



ADC-Bank: Detecting Acoustic Out-of-Band Signal Injection on Inertial Sensors

Jianyi Zhang^{1,3(✉)}, Yuchen Wang², Yazhou Tu³, Sara Rampazzi⁴,
Zhiqiang Lin⁵, Insup Lee⁶, and Xiali Hei³

¹ Beijing Electronic Science and Technology Institute, Beijing 100070, BJ, China
zjy@besti.edu.cn

² Academy of Information and Communications Technology, Beijing 100191, BJ, China

³ University of Louisiana at Lafayette, Lafayette, LA 70503, USA
{yazhou.tu1,xiali.hei}@louisiana.edu

⁴ University of Florida, Gainesville, FL 32611, USA
srampazzi@ufl.edu

⁵ Ohio State University, Columbus, OH 43210, USA
zlin@cse.ohio-state.edu

⁶ University of Pennsylvania, Philadelphia, PA 19104, USA
lee@cis.upenn.edu

Abstract. Inertial sensors are widely used in navigation, motion tracking, and gesture recognition systems. However, these sensors are vulnerable to spoofing attacks, where an attacker injects a carefully designed acoustic signal to trick the sensor readings. Traditional approaches to detecting and mitigating attacks rely on module redundancy, i.e., adding multiple sensor modules to increase robustness. However, this approach is not always feasible due to the limited space and increased complexity of current printed circuit boards.

This paper proposes a new method, ADC-Bank, to detect inertial sensor spoofing attacks via acoustic out-of-band signals. Unlike other multiple-sensor-based solutions, it is based on component redundancy within one sensor, using multiple analog-to-digital converters (ADCs) with different sampling rates to simultaneously sample the output of the sensors. The different sample rates result in different aliasing frequencies for out-of-band signals that can be used to detect attacks. The proposed method is evaluated on off-the-shelf inertial sensors with commercial ADCs, demonstrating its ability to detect the attacking signals with relatively low cost and computation overhead.

Keywords: Spoof attack · Out-of-band acoustic signal injection · Inertial sensor · Multiple ADCs · Detection and correction

1 Introduction

Micro-electro-mechanical systems (MEMS) inertial sensors are known to be susceptible to acoustic out-of-band signal injections [4–6, 9, 30, 36–38, 40, 41]. These

attacks used acoustic signals at frequencies close to the sensor’s resonant frequency to induce high-frequency analog signals in the sensor circuits. Ideally, the injected signals should be filtered out because they are out-of-band signals. However, attackers can still inject these signals into the system. The essential feature of out-of-band signal injections is that the induced analog signals will be undersampled, resulting in signal aliasing. When aliasing occurs, attackers can change the output of sensors by maliciously generated stimuli, then deceive the sensing and actuation systems into executing malicious actions accordingly [11]. For example, a self-balancing scooter can adjust direction and speed according to its lean angles, which are described by inertial sensors. However, an attacker can induce an intentional sound at the resonant frequencies of the gyroscope; the output of the inertial sensor will be distorted, and the attacker can make the scooter move in a corresponding opposite direction [38].

In recent years, several defense strategies have been studied to solve the problem of acoustic-based spoofing attacks. For example, shielding [3, 16, 30, 41] was recommended to mitigate out-of-band injections into inertial sensors. However, shielding can cause heat dissipation, cost, size, and usability issues. Another defense consists of low-pass filters that can filter out malicious high-frequency signals and mitigate attack at inertial sensors [16, 37, 45]. In practice, implementing ideal anti-aliasing filters that eliminate all out-of-band signals is trivial. For example, a high-order filter that eliminates all signals above the cutoff frequency will cause signals that change rapidly to ring on for a long time. Moreover, analog filters lead to an unequal time delay as a function of frequency [33]. If the phase delay introduced by filters is large, it is difficult to minimize this delay or compensate for it in software [8]. Moreover, the integrated low-pass filter does not have clear cut-offs [25, 32]. An additional defense approach consists of using high-frequency sampling of the analog signal. For instance, the inertial sensor signal frequency induced by movement is generally below 20 Hz. If the sensor designers choose ADCs with sampling rates high enough to handle the resonant frequencies, it will increase the production costs and decrease the sampling resolution and the processing speed due to the over-wide bandwidth. Recent work has studied purely software-based detection methods [35] and module redundancy methods (multisensors for sensor fusion) [3, 19, 28, 41–43]. However, false positives/negatives can occur when external factors or injected data differ from the assumed patterns. The researchers also noted that attacks with a directed magnetic field that can precisely control both the magnetometer and the gyroscope would cause their sensor fusion-based detection method to fail [35].

In this paper, we present ADC-Bank, a novel out-of-band signal defense method using component redundancy within a sensor in contrast to the work mentioned above. Compared to other defense strategies, our method is easy to manufacture and has fewer attack surfaces than module redundancy strategies, such as multiple sensor-based methods. After implementing multiple circuit components that simultaneously elaborate the physical stimulus under different configurations and settings, we provide multiple metrics on the legitimacy of the measurement at the software layer. This information is then used to detect

the system from processing an altered signal. We evaluate our method on off-the-shelf inertial MEMS sensors from three different vendors. Our experimental results show that ADC-Bank can detect physical injection attacks via out-of-band acoustic signals on all models of inertial MEMS sensors that we tested.

Despite many existing defense mechanisms against acoustic physical injection attacks at MEMS sensors, there is no fundamental solution to detect these malicious transmissions and prevent vulnerabilities in the physics of a MEMS sensor. Our work fills this gap through the following contributions:

1. We propose a component redundancy scheme to detect acoustic out-of-band signal injection by elaborating and comparing the physical stimulus in different settings.
2. We investigate how to extract the real physical stimulus from different results of the redundant components.
3. We deploy our defense method on off-the-shelf inertial sensors with commercial ADCs to evaluate our method.
4. We discuss how our strategy can be used in the design and manufacturing of future sensors.

2 Background

2.1 MEMS Inertial Sensors

Almost all MEMS inertial sensors have a mass and a support spring, and they use this mechanical structure to detect motion stimuli [26]. MEMS accelerometers sense linear accelerations by displacement of the mass supported by springs and measure the capacitance change between the mass and fixed electrodes [17, 44]. MEMS gyroscopes are relatively complex. They have a continuously vibrating mass that, like accelerometers, is supported by springs. They measure the Coriolis force generated by the applied angular velocity on the vibrating mass [29].

After transduction, the sensor output needs a series of additional processing to interface with external components such as microcontrollers. In general, the change in capacitance causes a change in voltage. For an analog sensor, this analog signal is typically amplified and outputted directly from the amplifier. For the digital sensor, the amplified signal is digitized via an analog-to-digital converter (ADC) and then transferred to the control system by standard digital interfaces like SPI, I²C, and UART. In this work, we consider analog inertial sensors to explain our approach.

2.2 ADCs and Aliasing

After the sensor transforms the physical measurement into an analog signal, a built-in ADC digitizes the sensor's output. The analog signal that is continuous in time should be converted at a certain rate by ADC, and this rate is defined as the sampling rate or sampling frequency of the converter. According to the Nyquist-Shannon sampling theorem, the sampling rate should be at least twice

the signal’s maximum frequency when the original physical measurement can be reconstructed from the discrete data by the mitigation filter. If the system acquires data at an insufficient rate, called undersampling, the signal will be incorrectly detected at a specified interval as a lower frequency [18]. Then aliasing will occur.

2.3 Acoustic Injection

Because of its miniaturized mechanical and integrated electronic structure, these sensors’ output could be changed to incorrect values by resonant acoustic interferences [37]. The successful modification relies on two vulnerabilities of the MEMS inertial sensor: the mass-spring structure that works as the receiving system for resonant acoustic signals and the non-linearity of electronic components like the overdriven amplifier or under-sampling of an ADC. According to the second vulnerability, the acoustic injection attack can be categorized into two classes: output control attack and output biasing attack [37]. The output control attack leverage signal clipping at the insecure amplifier to introduce a DC component into the acceleration signal, which slips through any subsequent LPF [15, 27, 39]. However, triggering this kind of attack requires a signal beyond the amplifier’s capability, which means high power and deafening volume. Therefore, it becomes impractical to generate the required loudness and attack the sensor from a long distance [10].

3 Threat Model

We assume that the attackers’ objective is to spoof and manipulate the MEMS inertial sensors’ output. To achieve this, attackers need to transmit specific acoustic signals at the resonant frequencies to deceive the sensor and trigger the control system’s actuation.

Attack Scenarios. We assume that attackers can use an off-the-shelf speaker or transducer to generate the sound waveforms for the injection. Also, we assume that they are able to induce the sound, at the resonant frequencies of these sensors, at any position, distance, or angle. This might be done via means of amplifiers and constant directivity horns. We assume that the attackers have sufficient resources to optimize the power, directivity, and emitting area. More powerful attackers may utilize customized acoustic equipment to improve the effect. The signal source of attacks can be a built-in speaker, a function generator, an MCU board like Arduino, mini signal generator boards [24, 31], or even malicious codes in an email or webpage with JavaScript and autoplay audio enabled. The attacker can also use long-distance acoustic devices to play the sound waves as described by Tu et al. [38].

Attack Goals. We assume that attackers utilize the resonant acoustic signal to inject the sensor output and deliver adversarial control to the system. Such




Fig. 1. Scheme of an ADC-Bank’s signal processing. The measured output of the sensor is a linear combination of the original signal, $s(t)$, and the injection signal, $s_a(t)$. Unlike legitimate signals, malicious out-of-band signals sampled by the ADCs generate multiple frequency peaks. We can employ such an observation to detect and analyze various attack scenarios.

attacks on the IMU sensors will pose security and safety risks to cyber-physical systems like robots, stabilization systems, self-balancing scooters, drones, etc.

System Accessibility. We assume the attackers know the exact type and model of the MEMS inertial sensors and can easily access the datasheet to know the sensors’ components and structures.

4 Defense Approach

4.1 System Model

The system model of our proposed protection scheme is presented in Fig. 1. It has two blocks, including a multi-ADCs part and a signal analysis module.

The multi-ADCs part consists of more than two ADCs whose sampling rates have certain constraints like pairwise relatively prime. After sampling the sensor output synchronously, these ADCs send their respective measurement results of the same sensor to the signal analysis module.

The signal analysis module for spectral analysis consists of three parts: frequency analysis, peak detection, and a calculation and reasoning phase. The frequency analysis performs a Fast Fourier transform (FFT) on each measurement result of the different ADCs and transfers the detection into the frequency domain. According to the results of peak detection, ideally, there will be one overlapped peak in the frequency domain, which means that the signal has normal behavior. Otherwise, multiple separated peaks would suggest the presence of out-of-band physical signal injection attacks. In the calculation and reasoning phase, when no multiple peaks are found, which means that there is no injection signal, the A/D conversion and the measurement value are considered trustworthy. Hence, the actuation system knows the result is digitized from the original sensor’s output. However, when the signal injection attack is detected, we can calculate the approximate frequencies of the injection signal based on prior knowledge of the intended frequency range of the attack (more details in [38]).

From the signal perspective, the sensor output generated by the real movement is $s(t)$ in the absence of attackers. The attack signal is $s_a(t)$ generated by acoustic injection. We model the measured output of the sensor as a linear combination of the original signal $s(t)$ and the injection signal $s_a(t)$. Hence the measured signal $\tilde{s}(t)$ is:

$$\tilde{s}(t) = s(t) + s_a(t) \quad (1)$$

Since the mechanical structure of the sensor under resonance oscillates at the same frequency as the attacking signal, we model the resulting signal with a resonant frequency F_a and an initial phase ϕ as:

$$s_a(t) = A \cdot \sin(2\pi F_a t + \phi) \quad (2)$$

where coefficients $A = A_0 k_a k_s$. A_0 is the amplitude of the attacking signal, k_a is the acoustics attenuation when the attacking signal is transmitted to the target sensor, and k_s is the sensitivity of the sensing mass. Substitute Eq. (2) into Eq. (1), we have the measured value:

$$\tilde{s}(t) = s(t) + A \cdot \sin(2\pi F_a t + \phi) \quad (3)$$

Then, the combination signal will be sampled by multiple ADCs. Typically, the sampling rate of the ADC in the inertial sensor system is designed to be high enough to sample the movement signal, so the true sensor measurement $s(t)$ will be normally converted. However, the frequency of attacking signals injected through resonance is usually much higher than the sampling rate. Therefore, sampling these out-of-band high-frequency signals will cause aliasing. A sinusoidal analog signal with frequency F will be aliased to a digital signal with a frequency of ε when $F > 2F_S$, where F_S is the sampling rate. We have

$$F = n \cdot F_S + \varepsilon \quad \left(-\frac{1}{2}F_S < \varepsilon \leq \frac{1}{2}F_S, n \in \mathbb{Z}^+\right) \quad (4)$$

Therefore, assuming that F_a is the resonant frequency of the sensor, the adversary uses it as the frequency of injection signals. For multiple ADCs, based on Eq. (4), we have:

$$F_a = n_i \cdot F_{S_i} + \varepsilon_i \quad \left(-\frac{1}{2}F_{S_i} < \varepsilon_i \leq \frac{1}{2}F_{S_i}, n_i \in \mathbb{Z}^+\right) \quad (5)$$

where F_{S_i} is the sampling rate of the i -th ADC, and ε_i is the resulted frequency of the corresponding ADC output. For simplicity, we assume that n in Eq. (4) and Eq. (5) is the integer multiple of the sampling rate F_S . Therefore, n_i stays the same when ε , F_S changes slightly.

According to Eq. (5), these multiple ADCs with different sampling rates will generate different results ε_i for the same input signal F_S . Out-of-band signal injections can be detected on the basis of this separation. Meanwhile, based on these sampling rates, the possible n_i can be traversed according to the reading, and the approximate range of F_a in Eq. (5) can be found according to multiple F_{S_i} and ε_i .

5 Detection

Multi-ADCs. In the out-of-band signal injection attack against the MEMS inertial sensor, under acoustic injection, malicious sound waves are transmitted to the mechanical structure of the inertial sensor, forcing the sensing mass to resonate.

In the analog-to-digital conversion process, the input signal is sampled. Only when the sampling rate F_S is greater than twice the highest frequency $2F_{Max}$ in the analog signal spectrum, the analog signal can be recovered without distortion. Therefore, the ADC sampling rate in the inertial sensor system is designed to be high enough to sample the movement signal. However, when inertial sensors face ultrasound resonant signal injection, also known as out-of-band signal injection attacks, the frequency of attacking signals is usually much higher than the sampling rate.

The sampling rate in the inertial sensor system is usually in the tens or hundreds, while the resonant frequency is usually higher than 2 kHz for the accelerometer and 19 kHz for the gyroscope. Since the resonant frequency is much higher than the sampling rate, signal aliasing will occur and be reconstructed into a new low-frequency in-band signal.

To detect suspicious out-of-band signal injection attacks, we take advantage of the phenomenon of undersampling. Specifically, the multi-ADC part consists of more than two ADCs that sample the input signal, respectively. Then, reconstructing these undersampled signals from the digital samples will cause signal aliasing. Our defense solution consists of comparing such aliased signals to determine the reconstructed original signal.

The microcontroller of the control system can then be used to measure the physical quantity and hence can detect the attack based on the outputs of the ADCs. We suppose that the attacker remotely injects the malicious waveforms into the inertial sensor circuit. After sampling and digitizing the stimulus by multiple ADCs with different sampling rates, the control system can spot the attack immediately since the results in the frequency domain are totally different. With the help of well-designed parameters, we can not only detect the existence of malicious signals, but also recover the real signal from the measurement of the sensor's outputs. In particular, if the sampling rates of multiple ADCs we selected are pairwise relatively prime, according to the Chinese remainder theorem [21], the microcontroller can easily calculate the range of attack frequencies. After that, we can easily filter the frequencies induced by attack signals and provide reliable measurements to the control system.

Frequency Analysis. In our defense approach, a key part is to analyze the frequency of the reconstructed signal. When multi-ADCs sample and digitize the input signal respectively, each measurement result of the different ADCs will be performed frequency analysis via Fast Fourier transform (FFT) and transferred the detection into the frequency domain. With the help of frequency domain analysis, we can obtain the frequency information of the input signal immediately.

If the input signal is generated by normal motion, the maximum frequency of the input signal must be within half of the sampling rate of the ADCs, and it will be able to be digitized normally. This means that the measurement results of different ADCs will produce the same frequency component after FFT. When faced with the injection of malicious acoustic signals, the situation will become different. The input signals far beyond the sampling rate of ADCs will cause aliasing. Because the sampling rates of ADCs are different from each other and relatively prime to each other, the measurement results will produce different frequencies after FFT. We use a simple peak detection algorithm to determine the credibility of the measured sensor value.




Fig. 2. Scheme of an ADC-Bank’s attack detection. In contrast with legitimate signals, malicious out-of-band signals sampled by the ADCs generate multiple frequency peaks. This technique can be used to detect and analyze various attack scenarios.

Peak Detection. The peak detection algorithm is used to quickly measure the results after FFT. Figure 2 shows the main peak detection process in the frequency domain; we detect the peak of the FFT results of ADC measurements, respectively. If only one overlapping peak is detected, it indicates that the signal is not attacked and is credible for the subsequent actuation system. If there are multiple separate peaks, that means that there is a potential attack. These signals will not be able to be transmitted directly to the actuation system and will need to be corrected.

6 Experiments

In this section, we try to prove the effectiveness of our method in a real-world case study. To prove the effectiveness of our signal process scheme, we designed a series of experiments. We have built an acoustic injection attack environment to collect raw data and perform signal processing and analysis.

We evaluate our approach from the following two situations. 1) To simulate a real attack environment, we use a signal generator output to drive the speaker and then interfere with the inertial sensor. Then we collect the motion signal and the injection signal from the inertial sensor, respectively, using an NI USB-4431 Data Acquisition (DAQ) [14]. 2) We use ADCs and microcontroller units (MCUs) to build a set of data acquisition environments. We use a signal generator output to drive the speaker, and then interfere with the inertial sensor. Then we collect and upload the sensor data to a PC for further signal processing.

6.1 Experimental Setup

Figure 3 shows the experimental setup. DG5300 signal generator is used to generate an acoustic signal [23]. Here, the output amplitude is set to 5 v. A power amplifier is used to enhance signal power, and the Vifa speaker [2] is responsible for outputting acoustic waves. The inertial sensor chip is mounted on an evaluation board and driven by 3.3 v/5 v DC provided by the external Arduino [1].




Fig. 3. Schematic of the experimental setup. The inertial sensor chip is mounted on an evaluation board placed on the experimental platform. The attack range can be between one and three meters in a real attack scenario [38].

After determining the resonant frequency of each inertial sensor chip, we select the appropriate frequency to carry out an acoustic injection attack on each

chip and then use the NI USB-4431 DAQ module and the A/D Data Acquisition System we built to collect the sensors' output, respectively.

In the following experiments, ADXL335 is used as the target accelerometer [7], and LPY550AL is used as the target gyroscope [34]. We also selected some other inertial sensor models, as shown in Table 1.

Table 1. Resonant frequency and aliasing frequency results of inertial sensors in the experiment

Chip Enterprise	Chip Model	Type	Axis	Resonant Frequency	Attack Frequency	Aliasing Frequency				
						NI USB-4431 DAQ Results			A/D Results	
						280 Hz	700 Hz	1000 Hz	250 Hz	920 Hz
Murata	ENC-03MB	Gyro	x	22 kHz-25.2 kHz	25135 Hz	215 Hz	65 Hz	135 Hz	115 Hz	295 Hz
Murata	ENC-03RC	Gyro	x	30 kHz-33 kHz	32295 Hz	185 Hz	95 Hz	295 Hz	45 Hz	95 Hz
STMicroelectronics	LPY550AL	Gyro	x	22 kHz-23 kHz	22785 Hz	105 Hz	315 Hz	215 Hz	35 Hz	215 Hz
STMicroelectronics	LPY550AL	Gyro	y	22 kHz-23 kHz	-	-	-	-	-	-
ADI	ADXL335	Acce	x	4 kHz-5.5 kHz	4490 Hz	10 Hz	290 Hz	490 Hz	10 Hz	110 Hz
ADI	ADXL335	Acce	y	4 kHz-5.5 kHz	-	-	-	-	-	-
ADI	ADXL335	Acce	z	4 kHz-5.5 kHz	-	-	-	-	-	-

At the same time, in order to simulate the signal output generated by real motion, we place the inertial sensor chip mounted on an evaluation board on top of a vibration platform, where we set the vibrating frequency below 50 Hz, then we use the above two acquisition systems to collect the signal output, respectively.

6.2 Evaluation Experiment

Inertial Sensors with DAQ. In this set of experiments, we first put the chip on the vibration platform and then set the vibration platform frequency to 16 Hz to simulate a true motion. For each time sensor output, we sample the output using three different sampling rates and analyze it in the frequency domain.

Firstly, we carried out experiments on an accelerometer, ADXL335. Figure 4 shows the detailed results of ADXL335. Figure 4a shows the raw data sampled by the system in the time domain. We process the data before transforming them to the frequency domain. First, we remove the DC component of the signal because it has no significance for us in detecting the frequency of the sensor output signal. Second, we normalize the data to make their amplitudes close. Then we transform the data sampled at three different sampling rates into the frequency domain, and the results are shown in Fig. 4b. We found that there are overlapping peaks at 16 Hz in the spectrum. This is also consistent with the frequency that we generate through the vibrating platform, which means that this signal is a normal motion signal. Our signal processing scheme will give the signal high confidence and let the signal enter the control system without affecting the response of the actuation system.

Then, we conducted a similar experiment on a gyroscope, LPY550AL. The environment settings are the same as ADXL335. The results are shown in Fig. 5. Similarly, we can also see that in Fig. 5b), there is an overlapping peak at 16 Hz.

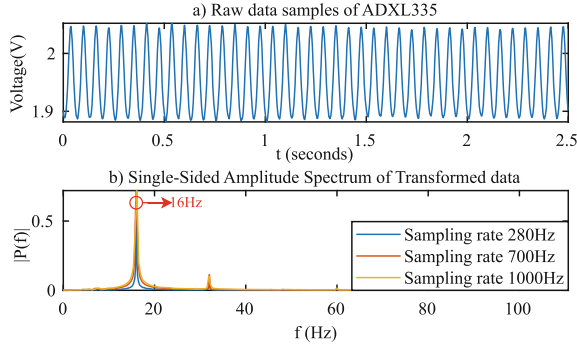


Fig. 4. The testing results of ADXL335 with NI USB-4431 DAQ. a) shows the sampled time domain raw data. In b), the raw time domain data is converted to the frequency domain after data processing.




Fig. 5. The testing results of LPY550AL with NI USB-4431 DAQ. a) shows the sampled time domain raw data. In b), the raw time domain data is converted to the frequency domain after data processing.

Since the NI USB-4431 DAQ module has a good anti-aliasing filter, the resonant signal beyond the sampling rate will not be collected. Therefore, we choose a sampling rate of 70,000 Hz, which is much higher than the resonant frequency and can sample normally, and then we simulate the aliasing process under different sampling rates by down-sampling.

For ADXL335, we use the signal generator to generate a 3,525 Hz signal, which is also the resonant frequency of the inertial sensor. The signal is output through the speaker to interfere with the accelerometer. By down-sampling, we get the raw data sampled at three sampling rates. After the signal is processed, we convert it to the frequency domain.

The data acquisition and frequency analysis results are shown in Fig. 6. Figure 6 a), b), and c) are the raw time-domain data of resonant signal down-sampled to 280 Hz, 700 Hz, and 1,000 Hz (due to the limitation of down-sampling), respectively. Figure 6d) is data converted to the frequency domain

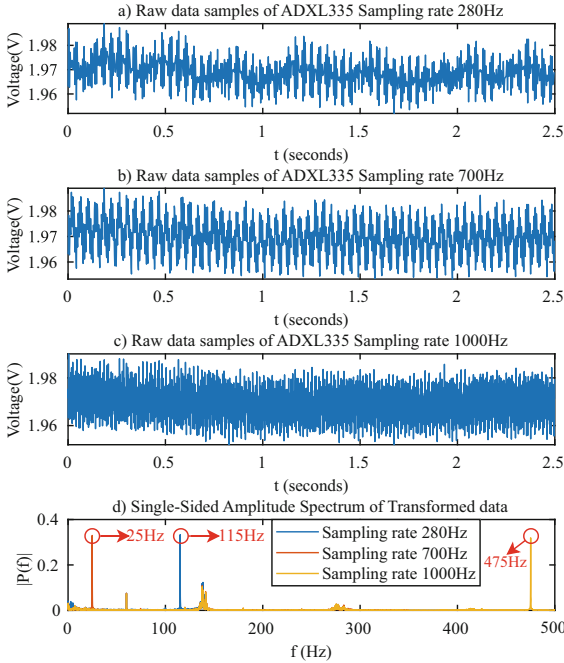


Fig. 6. The testing results of ADXL335 with NI USB-4431 DAQ. Figure a, b and c show the sampled time domain raw data. In Fig. d, the raw time domain data is converted to the frequency domain after data processing.

after signal processing. We can clearly see in the spectrum that there are three separate peaks because different sampling rates lead to different aliasing frequencies. Therefore, we can determine that this abnormal signal needs filtering.

Regarding the gyroscope, we also performed experiments on an LPY550AL. We generated a 22,700 Hz signal, which is also the resonant frequency of the inertial sensor, through the signal generator and output by a speaker. The results are shown in Fig. 7. In the spectrum diagram, we can see that there are only two separated peaks, one of which is the overlapping peak, which is due to the same aliasing frequency of the two sampling rates. However, it can still be determined that the signal is abnormal.

In this set of experiments, we also tested other types of inertial sensor chips, and the results are shown in Table 1. In fact, on all types of inertial sensor chips, we can clearly distinguish whether there is abnormal signal input. This also preliminarily proves that our detection method is applicable to real-world actuation systems.

Inertial Sensors with commercial ADCs. The main difference between this group of experiments and the previous group of experiments lies in the sampling method. We did not use a professional DAQ module to collect the inertial sensor

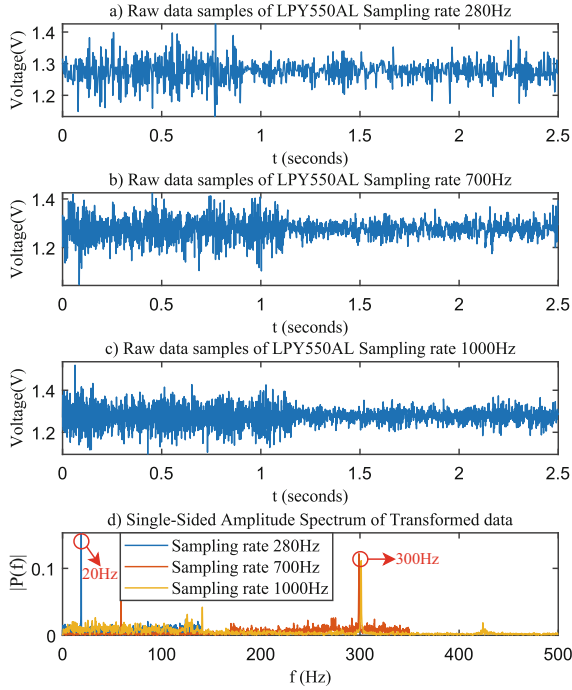


Fig. 7. The testing results of LPY550AL with NI USB- 4431 DAQ. Figure a, b and c show the sampled time domain raw data. In Fig.d, the raw time domain data is converted to the frequency domain after data processing.

output like before; instead, we will use a commercial ADC to collect the inertial sensor output to fully simulate the situation in a real actuation system.

In this set of experiments, we also tested four different inertial sensor models, as shown in Table 1, including the acquisition and analysis of normal motion signals and abnormal resonant signals. The ADC model we use is ADS1015, which has seven optional sampling rates. For normal motion signals, the sampling rate is usually several hundred Hz. At the same time, to evaluate whether the two ADCs can completely detect abnormal signals, we use two ADS1015, which are connected to the Arduino microcontroller and configured with different sampling rates. Here, we set the sampling rates of ADC to 250 Hz and 920 Hz (which are two optional sampling rates for ADS1015), respectively. We use these two ADCs with different sampling rates to sample the inertial sensor output simultaneously, then upload the sampling data to a PC for further data processing and analysis.

Here, we take ADXL335 as an example, as shown in Fig. 8. For the normal motion signal, we also select 16 Hz as the vibration frequency to simulate normal motion. We observed overlapping peaks at 16 Hz in Fig. 8c. It can be seen that the in-band, normal motion signals can be determined to be trustworthy by two ADCs' simultaneous sampling.

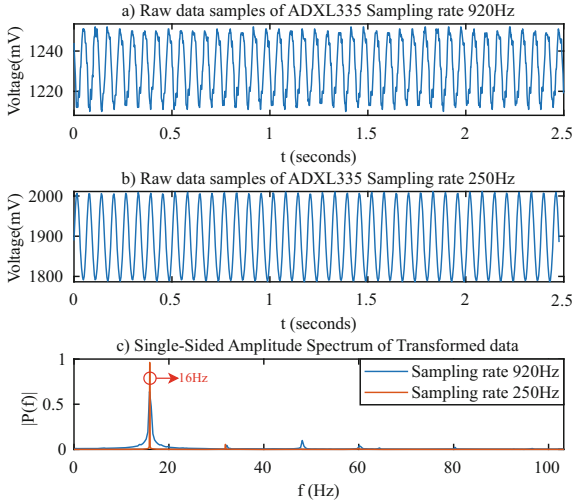


Fig. 8. The testing results of ADXL335 with commercial ADCs. Figure a and b show the sampled time domain raw data. In Fig. c, the raw time domain data is converted to the frequency domain after data processing.

For an abnormal resonant signal, the resonant frequency is 4,485 Hz. The results of signal sampling are shown in Fig. 9. As shown in Fig. 9c, two different peaks are generated at two different sampling rates. In fact, there is a certain deviation between the peak frequency and the theoretical aliasing frequency, but we can still determine that there is an abnormal signal to be filtered.

6.3 Attack Frequency Analysis

During the previous data acquisition and processing, we set the attack signal frequency through the signal generator, and then obtain the ADC sampling rate and the frequency of an aliased in-band signal. According to the prior knowledge, we have a known range of possible attack frequencies. For the sampling rate of each ADC, we can traverse the possible small frequency ranges within the possible attack frequency range according to the in-band signal frequency. Then we find the intersection of the frequency ranges determined by different ADCs, and can obtain a calculated attack frequency range. Through the previous experiments, we have collected data from some models of inertial sensors. Next, we will calculate and analyze the specific data for example.

According to Eq. (5), the possible attack frequency ranges of several segments can be calculated according to the peak frequency obtained from ADC of a certain sampling rate. We have a prior range of attack frequency, 2–5 kHz for accelerometers and 19–27 kHz for gyroscopes. For multiple ADCs, we can find different ranges of their peak frequencies, and get the intersections of these ranges. Under the above experimental configuration, we try to calculate the




Fig. 9. The testing results of ADXL335 with commercial ADCs. Figure a and b show the sampled time domain raw data. In Fig. c, the raw time domain data is converted to the frequency domain after data processing.

attack frequency range according to the aliasing frequency and ADC sampling rate. We show the results as shown in Fig. 10.




Fig. 10. Frequency range determined by calculation. The black part in the figure is the possible attack frequency range calculated within the range of 2–5 kHz. The red line indicates the overlapping range calculated under the sampling rates of 250 Hz and 920 Hz. We regard the overlapping range as the potential attack frequency range. (Color figure online)

As shown in Fig. 10, we can see that the attack frequency ranges determined by multiple ADCs form an intersection. We compare the attack frequency set by the signal generator with the calculated results, and it can be seen that the actual attack frequency falls within the frequency range we calculated.

7 Discussion

7.1 Adaptive Attacks and Frequency Drift

In acoustic-based spoofing attacks, slight frequency drift or sample rate jitter could be amplified and cause significant deviation in the digital output of the sensors [38]. Due to this drift, the frequency of the aliased output is not constant.

During adaptive attacks, if the attacker knows the details of the algorithm and the sampling rate, and hopes to attack with acoustic signals whose frequency is the common integer multiple of the sampling rates, it will be difficult to implement because the sampling rate is not completely accurate. Even if the attack frequency is an integer multiple of the sampling rate, it will be recognized due to frequency drift in a short time. Additionally, the common integer multiple of the sampling rates may not fall within the resonant frequency range of the inertial sensor. Increasing the number of ADCs will greatly reduce this possibility.

If the attacker wants to attack through frequency sweep and frequency hopping, the attack cannot be implemented because the accurate sampling rate cannot be known. In addition, in our defense method, we do not need to obtain a certain constant frequency output. We focus on whether different ADCs at the same time have the same output and then determine whether there is an abnormal signal. As long as an attack occurs, there will be multiple different peaks in the spectrum.

7.2 Consistency of ADC

In the experiment, we found that under the same sampling rate setting, the raw data obtained beyond the sampling rate were different for the two ADCs with the same model. Therefore, we have reason to believe that different ADCs of the same models have consistent differences. This will also cause errors in the aliasing frequency, which will affect the estimation of the attack frequency.

At the same time, we believe that, at the beginning of future sensor design, ADC with integrated component redundancy will have better consistency and help to reduce errors.

7.3 Future Design and Manufacturing

In our simulation and experiment, as the system is closer to reality, we built our defense system using existing commercially available modules, and the systematic error of the data has increased. We believe this is due to the noise generated by the connection between the modules. In the future sensor design and manufacturing process based on our method, we believe that the integration of various parts will help reduce the generation of systematic errors and improve the accuracy of our defense methods. On the other hand, the manufacturing cost does not increase linearly with the increase of components, which also makes us believe that our defense method based on ADC redundancy is feasible.

8 Related Work

Designing a secure sensor and actuator system is not an easy task. Ever since the Ghost talk proposed by F. Kune et al. in 2013, which demonstrates that medical devices can be inhibited pacing and induce defibrillation shocks by intentional electromagnetic (EM) signals [16], attempts have been made to find defense methods. In this section, we divide existing work into three categories: surrounding defense, module defense, and component defense. The surrounding defense mainly depends on the shielding to mitigate injection. Sometimes, the researchers may add specific materials as a physical barrier to attenuate the malicious signal. In previous studies, barriers were built in conductor wires [16] optical EMI shielding [20], or as sound damping [3, 30, 41]. Sometimes, researchers can increase the difficulty of injection by selectively reducing the attack surface, increasing the directivity [28], or limiting the duration of sensor exposures [22]. However, some sensors are placed on high-density interconnect printed circuit boards (HDI PCBs), and some sensors must be exposed to the external environment. Thus, surrounding defenses may not always be applicable.

Regarding module defense, additional modules such as receivers, sensors, or actuators are used to detect or dampen the targeted out-of-band signals. Z. Wang [41] and C. Bolton [3] proposed the adoption of additional microphones to detect resonating sounds, which are out-of-band signals, against MEMS inertial sensors. In the same line, as suggested by Kune et al. [16], adopting the cardiac probe and comparing the result of actuation can distinguish between induced and measured signals. Furthermore, researchers utilized sensor fusion to enhance resiliency against these injection attacks. Many prior work adopted redundant sensors as a defense method when we can bear the cost and space of these sensors [3, 28, 41–43].

The component defense is a more common strategy in the previous work. New, modified, or improved components may be introduced into the signal conditioning chain to reduce an attacker's ability to exploit the injection. For example, researchers can augment the circuit with an additional low-pass filter to attenuate the signal outside the sensor's baseband and hence cancel out the aliasing by blocking the high-frequency, which possibly induces such problem [16, 37, 45]. Meanwhile, an adaptive filter can be used when a simple low-pass filter is not applicable. Y. Son et al. employed differential signaling to filter the signal injected in the sensing pathway by referring to a dynamically measured frequency [30]. However, some previous work demonstrated that the parasitic characteristics caused by the surface mount components might convert the low-pass filter into a band-stop filter. Attack signals above the cutoff frequency can still be coupled to the circuit and cause aliasing [12, 13]. Furthermore, some researchers may choose to use a particular sampling pattern called out-of-phase sampling to mitigate malicious out-of-band signals that are converted to in-band frequencies after ADC [37]. Meanwhile, some researchers may improve the performance of specific components. Trippel et al. proposed a secure amplifier whose dynamic range is wide enough to cope with the exploited saturation [37]. Wang et al. [41] and Son et al. [30] both proposed the redesigned MEMS gyroscopes,

although they do not give specific approaches to move the resonant frequencies to noncritical frequency bands. Furthermore, researchers may be able to apply randomness in the receiver pathway to mitigate the influence of the attacker on sensor output. Trippel et al. [37] suggested that using ADC with a random sampling rate can effectively deal with DC aliasing since attackers often utilize the predictable property, such as sampling rate, to bias and control the accelerometer and gyroscope output.

9 Conclusion

We have presented a new solution, ADC-Bank, to address the issue of inertial sensor spoofing attacks in embedded systems. Our method successfully detects these attacks by identifying the aliasing frequency of the attack signal. Our experiments and evaluations, conducted on various types of inertial sensors, demonstrate the effectiveness of ADC-Bank in protecting against spoofing attacks.

Acknowledgment. The authors thank the anonymous reviewers for their valuable comments that improved this paper. This work is supported in part by the US NSF under grants CNS-1812553, CNS-2117785, OIA-2229752, CNS-2231682, and two gifts from Meta.

References

1. Arduino: Arduino Uno Rev3. <https://store-usa.arduino.cc/products/arduino-uno-rev3>. Accessed 16 Aug 2022
2. Avisoft. Ultrasoundgate. <http://www.avisoft.com/ultrasoundgate/>. Accessed 16 Aug 2022
3. Bolton, C., Rampazzi, S., Li, C., Kwong, A., Xu, W., Fu, K.: Blue note: how intentional acoustic interference damages availability and integrity in hard disk drives and operating systems. In: 2018 IEEE Symposium on Security and Privacy (SP), pp. 1048–1062. IEEE (2018)
4. Castro, S., Dean, R., Roth, G., Flowers, G.T., Grantham, B.: Influence of acoustic noise on the dynamic performance of mems gyroscopes. In: ASME International Mechanical Engineering Congress and Exposition, vol. 43033, pp. 1825–1831 (2007)
5. Dean, R.N., et al.: On the degradation of mems gyroscope performance in the presence of high power acoustic noise. In: 2007 IEEE International Symposium on Industrial Electronics, pp. 1435–1440. IEEE (2007)
6. Dean, R.N., et al.: A characterization of the performance of a mems gyroscope in acoustically harsh environments. *IEEE Trans. Ind. Electron.* **58**(7), 2591–2596 (2010)
7. Analog Devices. Adxl335. <https://www.analog.com/cn/products/adxl335.html>. Accessed 16 Aug 2022
8. Sami Fadali, M., Visioli, A.: Digital Control Engineering: Analysis and Design, Chapter 12 Practical Issues. Academic Press, Cambridge (2012)
9. Gallego-Juárez, J.A., Rodríguez-Corral, G., Gaete-Garretón, L.: An ultrasonic transducer for high power applications in gases. *Ultrasonics* **16**(6), 267–271 (1978)

10. Gao, M., et al.: KITE: exploring the practical threat from acoustic transduction attacks on inertial sensors. In: 20th ACM Conference on Embedded Networked Sensor Systems (2022)
11. Giechaskiel, I., Rasmussen, K.: Taxonomy and challenges of out-of-band signal injection attacks and defenses. *IEEE Commun. Surv. Tutor.* **22**(1), 645–670 (2019)
12. Giechaskiel, I., Zhang, Y., Rasmussen, K.B.: A framework for evaluating security in the presence of signal injection attacks. In: Sako, K., Schneider, S., Ryan, P.Y.A. (eds.) *ESORICS 2019, Part I. LNCS*, vol. 11735, pp. 512–532. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29959-0_25
13. Hurley, R.: Design considerations for ESD/EMI filters: II low pass filters for audio filter applications. ON Semiconductor (2007)
14. National Instruments. NI USB-4431. <https://www.ni.com/pdf/manuals/376767a.pdf>. Accessed 16 Aug 2022
15. Kitchin, C.: Avoiding Op Amp instability problems in single-supply applications. Analog Devices, Tech. Rep (2001)
16. Kune, D.F., et al.: Ghost talk: mitigating EMI signal injection attacks against analog sensors. In: 2013 IEEE Symposium on Security and Privacy, pp. 145–159. IEEE (2013)
17. Laermer, F.: Mechanical Microsensors. *MEMS: A Practical Guide to Design, Analysis, and Applications*, pp. 523–566 (2006)
18. Mishali, M., Eldar, Y.C.: From theory to practice: sub-nyquist sampling of sparse wideband analog signals. *IEEE J. Sel. Top. Signal Process.* **4**(2), 375–391 (2010)
19. Nashimoto, S., Suzuki, D., Sugawara, T., Sakiyama, K.: Sensor confusion: defeating Kalman filter in signal injection attack. In: *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, pp. 511–524 (2018)
20. Park, Y., Son, Y., Shin, H., Kim, D., Kim, Y.: This ain't your dose: sensor spoofing attack on medical infusion pump. In: 10th USENIX Workshop on Offensive Technologies. USENIX (2016)
21. Pei, D., Salomaa, A., Ding, C.: *Chinese Remainder Theorem: Applications in Computing, Coding, Cryptography*. World Scientific (1996)
22. Petit, J., Stottelaar, B., Feiri, M., Kargl, F.: Remote attacks on automated vehicles sensors: experiments on camera and lidar. *Black Hat Europe* **11**(2015), 995 (2015)
23. Rigol. DG5352 function/arbitrary waveform generator. <https://rigol.com/products/DGdetail/DG5000>. Accessed 16 Aug 2022
24. SainSmart. UDB1002S DDS signal generator. <https://www.amazon.com/SainSmart-UDB1002S-Signal-Generator-Function/dp/B00JTR66CG/>. Accessed 10 July 2022
25. Sawigun, C., Thanapitak, S.: A 0.9-nW, 101-Hz, and 46.3-uv IRN low-pass filter for ECG acquisition using FVF biquads. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **26**(11), 2290–2298 (2018)
26. Shaeffer, D.K.: Mems inertial sensors: a tutorial overview. *IEEE Commun. Mag.* **51**(4), 100–109 (2013)
27. Shaw, M.L.: Accelerometer overload considerations for automotive airbag applications. *SAE Trans.*, 344–350 (2002)
28. Shin, H., Kim, D., Kwon, Y., Kim, Y.: Illusion and dazzle: adversarial optical channel exploits against lidars for automotive applications. In: Fischer, W., Homma, N. (eds.) *CHES 2017. LNCS*, vol. 10529, pp. 445–467. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66787-4_22
29. Söderkvist, J.: Micromachined gyroscopes. *Sens. Actuators, A* **43**(1–3), 65–71 (1994)

30. Son, Y., et al.: Rocking drones with intentional sound noise on gyroscopic sensors. In: 24th USENIX Security Symposium (USENIX Security 2015), pp. 881–896 (2015)
31. SparkFun. MiniGen pro mini signal generator shield. <https://www.sparkfun.com/products/11420>. Accessed 10 July 2022
32. Sreenivasulu, P., Hanumantha Rao, G., Rekha, S., Bhat, M.S.: A 0.3 V, 56 db DR, 100 Hz fourth order low-pass filter for ECG acquisition system. *Microelectron. J.* **94**, 104652 (2019)
33. Stilson, T.: Problems with the anti-aliasing filter. <https://ccrma.stanford.edu/CCRMA/Courses/252/sensors/node35.html>. Accessed 11 Dec 2022
34. STMICROELECTRONICS. Lpy550al. <https://pdf1.alldatasheetcn.com/datasheet-pdf/view/346169/STMICROELECTRONICS/LPY550AL.html>. Accessed 16 Aug 2022
35. Tharayil, K.S., et al.: Sensor defense in-software (SDI): practical software based detection of spoofing attacks on position sensors. *Eng. Appl. Artif. Intell.* **95**, 103904 (2020)
36. Tian, J., et al.: Mobile device fingerprint identification using gyroscope resonance. *IEEE Access* **9**, 160855–160867 (2021)
37. Trippel, T., Weisse, O., Xu, W., Honeyman, P., Fu, K.: Walnut: Waging doubt on the integrity of mems accelerometers with acoustic injection attacks. In: 2017 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 3–18. IEEE (2017)
38. Tu, Y., Lin, Z., Lee, I., Hei, X.: Injected and delivered: fabricating implicit control over actuation systems by spoofing inertial sensors. In: USENIX Security Symposium, pp. 1545–1562 (2018)
39. Tu, Y., Rampazzi, S., Hao, B., Rodriguez, A., Fu, K., Hei, X.: Trick or heat? Manipulating critical temperature-based control systems using rectification attacks. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pp. 2301–2315 (2019)
40. Tu, Y., Rampazzi, S., Hei, X.: Towards adversarial control loops in sensor attacks: a case study to control the kinematics and actuation of embedded systems. arXiv preprint [arXiv:2203.07670](https://arxiv.org/abs/2203.07670) (2022)
41. Wang, Z., Wang, K., Yang, B., Li, S., Pan, A.: Sonic gun to smart devices. Black Hat USA (2017)
42. Wenyuan, X., Yan, C., Jia, W., Ji, X., Liu, J.: Analyzing and enhancing the security of ultrasonic sensors for autonomous vehicles. *IEEE Internet Things J.* **5**(6), 5015–5029 (2018)
43. Yan, C., Wenyuan, X., Liu, J.: Can you trust autonomous vehicles: contactless attacks against sensors of self-driving vehicle. *Def Con* **24**(8), 109 (2016)
44. Yazdi, N., Ayazi, F., Najafi, K.: Micromachined inertial sensors. *Proc. IEEE* **86**(8), 1640–1659 (1998)
45. Zhang, G., Yan, C., Ji, X., Zhang, T., Zhang, T., Xu, W.: DolphinAttack: inaudible voice commands. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 103–117 (2017)