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?

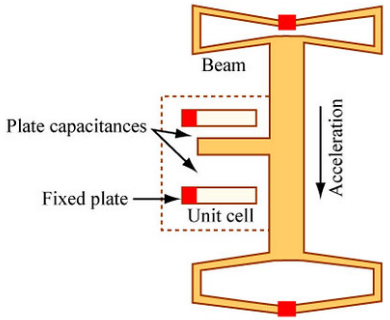

Accelerometers: [1] Marquardt et al., CCS '11, “(sp)iPhone..."
Gyroscopes: [2] Michalevsky et al., Usenix Security '14, “Gyrophone...”

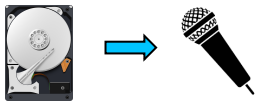

Hard Drive as a Microphone?

Challenges:

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

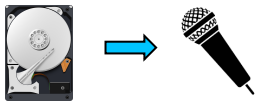

Contributions

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

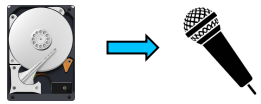

Threat Model

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

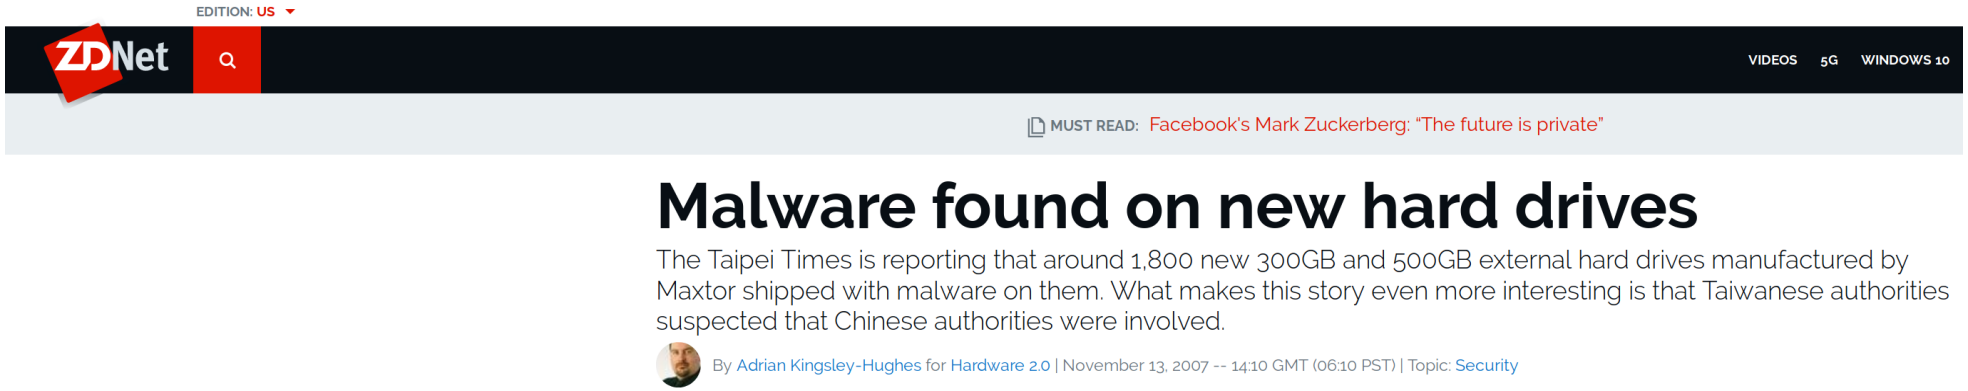

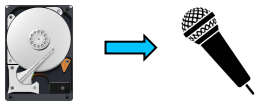

2018

EDITION: US

ZDNet | VIDEOS | 5G | WINDOWS 10 | CLOUD | AI | INNOVATION | SECURITY | MORE | NEWSLETTERS | ALL WRITERS

MUST READ: IT strategy: How an investment in diversity can boost your business

Apple's T2 security chip disconnects a MacBook's microphone when users close the lid

Feature only available for MacBook Pro and MacBook Air models released in 2018.

By Catalin Cimpanu for Zero Day | October 30, 2018 -- 20:00 GMT (13:00 PDT) | Topic: Security

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

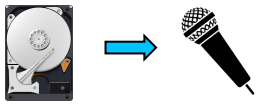

HDD as a microphone

- 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

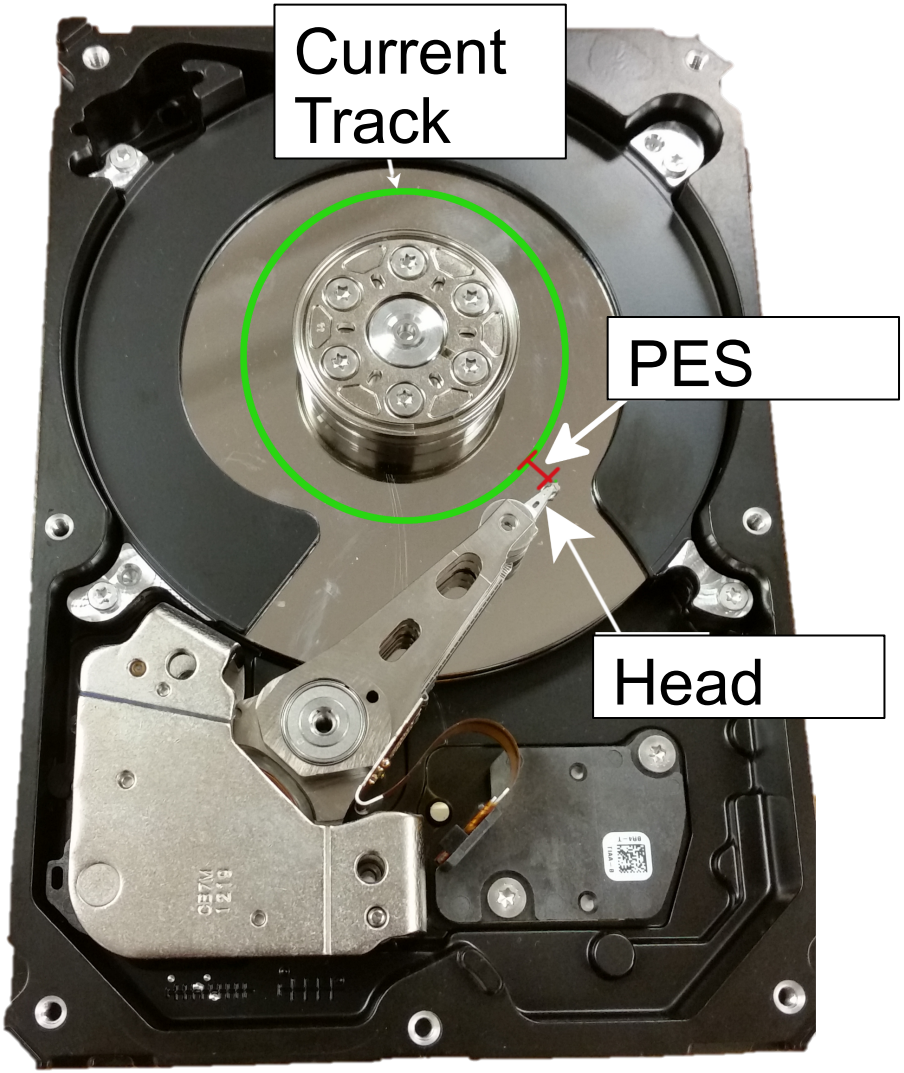

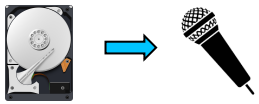

Head Tracking

- 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

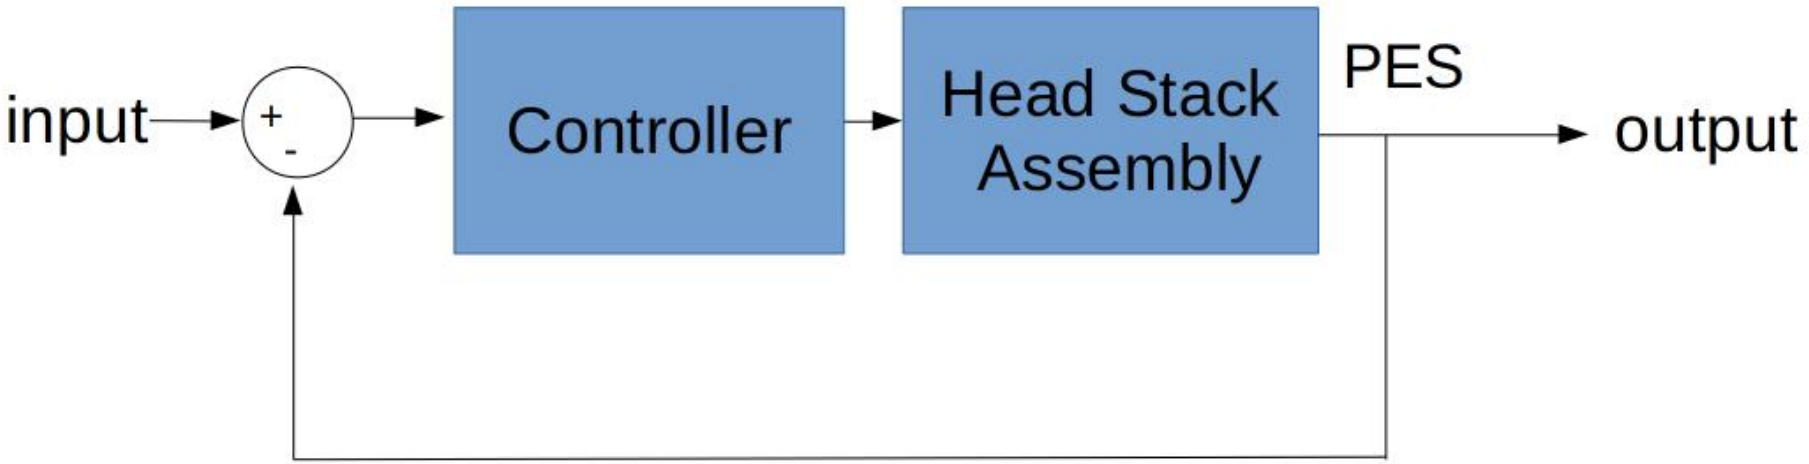

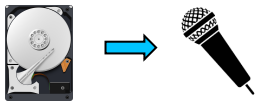

Similarities to Microphone

Microphone:

- Output measures diaphragm displacement
- Sound waves displace diaphragm

HDD:

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

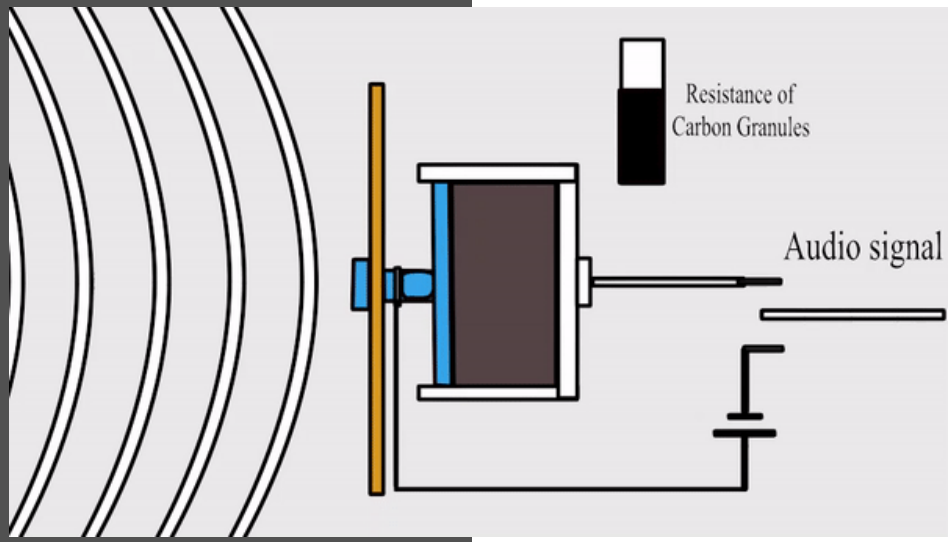

https://www.instructables.com/id/Simplified-Electronics-Microphone-DIY-How-It-Works/

PES approximates microphone output??

Andrew Kwong (https://andrewkwong.org)

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

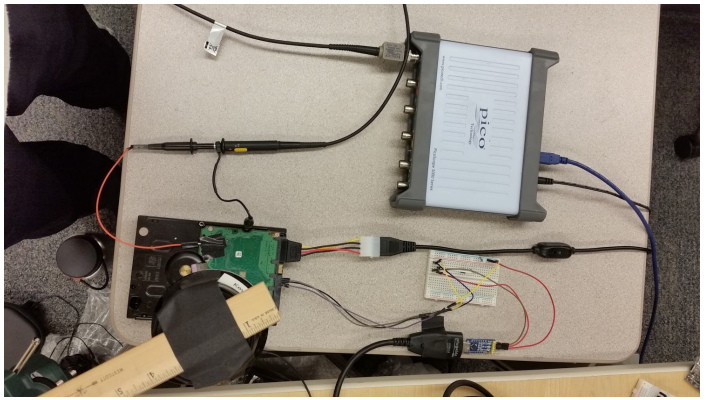

HDD Malware: [3] Zaddach et al., ACSAC '13

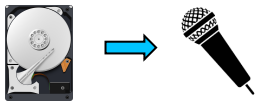

Sampling Rate

$$\text{frequency}_{\text{sampling}} = \text{frequency}_{\text{rotation}} * \text{num\_servo\_sectors\_per\_track}$$

$$= 120 \text{ Hz} * 288$$

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

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

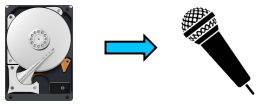

demo

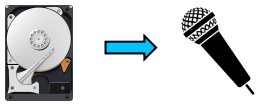

Experimental Setup

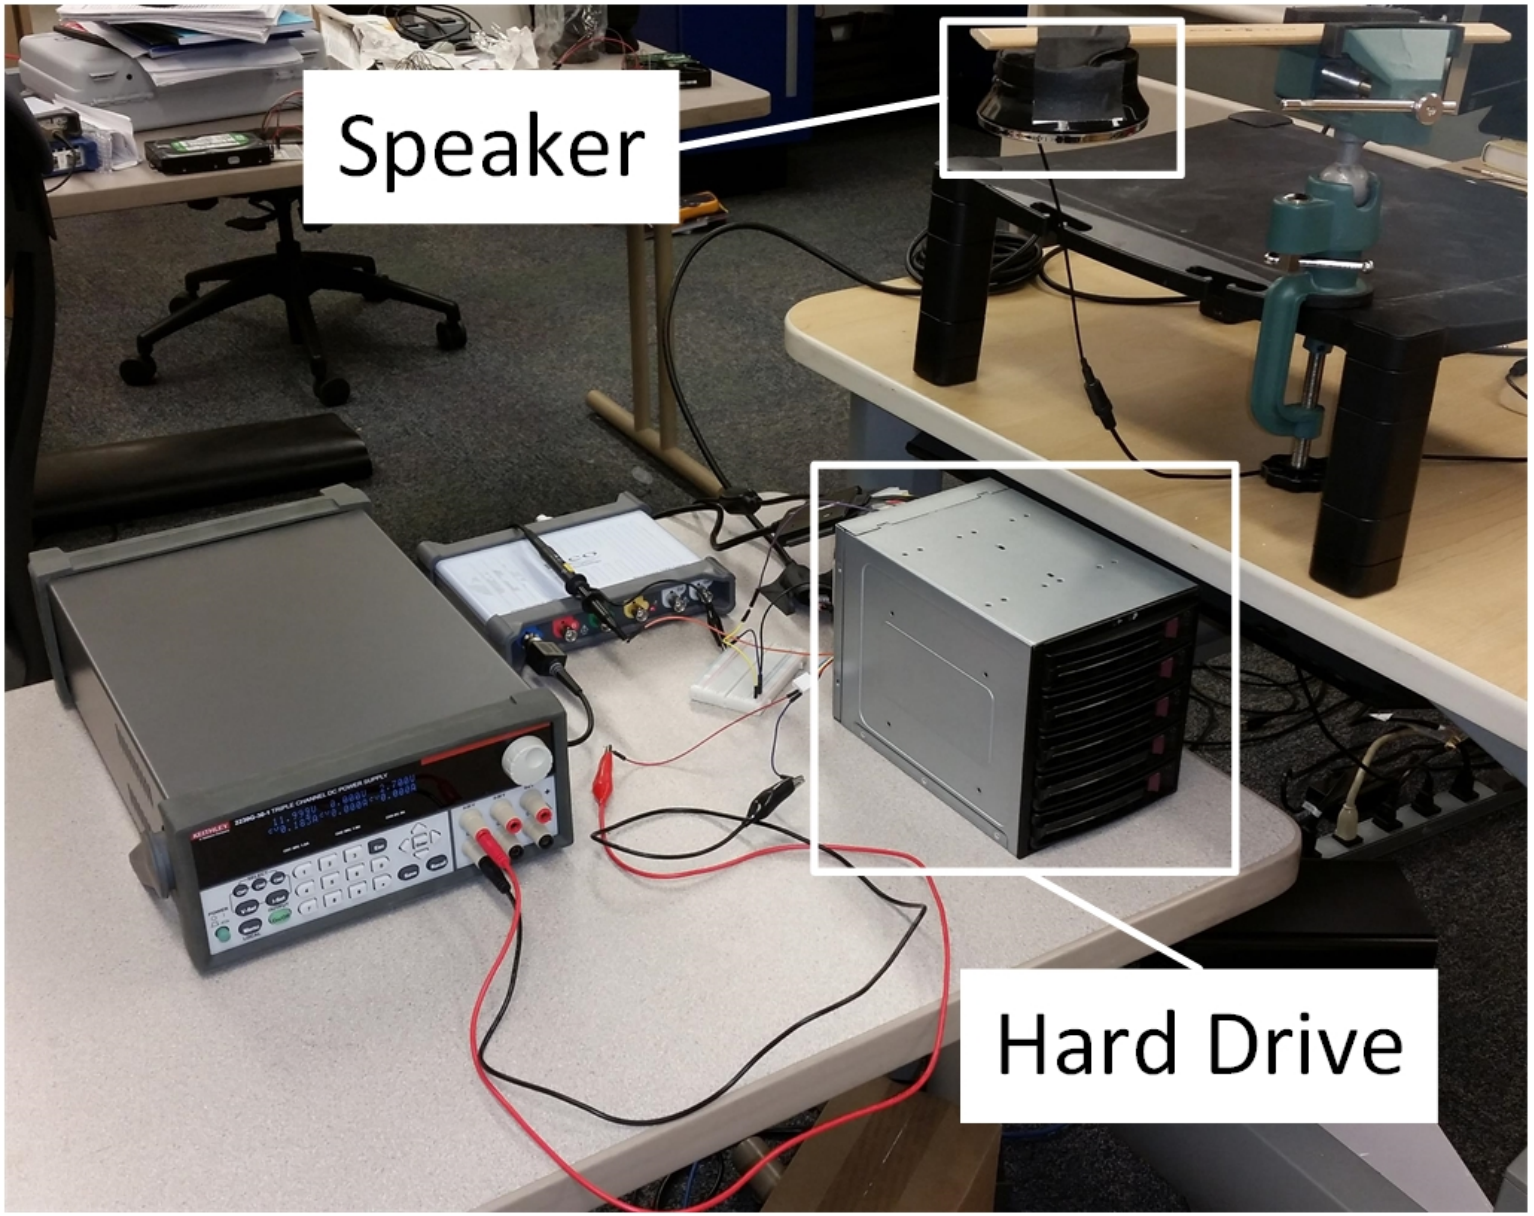

Speech Recovery

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

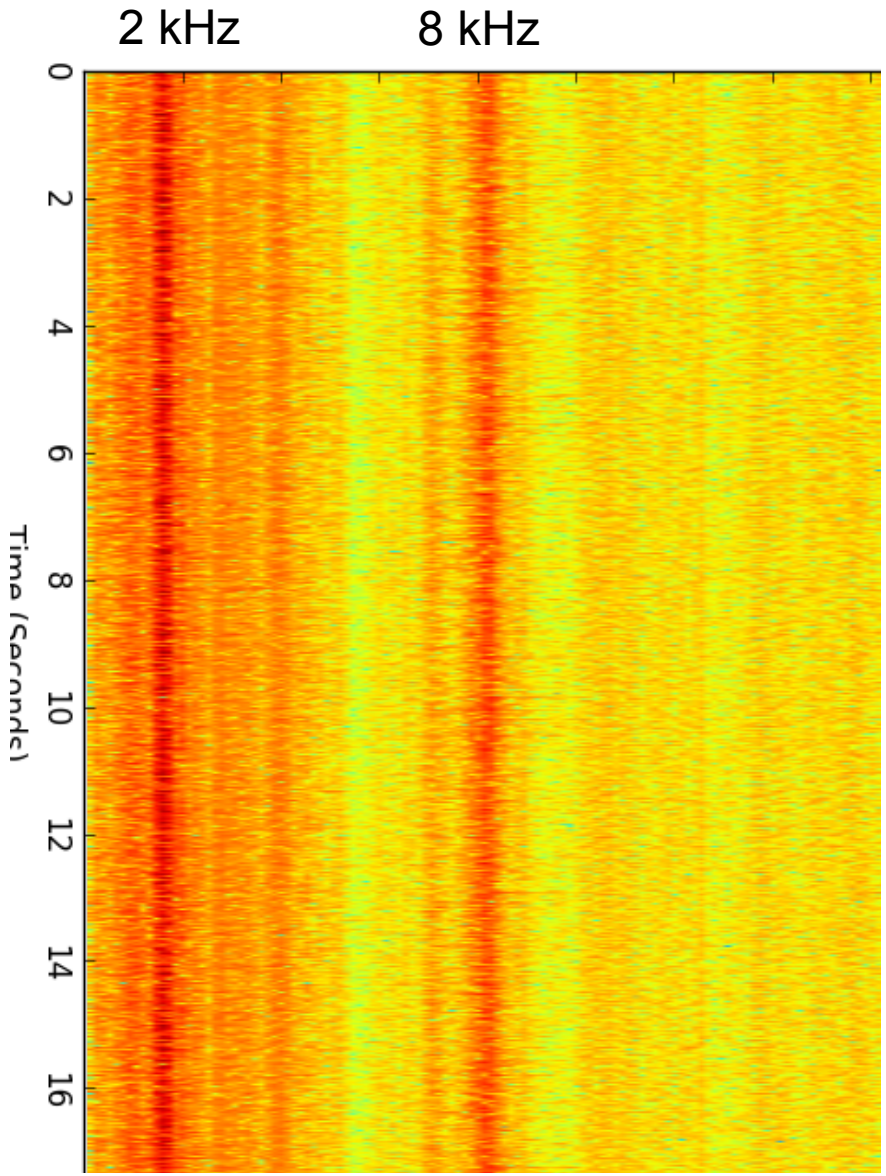

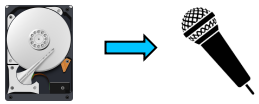

Signal Analysis

From HDD:

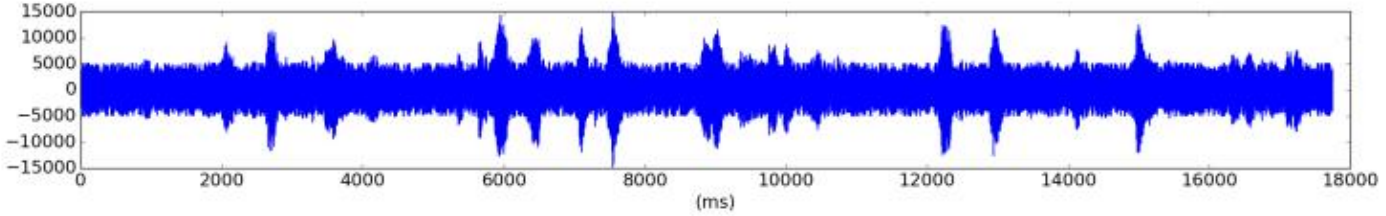

Original:

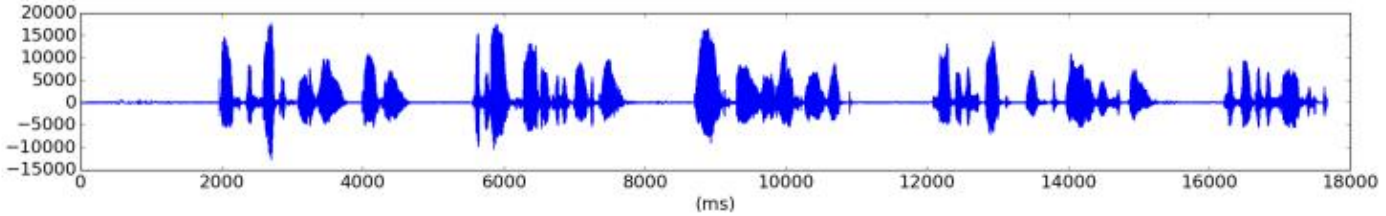

Cross Correlation:

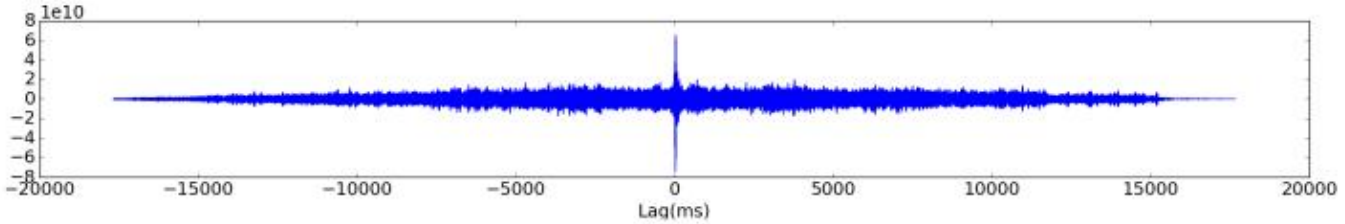

- Harvard Sentence male speaker with drive enclosed in case and fan powered at max (42W)

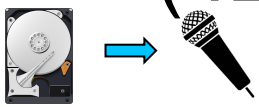

Quantitative Measures

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

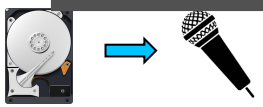

Speech Sample

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.

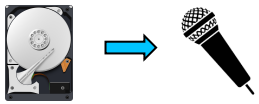

Shazam Recognition

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

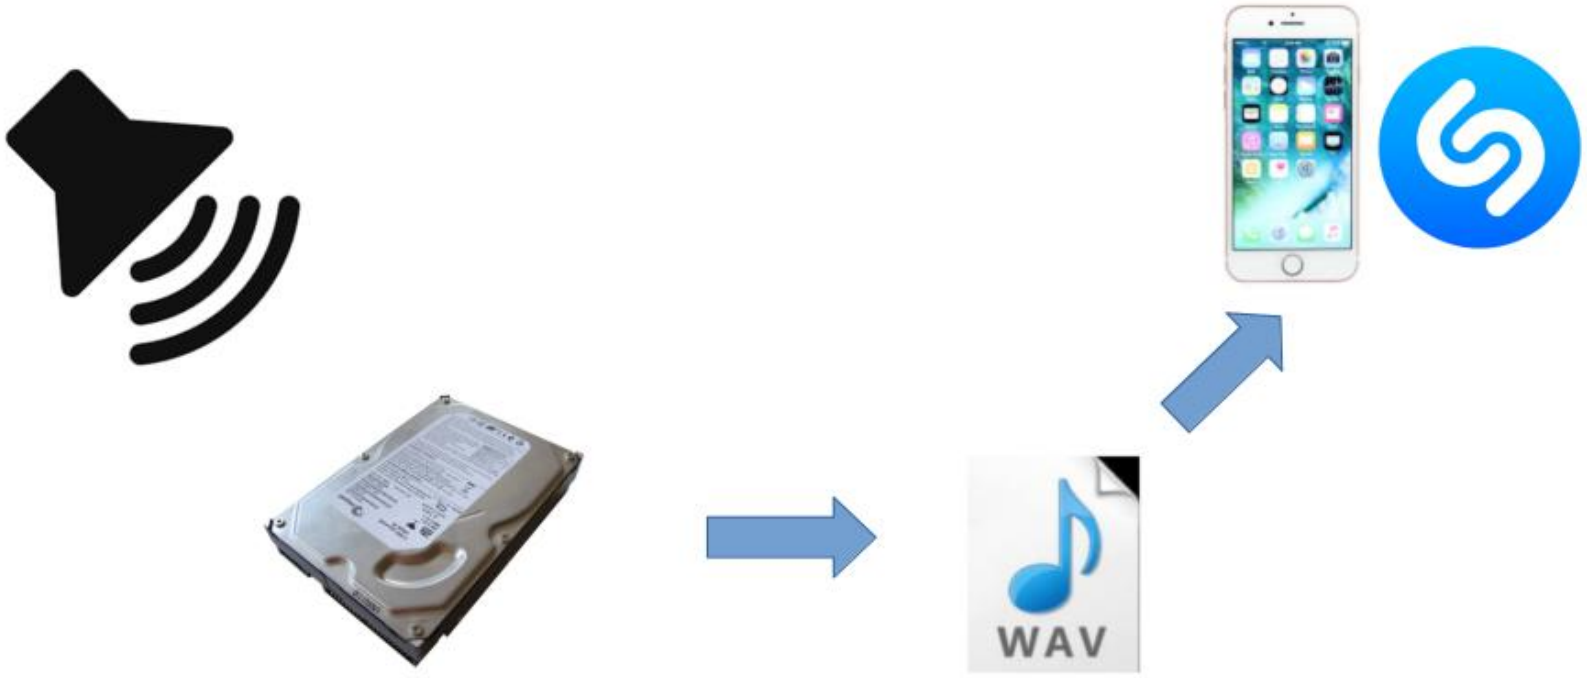

Success, but ...

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

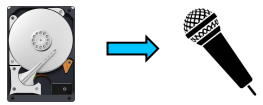

Potential Improvements

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

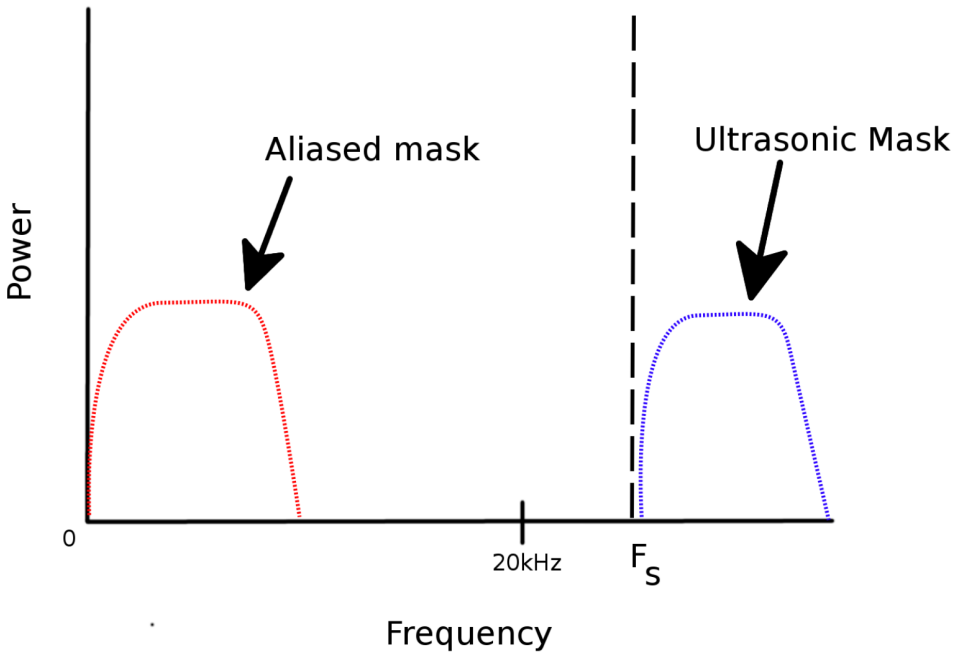

[3] [HDD Malware, ACSAC '03]

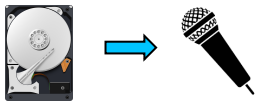

Conclusion

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

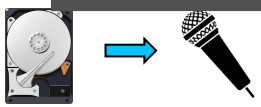

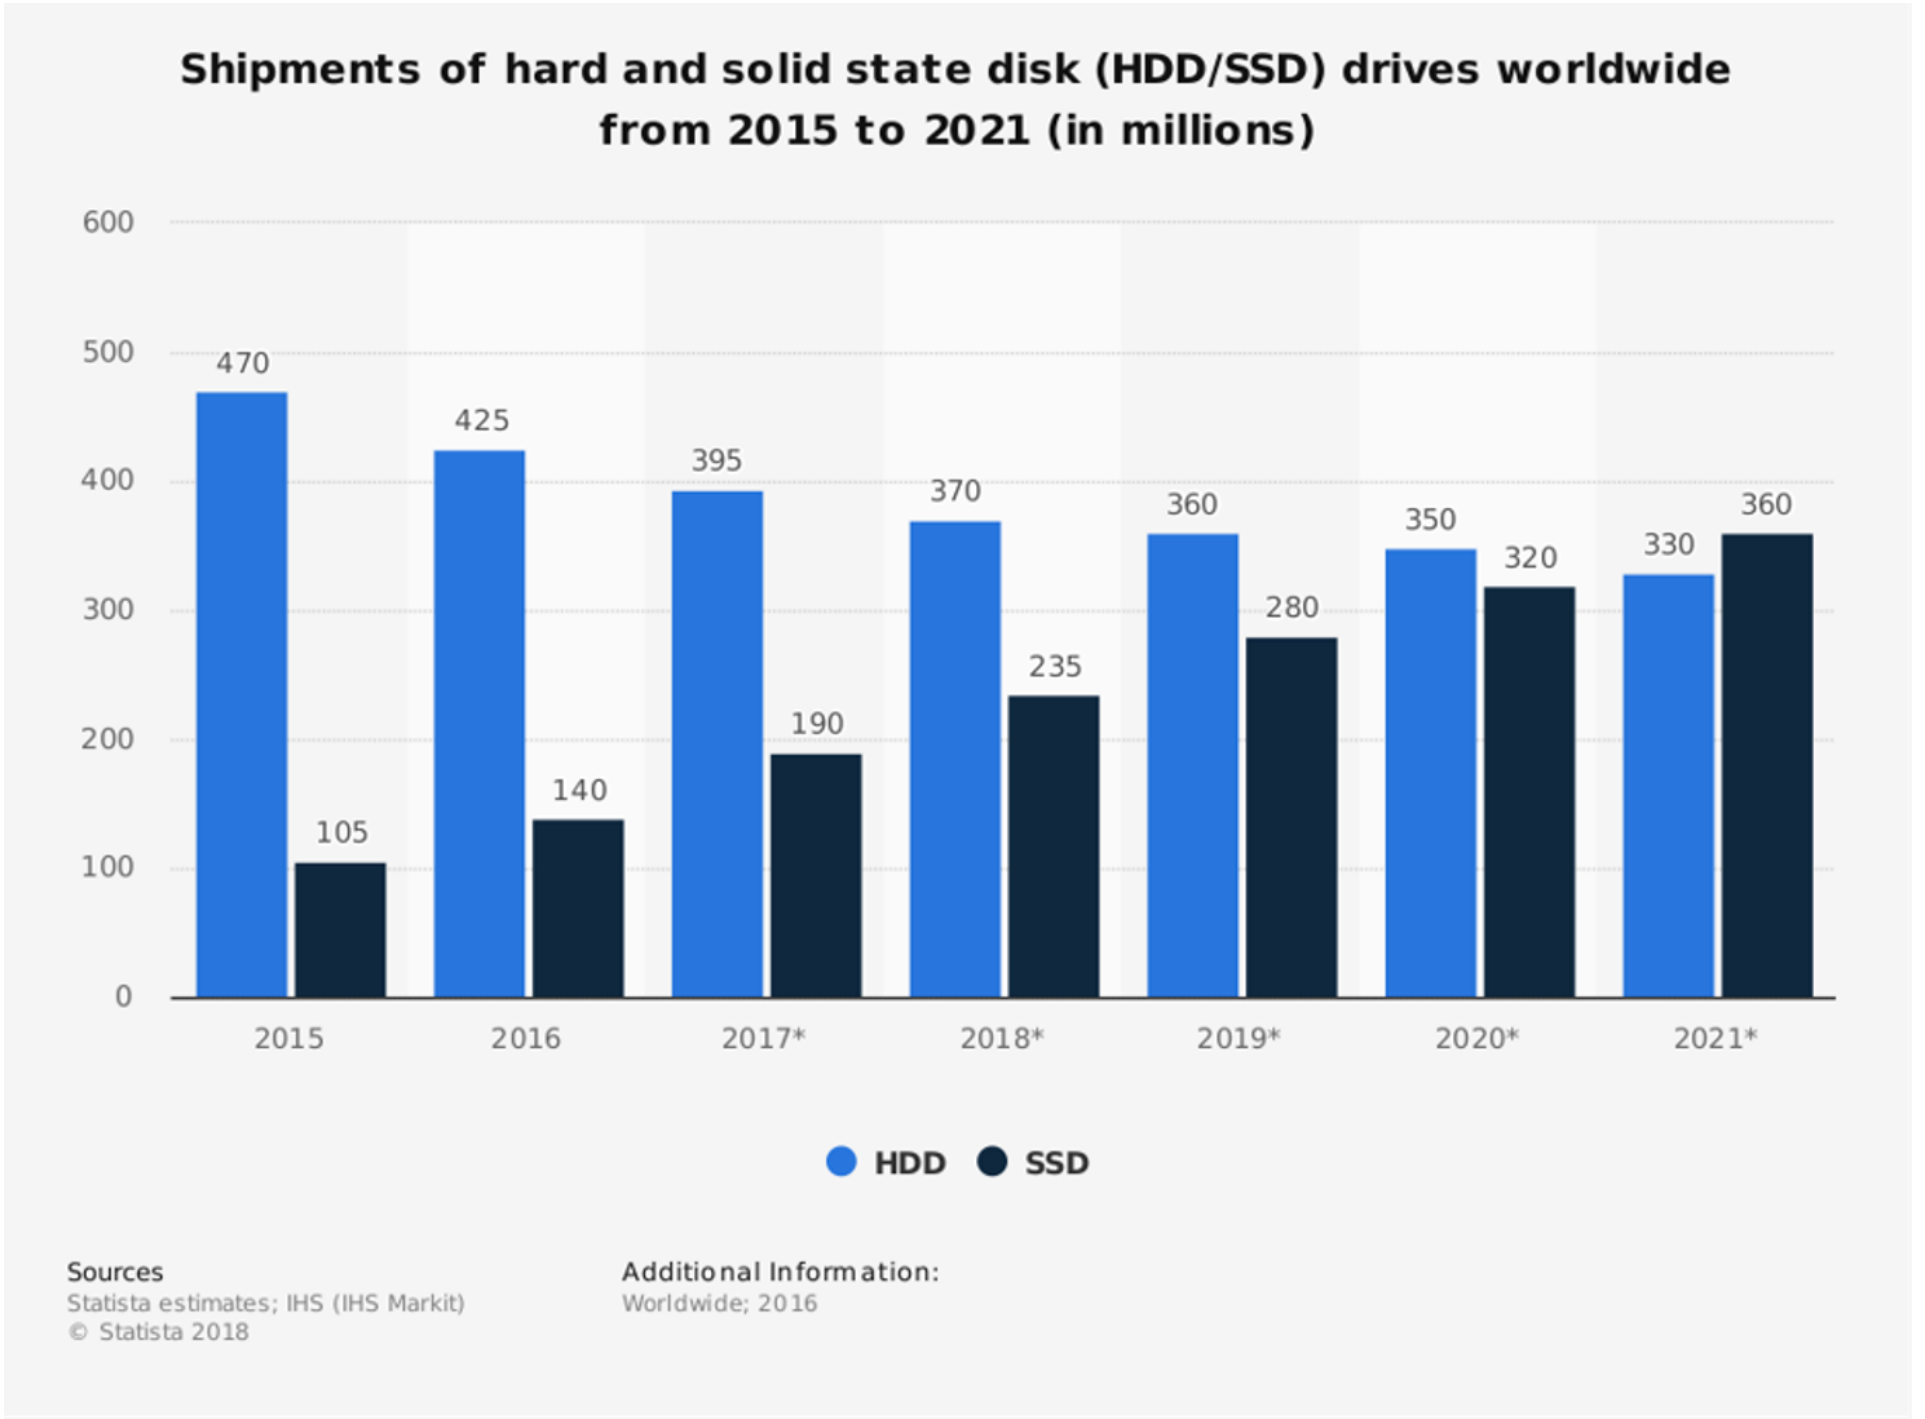
Shipments of hard and solid state disk (HDD/SSD) drives worldwide from 2015 to 2021 (in millions)
600
500
400
300
200
100
0
470
105
425
140
395
190
370
235
360
280
350
320
330
360
2015
2016
2017*
2018*
2019*
2020*
2021*
HDD
SSD
Sources
Statista estimates; IHS (IHS Markit)
© Statista 2018
Additional Information:
Worldwide; 2016

Granularity

- PES is a 16-bit value
- Granularity: 1/(2^12) of a track
- Only get 8 bits from AMUX pin
  - Chose bits 3-10

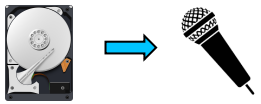

Accessibility to MCU

- Proof-of-Concept attack demonstrates what an attacker with firmware-resident malware can do
- First confirmed MCU's access to PES

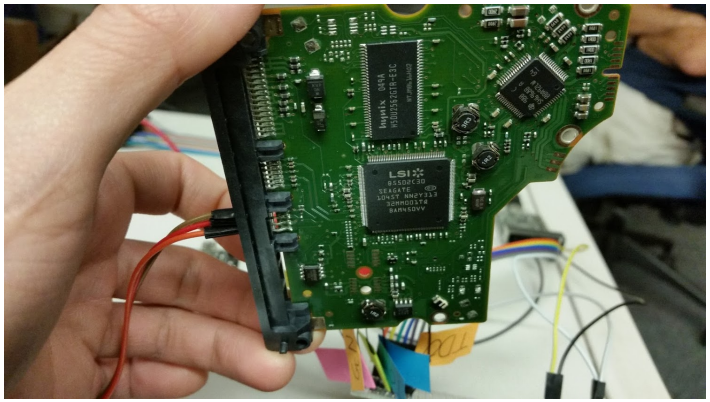

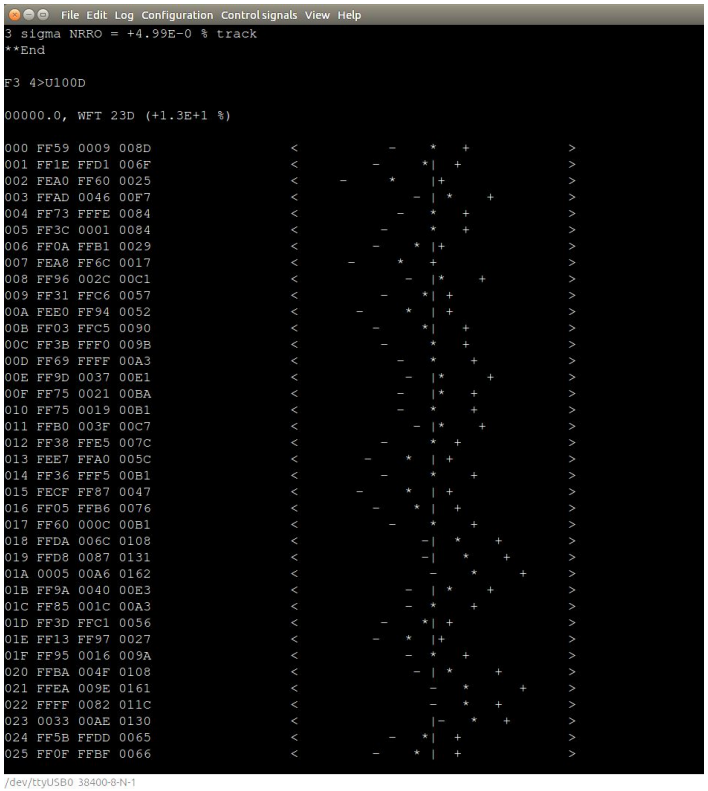

Frequency Response

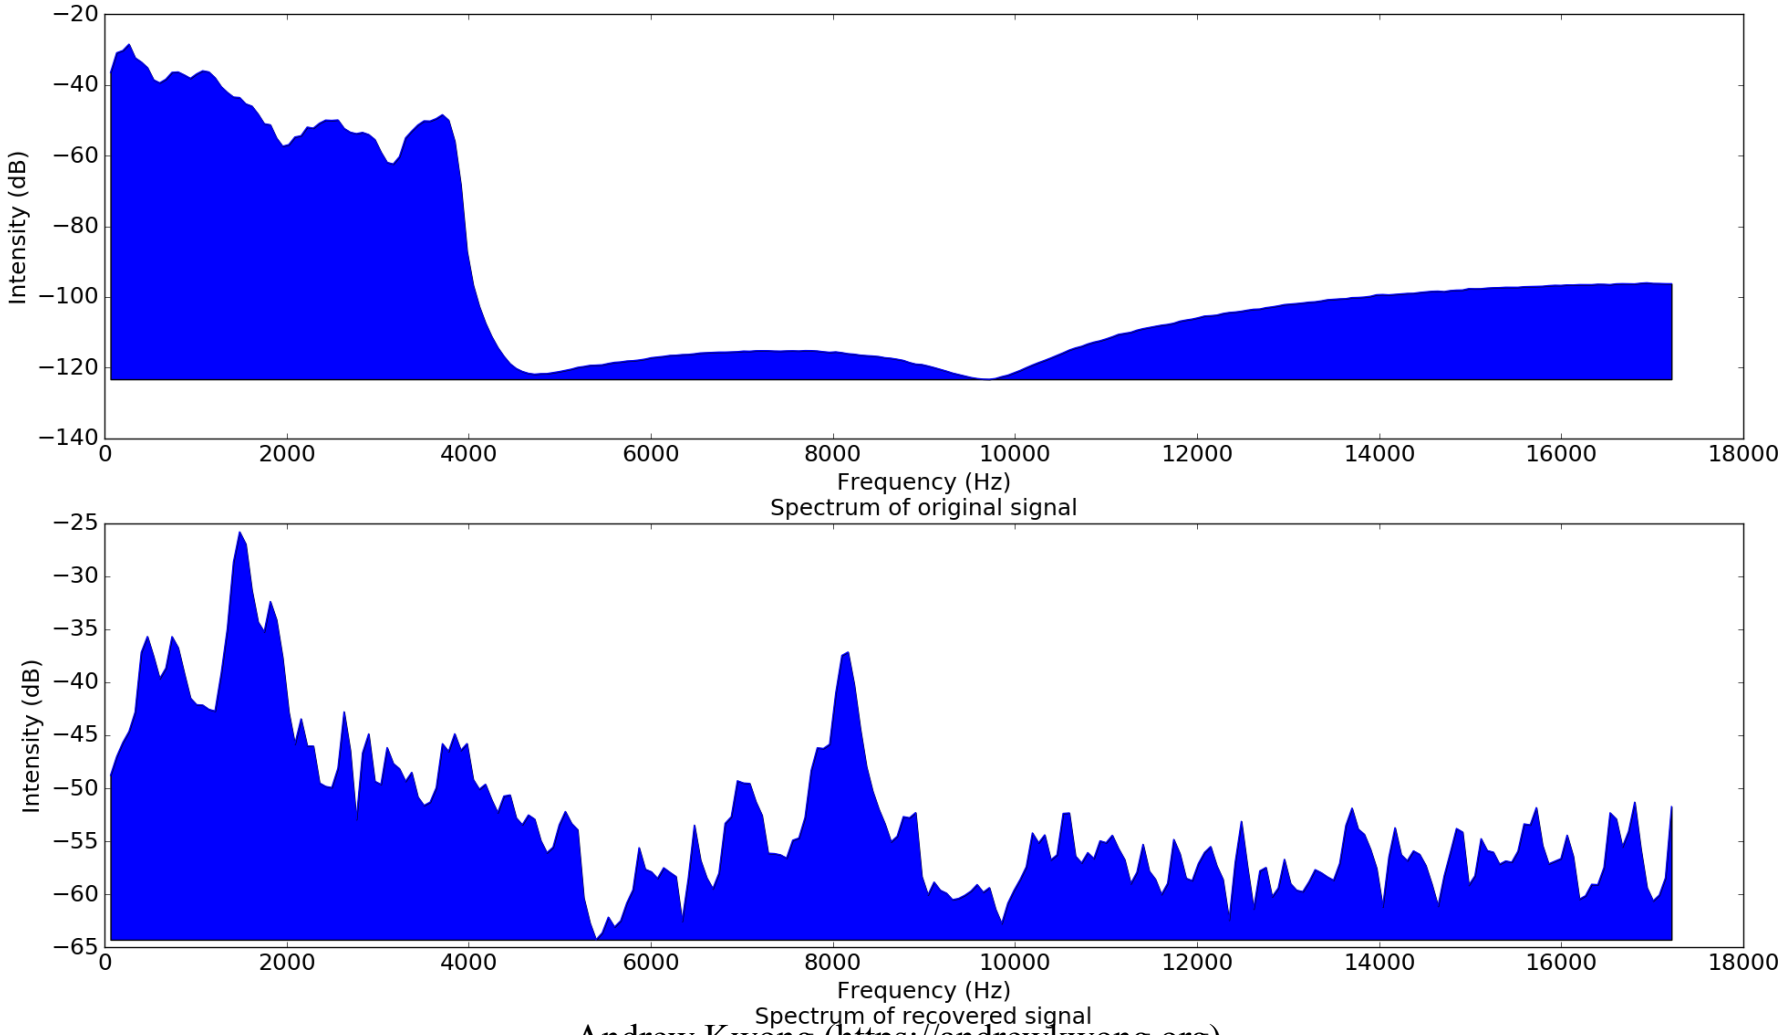

Spectral Analysis

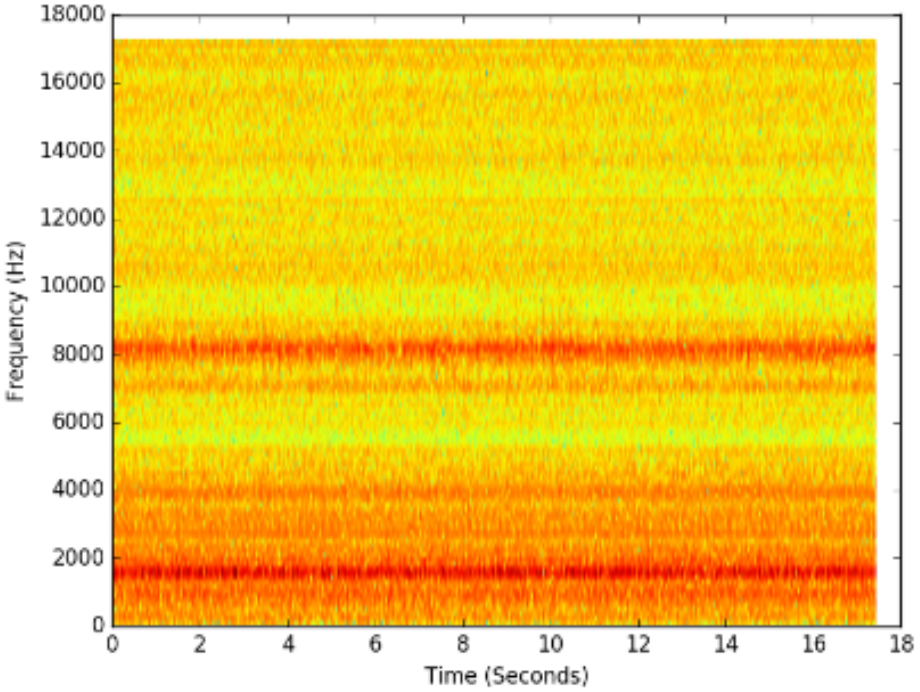

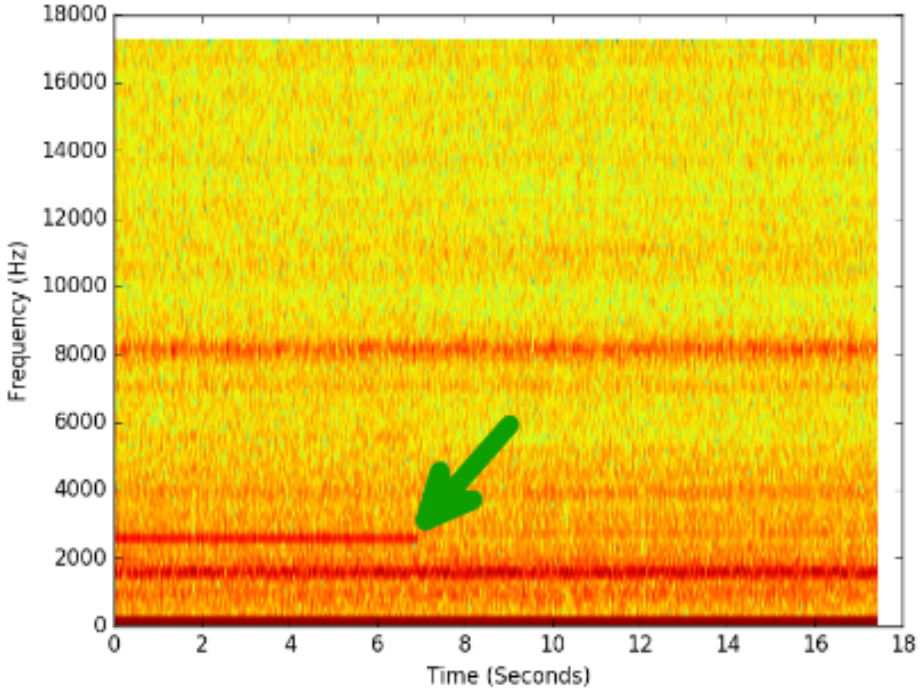

- Heavy bands of persistent noise around 8 kHz and 1900 kHz
- Responds well to 2.5 kHz tone

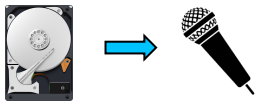

Reading PES

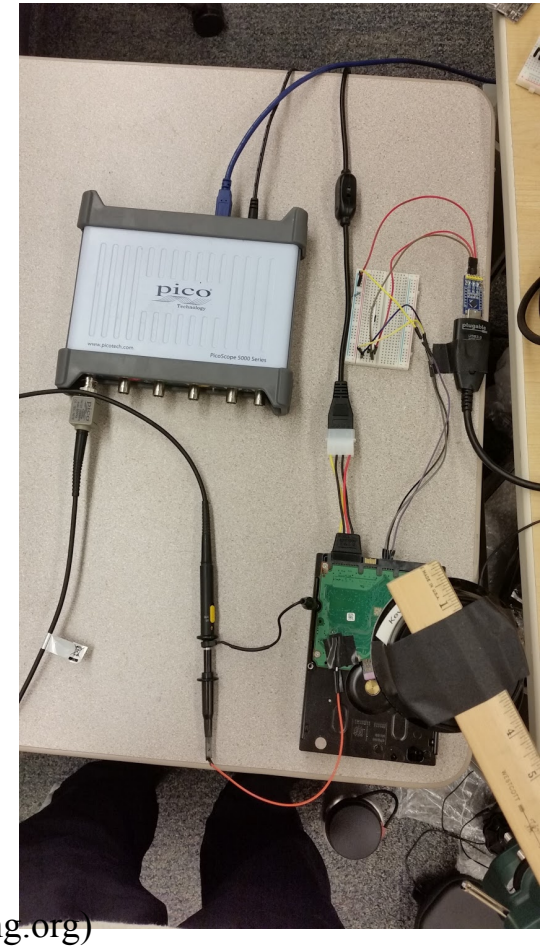

Digital Signal Processing

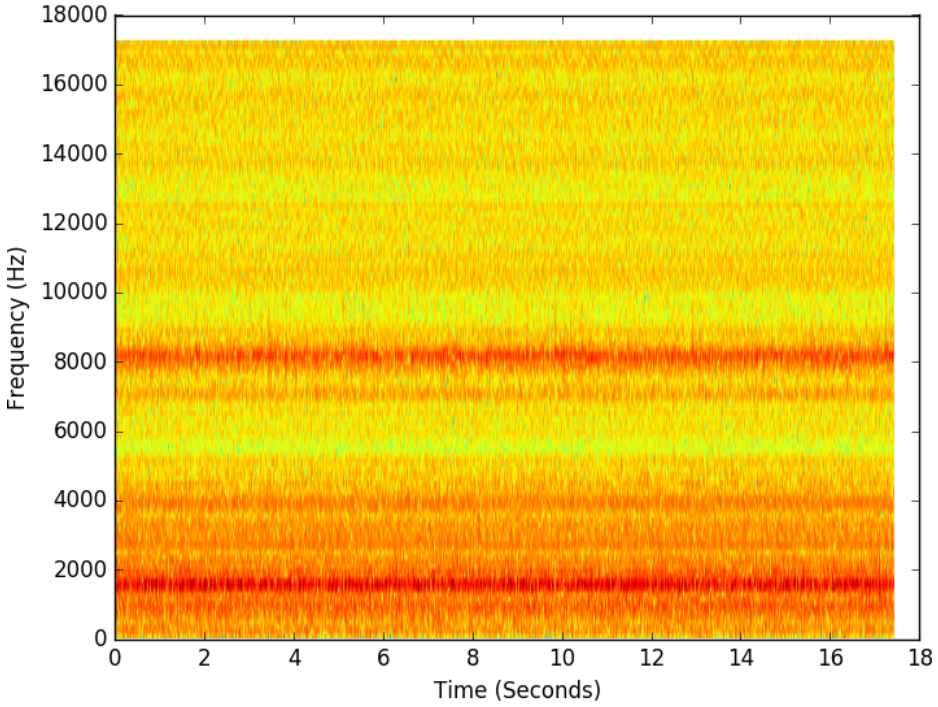

- Linearly filtering out 8 kHz and 1.9 kHz removes the heaviest bands of noise
- Made use of spectral noise gating for further filtering
  - Find noise thresholds at smaller sub-bands, only pass frequencies above the threshold