AFRL-AFOSR-JP-TR-2022-0094

Cyber Physical Analysis of System Software Survivability by Stimulating Sensors on Drones

Yongdae Kim
Korea Advanced Institute of Science and Technology
291 Daehak-ro, Yuseong-gu
Taejon, ,
KR

09/29/2022
Final Technical Report

Air Force Research Laboratory
Air Force Office of Scientific Research
Asian Office of Aerospace Research and Development
Unit 45002, APO AP 96338-5002

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
20220929	Final	START DATE 20200911	END DATE 20220810

4. TITLE AND SUBTITLE
Cyber Physical Analysis of System Software Survivability by Stimulating Sensors on Drones

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
	FA2386-20-1-4041	
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Yongdae Kim

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Korea Advanced Institute of Science and Technology
291 Daehak-ro, Yuseong-gu
Taejon
KR

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AOARD UNIT 45002 APO AP 96338-5002	AFRL/AFOSR IOA	AFRL-AFOSR-JP-TR-2022-0094

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
• We developed an acoustic injection testbed for MEMS gyroscopes and accelerometers; this testbed enables automated testing of the influence of compromised sensor values on drones, without the risk of physical damage to the drones.
• Using this testbed, we conducted rigorous experiments and discovered that sampling jitter is the essential factor influencing drone crashes during attacks. Notably, sampling jitter has not been discussed in previous studies.
• During our investigations, we discovered that sampling jitter produces noise-like signals. Based on this finding, we propose a novel prototype recovery system, UNROCKER, and demonstrated its capability through various experiments including real-world scenarios on physical sensors.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	SAR	4

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
AKIRA NAMATAME	3152277010

Cyber Physical Analysis of System Software Survivability by Stimulating Sensors on Drones

Yongdae Kim (PI), Dongkwan Kim, Dohyun Kim, Jaehoon Kim
KAIST
{yongdaek, dkay, dohyunjk, takeliberty}@kaist.ac.kr

Summary

- We developed an acoustic injection testbed for MEMS gyroscopes and accelerometers; this testbed enables automated testing of the influence of compromised sensor values on drones, without the risk of physical damage to the drones.
- Using this testbed, we conducted rigorous experiments and discovered that sampling jitter is the essential factor influencing drone crashes during attacks. Notably, sampling jitter has not been discussed in previous studies.
- During our investigations, we discovered that sampling jitter produces noise-like signals. Based on this finding, we propose a novel prototype recovery system, UNROCKER, and demonstrated its capability through various experiments including real-world scenarios on physical sensors.

Introduction

Several studies have been conducted to analyze security risks associated with such sensors. For instance, Son *et al*. [6] and Yazhou *et al*. [7] demonstrated that an acoustic sound noise can resonate the micro-electromechanical (MEMS) gyroscope sensors of IMUs, resulting in the crash of a drone to the ground. Similarly, Noh et al. *et al*. [3] proposed an adaptive approach to spoof GPS sensors. On the other hand, Kim *et al*. [2] analyzed misconfigured control parameters by adopting software fuzzing techniques. However, none of these studies clearly justifies the attack's underlying mechanism. One might imagine that the values of resonated sensors affect the values of the rotors, eventually crashing the drone. Is the attack always possible, regardless of the operational logic of sensors implemented on a target system? Unfortunately, it is difficult to answer this because existing studies have not deeply considered 1) the precise reasoning of the attack, 2) the practical conditions required for an end-to-end attack scenario, and 3) practical mitigation to prevent drones from crashing against the attack.

To address this, we 1) conduct a comprehensive study to understand the details of underlying operational logic of sensors and their associated attacks, and 2) propose practical mitigation against those attacks. We chose commercial drones as our analysis target in this study because they are widely-used CPSs in these days, from delivery services to recreational activities [5]. We focus on MEMS sensors, such as gyroscopes and accelerometers, because they have been extensively studied in previous studies [6], [7], [8], [1].

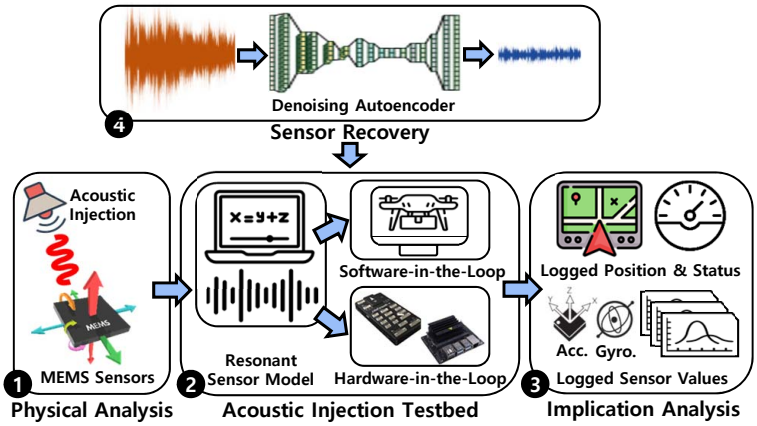

Fig. 1: Overview of our testbed and approach

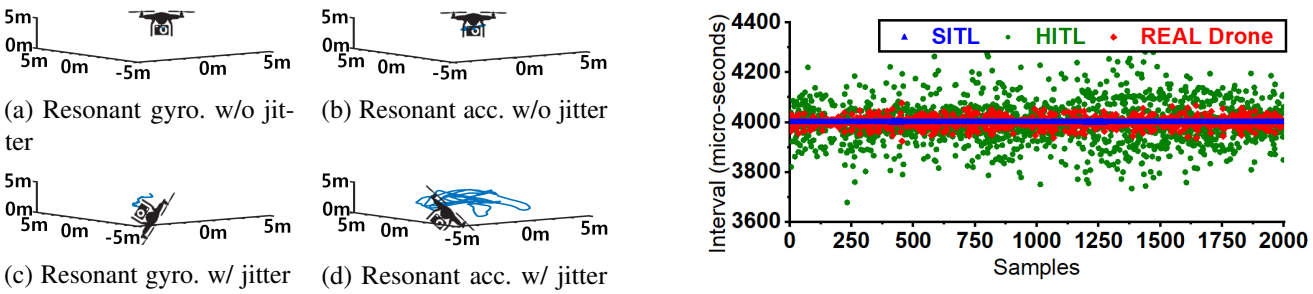

(a) Resonant gyro. w/o jitter (b) Resonant acc. w/o jitter

(c) Resonant gyro. w/ jitter (d) Resonant acc. w/ jitter

(e) Measured sampling intervals for each experiment

Fig. 2: Analyzing operation logic of sensors implemented in PX4.

Uncovering the Reasoning of Acoustic Noise-based Spoofing Attacks on Drones

First, to investigated the operational logic of MEMS sensors, we built a testbed based on PX4, the industry-leading open source drone project. Figure 1 depicts our testbed and overview of our approach. To simulate the acoustic attack, we conducted an empirical analysis of the resonated MEMS sensors to construct their software models. Then, we developed an automated analysis framework that can inject acoustic attack (*i.e.*, resonated) inputs into a simulated drone and monitor real-time sensor values. For this, we leveraged PX4's software-/hardware-in-the-loop (SITL/HITL) simulators.

We conducted four experiments to determine how a simulated drone (in the SITL mode) behaves when the resonated (*i.e.*, attacked) gyroscope/accelerometer signals are sampled at uniform/hardware intervals using our testbed. Interestingly, we discovered that the ideal uniform sampling scenario does not result in the crash of the simulated drone for both sensors, as shown in Figure 2a and Figure 2b. In contrast, when we conducted the same tests using real hardware, we discovered that the drone crashed to the ground, as illustrated in Figure 2c and Figure 2d. We used the HITL mode for these tests because conducting them directly with real drones would be prohibitively expensive; the drones can be broken apart. This result demonstrates that sampling from real hardware has another distinct factor that can cause the drone to crash other than the resonated sensor values. Figure 2e illustrates the measured sampling intervals in each experiment. While standard deviation of the sampling intervals was nearly 0 in the SITL experiments, that in the HITL experiments was approximately 457 μs and 3DR Solo showed approximately 103 μs standard deviation. This result demonstrates that sampling jitter is indeed present in real-world scenarios. However, is this sampling jitter sufficient to cause drone crashes?

To address this question, we performed additional in-depth analyses for the sampling jitter. We conducted the same SITL experiments, but by adding an intentional random jitter during the retrieval of sensor values in the sensor drivers. Thereafter, we investigated 1) drone crashes in the presence of jitter, and 2) the threshold of the jitter that caused the drone crashes for the sampled resonant sensor signals. For the gyroscopes, sampling jitter that was greater than 80 µs could cause the drone to crash. In the case of accelerometers, 6-µs sampling jitter was sufficient to crash the drone. For both the sensors, the threshold sampling jitter that crashed the drone was considerably smaller than the those in the HITL and actual drone experiments. Consequently, we concluded that the sampling jitter during hardware operation is a key factor influencing drone crashes during attacks.

Through careful analysis, we discovered that sampling jitter plays a critical role in drone crashes by spreading the resonant signals across various frequencies, including the in-band frequency range of the drones.

Mitigating Spoofing Attacks by Recovering Benign Sensor Values

Based on this knowledge, we propose a practical mitigation strategy that is capable of neutralizing sensor spoofing attacks. As shown in Figure 3, sampling jitter causes the dispersion of the resonant signal so even LPF cannot fully mitigate acoustic noise. Therefore, we propose UNROCKER, a sensor data recovery system using a denoising autoencoder [4], which is a DNN-based denoising filter, to neutralize compromised (*i.e.*, resonated) sensor values.

Figure 4b shows the overall architecture of our recovery system UNROCKER. UNROCKER is designed to operate in the middle of the sensor drivers and an flight controller; thus it directly recovers the digitized values from the sensors. UNROCKER receive the sensor data and transfers the value to its pre-trained DAE model. The architecture of a DAE model we used is shown in Figure 4a. We utilized the resonant sensor models and our testbed, described

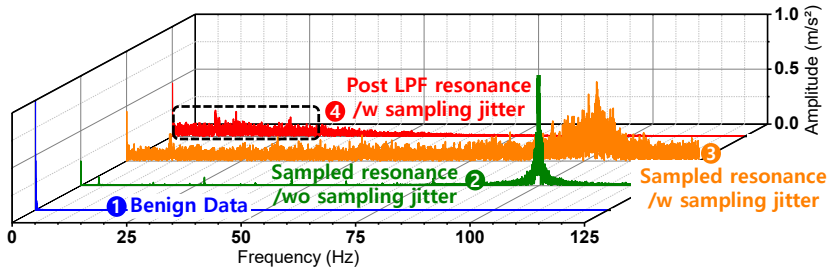

Fig. 3: Spectrum analysis of the sampled resonance signals

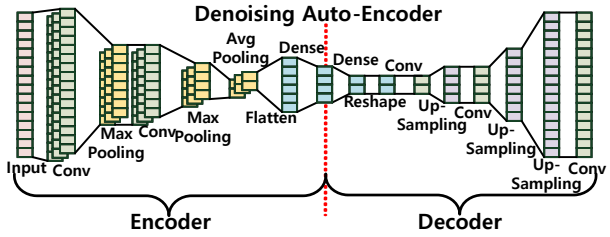

(a) Architecture of a denoising autoencoder (DAE)

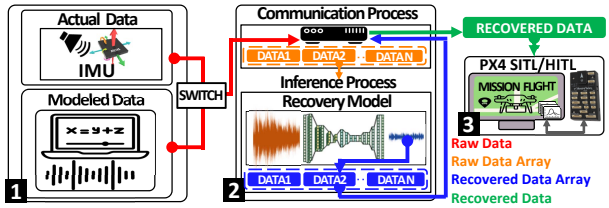

(b) Overview of our recovery system

Fig. 4: Architecture of our recovery system and DAE

in Figure 1, for the construction of the training dataset. For the HITL simulation and real-world evaluation, we used two target drones, namely 3DR Iris and Solo; these are based on the Pixhawk board.

The evaluation results are illustrated in Table I. As shown in the table, UNROCKER successfully recovered the compromised sensor signals for all the amplitudes by significantly reducing the errors. We also evaluated whether UNROCKER can recover the sensor data reflecting the dynamics of real flights and UNROCKER successfully reduced the standard deviation of errors for all cases. Moreover, we evaluated with actual sensors with actual acoustic noise in stationary cases and UNROCKER successfully reduced the standard deviation of the errors from 2.361 to 0.935 for the accelerometers, and from 1.469 to 0.022 for the gyroscopes.

References

[1] S. Khazaaleh, G. Korres, M. Eid, M. Rasras, and M. F. Daqaq. Vulnerability of mems gyroscopes to targeted acoustic attacks. *IEEE Access*, pages 89534–89543, 2019.

[2] Taegyu Kim, C. Kim, J. Rhee, Fan Fei, Z. Tu, Gregory Walkup, X. Zhang, X. Deng, and D. Xu. "rvfuzzer: Finding input validation bugs in robotic vehicles through control-guided testing". In *USENIX Security Symposium*, 2019.

[3] Juhwan Noh, Yujin Kwon, Yunmok Son, Hocheol Shin, Dohyun Kim, Jaeyeong Choi, and Yongdae Kim. Tractor beam: Safe-hijacking of consumer drones with adaptive gps spoofing. *ACM Transactions on Privacy and Security*, 22:1–26, 04 2019.

[4] Adrian Rosebrock. Denoising autoencoders with keras, tensorflow, and deep learning. *pyimagesearch*, 42(1):85–97, 2020.

[5] LUKAS SCHROTH. The drone market size 2020.

[6] Yunmok Son, H. Shin, D. Kim, Y. Park, Juhwan Noh, Kibum Choi, Jungwoo Choi, and Y. Kim. "rocking drones with intentional sound noise on gyroscopic sensors". In *USENIX Security Symposium*, 2015.

[7] Timothy Trippel, Ofir Weisse, Wenyuan Xu, Peter Honeyman, and Kevin Fu. "walnut: Waging doubt on the integrity of mems accelerometers with acoustic injection attacks". pages 3–18, 04 2017.

[8] Yazhou Tu, Zhiqiang Lin, Insup Lee, and Xiali Hei. "injected and delivered: Fabricating implicit control over actuation systems by spoofing inertial sensors". In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 1545–1562, 2018.

TABLE I: Standard deviation of errors against benign signals

Amplitude of Resonance Signals		Accelerometer ($\sigma, m/s^2$)				Gyroscope ($\sigma, rad/s$)			
		20	40	60	80	1	2	3	4
Compromised Signals		14.13	28.28	42.38	56.64	0.707	1.414	2.12	2.83
				↓			↓		
Recovery of 3DR Iris	X-axis	0.270	0.281	0.282	0.296	0.012	0.012	0.014	0.013
	Y-axis	0.086	0.081	0.097	0.083	0.021	0.022	0.022	0.026
	Z-axis	0.871	0.923	0.903	0.930	0.011	0.012	0.014	0.014
Recovery of 3DR Solo	X-axis	0.311	0.308	0.342	0.322	0.540	0.527	0.542	0.532
	Y-axis	0.125	0.118	0.131	0.143	0.057	0.057	0.057	0.064
	Z-axis	0.973	0.971	0.970	1.037	0.043	0.040	0.045	0.047
Recovery of Flight Data	X-axis	3.461	3.473	3.476	4.078	0.097	0.109	0.113	0.116
	Y-axis	2.212	2.213	2.213	2.228	0.051	0.052	0.054	0.055
	Z-axis	2.678	2.712	2.738	2.827	0.042	0.042	0.042	0.042

[1] The red cells indicate that the standard deviation of the errors was destructive to drones. Note that the minimum value to crash both Iris and Solo drones was 35.4 and 1.41 for accelerometers and gyroscopes, respectively.

[2] The amplitudes of the benign sensor signals were ranged from -18 to +3 m/s^2 for accelerometers, and from -5 to +5 rad/s for gyroscopes