F \frac{i}{F_S} + \phi_0) \quad (i \in \{0, 1, 2, 3, \dots\}) \quad (2)$$

The frequency of analog signals injected through resonance is usually much higher than the sampling rate. For instance, the typical resonant frequency is several kHz for accelerometers and more than 19 kHz for gyros, while the sample rate is usually in tens or hundreds. According to the Nyquist theorem, when $F > \frac{F_S}{2}$, there would be a problem of aliasing. We have,

$$F = n \cdot F_S + \varepsilon \quad (-\frac{1}{2}F_S < \varepsilon \leq \frac{1}{2}F_S, n \in \mathbb{Z}^+) \quad (3)$$

Substitute (3) into (2), we have:

$$V[i] = A \cdot \sin(2\pi \varepsilon \frac{i}{F_S} + \phi_0) \quad (i \in \{0, 1, 2, 3, \dots\}) \quad (4)$$

These equations describe the basic relationship between the out-of-band analog signal and the digitized signal: a sinusoidal analog signal with a frequency F will be aliased to a digital signal with a frequency of ε .

Our discussions in this section mainly focus on signals with frequencies close to the same integer multiple of sample rate. Therefore, we assume that n in (3) stays the same when ε, F or F_S slightly changes.




Figure 2: The output of the gyroscope (X-axis) in a stationary iPhone 5S when we inject acoustic signals with a fixed frequency (19,471 Hz). Due to sample rate drifts, the frequency of the induced output is not a constant.

Amplification Effect of Sample Rate Drifts. ADC is designed to sample the voltage of the analog signal at specific intervals. Theoretically, each interval should be exactly $\frac{1}{F_S}$. Therefore, given F , the value of ε should be determined (Equation 3). However, due to drifts in F_S , when we inject acoustic signals at a fixed frequency into a smartphone’s gyroscope, we find that the frequency of the digital output is deviating, as shown in Figure 2. We formalize the following theorem to explain why slight sample rate drifts could result in observable deviations in the frequency of the digital signal.

Theorem 1. *When a signal with a frequency F is sampled with an insufficient sample rate F_S ($F_S < 2F$), a drift ΔF_S in the sample rate will be amplified to a deviation of $-n \cdot \Delta F_S$ in the frequency (ε) of the sampled signal and $n = \frac{F-\varepsilon}{F_S}$ ($-\frac{1}{2}F_S < \varepsilon \leq \frac{1}{2}F_S, n \in \mathbb{Z}^+$).*

Proof. Let $\hat{\varepsilon}$ be the frequency of the sampled signal after sample rate drifts. We have,

$$\begin{aligned} F &= nF_S + \varepsilon \\ F &= n(F_S + \Delta F_S) + \hat{\varepsilon} \end{aligned} \quad (5)$$

Therefore, the deviation in the frequency of the sampled signal is,

$$\hat{\varepsilon} - \varepsilon = -n \cdot \Delta F_S \quad (6)$$

□

For instance, the resonant frequency of gyros could range from 19 kHz to above 30 kHz. If $F = 20,000$ Hz and $F_S = 200$ Hz, a tiny drift of 0.01 Hz in the sample rate would result in a deviation of -1 Hz in the frequency of the sampled signal. Due to the amplification effect of sample rate drifts, it is difficult to induce and maintain a DC output especially when the sensor is embedded.

4.2 Digital Amplitude Adjusting

The injected analog signal caused by mechanical resonance of the sensing mass is an oscillating sinusoidal signal. According to (4), the resulting digital signal will also be oscillating (when $\varepsilon \neq 0$). However, an oscillating digital output induced in the sensor could be interpreted as noises or environmental interferences by the system,




Figure 3: When an oscillating analog signal is sampled correctly, the digital signal is oscillating (left). When an oscillating analog signal is undersampled, amplitudes of different digital samples could be adjusted to modify the shape of the digital signal (right).

and its effect could be limited to disturbances or denial of service (DoS). In this subsection, we investigate the possibility to modify the oscillating pattern of the digital signal by modulating the amplitude of analog signals.

An essential feature of out-of-band signal injections is that the induced analog signal will be undersampled, resulting in distortions of the signal. While aliasing is a well-known effect of signal distortions caused by undersampling, it mainly focuses on changes of the signal in the frequency domain, and how to utilize such distortions to intentionally modify the ‘shape’ of an oscillating digitized signal has rarely been discussed.

Due to undersampling, the pattern of the analog signal may not be preserved in the digital signal. As illustrated in Figure 3, when an amplitude modulated oscillating analog signal is sampled correctly, the digital signal has an amplitude that changes gradually and still presents an oscillating pattern. However, when an oscillating analog signal is undersampled, amplitudes of different digital samples within one cycle of oscillation ($T = \frac{1}{\varepsilon}$) could be adjusted to modify the shape of the digital signal. In fact, when $F > \frac{F_S}{2}$, the continuity in the amplitude of the oscillating analog signal kept in digitized samples begins to decrease. As $\frac{2F}{F_S}$ grows, amplitudes of adjacent samples become less dependent on each other. When F is considerably larger than $\frac{F_S}{2}$, each digital amplitude can be adjusted independently. We have,

$$V[i] = A[i] \cdot \sin(2\pi\varepsilon \frac{i}{F_S} + \phi_0) \quad (i \in \{0, 1, 2, 3, \dots\}) \quad (7)$$

where $A[0], A[1], A[2], \dots$ could be adjusted by modulating the amplitude of the oscillating analog signal. In this way, during out-of-band signal injections, a digital signal with specific waveforms (such as a one-sided waveform in Section 5.1) instead of an oscillating signal could be fabricated.

4.3 Phase Pacing

In this subsection, we propose a novel approach to manipulate the phase of the oscillating digitized signal by changing the frequency of out-of-band analog signals.

Assuming the frequency of the analog signal changes from F_1 to F_2 at time t_c , and

$$\begin{aligned} F_1 &= n \cdot F_S + \varepsilon_1 & \left(-\frac{1}{2}F_S < \varepsilon_1 \leq \frac{1}{2}F_S, n \in \mathbb{Z}^+\right) \\ F_2 &= n \cdot F_S + \varepsilon_2 & \left(-\frac{1}{2}F_S < \varepsilon_2 \leq \frac{1}{2}F_S, n \in \mathbb{Z}^+\right) \end{aligned} \quad (8)$$

the analog signal will be:

$$V(t) = \begin{cases} A \cdot \sin(2\pi F_1 t + \phi_0) & 0 \leq t \leq t_c \\ A \cdot \sin(2\pi F_2 (t - t_c) + \phi_1) & t > t_c \end{cases} \quad (9)$$

where ϕ_0 is the initial phase of the analog signal, and ϕ_1 is the phase of the analog signal when we change its frequency at t_c . We have:

$$\phi_1 = 2\pi F_1 t_c + \phi_0 \quad (10)$$

From (9) and (10), we have,

$$V(t) = \begin{cases} A \cdot \sin(2\pi F_1 (t - t_c) + \phi_1) & 0 \leq t \leq t_c \\ A \cdot \sin(2\pi F_2 (t - t_c) + \phi_1) & t > t_c \end{cases} \quad (11)$$

For simplicity, assuming $t_c = \frac{i_c}{F_S}$, the digitized signal will be,

$$V[i] = A \cdot \sin(\Phi[i]) \quad (i \in \{0, 1, 2, 3, \dots\}) \quad (12)$$

where $\Phi[i]$ is the phase of the digital signal. We have,

$$\Phi[i] = \begin{cases} 2\pi\varepsilon_1 \left(\frac{i - i_c}{F_S}\right) + \phi_1 & i \in \{0, 1, \dots, i_c\} \\ 2\pi\varepsilon_2 \left(\frac{i - i_c}{F_S}\right) + \phi_1 & i \in \{i_c + 1, i_c + 2, \dots\} \end{cases} \quad (13)$$

Since $t_i = \frac{i}{F_S}$ is the sampling time, the derivative of the signal's phase will be

$$\Phi'[i] = \begin{cases} 2\pi\varepsilon_1 & i \in \{0, 1, \dots, i_c\} \\ 2\pi\varepsilon_2 & i \in \{i_c + 1, i_c + 2, \dots\} \end{cases} \quad (14)$$

Therefore, when the frequency of the analog signal changes at t_c , the phase of the signal is still ϕ_1 , but the derivative of the phase changes from $2\pi\varepsilon_1$ to $2\pi\varepsilon_2$. Especially, when

$$\varepsilon_1 \cdot \varepsilon_2 < 0, \quad (15)$$

the moving direction of the signal at t_c will be inverted because of the flipped sign of the phase derivative, as illustrated in Figure 4.

In fact, both parts of the digital signal can be represented in terms of positive frequencies. Assuming $\varepsilon_1 > 0$, $\varepsilon_2 < 0$, from (12), (13) and $\sin(x) = \sin(\pi - x)$, we have

$$V[i] = \begin{cases} A \cdot \sin(2\pi\varepsilon_1 \left(\frac{i - i_c}{F_S}\right) + \phi_1) & i \in \{0, 1, \dots, i_c\} \\ A \cdot \sin(2\pi(-\varepsilon_2) \left(\frac{i - i_c}{F_S}\right) + \pi - \phi_1) & i \in \{i_c + 1, \dots\} \end{cases} \quad (16)$$

We can see clearly there is a phase change of $\pi - 2\phi_1$ in the digital signal because of frequency switching at time t_c . We refer to the method that induces a phase offset in the digital signal by switching the frequency of out-of-band analog signals as *Phase Pacing*.

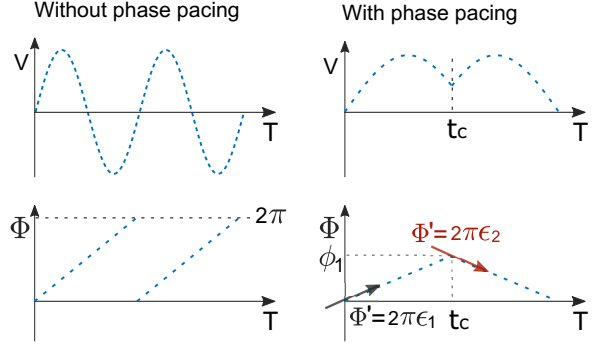


Figure 4: Without phase pacing, the digital signal is oscillating (left). With phase pacing at t_c , the moving direction of the digital signal is inverted due to the flipped sign of its phase derivative (right).

4.4 Out-of-band Signal Injection Model

In summary, during out-of-band signal injections, the digitized signal can be represented by,

$$V[i] = A[i] \cdot \sin(\Phi[i]) \quad (i \in \{0, 1, 2, 3, \dots\}) \quad (17)$$

Where,

$$\Phi[i] = 2\pi\varepsilon \frac{i}{F_S} + \phi_0 \quad (i \in \{0, 1, 2, 3, \dots\}) \quad (18)$$

The parameters that could be manipulated in this model are $A[i]$ and ε . By adjusting $A[i]$, the value of each digitized sample $V[i]$ can be manipulated proportionally. In addition, ε can be altered by changing the frequency of the analog signal. Especially, when the sign of ε is flipped, the moving direction of the digital signal will be inverted because of the phase offset.

5 Attack Methods

Inertial sensors are often used by control systems to ascertain the state of motion. One critical property derived from inertial measurements is the heading angle. A different heading angle detected by the control system often triggers different automated decisions and actuations. Therefore, in this section, we investigate attack methods on embedded inertial sensors to manipulate sensor readings as well as the derived heading angle.

5.1 Side-Swing Attack

The basic idea of Side-Swing attacks is to proportionally amplify the induced output in the target direction and attenuate the output in the opposite direction.

In DoS attacks, the potential accumulative inertial information induced is often limited because an oscillating signal contributes to about the same amount of inertial measurements in both directions. As illustrated in Figure 5, when an oscillating sensor output is induced in a

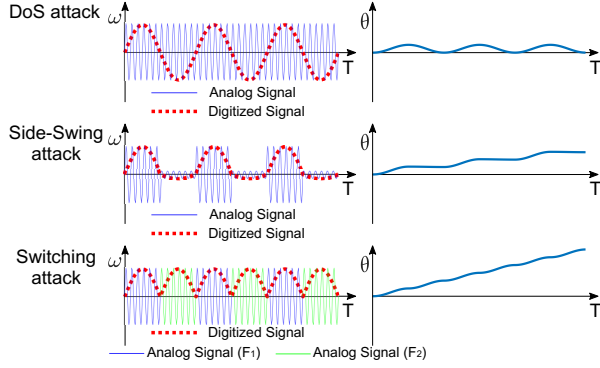


Figure 5: For an oscillating signal, the accumulative heading degree (θ) fluctuates and falls back to 0 after each cycle (top). Under Side-Swing attacks, the derived heading degree grows but only in half of each period of the signal (middle). The derived heading degree under Switching attacks keeps growing (bottom).

gyro, the heading angle θ accumulated in each cycle of oscillation is 0.

To address this problem, in Side-Swing attacks, the attacker can increase the amplitude when the digitized sample is in the target direction and decrease the amplitude otherwise. Recall in (17), we have $V[i] = A[i] \cdot \sin(\Phi[i])$. Assuming that the target direction is the positive direction, the attacker would increase $A[i]$ when $\sin(\Phi[i]) > 0$, otherwise decrease $A[i]$ to 0 or a very small value. In this way, the derived heading angle can be accumulated in the target direction.

Assuming that the injected analog signals are modulated with a high amplitude A_h and a low amplitude A_l alternatively, the heading angle accumulated in each cycle of the signal will be,

$$\theta = \int_0^{\frac{1}{2\varepsilon}} A_h \cdot \sin(2\pi\varepsilon t) + \int_{\frac{1}{2\varepsilon}}^{\frac{1}{\varepsilon}} A_l \cdot \sin(2\pi\varepsilon t) = \frac{A_h - A_l}{\pi\varepsilon} \quad (19)$$

The average angular speed during one cycle is:

$$\bar{\omega} = \varepsilon\theta = \frac{A_h - A_l}{\pi} \quad (20)$$

When $A_l = 0$, the heading angle accumulated in one cycle would be $\frac{A_h}{\pi\varepsilon}$, and the average angular velocity would be $\frac{A_h}{\pi}$. Attackers can adjust these values by adopting different values of A_h . The principle of Side-Swing attacks is illustrated in Figure 5.

We conduct Side-Swing attacks on the gyroscope of an iPhone 5. As shown in Figure 6, while the phone is stationary, the collected gyroscope data shows that it has rotated to the positive direction of X-axis for 17.6 rads (1008°) in about 25 seconds. The peak angular speed ω_{max} is 4.73 rad/s and the average angular speed $\bar{\omega}$ is 0.70 rad/s. The ratio of $\bar{\omega}$ to ω_{max} is 0.15.

In summary, Side-Swing attacks induce the outputs mainly in the target direction and allow the derived heading angle to be manipulated. In control systems, the mov-

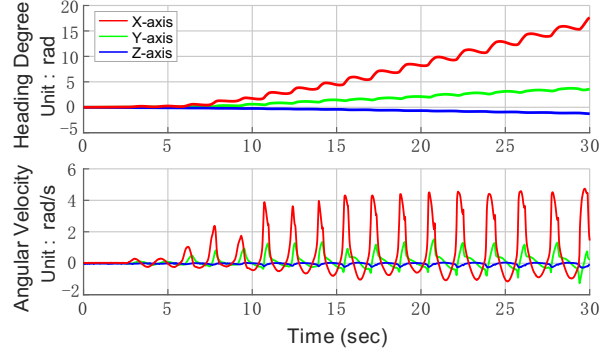


Figure 6: Output of the gyroscope in an iPhone 5 and the derived heading angle under Side-Swing attacks in X-axis. The phone is 0.5 m away from a 50-Watt sound source. The sound frequency is 19,976 Hz.

ing direction and speed of actuators are often determined by the measured angular velocity and the derived heading angle. Therefore, Side-Swing attacks could provide attackers a more direct way to manipulate the control system by modulating the amplitude of acoustic signals. However, during Side-Swing attacks, the derived heading angle increases in only half of each period of the signal and stops growing when the signal is in the opposite direction. This may limit the maximum heading angle accumulated in a certain amount of time.

5.2 Switching Attack

The principle of Switching attacks is to control the induced output by manipulating the phase of the digital signal with repetitive phase pacing.

Recall (8) and (15) in Section 4.3, when $\varepsilon_1 \cdot \varepsilon_2 < 0$ and the frequency of the analog signal changes from F_1 to F_2 , the moving direction of the digital signal will be inverted. Similarly, if the frequency of the analog signal changes from F_2 to F_1 , the condition of phase pacing ($\varepsilon_2 \cdot \varepsilon_1 < 0$) also holds. Therefore, in Switching attacks, the attacker uses two frequencies (F_1 and F_2) and switches the frequency of acoustic signals between them to induce phase pacing repeatedly. Different from Side-Swing attacks, the accumulated heading angle in Switching attacks keeps growing under the sustained influence of the induced angular speed in the target direction, as illustrated in Figure 5.

Assuming the target direction is the positive direction and the attacker switches the frequency when the signal drops from the target direction to the opposite direction, the heading degree accumulated in one period would be:

$$\theta = \int_0^{\frac{1}{2\varepsilon}} A \cdot \sin(2\pi\varepsilon t) + \int_{\frac{1}{2\varepsilon}}^{\frac{1}{\varepsilon}} A \cdot \sin(-2\pi\varepsilon t + \pi) = \frac{2A}{\pi\varepsilon} \quad (21)$$

where we assume $\varepsilon_1 > 0$, $\varepsilon_2 < 0$ and $|\varepsilon_1| = |\varepsilon_2| = \varepsilon$ to simplify the discussion. The average angular speed in one period of the signal is




Figure 7: Output of the gyroscope in an iPhone 7 and the derived heading angle under Switching attacks in Y-axis. The phone is 0.3 m away from a 50-Watt sound source. The sound frequencies are 27,378 and 27,379 Hz.

$$\bar{\omega} = \varepsilon\theta = \frac{2A}{\pi} \quad (22)$$

The values of θ and $\bar{\omega}$ can be adjusted by adopting different amplitudes. In fact, the attacker can switch the frequency more frequently to keep the signal at a higher level and induce a larger heading angle. As shown in Figure 7, we conduct Switching attacks on the gyroscope of an iPhone 7. While the phone is stationary, the collected gyroscope data shows that it has rotated to the positive direction of Y-axis for 6.5 rads (372.4°) in about 25 seconds. The peak angular speed ω_{max} is 0.45 rad/s and the average angular speed $\bar{\omega}$ is 0.26 rad/s. The ratio of $\bar{\omega}$ to ω_{max} is 0.58, which is much larger than 0.15 in the previous experiment with Side-Swing attacks, implying that Switching attacks are more efficient than Side-Swing attacks and could be used to achieve a larger heading angle. However, acoustic frequencies used in Switching attacks should satisfy (8) and (15). We can assume $F_2 = F_1 + step$ ($F_1 < F_2$), and the parameter $step$ can be selected by the attacker to control the length of the interval $[F_1, F_2]$ that bounds the integer multiple of F_3 . In our settings, $step$ is set to 1.

In summary, both Side-Swing and Switching attacks could induce spoofed sensor outputs in the target direction and manipulate the derived heading angle. The target direction can be either positive or negative, determined by the attacker. Theoretically, these methods are not limited to controlling oscillating digitized signals with a very small $|\varepsilon|$. However, in practice, the value of $|\varepsilon|$ should be less than 0.5 or 1, depending on the reaction speed of an attacker. With a very large ε , the signal would oscillate rapidly and may allow not enough time to manually tune acoustic signals effectively. Since the frequency (ε) of the induced signal is closely related to the behavior of the device under attacks, we assume attackers could analyze the behavior of an identical device under acoustic effects to find suitable sound frequencies that could be used in the attack.

6 Evaluations

MEMS inertial sensors are widely used in consumer, industrial, and low-end tactical control systems [55, 58]. Depending on the application, the control algorithm and usage of inertial sensors might be different. Therefore, a key question is: *Can non-invasive spoofing attacks on embedded inertial sensors deliver adversarial control to various types or just one particular type of systems?* The answer to this question will give us a clearer understanding of the potential attack scope and facilitate the evaluation of vulnerabilities that might ubiquitously exist in control systems relying on MEMS inertial sensors.

We evaluate the non-invasive attacks on various types of real systems equipped with MEMS inertial sensors. The results of our attack experiments are summarized in Table 1 and Table 2. Among the 25 tested devices, 17 devices are susceptible to implicit control. In remaining devices, 2 of them can be controlled very limitedly due to insufficient sound strength and 4 of them are vulnerable to DoS attacks. Only 2 devices are not affected by acoustic signals. Our proof-of-concept attacks demonstrate implicit control over various systems including self-balancing, aiming and stabilizing, motion tracking and controlling, navigation systems, etc.

In our experiments, we find that attacks on gyros induce more responsive actuations in the system and demonstrate more adversarial control than attacks on accelerometers. Possible reasons could be that gyros are usually more sensitive, and in most control systems with both gyros and accelerometers, the heading angle of the device is mainly derived from angular velocities measured by gyros, while accelerometers are often used as a gravity sensor and could slowly calibrate the derived orientation information.

6.1 Attack Overview

Without accessing the real-time inertial sensor data, it could be difficult for attackers to decide when to change the amplitude or frequency of acoustic signals so that malicious sensor data is induced in the target direction. However, we find that decisions made by control systems could give away certain information about the induced digital signal, and such information could be observed and leveraged to guide the attack.

During attacks, the induced sensor output could influence actuation decisions of the system instantaneously. For instance, when positive sensor output is detected in the X-axis of the embedded gyro, a self-balancing human transporter would apply forward accelerations to the motor, while negative angular velocities would trigger accelerations to the opposite direction. The amount of the induced acceleration is related to the amount of the spoofed angular velocity. In turn, by observing conse-




Figure 8: An illustration of the reverse signal mapping method. Attackers could reversely infer the current direction and amount of the induced sensor output by observing the consequent actuations or accelerations.

quent actuations or accelerations in the system, attackers could estimate the current direction and amount of the induced sensor output, as illustrated in Figure 8. Another property that could be observed and estimated is the frequency ($|\varepsilon|$) of the induced signal, which could be reversely mapped from the frequency of oscillating movements induced in actuation systems. Such oscillating movements could be periodic accelerations and decelerations of a motor, shaking or circling movements of visual information in VR/AR systems, etc.

The reversely inferring method could be used in following steps to guide the attack:

1) *Profiling*. Before the attack, attackers could analyze the behavior of an identical device under acoustic effects to find the resonant frequency range and profile suitable attack frequencies of the embedded inertial sensor.

To find the resonant frequency range, attackers could generate single-tone sound and sweep a frequency range at an interval of 10 Hz. Attackers apply the sound to a device that is stationary or in a well-balanced status, and there is no other input to control or interfere with the target system. The range of sound frequencies that noticeably affect the motion sensing unit and induce actuations in the device can be recorded as the resonant frequency range. We notice that acoustic frequencies in the middle part of the range could affect the target device more significantly since they are closer to the natural frequency.

Attackers could then generate single-tone sound in the resonant frequency range and adjust the frequency with an interval of 1 Hz or smaller to find and profile attack frequencies. Acoustic frequencies used in our attacks are usually close to the integer multiple of the sensor's sample rate and we have $F = n_0 \cdot F_S + \varepsilon$ ($|\varepsilon| < 1, n_0 \in \mathbb{Z}^+$), where $n_0 F_S$ is an integer multiple of F_S that is in the resonant frequency range of the sensor. Attackers could observe the induced actuations and estimate $|\varepsilon|$. In our settings, when $|\varepsilon| < 1$, the corresponding acoustic frequencies (F) can be considered as suitable attack frequencies.

In practice, due to sample rate drifts, $n_0 F_S$ could fluctuate in a range. As a result, there could be a range of possible attack frequencies. Since we want to use frequencies near $n_0 F_S$, by tracking the range of $n_0 F_S$, the

range of possible attack frequencies can also be located. Attackers could try to make $|\varepsilon|$ as small as possible by adjusting F and estimate $n_0 F_S$ from $F = n_0 F_S + \varepsilon$.

Empirically, the drift of $n_0 F_S$ is usually less than 1 Hz in 1 or 2 minutes, but the accumulative drift in a long time could be larger and $n_0 F_S$ could fluctuate in a frequency range with a width of around 10 Hz. We track $n_0 F_S$ of the gyro in an iPhone 5 for 3 hours and find that it fluctuates in the range of 19,966 to 19,976 Hz. While it might be difficult to predict $n_0 F_S$ deterministically, we notice that $n_0 F_S$ tends to decrease as the target system is running, which could be caused by the increased temperature. For instance, when we just turn on a gyro-based application in an iPhone 5, $n_0 F_S$ is more likely to be close to 19,975 Hz. If the application has been running for a while, $n_0 F_S$ may become close to 19,970 Hz. If the application has been running for a long time such as an hour, $n_0 F_S$ could be between 19,966 to 19,970 Hz.

2) *Synchronizing*. Based on the profiled range of possible attack frequencies, attackers could select a frequency that is more likely to be close to $n_0 F_S$ and adjust the sound frequency to 'synchronize' to a suitable attack frequency to initiate the attack.

Attackers could observe changes in $|\varepsilon|$ while they are adjusting F . Based on $F = n_0 F_S + \varepsilon$, if the observed $|\varepsilon|$ decreases when F increases, attackers could infer $F < n_0 F_S$ and $\varepsilon < 0$. Otherwise, they could infer $\varepsilon > 0$ and F should be decreased to get closer to $n_0 F_S$. In this way, attackers could adjust F more effectively since they could infer the sign of ε and know whether the adjusted F is getting closer to or further away from $n_0 F_S$.

After synchronizing to a frequency F with $|\varepsilon|$ less than 0.5 or 1, attackers could start Side-Swing attacks. For Switching attacks, if attackers find a suitable F_1 with $-1 < \varepsilon_1 < 0$, they could find F_2 by $F_2 = F_1 + 1$. Similarly, they could also acquire $F_1 = F_2 - 1$ if they find a suitable F_2 with $0 < \varepsilon_2 < 1$. Usually, we make both $|\varepsilon_1|$ and $|\varepsilon_2|$ close to 0.5 so that $n_0 F_S$ is well bounded by $[F_1, F_2]$.

In our settings, this process involves manually tuning the acoustic frequency with an off-the-shelf function generator and observing consequent actuations of the target device. Usually, such interactions between attackers and the target system could take about 10 to 60 seconds.

3) *Manipulating*. In Side-Swing attacks, attackers can increase the amplitude when the induced actuation is in the target direction and otherwise decrease the amplitude. In Switching attacks, attackers can switch the frequency of acoustic signals when the induced actuation or acceleration in the target direction begins to attenuate.

4) *Adjusting (optional)*. After several minutes of manipulation, $n_0 F_S$ could deviate from F because of sample rate drifts. Attackers could accommodate the deviation by observing changes in ε and adjusting F . For exam-




Figure 9: Unweighted SPL measurements of different speakers we use. The speaker is placed 10 cm from the microphone and operated near its maximum amplitude.

ple, if attackers observe that $\varepsilon < 0$ and $|\varepsilon|$ increases, they could infer that $n_0 F_S$ has increased and could increase F to compensate for the deviation.

6.2 Experimental Setup

In our experiments, we use several types of consumer-grade tweeter speakers, including two electromagnetic (EM) speakers [20, 21] and one piezo speaker [17]. We measure the Sound Pressure Level (SPL) of the speakers with an NI USB-4431 sound measuring instrument and a GRAS 46AM free-field microphone that has a wide frequency range. The speaker plays single-tone sound from 1.5 kHz to 31.5 kHz with an interval of 100 Hz. We set the sample rate of the microphone to 96 kHz instead of 48 kHz to pick up ultrasonic signals correctly.

Figure 9 shows the average SPL values of the speakers, from which we can select a speaker that has the maximum SPL for each attack. The SPL of our sound source can be represented by $\max(SPL_{em1}, SPL_{em2}, SPL_{piezo})$. By selecting from multiple speakers, we avoid sharp performance degradations of one specific speaker in certain frequency bands and enhance the overall performance of the sound source. The resulting improvement of SPL can be crucial in attacks on embedded sensors since the actual sound pressure grows exponentially as the sound level increases; a gain of 6.02 dB in SPL doubles the amount of sound pressure. During attacks, we use a directivity horn, such as [16] and [19], to improve the directivity of the sound source. The speaker is powered by a 50-Watt Lepy LP-2051 audio amplifier and the signal source is an Agilent 33220A function generator. We conduct the experiments indoor and put acoustic foams in the environment to reduce potential sound reflections.

In Table 1 and Table 2, we measure the maximum horizontal distance D_{Max} between the sound source and the target device that an observable actuation or an inertial output with an amplitude of 0.1 rad/s can be induced under acoustic effects. Empirically, the possible attack distance with our sound source is about $\frac{D_{Max}}{4}$ for Side-




Figure 10: An illustration of Side-Swing attacks on a self-balancing scooter. The system is tricked to actuate its motor based on the spoofed angular speed. The attack is demonstrated in [6].

Swing attacks, and $\frac{D_{Max}}{3}$ for Switching attacks to achieve adversarial control. Manufacturer information of inertial sensors is collected for statistical purposes. We find sensor information of iPhones and VR devices in online disassembling reports [15]. Android devices provide APIs to retrieve sensor information. We disassemble other devices to reveal the information written on the package of the embedded inertial sensor, but some devices do not specify the sensor model explicitly even on the sensor’s package. Lastly, we record the alignments of affected and functional axes based on the orientation of the sensor when the embedded inertial sensing module is recognized. Otherwise, the alignments of axes are based on the orientation of the device.

6.3 Experiments on Closed-loop Systems

In a closed-loop control system, there is usually a goal state. The system continuously compares the goal state with its current state based on inertial measurements and tries to diminish the difference between them through actuations. We evaluate our attacks on different instances of four types of closed-loop systems, including self-balancing human transporters, robots, stabilizers, and anti-tremor devices. These systems present different features under acoustic effects. Nevertheless, we find that a large part of them are susceptible to implicit control.

(1) *Human transporters*. The goal state of self-balancing human transporters is a vertical position of the system with a tilt angle of 0° . Inertial sensors are used to detect tilts of the transporter. Based on the direction and amount of the tilt, the control system applies accelerations to motors to correct the position of the system.

We evaluate acoustic attacks on four instances of self-balancing transporters: a Megawheels TW01 scooter, a Veeko 102 scooter, a Segway one S1 unicycle, and a Segway Minilite scooter. We find that, by spoofing the angular speed measured by gyros, the moving direction and speed of the motor could be controlled, as illustrated in Figure 10.

Results. The Megawheels scooter and the Veeko scooter are vulnerable to adversarial control over the moving direction and speed of the motor through ultrasonic signals.

Table 1: Results of our attack experiments on closed-loop control systems

Device	Sensor		Resonant Freq. (kHz)	Affected/ Func. Axes	Max Dist. (m)	Control Level
	Type	Model [†]				
Megawheels scooter	Gyro	IS MPU-6050A	27.1~27.2	y/y	2.9	Implicit control
Veeko 102 scooter	Gyro	Unknown	26.0~27.2	x/x	2.5	Implicit control
Segway One S1	Gyro	Unknown	20.0~20.9	x/x	0.8	Implicit control
Segway Minilite	Gyro	Unknown	19.2~20.0	x/x	0.3	DoS
Mitu robot	Gyro	N/A SH731	19.0~20.7	x/x	7.8	Implicit Control
MiP robot	Acce	Unknown	5.2~5.4	x/x	1.2	DoS
DJI Osmo stabilizer	Gyro	IS MP65	20.0~20.3	x,y,z/x,y,z	1.2	Implicit control
WenPod SP1 stabilizer	Gyro	IS MPU-6050	26.0~26.9	z/y,z	1.8	Implicit control
Gyenno steady spoon	Gyro	Unknown	Not found	Unknown	N/A	Not affected
Liftware level handle	Acce	IS MPU-6050	5.1	x/x	0.1	DoS

[†] IS: InvenSense, N/A: Unknown manufacturer.

While the Segway One S1 unicycle can be manipulated by Switching attacks, the range of induced actuations is very small. The unicycle only tilts slightly to the target direction. The Segway Minilite scooter tends to lose control under acoustic effects. Our Side-Swing attacks and Switching attacks on smart human transporters are demonstrated in [6] and [11]¹. The transporter is in a relatively static experimental setting, and we lift the wheels of the transporter up from the ground during the experiments.

(2) *Robots*. Self-balancing robots work similarly to self-balancing human transporters but without a rider. We test two self-balancing robots equipped with MEMS gyros and accelerometers: a Mitu robot and a MiP robot.

Results. We find that the gyro of Mitu robot is susceptible to adversarial control. The robot would speed up to the same direction as the spoofed rotations under Side-Swing attacks, as demonstrated in [5]. While the gyro of MiP robot is not affected by acoustic attacks, its accelerometer is vulnerable to DoS attack, which makes it suddenly stop working and fall to the ground.

(3) *Stabilizers*. MEMS inertial sensors are widely used in aiming and stabilizing systems. The goal of such systems is to maintain a device or platform in a certain orientation despite external forces or movements. Therefore, when movements are detected by inertial sensors, the system would actuate in opposite directions to cancel the effect of external movements.

We evaluate our attacks on two camera stabilizers: a DJI Osmo stabilizer and a Wenpod SP1 stabilizer. Our results show that by spoofing the gyro and manipulating the derived heading angle, the pointing direction of a stabilizer could be controlled. However, fabricated heading angles in X and Y axes will be gradually calibrated by the system based on gravity information. As illustrated




Figure 11: An illustration of Switching attacks on a stabilizer. The stabilizer tries to correct the fabricated heading angle in Y-axis of the device by rotating to the opposite direction. The attack is demonstrated in [13].

in Figure 11, we can use Switching attacks to induce a maximum heading degree in the stabilizer. As the induced heading angle increases, the calibration effect also becomes stronger until the maximum heading angle is reached.

Results. Both instances of stabilizers are vulnerable to adversarial control through ultrasonic signals. The Osmo stabilizer is mainly affected in X-axis while the Wenpod stabilizer can only be manipulated in Y-axis of the device (which is the Z-axis based on the orientation of the embedded inertial sensor). Our Side-Swing attacks and Switching attacks on stabilizers are demonstrated in [8] and [13].

(4) *Anti-tremor Devices*. Inertial sensors can be used by anti-tremor gadgets in health-care applications, such as gyroscopic tablewares and gloves [32] that mitigate hand tremors and assist users to perform daily tasks. We evaluate acoustic attacks on a Liftware level handle and a Gyenno gyroscopic spoon.

Results. The Liftware handle is vulnerable to DoS attacks on its accelerometer. The handle under attacks would abnormally actuate its motor to one direction and become unusable. The Gyenno gyroscopic spoon is not affected by acoustic signals.

¹Precautions were used to ensure the safety of researchers.

6.4 Experiments on Open-loop Systems

Different from closed-loop systems that have a goal state, open-loop control systems simply take inertial measurements as inputs and actuate accordingly. We evaluate our attacks on various types of devices that use real-time inertial data for open-loop control. These devices use various MEMS inertial sensors from different vendors. Nevertheless, we find that most of them could be susceptible to implicit control.

(1) *3D mouses.* Inertial sensors can be used in input devices for remote control. 3D mouses use gyros to detect a user's hand movements and move the cursor accordingly. We evaluate our spoofing attacks on an IOGear 3D mouse and a Ybee 3D mouse.

Results. Both instances of 3D mouse are vulnerable to adversarial control through ultrasonic signals. By spoofing the gyroscope, attackers could point the cursor of the 3D mouse in a remote system to different targets. We demonstrate Side-Swing attacks and Switching attacks on 3D mouses in [4] and [9].

(2) *Gyroscopic screwdrivers.* The gyroscopic screwdriver is an industrial application that controls a mechanical system based on inertial measurements. The moving direction and speed of the motor in the screwdriver is decided by the heading angle derived from gyroscope data.

In gyroscopic screwdrivers, there is usually no mechanism to calibrate the heading angle. Therefore, the induced heading angle will not be eliminated even when the attack ceases. Based on this feature, we adjust our attack method to *Conservative Side-Swing Attacks*. The basic idea is that attackers emit acoustic signals only when changing the direction or speed of the motor. Once the motor is tricked to move with a desired speed in the target direction, attackers can turn off acoustic signals to keep the heading angle in the system, as illustrated in Figure 12. We evaluate our attacks on an E-design ES120 screwdriver, a B&D gyroscopic screwdriver, and a Dewalt gyroscopic screwdriver.

Results. By spoofing the gyro and manipulate the derived heading angle, both the moving direction and speed of the motor in the ES120 screwdriver can be controlled. The B&D screwdriver can be manipulated only after we remove its external panel and the Dewalt screwdriver is not affected by acoustic signals.

(3) *VR/AR devices.* Inertial sensors are used by Virtual/Augmented Reality (VR/AR) headsets and kinetic controllers to track the user's movements and control visual information in an image system. The user's view in VR systems or the position of augmented information displayed in AR systems is often determined by heading angles of the headset. In addition, the movements de-

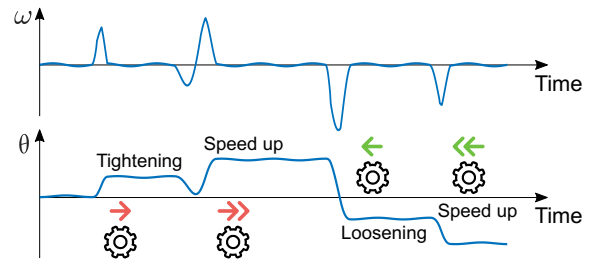


Figure 12: An illustration of Conservative Side-Swing attacks on a screwdriver. Both the moving direction and speed of the motor can be manipulated by spoofing the gyroscope. The attack is demonstrated in [2].

TECTED by the kinetic controller will directly be used to control an object in the image system. We evaluate our attacks on an Oculus Rift VR headset, an Oculus Touch controller, and a Microsoft HoloLens AR headset.

Results. By spoofing the gyros with ultrasonic signals, the user's view in Oculus Rift headset and the orientation of an object controlled by Oculus Touch can both be manipulated in X-axis. The HoloLens headset can only be affected very slightly by our sound source. Our Switching attacks on VR devices are demonstrated in [10] and [14]. Recent researches have shown that buggy or maliciously exploited visual information in an immersive environment might startle or mislead a user and cause unexpected consequences [50, 51]. Furthermore, a few prototype products use AR applications to assist critical real-world tasks [33, 31], and plenty of studies utilize inertial measurements to remotely control mechanical systems such as a robotic arm [38]. Our experimental results might help designers of these rapidly emerging applications to be aware of potential threats that might be caused by spoofing inertial sensors.

(4) *Smartphones.* Smartphones have become a platform that provides sensor data and computation resource for large amounts of applications. Inertial sensor data of smartphones is often used in mobile VR/AR applications and navigation systems. We evaluate our attacks on six smartphones in different models. Both iOS and Android devices are tested.

Results. The smartphones we test have different gyroscopes, which have different resonant frequency ranges. While their sensitivity to resonant sound differs, we find that all of them are vulnerable to adversarial control. Our Side-Swing attacks and Switching attacks on mobile VR applications are demonstrated in [7] and [12]. In the demos, we manipulate the VR user's view and aim several targets by spoofing the gyroscopic sensor.

(5) *Motion-aware devices.* Using inertial sensors to detect motions is a popular wake-up mechanism in smart devices. This mechanism can also be used to control

Table 2: Results of our attack experiments on open-loop control systems

Device	Sensor		Resonant Freq. (kHz)	Affected/ Func. Axes	Max Dist. (m)	Control Level
	Type	Model [†]				
IOGear 3D mouse	Gyro	IS M681	26.6~27.6	x,z/x,z	2.5	Implicit control
Ybee 3D mouse	Gyro	Unknown	27.1~27.3	x/x,z	1.1	Implicit control
ES120 screwdriver	Gyro	ST L3G4200D	19.8~20.0	y/y	2.6	Implicit control
B&D screwdriver	Gyro	IS ISZ650	30.3~30.6	z/z	0	Limited control
Dewalt screwdriver	Gyro	Unknown	Not found	none/y	N/A	Not affected
Oculus Rift	Gyro	BS BMI055	24.3~25.6	x/x,y,z	2.4	Implicit control
Oculus Touch	Gyro	IS MP651	27.1~27.4	x/x,y,z	1.6	Implicit control
Microsoft Hololens	Gyro	Unknown	27.0~27.4	x/x,y,z	0	Limited control
iPhone 5	Gyro	ST L3G4200D	19.9~20.1	x,y,z/x,y,z	5.8	Implicit control
iPhone 5S	Gyro	ST B329	19.4~19.6	x,y,z/x,y,z	5.6	Implicit control
iPhone 6S	Gyro	IS MP67B	27.2~27.6	x,y,z/x,y,z	0.8	Implicit control
iPhone 7	Gyro	IS 773C	27.1~27.6	x,y,z/x,y,z	2.0	Implicit control
Huawei Honor V8	Gyro	ST LSM6DS3	20.2~20.4	x,y,z/x,y,z	7.7	Implicit control
Google Pixel	Gyro	BS BMI160	23.1~23.3	x,y,z/x,y,z	0.4	Implicit control
Pro32 soldering iron	Acce	NX MMA8652FC	6.2~6.5	Unknown	1.1	DoS

[†] IS: InvenSense, ST:STMicroelectronics, BS: Bosch, NX: NXP Semiconductors.

critical functions of an embedded system. The Pro32 soldering iron uses an accelerometer to detect movements. If there is no movement for a long time, the system will cool down the iron tip and go into the sleep mode. This protects the iron from overheating and reduces the risk of accidental injuries or fire. However, we find that this mechanism could be compromised by resonant acoustic interferences. Our experiments show that attackers can wake the Pro32 soldering iron up from the sleep mode through DoS attacks on the accelerometer, and make the iron tip heat up to a high working temperature repetitively. The attack is demonstrated in [3].

7 Automatic Attack

In this section, we present a novel automatic attack method and implement a proof-of-concept spoofing attack on a mobile navigation system. We find that in both iOS and Android smartphones, inertial sensor data can be accessed through a script in a web page or an application without any permission. In our scope, a key question is: *Can an attack program facilitate spoofing attacks on inertial sensors by leveraging the real-time sensor data?* To answer this question, we investigate automatic methods to implement Switching attacks.

Automatic Method. In automatic attacks, the attack program modulates acoustic signals automatically based on parameters set by the attacker. These parameters include initial sound frequencies, threshold, target direction, etc. The attacker can set the initial sound frequencies F_1 and F_2 based on the real-time feedback of the sensor. The threshold is used by the attack program to decide when to switch the sound frequency. During attacks, the at-

tacker can send commands to the program to change the target direction, to stop or restart the attack.

The attack program monitors the output of the sensor and switches the frequency of acoustic signals between F_1 and F_2 when the induced signal drops to the opposite direction and falls below a threshold. However, we find that this setting only allows the program to attack automatically for one or two minutes. After two minutes, the integer multiple of the sensor's sample rate might fall outside (F_1, F_2) because of drifts in F_S and the condition of phase pacing ($\varepsilon_1 \cdot \varepsilon_2 < 0$) would no longer hold. As a result, the attacker would need to manually adjust the sound frequencies every one or two minutes.

A method to address this issue is to *actively adapt to the drifts in the sample rate*. Due to drifts in F_S , the value of $n_0 F_S$ may become $n_0 \hat{F}_S$. If $n_0 \hat{F}_S$ falls outside (F_1, F_2) , the condition of phase pacing will no longer be satisfied. Therefore, the goal of adaptation is to actively adjust the sound frequencies to \hat{F}_1 and \hat{F}_2 so that $n_0 \hat{F}_S$ is at the midpoint of (\hat{F}_1, \hat{F}_2) . Assuming $\varepsilon_1 < 0, \varepsilon_2 > 0$, we have,

$$F_1 - \varepsilon_1 = n_0 \hat{F}_S = F_2 - \varepsilon_2 \quad (23)$$

After adaptation, we would have,

$$\hat{F}_1 + \frac{\varepsilon_2 - \varepsilon_1}{2} = n_0 \hat{F}_S = \hat{F}_2 - \frac{\varepsilon_2 - \varepsilon_1}{2} \quad (24)$$

Therefore,

$$\Delta F = \hat{F}_1 - F_1 = \hat{F}_2 - F_2 = -\frac{\varepsilon_1 + \varepsilon_2}{2(\varepsilon_2 - \varepsilon_1)}(\varepsilon_2 - \varepsilon_1) \quad (25)$$

Since $\varepsilon_2 - \varepsilon_1 = F_2 - F_1$, we have,

$$\Delta F = \frac{r-1}{2(r+1)}(F_2 - F_1) \quad (26)$$

where $r = \frac{|\varepsilon_1|}{|\varepsilon_2|} = \frac{-\varepsilon_1}{\varepsilon_2}$, and can be derived from

$$r = \frac{T_2}{T_1} \approx \frac{T_2'}{T_1'} \quad (27)$$




Figure 13: Controlling the orientation of a mobile navigation system with automatic Switching attacks on the gyroscopic sensor. The attack is demonstrated in [1].

T_1 and T_2 are periods of the induced signals. The ratio $\frac{T_2}{T_1}$ can be estimated by $\frac{T'_2}{T'_1}$, where T'_1 and T'_2 correspond to the time intervals between adjacent frequency switching operations. During attacks, T'_1 and T'_2 can be recorded by the program. The program computes ΔF and adapts the frequencies after every two times of frequency switching.

Evaluation. We evaluate our attacks on a Huawei Honor V8 smartphone and demonstrate the attack effects with a mobile navigation system (Google Maps). In mobile navigation systems, inertial sensors are often used to aid the GPS system to provide a more timely and accurate positioning service. The gyroscope is often used to determine the orientation of the system.

We implement the automatic attack method in an Android application. The application utilizes the smartphone's built-in speaker to generate ultrasonic signals and surreptitiously manipulate the gyroscope data while running in the background. As shown in Figure 13, we first induce positive outputs in the Z-axis of gyro and the navigation system is tricked to rotate its orientation counter-clockwisely. The accumulated heading angle is 6.85 rads in 32 seconds. After we change the target direction, the navigation system is deceived by negative outputs and rotates the orientation clockwisely. The accumulated heading angle is -6.82 rads in about 31 seconds.

Our results show that, with real-time sensor data, spoofing attacks on inertial sensors could manipulate the orientation of a navigation system. When the displayed orientation of a navigation system is manipulated, users or systems guided by the navigation information could be led to a wrong path. Additionally, for areas not well covered by GPS or situations when the GPS signal is jammed or spoofed [56, 60], errors in the orientation information will not be effectively calibrated and could cause more troubles to the positioning service.

Several recent approaches have been proposed to control the access to inertial sensors in smartphones, but with a focus on privacy issues [59, 63]. Our automatic attack also demonstrates that unprotected inertial sensor data could be leveraged to manipulate the sensor output. Our results confirm that protection mechanisms over in-

ertial sensor data are necessary. Devices should control the access to the sensor data. In addition, when a remote autonomous agent transmits real-time inertial sensor data for navigation purposes, the data should be encrypted.

8 Discussion

8.1 Countermeasures

It is important to protect control systems from sensor spoofing attacks, however, feasible countermeasures to be deployed in embedded systems should not cause too much expenses in cost and size or compromises in designs. Therefore, the countermeasures we discuss mainly focus on two aspects: (1) Damping and isolation. These approaches mitigate acoustic or vibrational noises physically. (2) Filtering and sampling. These approaches eliminate or mitigate malicious signals in the signal conditioning circuits.

Damping and Isolation. Early mitigation approaches against acoustic interferences include using isolating boxes and acoustic foams to surround the sensor [41]. The simple strategy could achieve substantial protection from acoustic noises, but issues in size and design concerning an embedded environment were not addressed.

To protect MEMS inertial sensors without compromising their advantages in size, weight, power, and cost (SWaP-C [48]), recent studies have been dedicated to using micro-level techniques for acoustic isolation. Dean et al. proposed the use of microfibrus metallic cloth as an acoustic damping material to protect MEMS gyroscopes [43]. Soobramaney et al. evaluated the mitigation effects of microfibrus cloth on noise signals induced in MEMS gyros under acoustic interferences [65]. They tested 7 MEMS gyros and showed that, by surrounding the sensor with 12 mm of the media, 65% reduction in the amplitude of noise signals can be easily obtained and up to 90% reduction could be achieved [65]. Additionally, Yunker et al. suggested to use MEMS fabricated acoustic metamaterial to mitigate acoustic signals at frequencies close to the resonant frequency of the MEMS gyroscope [76]. Furthermore, Kranz et al. showed that a MEMS-fabricated micro-isolator can be applied within the sensor packaging but their work mainly focused on isolating mechanical vibrations [48].

Filtering. As suggested in [68], a low-pass filter (LPF) should be used to eliminate the out-of-band analog signals. According to the datasheets [30, 28], we find that many inertial sensors have an analog LPF in their circuits, but are still vulnerable to acoustic attacks, which could be due to a cut-off frequency that is set too high. We also find that most programmable inertial sensors use a digital LPF for bandwidth control [27, 29]. However,

filters in digital circuits will not alleviate the problem because out-of-band analog signals have already been aliased to in-band signals after sampling.

Sampling. Trippel et al. proposed randomized sampling and 180° out-of-phase sampling methods for inertial sensors with analog outputs and software controlled ADCs [68]. These approaches were designed to eliminate an attacker’s ability to achieve a DC signal alias and limit potential adversarial control. However, adding a randomized delay to each sampling period or computing the average of two samples at a 180° phase delay could degrade the accuracy of inertial measurements. Small errors in the measurements could accumulate in a long time and might affect the performance of the system.

We think an alternative sampling method to mitigate potential adversarial control without degrading the performance is to use a *dynamic sample rate*. Recall in (3) and (4), the frequency ε of the induced digital signal depends on both F and F_S . With a dynamic F_S , attackers may not be able to induce a digital signal with a predictable frequency pattern. In this case, the ability of attackers will be limited and it could be difficult for attackers to accumulate a large heading angle in a target direction. This might be a general mitigation method for ADCs subject to out-of-band signal injections.

Additionally, redundancy-based approaches could enhance the resilience of the system. For example, multiple sensors could still provide trustworthy information when one of them is under attack. It might still be possible to attack or interfere several sensors simultaneously to affect the functioning of the system, but such attacks could be more difficult to implement.

In summary, acoustic attacks on inertial sensors are enabled by two weaknesses in the analog domain: (1) Susceptibility of the micro inertial sensing structure to resonant sound. (2) Incapability of signal conditioning circuits to handle out-of-band analog signals properly. Employing both acoustic damping and filtering approaches in the designs of future sensors and systems can address these weaknesses. Additionally, acoustic damping can also be used to mitigate the susceptibility of currently deployed sensors and systems to acoustic attacks.

8.2 Sound Source

Applications of sonic weapons [34], ultrasonic transducers [47], and long-range acoustic devices [18, 26] have already shown the capability of specialized devices to generate more powerful sound with a further transmitting distance than common audio devices. In addition, we find several consumer-grade techniques that could be used to optimize a sound source.

The most direct acoustic amplification method is to

use speakers and amplifiers with better performance and output capabilities. Besides, the sound played by common audio speakers usually diffuses into the air with little directivity, leading to losses of acoustic energy. With directivity horns [16, 19], the sound waves can be focused into a certain emitting area and transmit through a longer distance. Another important approach is to use multiple speakers to form a specialized speaker array. With appropriate arrangement of speakers and directivity horns to focus the sound waves, the sound strength, transmitting distance, and emitting area of the sound source could be customized and improved. Moreover, ultrasonic transducers [73, 72] could have small sizes, variable resonant frequencies, and high efficiency. It might be possible to build a more powerful and efficient sound source by selecting and using a large number of transducers.

With multiple speakers or transducers, the performance of a sound source could be improved. If the sound waves are in phase, the add-up SPL of n coherent sources could be [25],

$$L_{\Sigma} = 20 \log_{10} \left(10^{\frac{L_{p1}}{20}} + 10^{\frac{L_{p2}}{20}} + \dots + 10^{\frac{L_{pn}}{20}} \right) \quad (28)$$

Assuming each coherent source is identical, we have

$$L_{\Sigma} = 20 \log_{10}(n) + L_{p1} \quad (29)$$

Theoretically, with 8 identical sources, the level increase could be $L_{\Sigma} - L_{p1} \approx 18.0$ dB. In practice, the performance could also depend on arrangements of multiple sources, designs of the enclosure and horns, and differences in phases need to be considered and accommodated. The distance attenuation of SPL can be quantified by [23]: $L'_p = L_p + 20 \log_{10}(\frac{D}{D'})$, where D and D' are distances. Therefore, a level increase of 18.0 dB could increase the possible attack distance by a factor of 8.

8.3 Limitations

Moving targets. Depending on the speed and range of movements, it could be difficult for attackers to follow and aim a moving target while manually tuning acoustic signals. It could be helpful to predict the movements and align the sound beam with the trajectory of the target. In certain circumstances, it might be possible to attach a sound source to the victim device or exploit a sound source in close proximity to the device. Additionally, it might be possible to carry the sound source with a vehicle or drone that follows the target.

Ideally, an automatic tracking and aiming system might be implemented to aim the target. It might use cameras or radar sensors to track the position of a target and use a programmable 3-way pan/tilt platform to aim.

Timing. In our experimental settings, attackers observe actuations of a target and manually tune acoustic signals

with off-the-shelf devices. In certain circumstances, however, such settings could be slow and ineffective; it might be difficult for attackers to analyze the observed movements and modulate signals timely and correctly.

To reduce potential delays caused by hand tuning and observing, it might be possible to use more customized devices, tools, and programs. As we have investigated in Section 7, a program could help attackers to modulate acoustic signals more timely and accurately. Moreover, it might also be possible to use systems with cameras or radar sensors to help attackers observe and analyze the behavior of a target more automatically.

In addition, the pattern of a closed-loop system could be more complex than the simple signal mapping model in Section 6.1. For example, when a user is riding the self-balancing scooter, user involvements (including unintentional involvements) could counter or disrupt attack effects. Attackers might need a more specific model to analyze and predict the movement patterns.

8.4 Generalization

Acoustic attacks on inertial sensors exploit resonance and inject analog signals with very high frequencies. To explore the generalizability of the out-of-band signal injection model and attack methods, we investigate whether the oscillating digitized signal could be manipulated when analog signals are sent at relatively low frequencies through a more common sensing channel.

We use a vibrating platform to generate mechanical vibration signals and implement Side-Swing and Switching attacks on the accelerometer of a smartphone, as shown in Figure 14. We place the Google Pixel smartphone on the platform. In Side-Swing attacks, we generate sinusoidal vibration signals at 19.6 Hz. While the phone remains on the platform, the collected accelerometer data shows that the phone is launched to the sky and has accumulated a speed of 73.9 m/s in about 25 seconds. In Switching attacks, we switch the frequency of the sinusoidal vibration signal between 19.4 Hz and 20.4 Hz. While the phone is still placed on the platform, the accelerometer data shows that it has accumulated an upward speed of 74.5 m/s in about 25 seconds.

We try to find the approximate sample rate of the embedded accelerometer by inducing an aliased DC-like signal. We increase the vibration frequency with an interval of 0.1 Hz and observe the induced output. The first DC-like signal is induced at $F = 19.9$ Hz, the second at 39.8 Hz, and the third at 59.7 Hz. Based on $F = nF_S + \epsilon_0$ ($\epsilon_0 \approx 0$), we infer that the sample rate of the ADC is approximately 19.9 Hz.

Our experimental results show that, when analog signals are sent at relatively low frequencies, such as frequencies close to F_S , the oscillating digitized signal could




Figure 14: The output of the accelerometer (Z-axis) in a Google Pixel smartphone. We implement Side-Swing (top) and Switching attacks (bottom) with low-frequency vibration signals to manipulate the sensor output. The phone is placed with the Z-axis pointing upward, and the default output in Z-axis is 1 g if the device is at rest.

still be manipulated. Moreover, instead of exploiting resonance, malicious signals could be injected and manipulated through the sensing channel as well.

As we have discussed, sensors without a correctly configured analog LPF could be vulnerable to out-of-band signal injections. Furthermore, some digital sensors could have a configurable sample rate and use a programmable digital LPF for bandwidth control. For example, the ADC sample rate of the MPU-6500 gyroscope is programmable from 8,000 samples per second, down to 3.9 samples per second [29]. In this case, assuming the cut-off frequency of the analog LPF is 4 kHz, which is the half of the maximum sample rate, if applications set F_S to 4 kHz or lower, out-of-band signals with relatively low frequencies (such as frequencies close to F_S) would not be eliminated by the analog LPF and could be exploited to manipulate the digitized signal.

9 Related Work

Since measurements of embedded sensors are often trusted by control systems to make critical decisions, the security of analog sensors has become an increasingly important concern. This section discusses security of inertial sensors and attacks against analog sensors.

Attacks on Inertial Sensors. MEMS inertial sensors have drawn the attention of recent security researches because of their criticality in control systems. Son et al. [64] proposed a DoS attack against MEMS gyroscopes and showed that a drone could be caused to crash by intentional resonant sound. Additionally, Wang et al. developed a sonic gun and showed that a range of smart devices could lose control under acoustic attacks on inertial sensors [71]. Furthermore, Trippel et al. [68] proposed output biasing attacks and output control attacks to com-

promise the integrity of MEMS accelerometers. However, output biasing attacks were only implemented on exposed sensors with an insufficiently realistic attack setting; while the output control attack method only works on sensors with an insecure amplifier and the generalizability could be limited in two aspects: (1) To trigger signal clipping in the amplifier, the amplitude of the induced analog signal needs to exceed the operating range of the amplifier. (2) The direction of induced outputs is determined by the asymmetry of signal clipping that occurs in the saturated amplifier and cannot be controlled. Different from prior works, this work shows that an oscillating digitized signal, which is often regarded as noises, could be manipulated to deliver adversarial control, and demonstrates implicit control over different kinds of real systems through non-invasive attacks against embedded inertial sensors.

Eavesdropping through Inertial Sensors. Inertial sensors have become ubiquitous in mobile devices. It is known that access to inertial sensors in both iOS and Android devices does not require permissions from the operating system [40, 53]. Therefore, attackers could surreptitiously read inertial sensor data through either a web script or a malicious application. The inertial sensing data in smartphones could be used to recover keystroke information [40, 37, 54]. Furthermore, the works of [53] and [35] showed that it might be possible to utilize inertial sensors in a smartphone to eavesdrop speech signals in certain scenarios. Additionally, a user's keystroke information could be recovered by exploiting inertial sensors in smart watches [52, 69, 70]. More recent studies showed that inertial sensors in mobile devices could be exploited to establish a covert channel due to their sensitivity to vibrations [46, 39]. All these works focused on utilizing inertial sensing data for eavesdropping or data exfiltration purposes. To our knowledge, the automatic attack we demonstrate is the first method that leverages inertial sensor data to manipulate the sensor output with a malicious program.

Analog Sensor Spoofing Attacks. Foo Kune et al. showed that bogus signals could be injected into analog circuits of a sensor through electromagnetic interference to trigger or inhibit critical functions of cardiac implantable electrical devices [49]. Park et al. studied a saturation attack against infrared drop sensors to manipulate the dosage delivered by medical infusion pumps [57]. In automotive embedded systems, Shoukry et al. presented non-invasive spoofing attacks on magnetic wheel speed sensors in anti-lock braking systems [62]. Yan et al. investigated contactless attacks against environment perception sensors in autonomous vehicles [74]. Recently, Shin et al. studied spoofing attacks on Lidar sensors in automotive systems to manipulate the distance of objects

detected by the system [61]. In addition, Davidson et al. investigated a sensor input spoofing attack against optical flow sensing of unmanned aerial vehicles [42]. Finally, Zhang et al. presented an inaudible attack on voice controllable systems that injects commands into a microphone through ultrasonic carriers [77].

10 Conclusion

Embedded sensors in a control loop play important roles in the correct functioning of control systems. A wide range of control systems depend on the timely feedback of MEMS inertial sensors to make critical decisions. In this work, we devised two sets of novel attacks against embedded inertial sensors to deceive the system. Our attack evaluations on 25 devices showed that it is possible to deliver implicit control to different kinds of systems by non-invasive attacks.

We characterized the out-of-band signal injection model and methods to manipulate an oscillating digitized signal, which was often considered as noises, to deliver adversarial control. To explore the generalizability of our methods, we showed that the oscillating digitized signal could also be manipulated by sending analog signals at relatively low frequencies through the sensing channel.

Acknowledgment

The authors would like to thank the anonymous reviewers and our shepherd Yongdae Kim for their numerous, insightful comments that greatly helped improve the presentation of this paper. This work is supported in part by ONR N000141712012 and US NSF under grants CNS-1812553, CNS-1834215, and CNS-1505799.

References

- [1] A video demonstration of automatic Switching attacks to spoof GoogleMaps. <https://youtu.be/dy6gm9ZLKuY>.
- [2] A video demonstration of Conservative Side-Swing attacks on a gyroscopic screwdriver. <https://youtu.be/SCAYbyMIJAc>.
- [3] A video demonstration of DoS attacks on a soldering iron. <https://youtu.be/itgm0121zoc>.
- [4] A video demonstration of Side-Swing attacks on a 3D mouse. <https://youtu.be/YoYpNeIJh5U>.
- [5] A video demonstration of Side-Swing attacks on a self-balancing robot. <https://youtu.be/oy3B1X41u5s>.
- [6] A video demonstration of Side-Swing attacks on a self-balancing scooter. <https://youtu.be/Y1LLiyhCn9I>.
- [7] A video demonstration of Side-Swing attacks on a smartphone. https://youtu.be/w16c_zBG1pU.
- [8] A video demonstration of Side-Swing attacks on a stabilizer. <https://youtu.be/FDxaLUtgaCM>.
- [9] A video demonstration of Switching attacks on a 3D mouse. <https://youtu.be/iWXTJ6We0UY>.

- [10] A video demonstration of Switching attacks on a kinetic controller. <https://youtu.be/MtXxcSzWcQA>.
- [11] A video demonstration of Switching attacks on a self-balancing scooter. <https://youtu.be/D-etuH04pms>.
- [12] A video demonstration of Switching attacks on a smartphone. <https://youtu.be/psu0hyUvDQk>.
- [13] A video demonstration of Switching attacks on a stabilizer. https://youtu.be/JcA_WXhrUEs.
- [14] A video demonstration of Switching attacks on a VR headset. <https://youtu.be/Jf9xHAW1PJY>.
- [15] Device teardown reports. <https://www.ifixit.com/> <https://www.chipworks.com/>.
- [16] Goldwood Sound directivity horns. <http://www.goldwoodparts.com/directivity-horns>. Accessed: 2018-05-05.
- [17] Goldwood Sound GT-1188 piezo tweeter speaker. <http://www.goldwoodparts.com/gt-1188.shtml>. Accessed: 2018-05-05.
- [18] L. Corporation, LRAD 2000X datasheet. https://www.dropbox.com/s/4qth9beayjx5gxr/LRAD_Datasheet_2000X.pdf. Accessed: 2018-04-25.
- [19] Myskunkworks 10" long-range horn. http://myskunkworks.net/index.php?route=product/product&path=61&product_id=63. Accessed: 2018-05-05.
- [20] Myskunkworks 130dB tweeter speaker. http://myskunkworks.net/index.php?route=product/product&path=61&product_id=79. Accessed: 2018-05-05.
- [21] Pyle PDBT78 tweeter speaker. <https://www.amazon.com/Pyle-PDBT78-2-Inch-Titanium-Tweeter/dp/B000JLB06E>. Accessed: 2018-05-05.
- [22] SainSmart UDB1002S DDS signal generator. <https://www.amazon.com/SainSmart-UDB1002S-Signal-Generator-Function/dp/B00JTR66CG/>. Accessed: 2018-05-05.
- [23] Sound pressure - Wikipedia. http://en.wikipedia.org/wiki/Sound_pressure. Accessed: 2018-06-01.
- [24] SparkFun MiniGen mini signal generator board. <https://www.sparkfun.com/products/11420>. Accessed: 2018-05-05.
- [25] Total SPL adding of coherent sound sources. <http://www.sengpielaudio.com/calculator-coherentsources.htm>. Accessed: 2018-06-01.
- [26] UltraElectronics HyperShield datasheet. https://www.ultra-hyperspike.com/Data/Pages/fe14c65c8b5fa0e0b19b46fca45fa01d-HyperShield_Dat_Sheet.pdf. Accessed: 2017-05-30.
- [27] STMicroelectronics L3G4200D datasheet. https://www.elecrow.com/download/L3G4200_AN3393.pdf, 2011. Accessed: 2017-06-12.
- [28] STMicroelectronics LSM330 datasheet. www.st.com/resource/en/datasheet/dm00037200.pdf, 2012. Accessed: 2018-06-14.
- [29] InvenSense MPU-6500 datasheet. https://store.invensense.com/datasheets/invensense/MPU_6500_Rev1.0.pdf, 2013. Accessed: 2017-06-12.
- [30] STMicroelectronics L3GD20 datasheet. <http://www.st.com/en/mems-and-sensors/l3gd20.html>, 2013. Accessed: 2017-06-12.
- [31] Future Vision. MINI augmented reality glasses make the future a reality. <http://www.bmwgroupdesignworks.com/work/mini-ar-glasses/>, 2015. Accessed: 2017-05-16.
- [32] Hope in a glove for Parkinson's patients. <https://www.technologyreview.com/s/545456/hope-in-a-glove-for-parkinsons-patients/>, 2016. Accessed: 2018-02-01.
- [33] Heads-up display to give soldiers improved situational awareness. <https://www.army.mil/article/188088>, 2017. Accessed: 2017-12-19.
- [34] ALTMANN, J. Acoustic weapons-a prospective assessment. *Science & Global Security* 9, 3 (2001), 165–234.
- [35] ANAND, S. A., AND SAXENA, N. Speechless: Analyzing the threat to speech privacy from smartphone motion sensors. In *IEEE Symposium on Security and Privacy* (2018).
- [36] ANTONELLO, R., AND OBOE, R. MEMS gyroscopes for consumers and industrial applications. In *Microsensors*. InTech, 2011.
- [37] AVIV, A. J., SAPP, B., BLAZE, M., AND SMITH, J. M. Practicality of accelerometer side channels on smartphones. In *Proceedings of the 28th Annual Computer Security Applications Conference* (2012), ACM.
- [38] BHUYAN, A. I., AND MALLICK, T. C. Gyro-accelerometer based control of a robotic arm using AVR microcontroller. In *9th International Forum on Strategic Technology (IFOST)* (2014), IEEE.
- [39] BLOCK, K., NARAIN, S., AND NOUBIR, G. An autonomic and permissionless android covert channel. In *Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks* (2017).
- [40] CAI, L., AND CHEN, H. On the practicality of motion based keystroke inference attack. In *International Conference on Trust and Trustworthy Computing* (2012), Springer.
- [41] CASTRO, S., DEAN, R., ROTH, G., FLOWERS, G. T., AND GRANTHAM, B. Influence of acoustic noise on the dynamic performance of MEMS gyroscopes. In *ASME International Mechanical Engineering Congress and Exposition* (2007).
- [42] DAVIDSON, D., WU, H., JELLINEK, R., SINGH, V., AND RISTENPART, T. Controlling UAVs with sensor input spoofing attacks. In *10th USENIX Workshop on Offensive Technologies (WOOT)* (2016).
- [43] DEAN, R., BURCH, N., BLACK, M., BEAL, A., AND FLOWERS, G. Microfibrous metallic cloth for acoustic isolation of a MEMS gyroscope. In *Proceedings of Industrial and Commercial Applications of Smart Structures Technologies* (2011), Society of Photo-Optical Instrumentation Engineers.
- [44] DEAN, R. N., CASTRO, S. T., FLOWERS, G. T., ROTH, G., AHMED, A., HODEL, A. S., GRANTHAM, B. E., BITTLE, D. A., AND BRUNSCH, J. P. A characterization of the performance of a MEMS gyroscope in acoustically harsh environments. *IEEE Transactions on Industrial Electronics* (2011).
- [45] DEAN, R. N., FLOWERS, G. T., HODEL, A. S., ROTH, G., CASTRO, S., ZHOU, R., MOREIRA, A., AHMED, A., RIFKI, R., GRANTHAM, B. E., ET AL. On the degradation of MEMS gyroscope performance in the presence of high power acoustic noise. In *IEEE International Symposium on Industrial Electronics* (2007).
- [46] FARSHTEINDIKER, B., HASIDIM, N., GROSZ, A., AND OREN, Y. How to phone home with someone else's phone: Information exfiltration using intentional sound noise on gyroscopic sensors. In *10th USENIX Workshop on Offensive Technologies (WOOT)* (2016).
- [47] GALLEGU-JUÁREZ, J. A., RODRIGUEZ-CORRAL, G., AND GAETE-GARRETON, L. An ultrasonic transducer for high power applications in gases. *Ultrasonics* 16, 6 (1978), 267–271.

- [48] KRANZ, M., WHITLEY, M., RUDD, C., CRAVEN, J. D., CLARK, S. D., DEAN, R. N., AND FLOWERS, G. T. Environmentally isolating packaging for MEMS sensors. In *International Symposium on Microelectronics* (2017), International Microelectronics Assembly and Packaging Society.
- [49] KUNE, D. F., BACKES, J., CLARK, S. S., KRAMER, D., REYNOLDS, M., FU, K., KIM, Y., AND XU, W. Ghost talk: Mitigating emi signal injection attacks against analog sensors. In *IEEE Symposium on Security and Privacy* (2013).
- [50] LEBECK, K., RUTH, K., KOHNO, T., AND ROESNER, F. Securing augmented reality output. In *IEEE Symposium on Security and Privacy* (2017).
- [51] LEBECK, K., RUTH, K., KOHNO, T., AND ROESNER, F. Towards security and privacy for multi-user augmented reality: Foundations with end users. In *IEEE Symposium on Security and Privacy* (2018).
- [52] LIU, X., ZHOU, Z., DIAO, W., LI, Z., AND ZHANG, K. When good becomes evil: Keystroke inference with smartwatch. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (2015).
- [53] MICHALEVSKY, Y., BONEH, D., AND NAKIBLY, G. Gyrophone: Recognizing speech from gyroscope signals. In *Proceedings of USENIX Security Symposium* (2014).
- [54] MILUZZO, E., VARSHAVSKY, A., BALAKRISHNAN, S., AND CHOUDHURY, R. R. Tappprints: your finger taps have fingerprints. In *Proceedings of the 10th international conference on Mobile systems, applications, and services* (2012), ACM.
- [55] NASIRI, S. A critical review of MEMS gyroscopes technology and commercialization status. InvenSense whitepaper.
- [56] NIGHSWANDER, T., LEDVINA, B., DIAMOND, J., BRUMLEY, R., AND BRUMLEY, D. GPS software attacks. In *Proceedings of the 2012 ACM conference on Computer and Communications Security* (2012).
- [57] PARK, Y., SON, Y., SHIN, H., KIM, D., AND KIM, Y. This ain't your dose: Sensor spoofing attack on medical infusion pump. In *10th USENIX Workshop on Offensive Technologies (WOOT)* (2016).
- [58] PASSARO, V., CUCCOVILLO, A., VAIANI, L., CARLO, M. D., AND CAMPANELLA, C. E. Gyroscope technology and applications: A review in the industrial perspective. *Sensors* 17, 10 (2017).
- [59] PETRACCA, G., REINEH, A.-A., SUN, Y., GROSSKLAGS, J., AND JAEGER, T. Aware: Preventing abuse of privacy-sensitive sensors via operation bindings. In *Proceedings of USENIX Security Symposium* (2017).
- [60] PSIAKI, M. L., O'HANLON, B. W., POWELL, S. P., BHATTI, J. A., WESSON, K. D., AND HUMPHREYS, T. E. GNSS spoofing detection using two-antenna differential carrier phase. In *Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation* (2014).
- [61] SHIN, H., KIM, D., KWON, Y., AND KIM, Y. Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications. In *International Conference on Cryptographic Hardware and Embedded Systems* (2017), Springer.
- [62] SHOUKRY, Y., MARTIN, P., TABUADA, P., AND SRIVASTAVA, M. Non-invasive spoofing attacks for anti-lock braking systems. In *Cryptographic Hardware and Embedded Systems* (2013), Springer.
- [63] SIKDER, A., AKSU, H., AND ULUAGAC, A. S. 6thSense: A context-aware sensor-based attack detector for smart devices. In *Proceedings of USENIX Security Symposium* (2017).
- [64] SON, Y., SHIN, H., KIM, D., PARK, Y., NOH, J., CHOI, K., CHOI, J., AND KIM, Y. Rocking drones with intentional sound noise on gyroscopic sensors. In *Proceedings of USENIX Security Symposium* (2015).
- [65] SOOBARAMANEY, P., FLOWERS, G., AND DEAN, R. Mitigation of the effects of high levels of high-frequency noise on MEMS gyroscopes using microfibrinous cloth. In *ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (2015).
- [66] TAKEDA, S., MORIOKA, I., MIYASHITA, K., OKUMURA, A., YOSHIDA, Y., AND MATSUMOTO, K. Age variation in the upper limit of hearing. *European journal of applied physiology and occupational physiology* 65, 5 (1992), 403–408.
- [67] TIAN, J., YANG, W., PENG, Z., TANG, T., AND LI, Z. Application of MEMS accelerometers and gyroscopes in fast steering mirror control systems. *Sensors* 16, 4 (2016).
- [68] TRIPPEL, T., WEISSE, O., XU, W., HONEYMAN, P., AND FU, K. Walnut: Waging doubt on the integrity of MEMS accelerometers with acoustic injection attacks. In *Proceedings of IEEE European Symposium on Security and Privacy* (2017).
- [69] WANG, C., GUO, X., WANG, Y., CHEN, Y., AND LIU, B. Friend or foe?: Your wearable devices reveal your personal pin. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security* (2016).
- [70] WANG, H., LAI, T. T.-T., AND ROY CHOUDHURY, R. Mole: Motion leaks through smartwatch sensors. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking* (2015), ACM.
- [71] WANG, Z., WANG, K., YANG, B., LI, S., AND PAN, A. Sonic gun to smart devices: Your devices lose control under ultrasound/sound. *Blackhat USA* (2017).
- [72] WANG, Z., ZHU, W., MIAO, J., ZHU, H., CHAO, C., AND TAN, O. K. Micromachined thick film piezoelectric ultrasonic transducer array. *Sensors and Actuators A: Physical* 130 (2006), 485–490.
- [73] WYGANT, I. O., KUPNIK, M., WINDSOR, J. C., WRIGHT, W. M., WOCHNER, M. S., YARALIOGLU, G. G., HAMILTON, M. F., AND KHURI-YAKUB, B. T. 50 kHz capacitive micromachined ultrasonic transducers for generation of highly directional sound with parametric arrays. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 56, 1 (2009), 193–203.
- [74] YAN, C., XU, W., AND LIU, J. Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle. *DEF CON 24* (2016).
- [75] YUNKER, W. N., SOOBARAMANEY, P., BLACK, M., DEAN, R. N., FLOWERS, G. T., AND AHMED, A. The underwater effects of high power, high frequency acoustic noise on MEMS gyroscopes. In *ASME 2011 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (2011).
- [76] YUNKER, W. N., STEVENS, C. B., FLOWERS, G. T., AND DEAN, R. N. Sound attenuation using microelectromechanical systems fabricated acoustic metamaterials. *Journal of Applied Physics* (2013).
- [77] ZHANG, G., YAN, C., JI, X., ZHANG, T., ZHANG, T., AND XU, W. Dolphinattack: Inaudible voice commands. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (2017).