Embedded Security

EECE 5698-08: Special Topics: Cyber-Physical Security of IoT Systems in the Age of AI

Lecture 6: Sound and Confidentiality
"Everything is a Microphone"

Prof. Kevin Fu

September 29, 2025

https://spqrlab1.github.io/emsec/

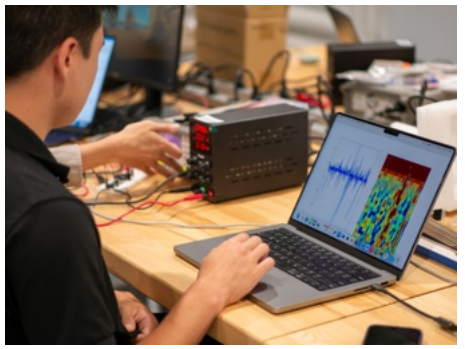

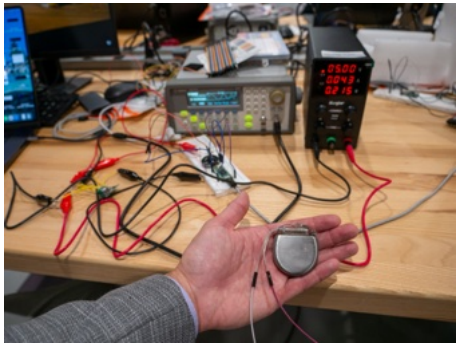

Today's Learning Goals

- Learn how everyday devices can be synthesized into microphones
  - Gyroscopes
  - Hard drives
  - Fiber optic cables
  - Anything that transduces acoustic energy into another form of energy
- Learn how one-way valves are usually two-way:
  - Microphones are speakers
  - Speakers are microphones

Review: Unintentional Demodulation

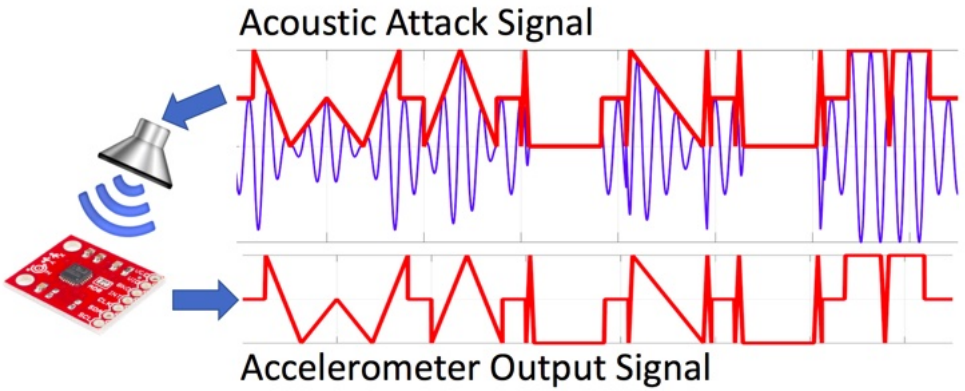

vs.

Both: Intentional signal modulation

Intentional
signal demodulation

<u>Unintentional</u>
signal demodulation

Review: MEMS Sensors

- Micro-Electro-Mechanical Systems
  - **Accelerometers**
  - **Gyroscopes**
  - **Clocks**
- Advantages
  - Low cost
  - Low power
    some < 1 mA
  - Small integrated circuit

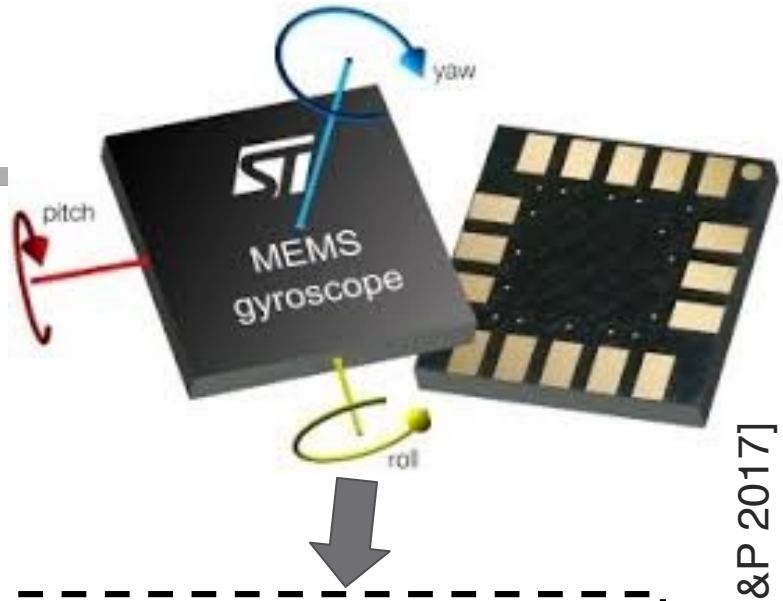

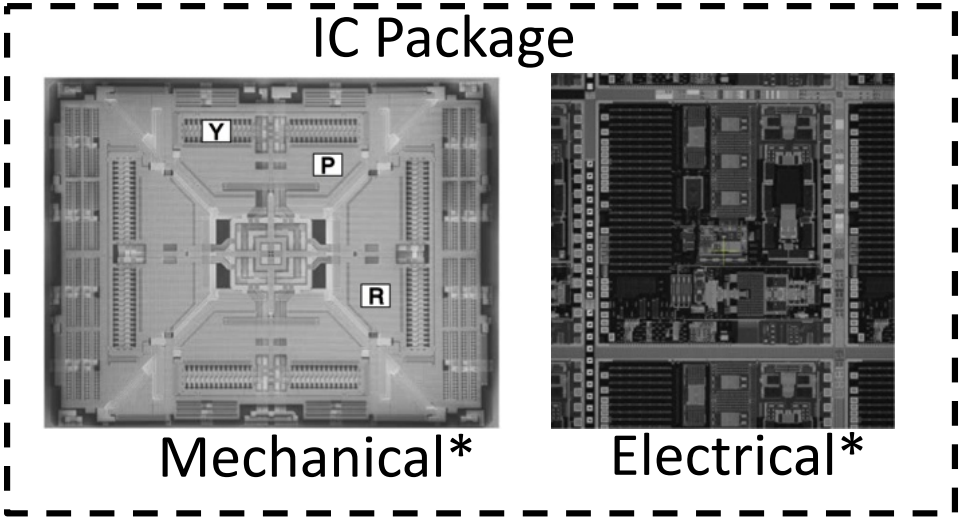

*Photos courtesy of "*Everything about STMicroelectronics' 3-axis digital MEMS gyroscopes – Technical Report*", by STMicroelectronics.

["WALNUT" by Trippel et al., IEEE Euro S&P 2017]

Review: Signal Distortion

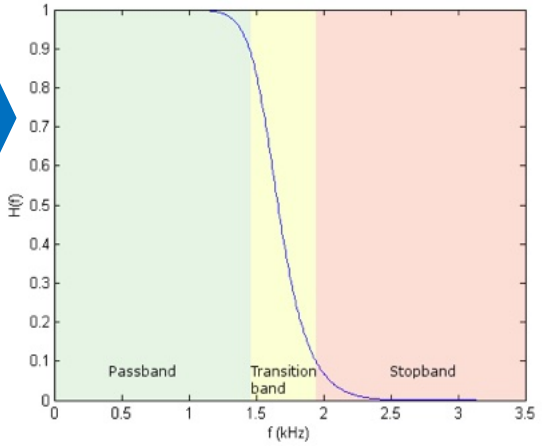

[Trippel et al., IEEE Euro S&P 2017]

Two types of spoofed acceleration

- Fluctuating accelerometer output
- Constant accelerometer output

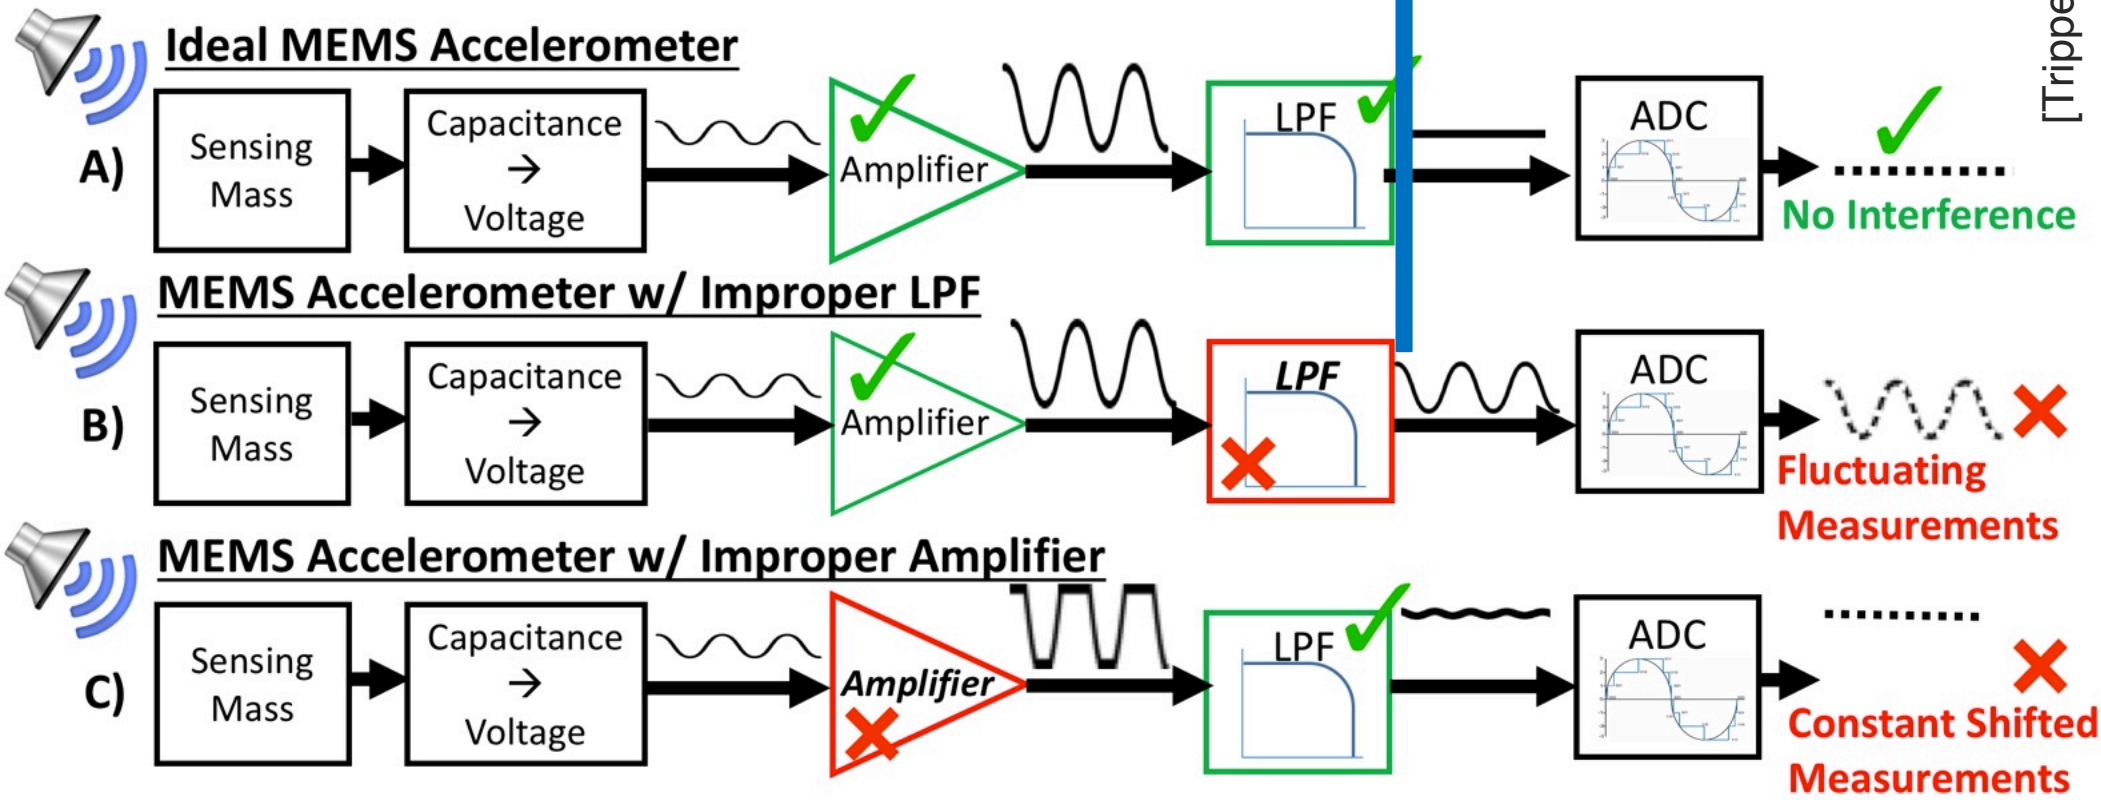

asymmetric clipping

Review: Output Biasing via Aliasing

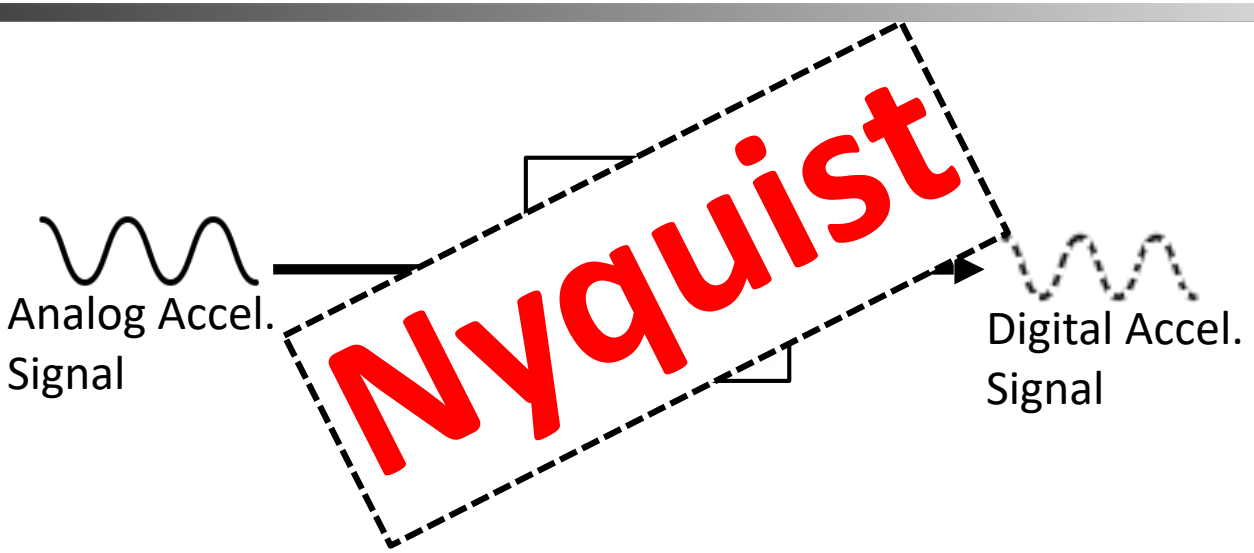

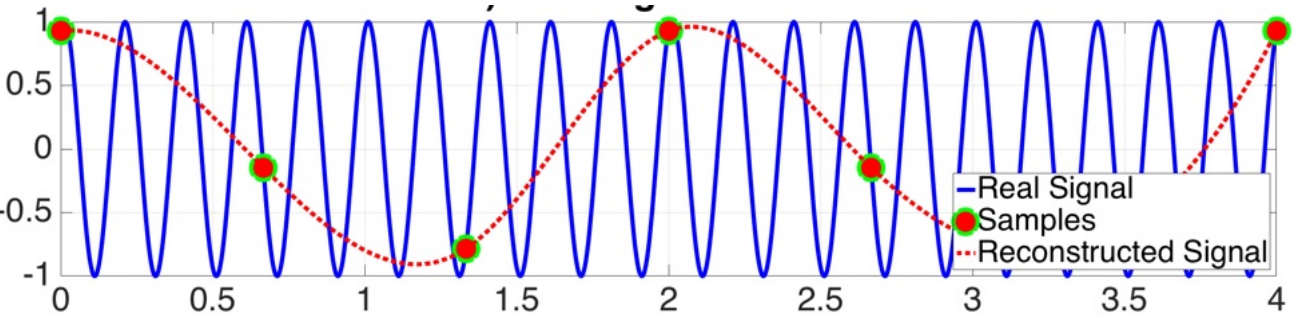

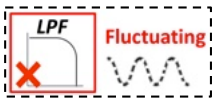

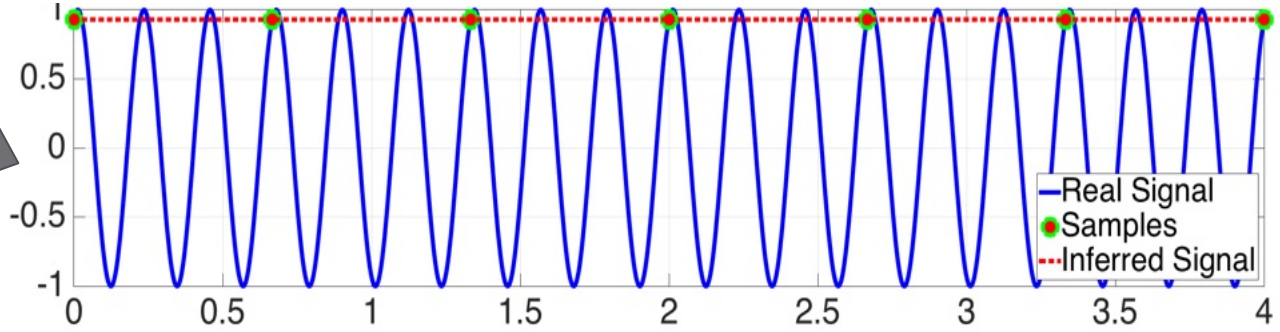

["WALNUT" by Trippel et al., IEEE Euro S&P 2017]

Everything is a Microphone

https://crypto.stanford.edu/gyrophone/

A Window into Stomach Anatomy

- 203 years ago in 1822, Alexis St. Martin was accidentally shot in stomach at Fort Mackinac, Michigan
- Dr. William Beaumont experimented on him
- Created a "window" and portal into the stomach
- Observe the stomach at work and even inject stuff into the portal
- Learned a lot, but ethically dubious

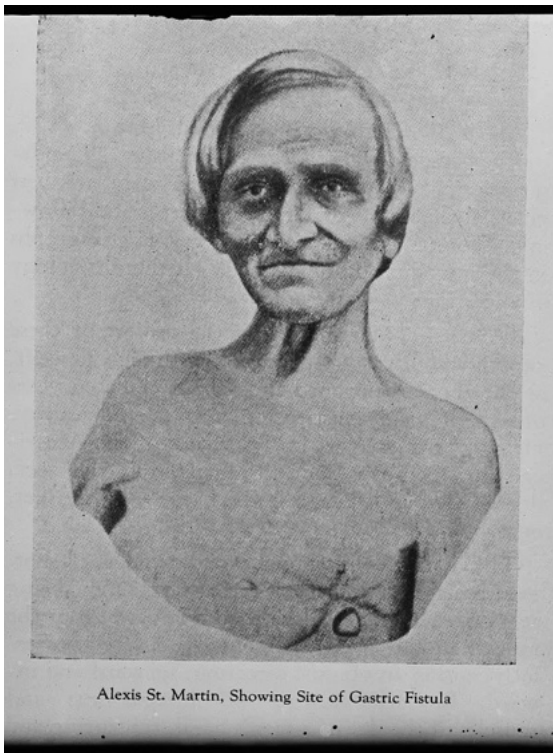

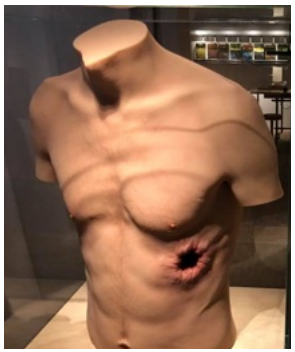

https://www.smithsonianmag.com/smart-news/grisly-story-human-guinea-pig-alexis-st-martin-180963520/
https://diseasesofmodernlife.web.ox.ac.uk/article/sneak-preview-of-victorian-speed

Sound Motivation

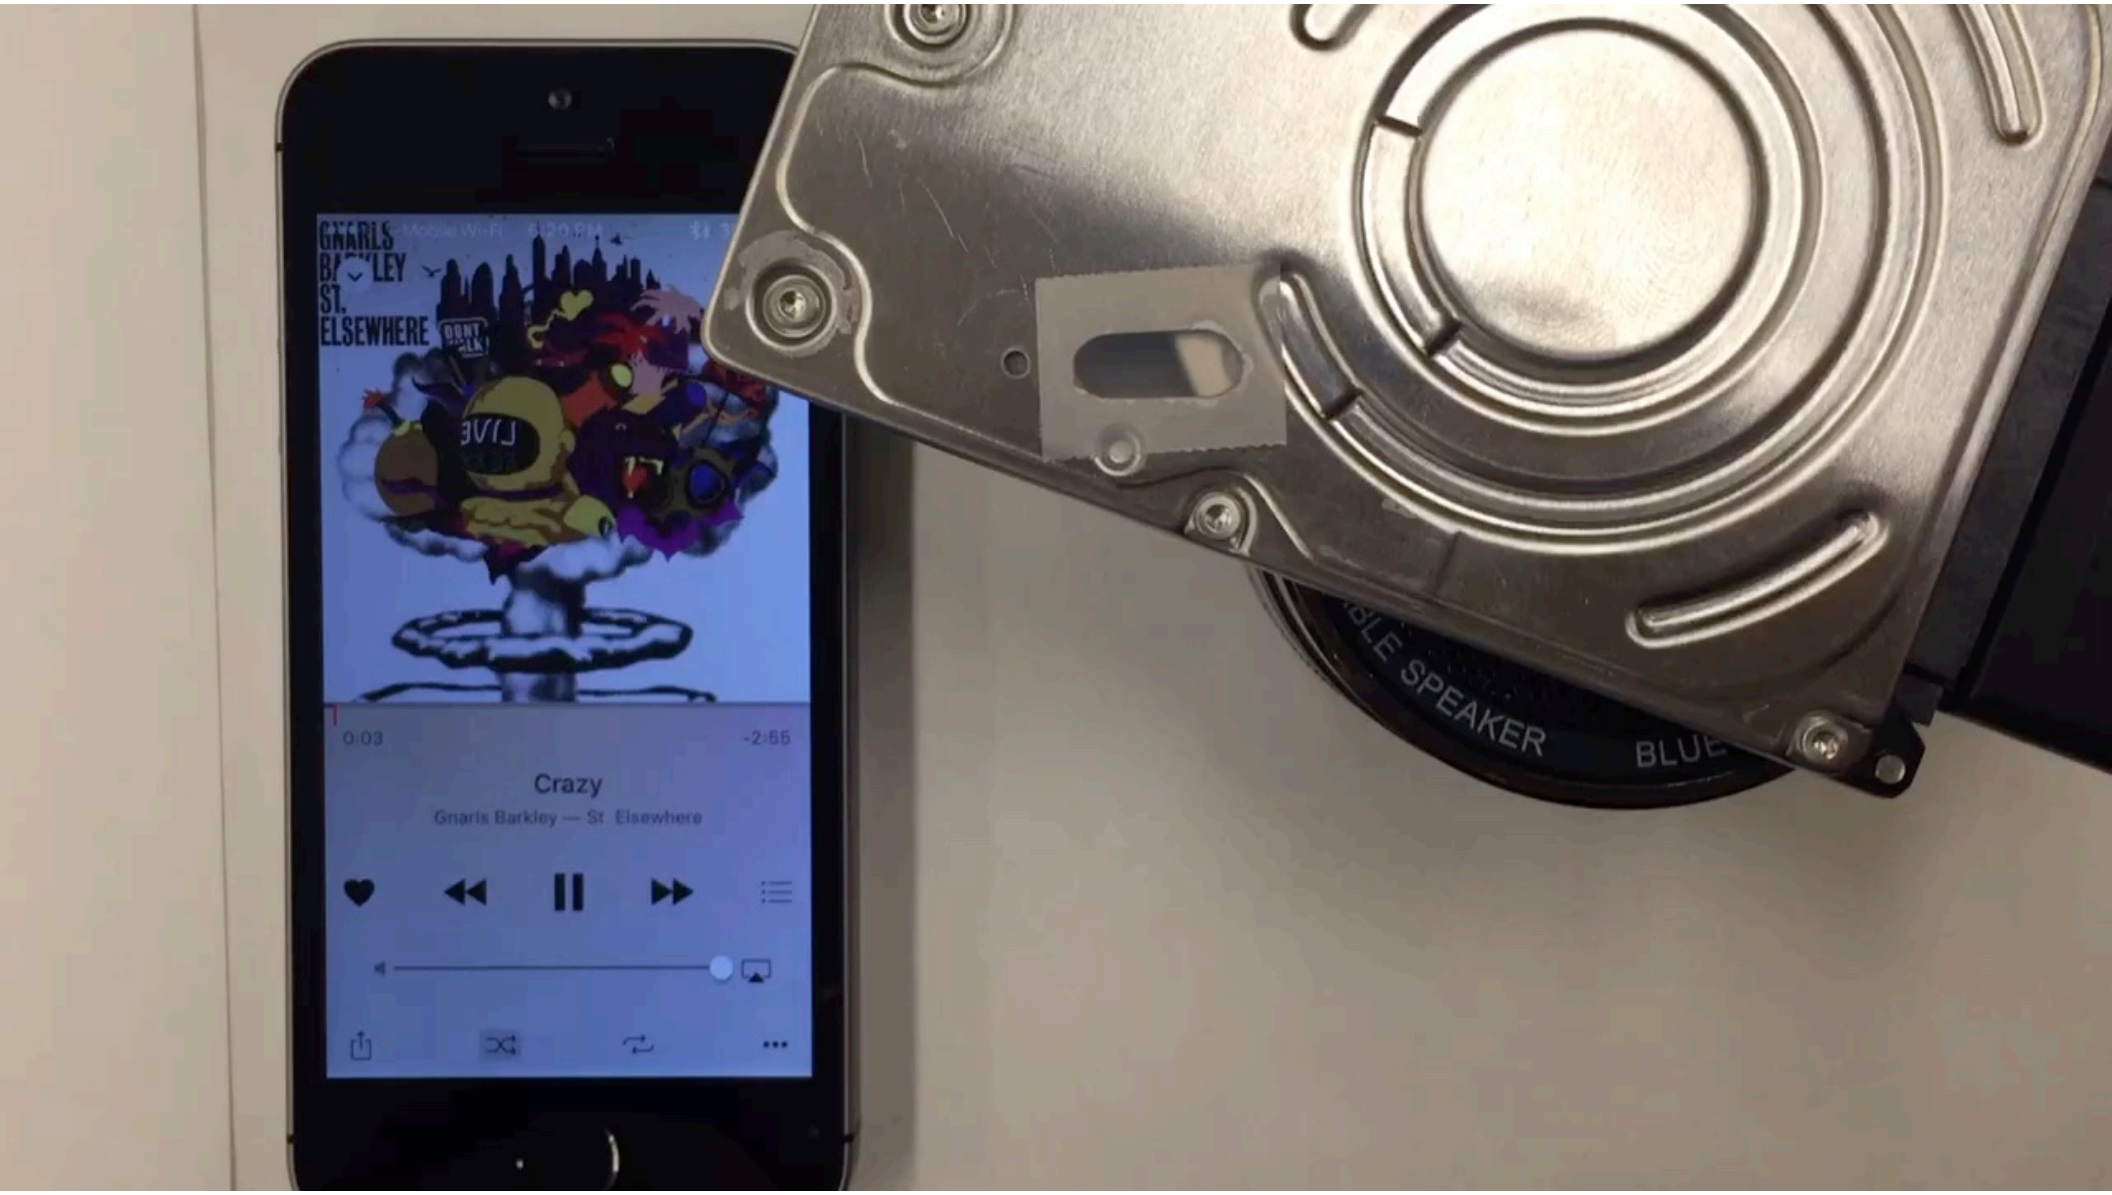

Hard Drive of Hearing [IEEE S&P '19]

Hard Drive of Hearing: Disks that Eavesdrop with a Synthesized Microphone

Andrew Kwong[1], Wenyuan Xu[2], and Kevin Fu[1]

[1]University of Michigan
[2]Zhejiang University

Abstract—**Security conscious individuals may take considerable measures to disable sensors in order to protect their privacy. However, they often overlook the cyberphysical attack surface exposed by devices that were never designed to be sensors in the first place. Our research demonstrates that the mechanical components in hard drives behave as microphones with sufficient precision to extract and parse human speech. These unintentional microphones sense speech with high enough fidelity for the Shazam service to recognize a song recorded through the hard drive. This proof of concept attack sheds light on the possibility of invasion of privacy even in absence of traditional sensors. We also present defense mechanisms, such as the use of ultrasonic aliasing, that can mitigate acoustic eavesdropping attacks by hard drives.**

1. Introduction

Magnetic hard disk drives continue to persist in everything from legacy laptops to server racks [1]. Because of their critical role in a wide variety of applications, hard drives make an annealing target for both cyber criminals

use this offset, known as the Position Error Signal (PES), in a feedback control loop; the microprocessor takes the PES as input for actuating the read/write head by use of a voice-coil motor (VCM) [4].

For both read and write operations, the read/write head can tolerate deviation from the center only on the order of nanometers. Accordingly, PES measurements are taken at a very fine granularity. These extremely precise measurements are sensitive to vibrations caused by even the slightest fluctuations in air pressure, such as those induced by human vocalizations.

Extracting speech from the PES, however, is complicated due to a weak signal-to-noise-ratio (SNR). Imperfections in the eccentricity of the platters, thermal drift, and turbulence from the rapid rotation of the disks all contribute to a large quantity of noise in the signal [5]. Through a mixture of digital filtering techniques in both the time domain and the frequency domain, however, we have managed to sufficiently clean the signal such that human speech can be completely reconstructed under certain conditions.

To prove the existence of this acoustic side-channel, we measured the PES directly from the hard drive under oper-

Turn hard drives into microphones with firmware

Ph.D. student: Andrew Kwong

Hard Drive of Hearing: Disks that Eavesdrop with a Synthesized Microphone

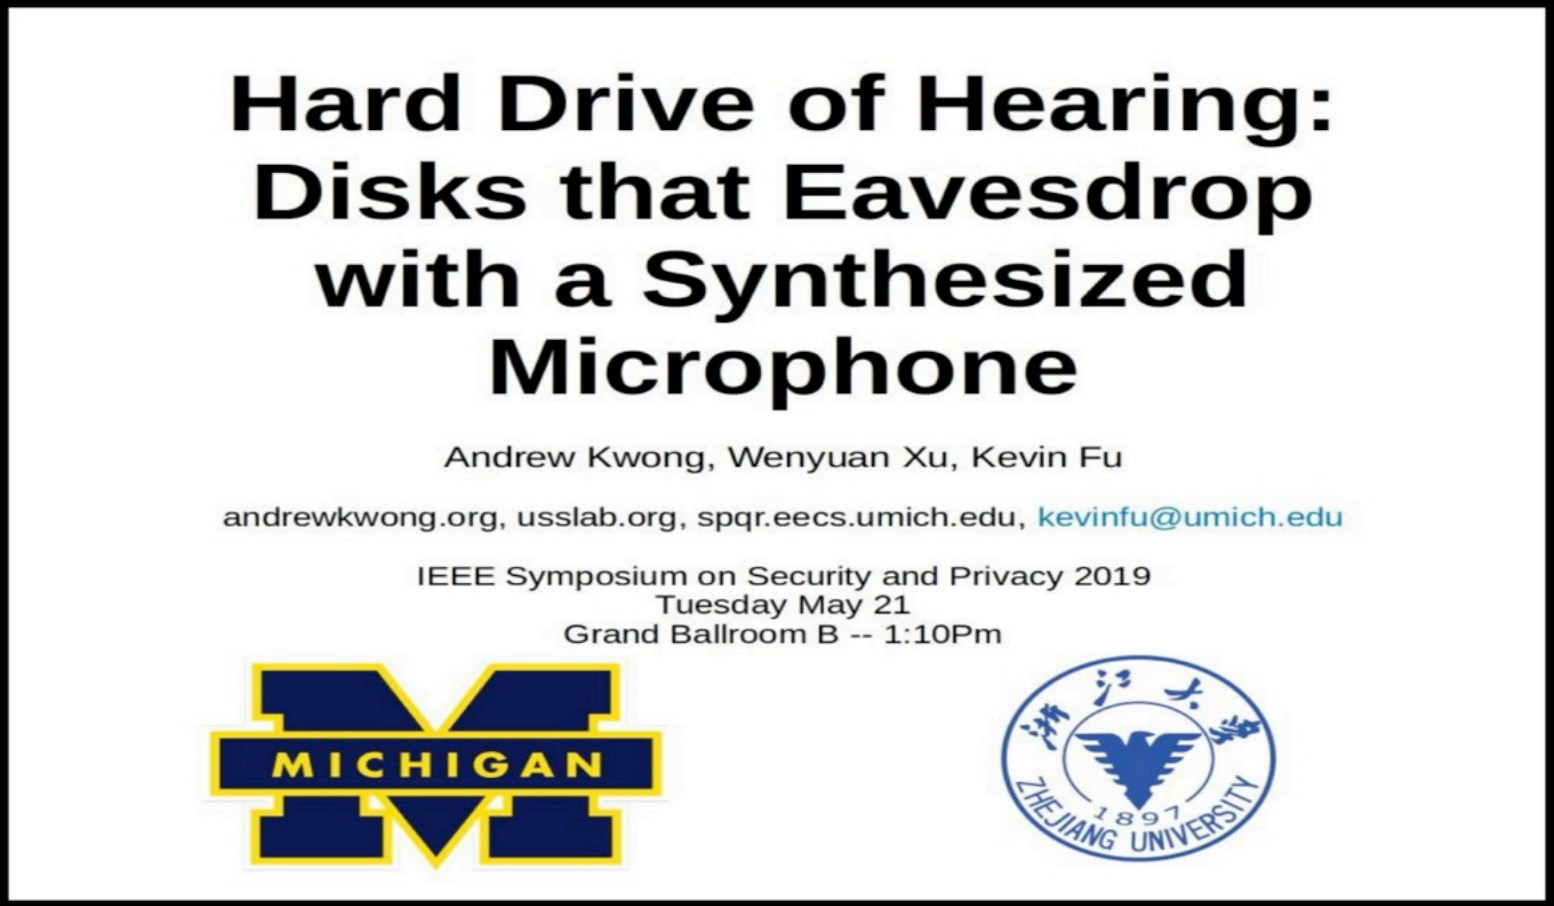

https://www.youtube.com/watch?v=2EHQCuWl2jg

Hard Drive of Hearing: Disks that Eavesdrop with a Synthesized Microphone

Andrew Kwong, Wenyuan Xu, Kevin Fu

andrewkwong.org, usslab.org, spqr.eecs.umich.edu, kevinfu@umich.edu

IEEE Symposium on Security and Privacy 2019
Grand Ballroom B -- 1:10Pm

Sensors Intrude on Privacy

- Accelerometers can leak keystrokes [1], gyroscopes can leak voice [2], etc.
- What is the threat from devices never intended to be sensors in the first place?

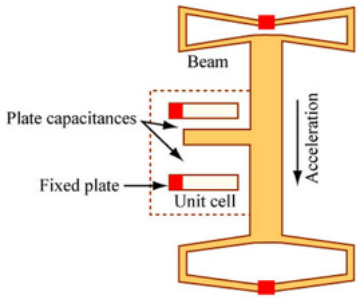

Accelerometers: [1] Marquardt et al., CCS '11, “(sp)iPhone..."
Gyroscopes: [2] Michalevsky et al., Usenix Security '14, “Gyrophone...”

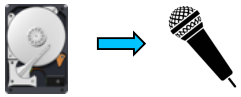

Hard Drive as a Microphone?

Challenges:

- HDDs are not designed as microphones
- Large quantity of self-noise
- Low signal-to-noise ratio

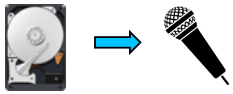

Contributions

HDD as a microphone

- Used SNReval measurements to evaluate extracted speech quality
- Used Shazam to recognize song recovered through HDD

Mitigations

- Ultrasonic aliasing
- Firmware signatures

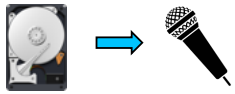

Threat Model

Firmware Resident Malware

- Drive firmware can be flashed from software

Flashing:

- MITM attacks (POODLE, LOGJAM, DROWN)
- Any compromise granting root access to a machine

2007

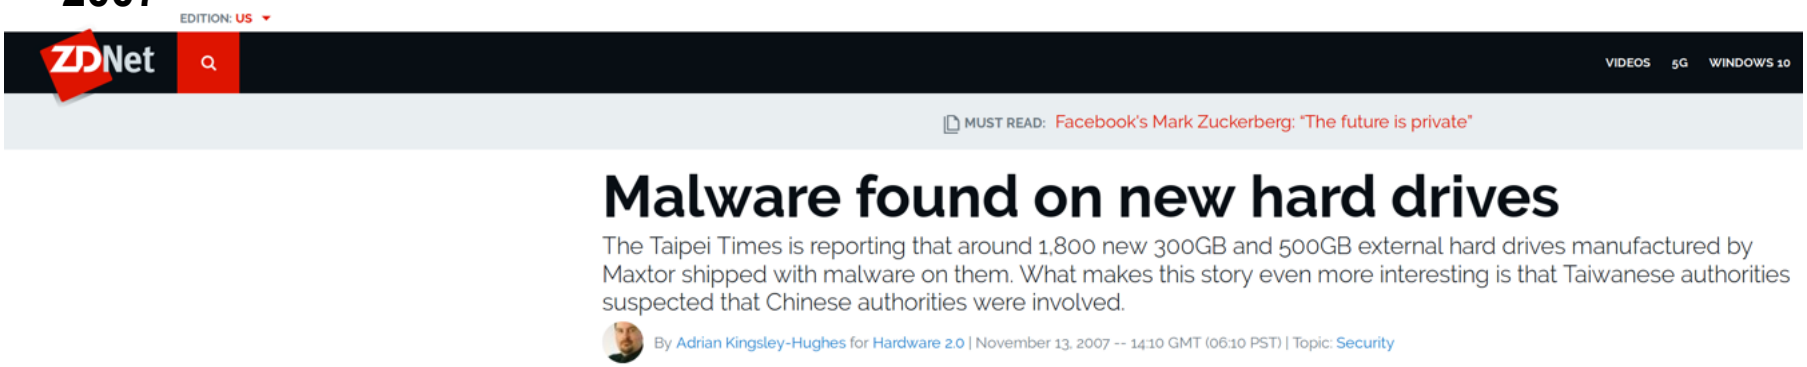

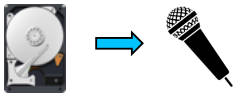

2018

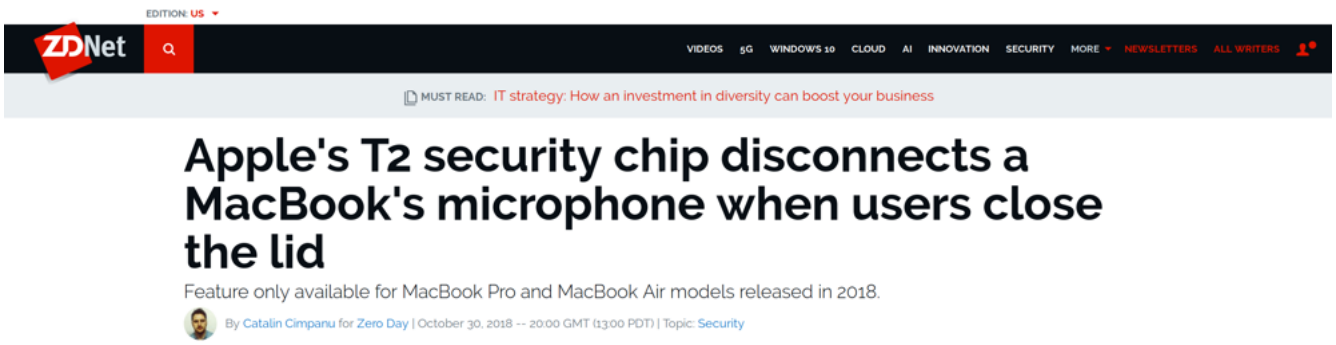

http://stahlke.org/dan/phonemute/

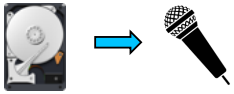

HDD as a microphone

- Head stack assembly actuates the read/write head as the disk spins beneath it
  - Head follows a track
  - can tolerate only tiny errors
- Position Error Signal(PES):
  - Head's offset from center of current track

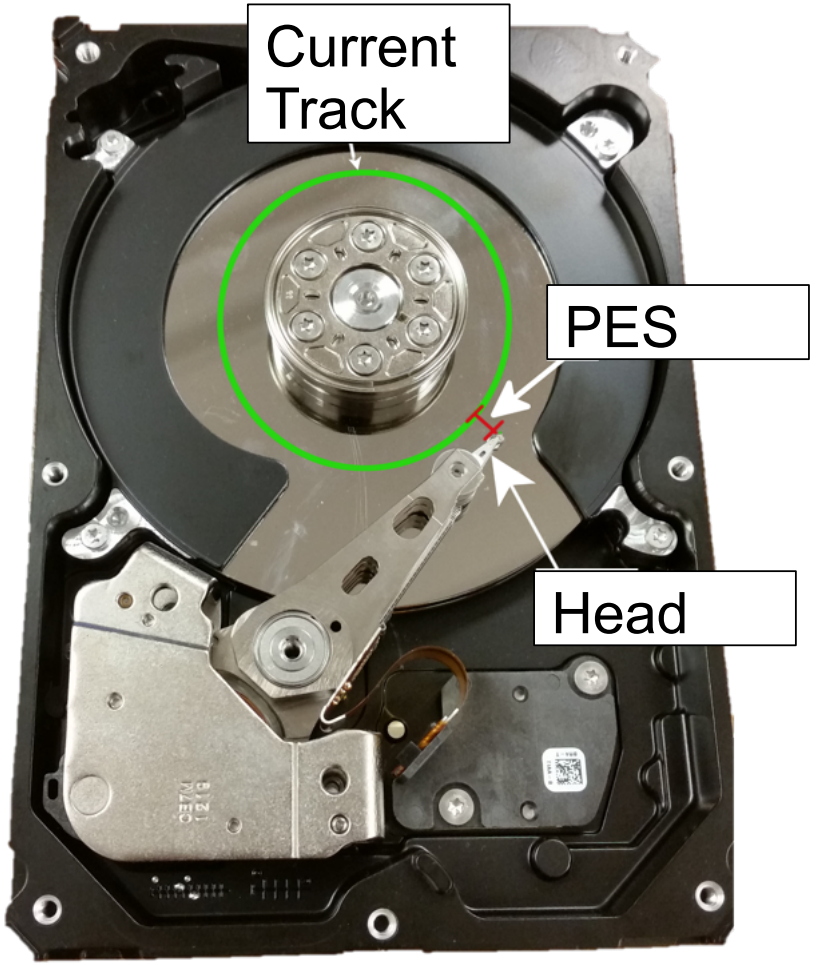

Head Tracking

- Utilizes Feedback-Control Loop to keep head on track
- Generates PES by reading out magnetic burst from servo sectors
  - Fixed number of servo sectors per track

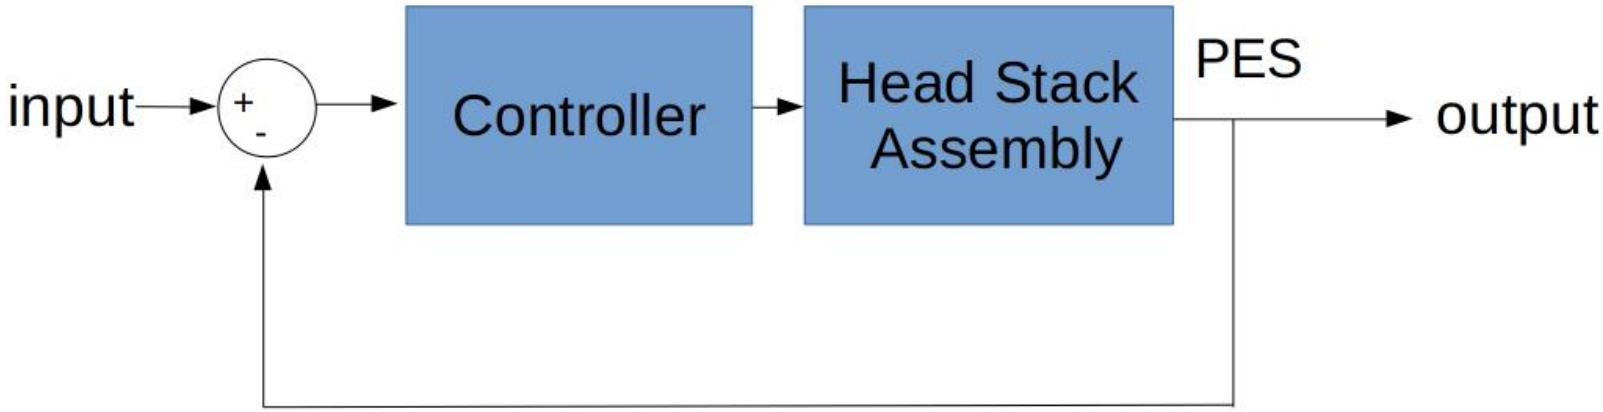

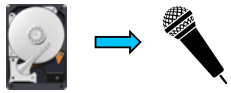

Similarities to Microphone

Microphone:

- Output measures diaphragm displacement
- Sound waves displace diaphragm

HDD:

- PES measures read/ write head displacement
- Sound waves displace write head?

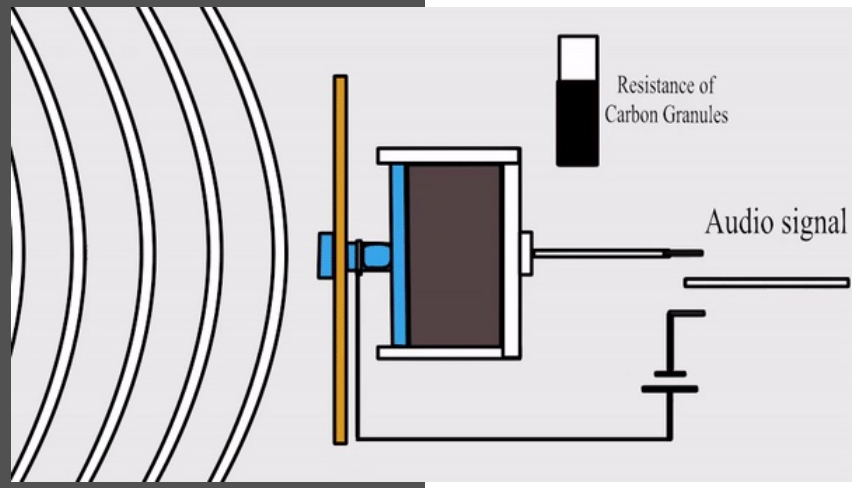

https://www.instructables.com/id/Simplified-Electronics-Microphone-DIY-How-It-Works/

PES approximates microphone output??

Measuring the PES

- Under our threat model, attacker would read it through firmware resident malware
  - Zaddach et al. [3] developed HDD firmware malware
- Proof of concept: suffices to read PES by tapping a debug pin
  - Used serial diagnostic port to output PES

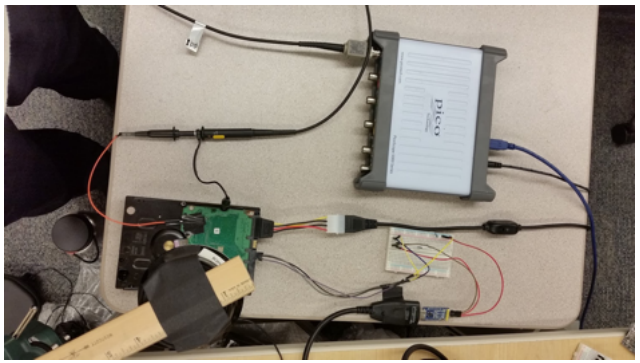

HDD Malware: [3] Zaddach et al., ACSAC '13

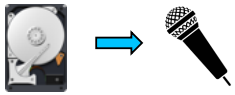

Sampling Rate

$$frequency_{sampling} = frequency_{rotation} * num\_servo\_sectors\_per\_track$$

$$= 120\ Hz * 288$$

$$= 34{,}560\ Hz$$

Nyquist-Shannon Sampling theorem:

- need sample at 2x the frequency of signal

Audible sound: 20 Hz-20 kHz

- Male fundamental: 85-180 Hz
- Female fundamental: 156-255 Hz
- POTS: 8 kHz

Experimental Setup

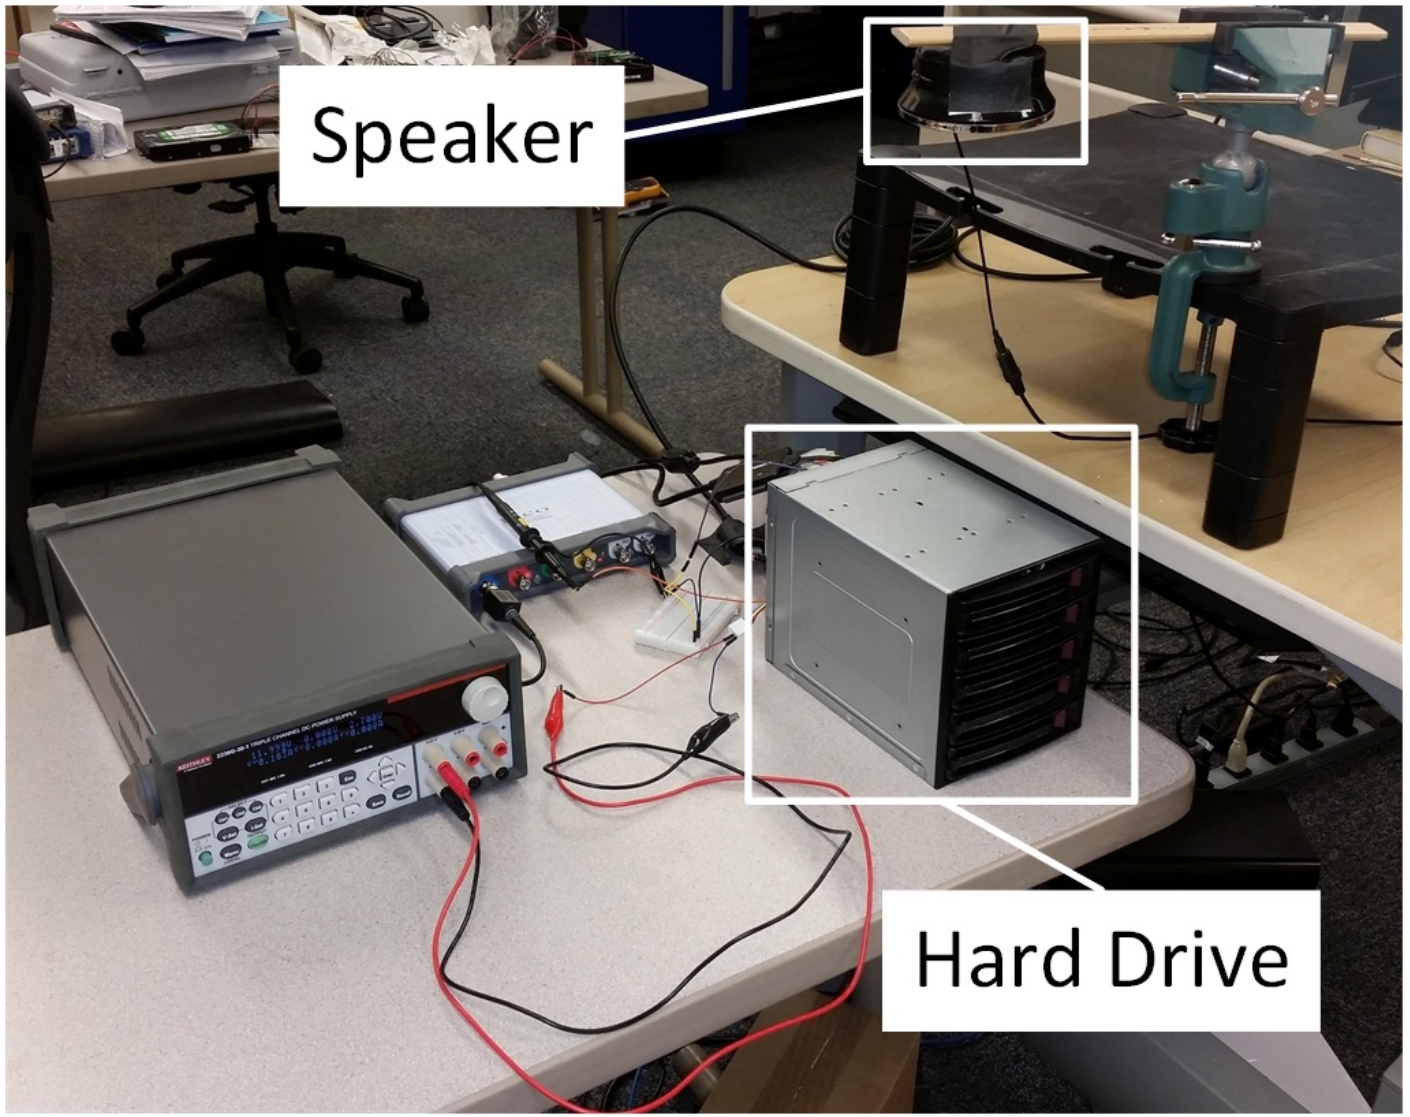

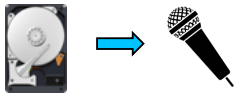

Speech Recovery

Must recover speech from PES readings

- PES values approximate instantaneous air pressure readings
- Wrote normalized PES values to WAV file

Noise from:

- Platter eccentricity
- Thermal drift
  - Errors 300X width of track
- turbulence

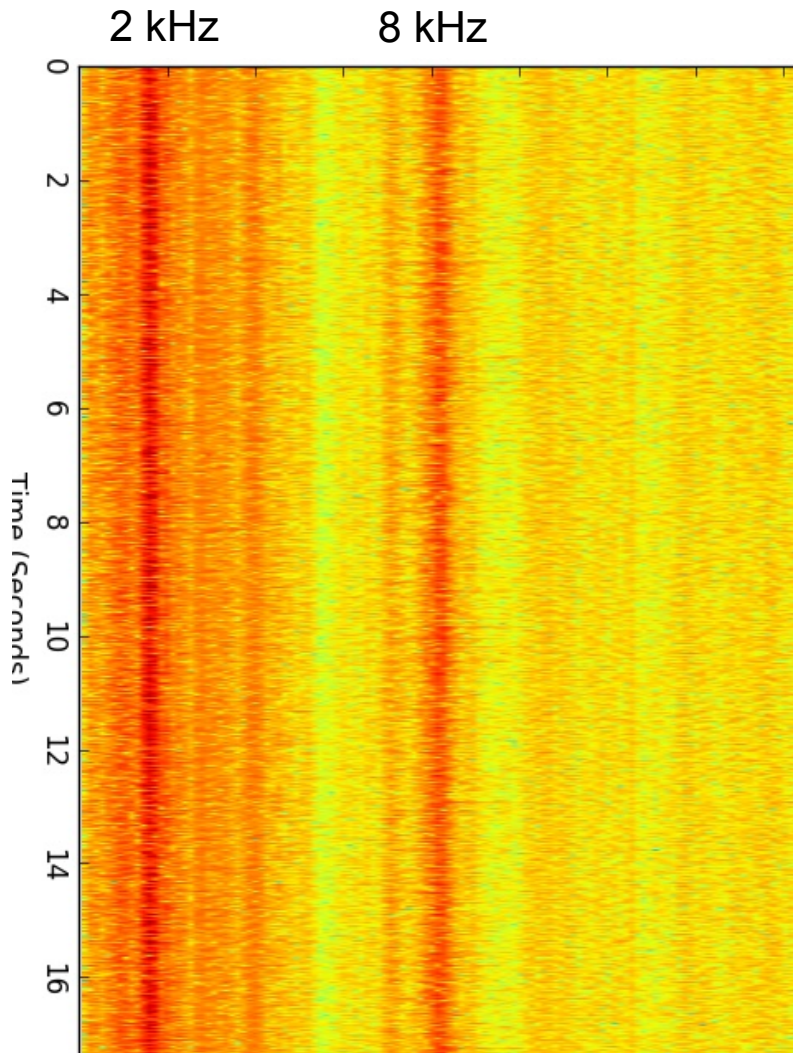

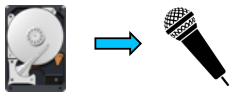

Signal Analysis

From HDD:

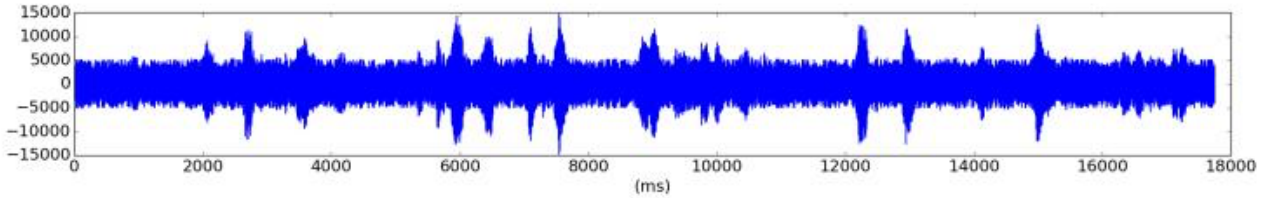

Original:

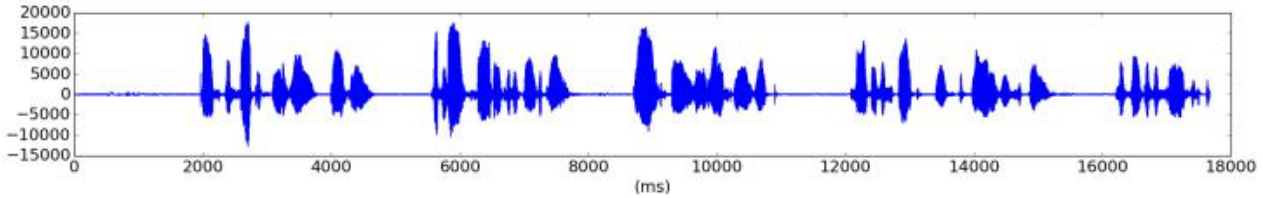

Cross Correlation:

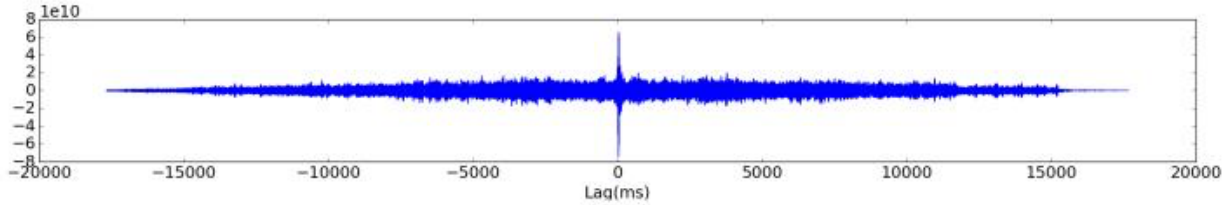

- Harvard Sentence male speaker with drive enclosed in case and fan powered at max (42W)

Frequency Response

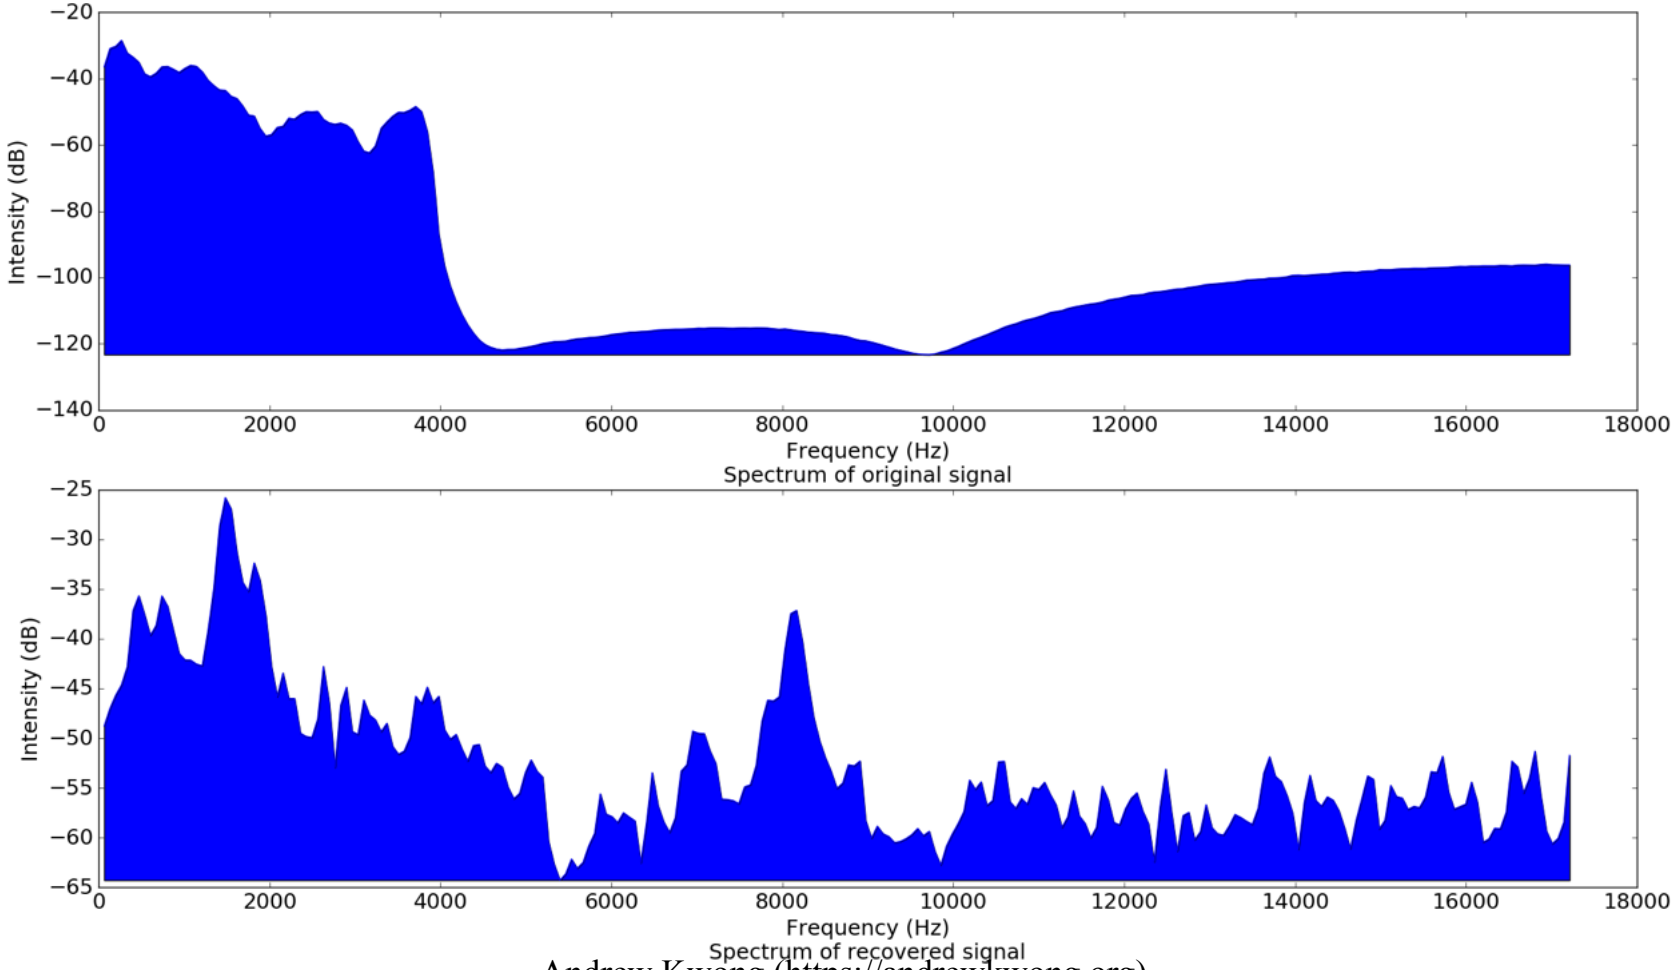

Spectral Analysis

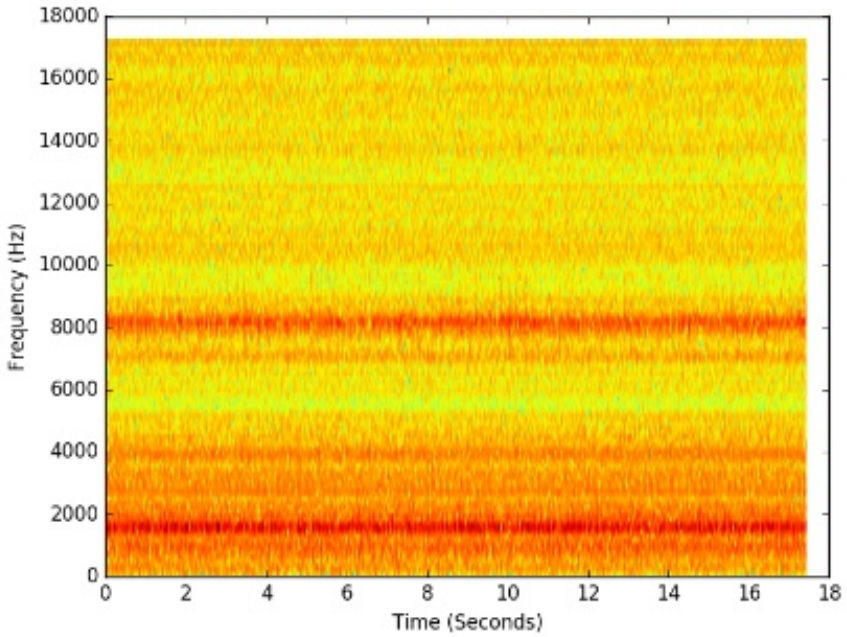

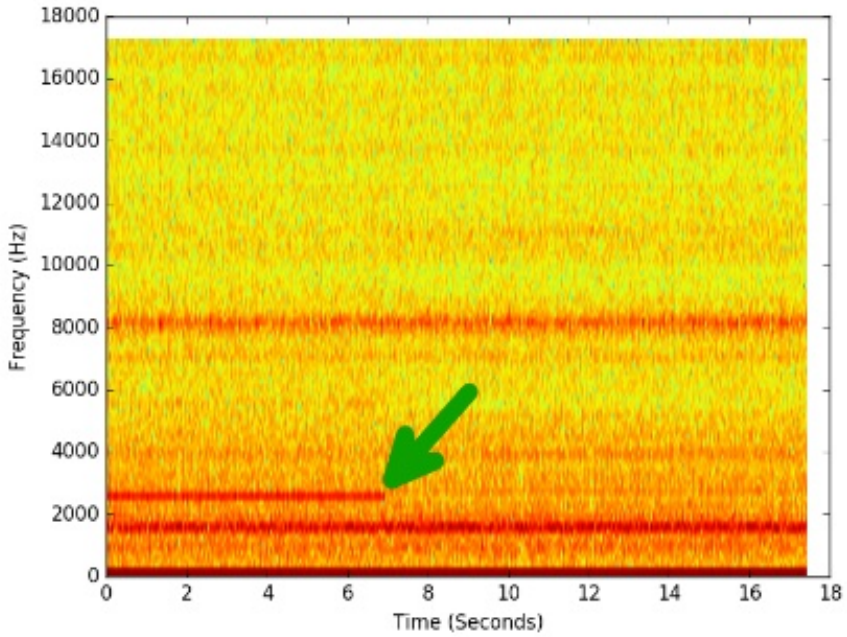

- Heavy bands of persistent noise around 8 kHz and 1900 kHz
- Responds well to 2.5 kHz tone

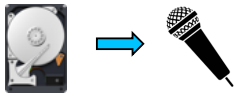

Reading PES

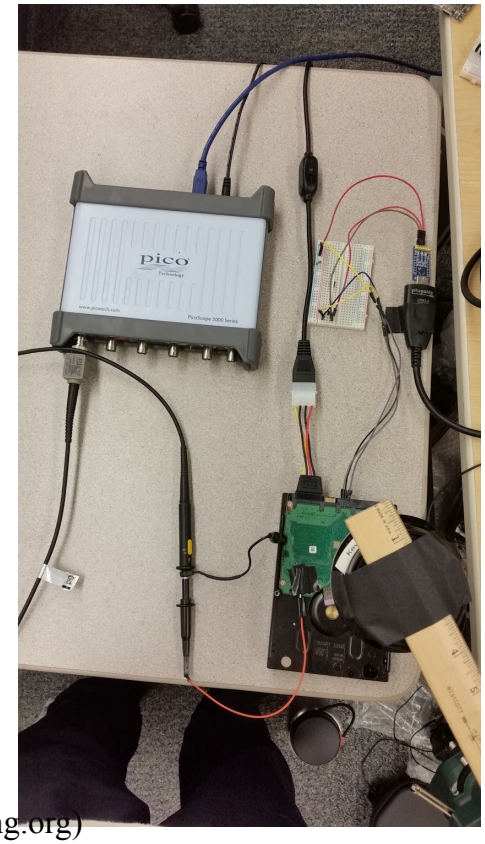

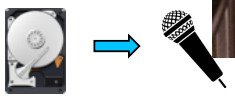

Digital Signal Processing

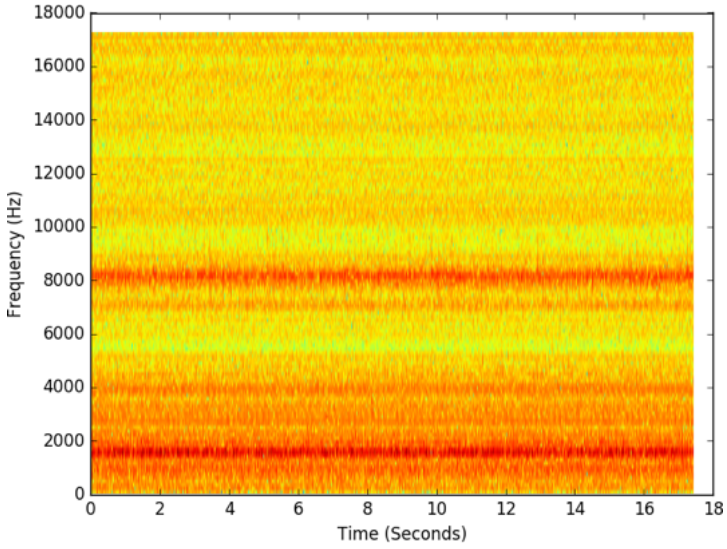

- Linearly filtering out 8 kHz and 1.9 kHz removes the heaviest bands of noise
- Made use of spectral noise gating for further filtering
  - Find noise thresholds at smaller sub-bands, only pass frequencies above the threshold

Quantitative Measures

PESQ MOS: Perceptual Evaluation of Speech Quality.

- Estimates intelligibility of speech
- Baseline: 1.7dB
- From exposed HDD: 1.4 dB
- Inside external hard drive enclosure: 1.6 dB

Enclosure actually improved results!

- Container presents a larger surface area to oncoming waves

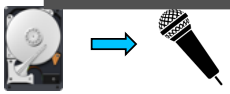

Speech Sample

Transcription:

- Paint the sockets in the wall dull green.
- The child crawled into the dense grass.
- Bribes fail where honest men work.
- Trample the spark, else the flames will spread.

Shazam Recognition

- Played Iron Maiden's "The Trooper" at hard drive

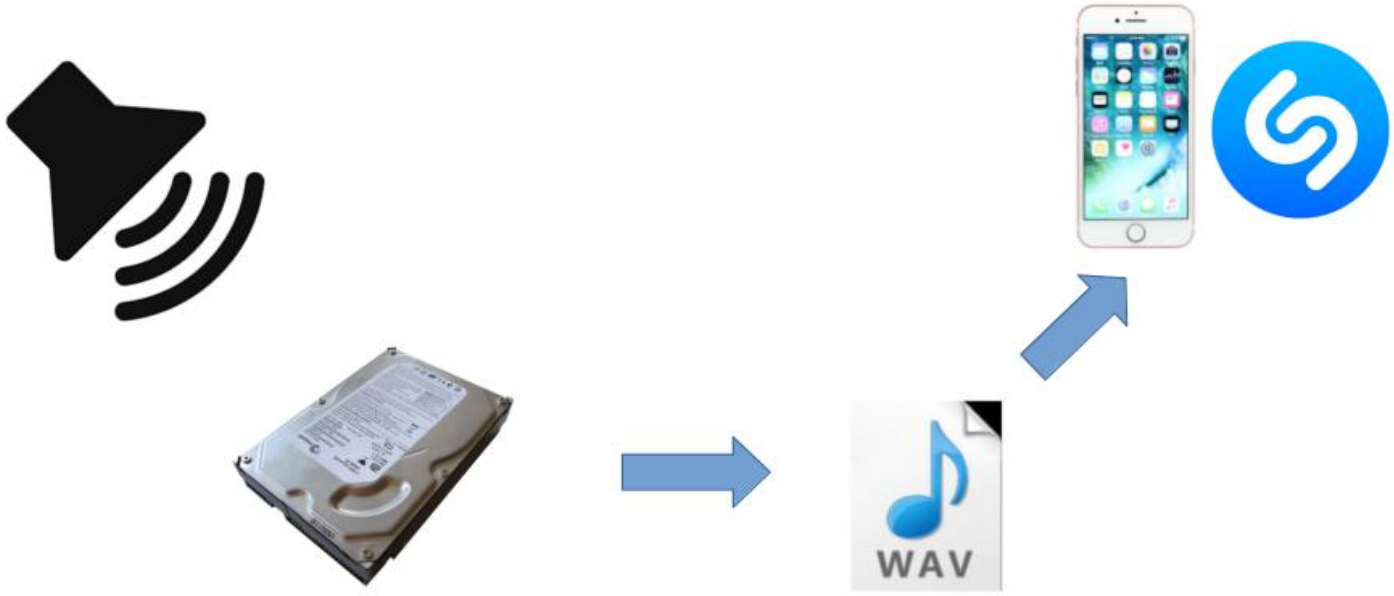

Success, but ...

Required higher volume (90 dBA), filtering didn't work

- Noise-gating discrimination errors ruined spectral fingerprint
- Recovered audio extremely poor
- Still enough information to be recognized

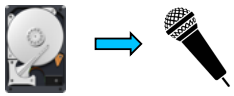

Potential Improvements

Multiple Hard drives

- Make use of signal averaging
- White noise averages to zero, signal averages to itself

Use auto-correlation to find repetitions of same utterance, average them

Mitigations

- Ultrasonic masking can protect deployed systems
- Sign firmware!
  - Zaddach et al. [3] didn't find signatures in use in any HDDs they examined

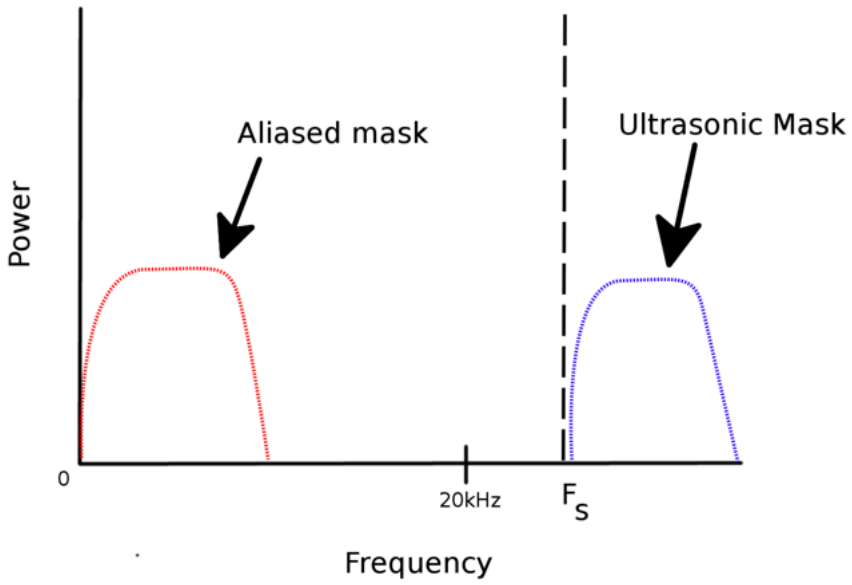

[3] [HDD Malware, ACSAC '03]

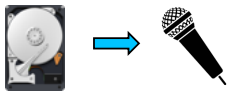

Conclusion

Our research sheds light on overlooked threat of devices that weren't designed as sensors

Defenses for already deployed systems are challenging

Hard drives can approximate crude microphones

Other Applications: other devices, such as printers; mechanical coupling

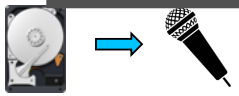

Everything is Microphonic

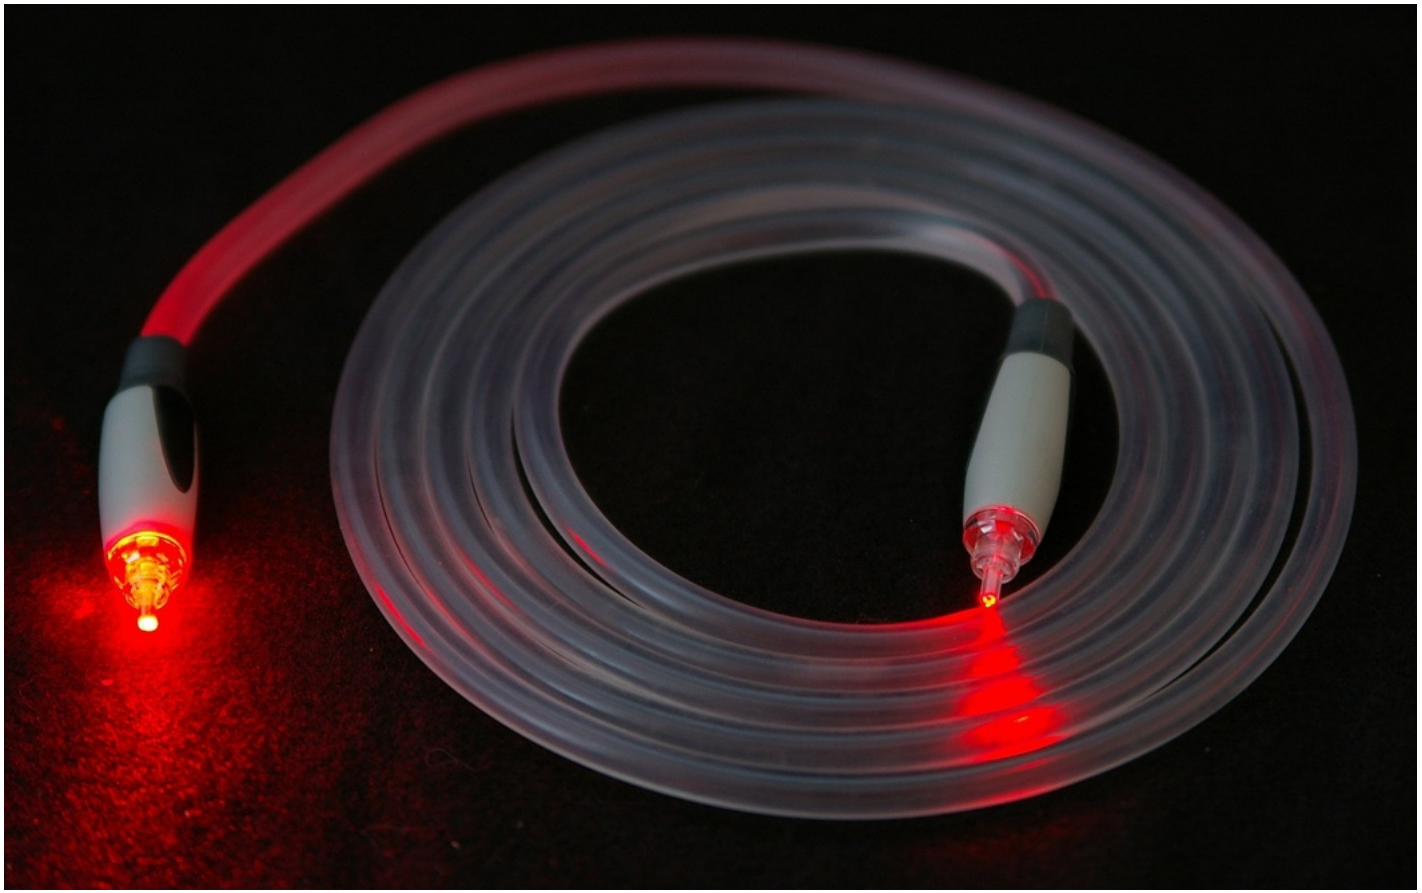

The first fiber-optic acoustic sensors published in 1977 [Cole et al. and Bucaro et al.]

https://en.wikipedia.org/wiki/Fiber-optic_cable

Everything is Microphonic

TWENTY-FIVE YEARS OF INTERFEROMETRIC FIBER OPTIC ACOUSTIC SENSORS AT THE NAVAL RESEARCH LABORATORY

James H. Cole

Clay Kirkendall

Anthony Dandridge

Gary Cogdell

T.G. Giallorenzi

Naval Research Laboratory
Washington, D.C.

Everything is Microphonic

PHOTONIC SENSORS / Vol. 11, No. 1, 2021: 109–122

Recent Progress in Fiber-Optic Hydrophones

Zhou MENG[1*], Wei CHEN[1], Jianfei WANG[1], Xiaoyang HU[1],
Mo CHEN[1], and Yichi ZHANG[1,2]

[1]*College of Meteorology and Oceanology, National University of Defense Technology, Changsha 410073, China*

[2]*Academy of Artillery and Air Defense, Nanjing 210000, China*

[*]Corresponding author: Zhou MENG E-mail: zhoumeng6806@163.com

Abstract: Fiber-optic hydrophone (FOH) is a significant type of acoustic sensor, which can be used in both military and civilian fields such as underwater target detection, oil and natural gas prospecting, and earthquake inspection. The recent progress of FOH is introduced from five aspects, including large-scale FOH array, very-low-frequency detection, fiber-optic vector hydrophone (FOVH), towed linear array, and deep-sea and long-haul transmission. The above five aspects indicate the future development trends in the FOH research field, and they also provide a guideline for the practical applications of FOH as well as its array.

Keywords: Fiber-optic hydrophone; large-scale array; very-low-frequency detection; fiber-optic vector hydrophone; towed linear array; deep sea; long-haul fiber transmission

Everything is Microphonic

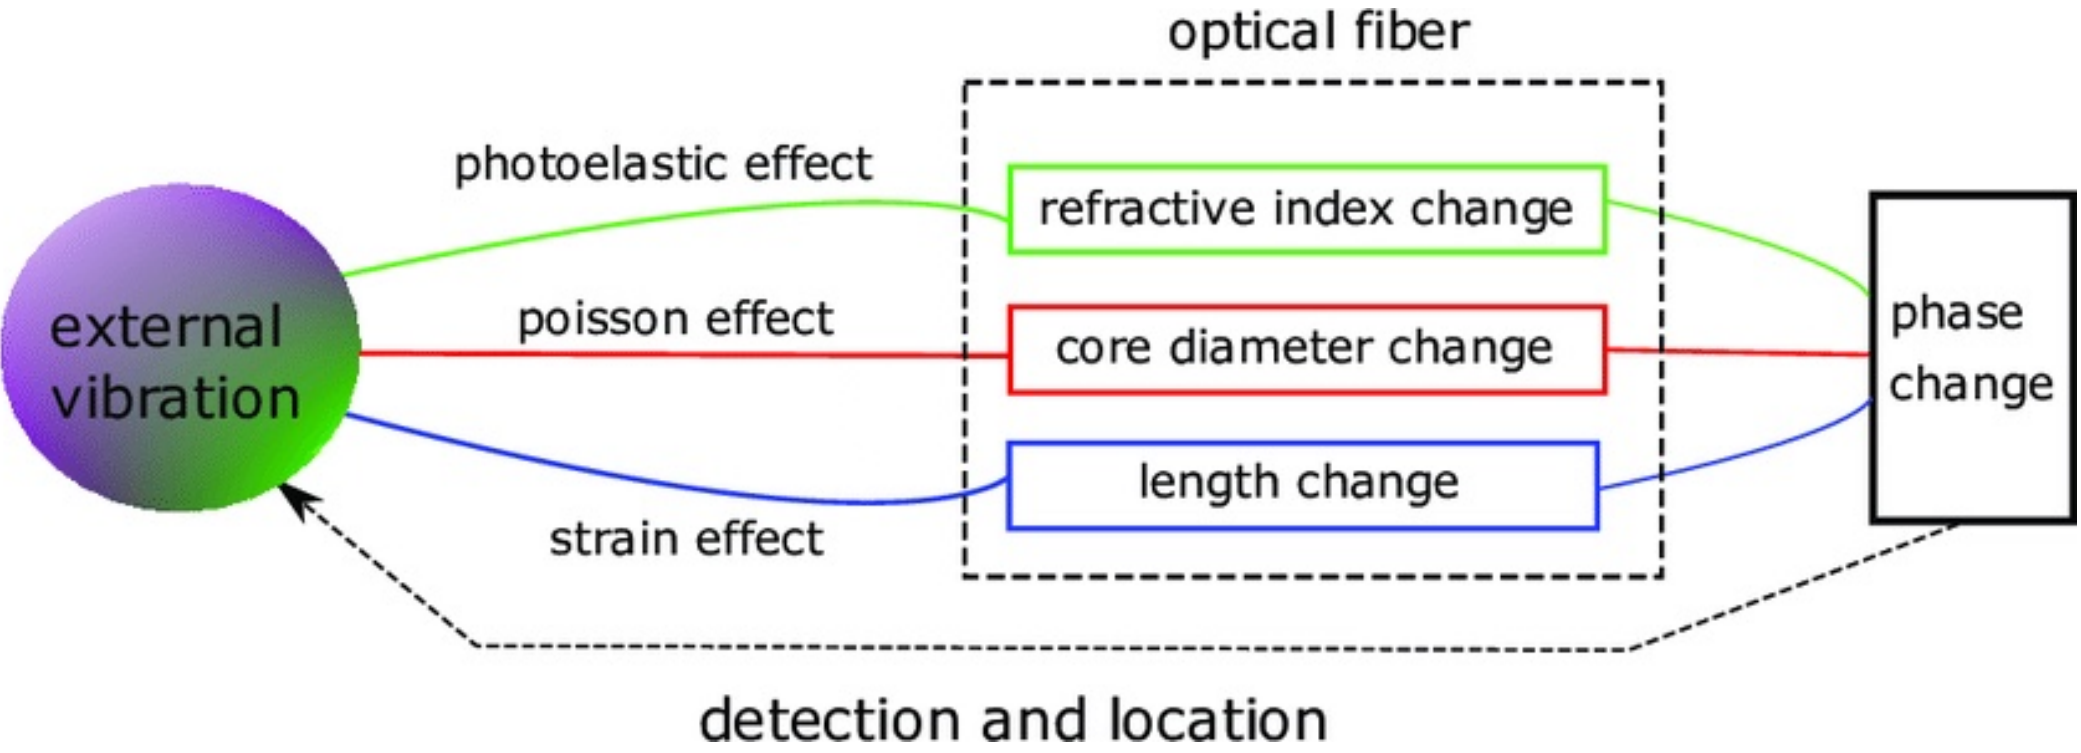

Microphonic Multi-Layer Ceramic Capacitors

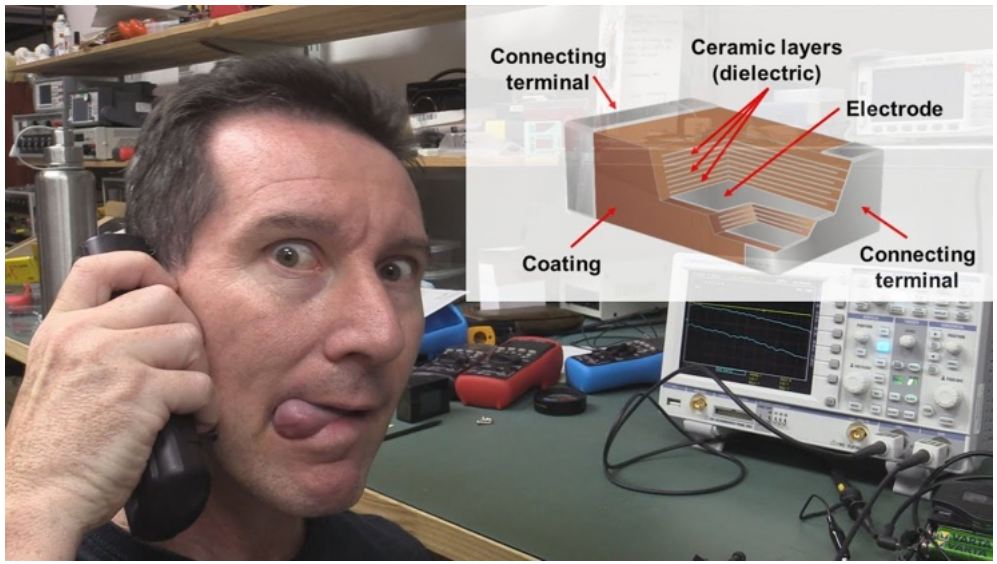

https://www.youtube.com/watch?v=RqEy8QekLDw

https://www.youtube.com/watch?v=F2gX-R1k7MM

Everything is Microphonic: Speake(a)r

- "Espionage malware that can covertly turn the head-phones, earphones, or simple earbuds connected to a PC in
- H

> Interestingly, the audio chipsets in modern motherboards and sound cards include an option to change the function of an audio port at software level, a type of audio port programming sometimes referred to as jack retasking or jack remapping. This option is available on Realtek's (Realtek Semiconductor Corp.) audio chipsets, which are integrated into a wide range of PC motherboards to-day. Jack retasking, although documented in applicable

https://www.usenix.org/system/files/conference/woot17/woot17-paper-guri.pdf

Everything is Microphonic: Lamphone

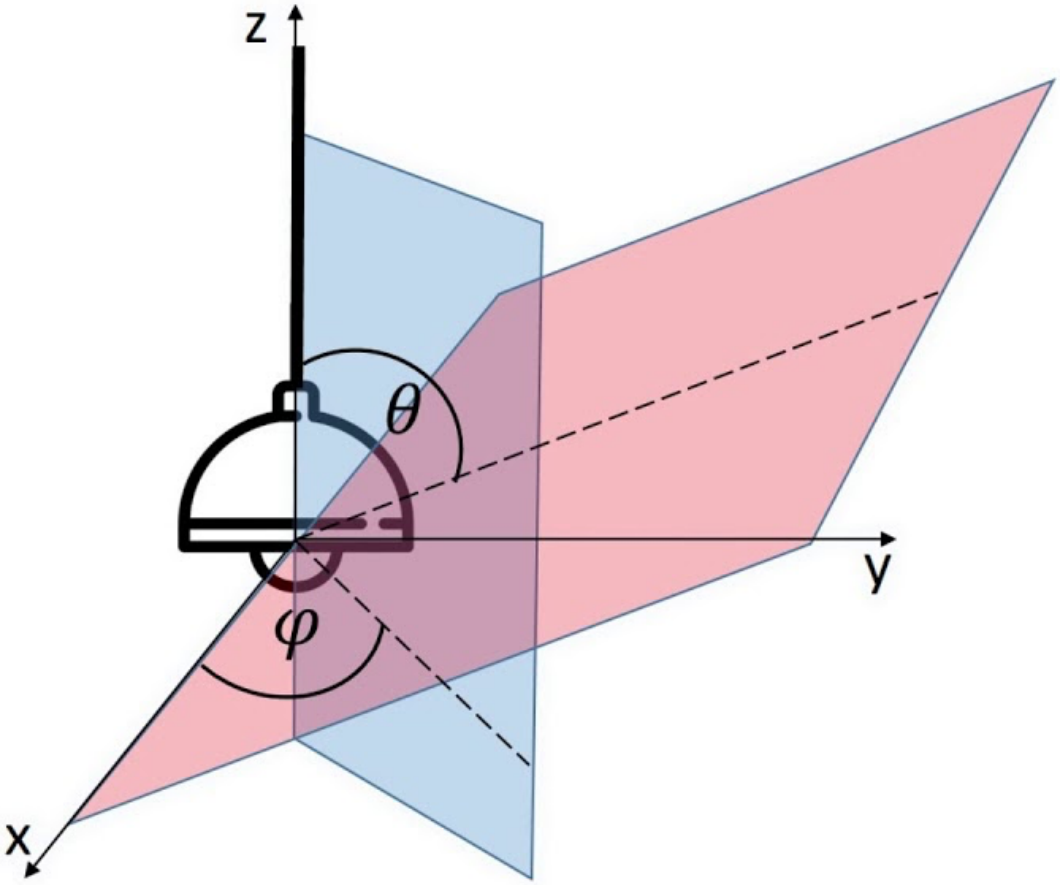

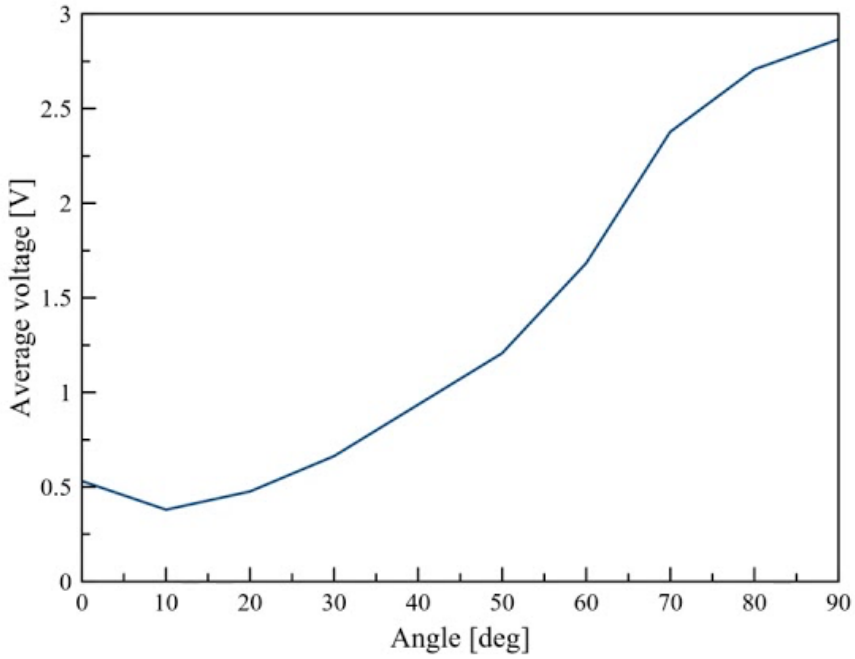

https://www.nassiben.com/lamphone

trigger warning

Homework and Next

- Homework
  - ✓ Essay #1: Done.
  - ✓ Prelab #2: Done.
  - ➡ Lab #2: Due Mon, October 6
- Next
  - Thursday: Lab #2 time in class
  - **Read for Monday:** BlueNote, Proc. IEEE Security & Privacy, 2018. https://spqrlab1.github.io/papers/bolton-blue-note.pdf