Who Let the Smart Toaster Hack the House?

Exploring Security and Privacy Risks in Connected Devices

Anna Maria Mandalari

annamandalari.com

UCL

What if we are exchanging privacy for gimmicks and minor convenience?

What is IoT exposing when it comes to privacy in a Smart Home?

What might this mean for the future?

Why were we interested in this?

They may listen to you
(e.g., smart speakers)

- They can (by definition) access the Internet and therefore may expose private information

- Lack of understanding on what information they expose, on when they expose it, and to whom

- Lack of understanding of regional differences (e.g., GDPR)

ey may know what
vatch (e.g., smart TVs)

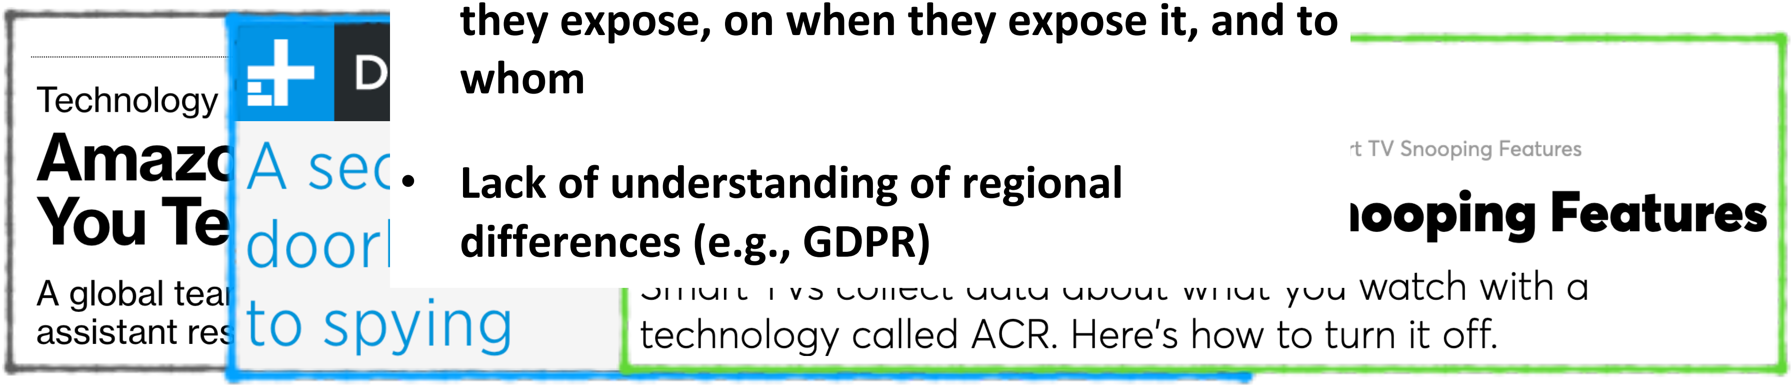

Course Overview

- ❑ Benchmarking privacy in IoT devices
- ❑ IoT devices identification
- ❑ Benchmarking security in IoT devices
- ❑ Benchmarking security solutions for IoT devices
- ❑ Privacy solutions for IoT devices at the edge
- ❑ Security solutions for IoT devices at the edge
- ❑ IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security

What is IoT exposing when it comes to privacy in a Smart Home?

Course Overview

- ❑ Benchmarking privacy in IoT devices
- ❑ IoT devices identification
- ❑ Benchmarking security in IoT devices
- ❑ Benchmarking security solutions for IoT devices
- ❑ Privacy solutions for IoT devices at the edge
- ❑ Security solutions for IoT devices at the edge
- ❑ IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security

Goal of Research

What is the destination of IoT network traffic?

What information is sent?

Does a device expose information unexpectedly?

Google, Amazon, and Apple have decided to collaborate on a universal smart home ecosystem.

210 devices in two different countries

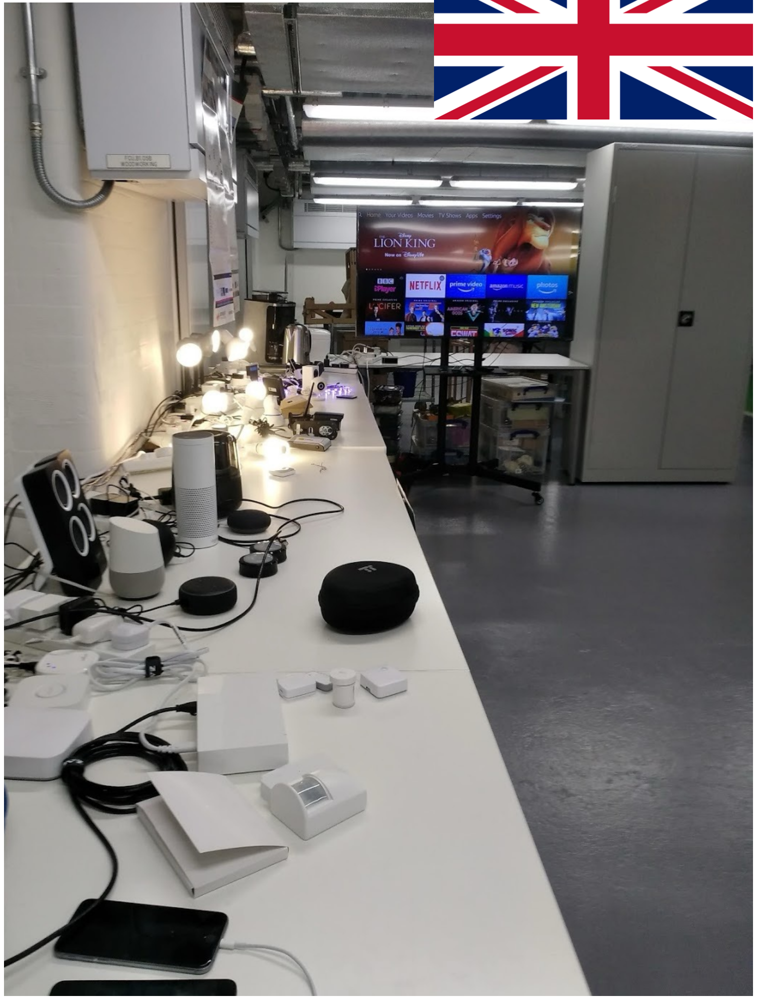

Design of Experiments

>200k Experiments

- **Controlled interactions**
  - Automated (repeated 30 times)
    - Text-to-speech to smart assistants (Alexa/Google/Cortana/Bixby)
    - Monkey instrumented control from Android companion apps
- **Idle: background traffic**

Activity	Description
Power	power on/off the device
Voice	voice commands for speakers
Video	record/watch video
On/Off	turn on/off bulbs/plugs
Motion	move in front of device
Others	change volume, browse menu

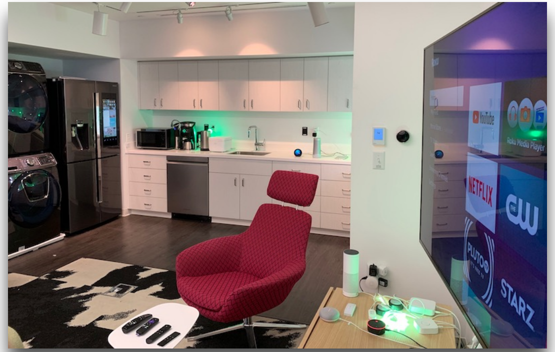

Data Collection Methodology

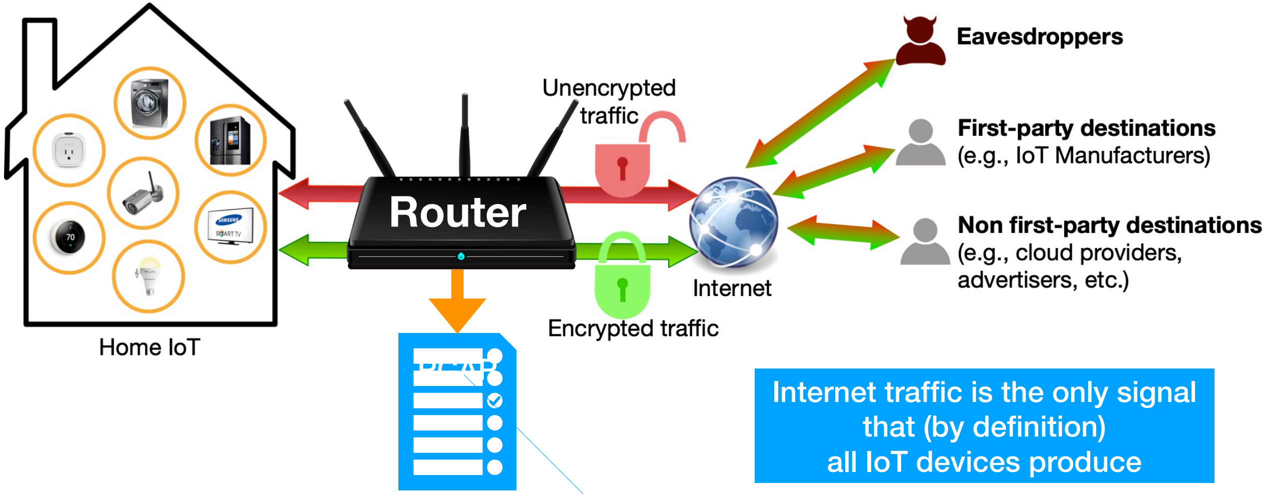

Internet traffic is the only signal that (by definition) all IoT devices produce

- Monitor all traffic at the **router**
    - per-device
    - per-experiment

What Is the Destination?

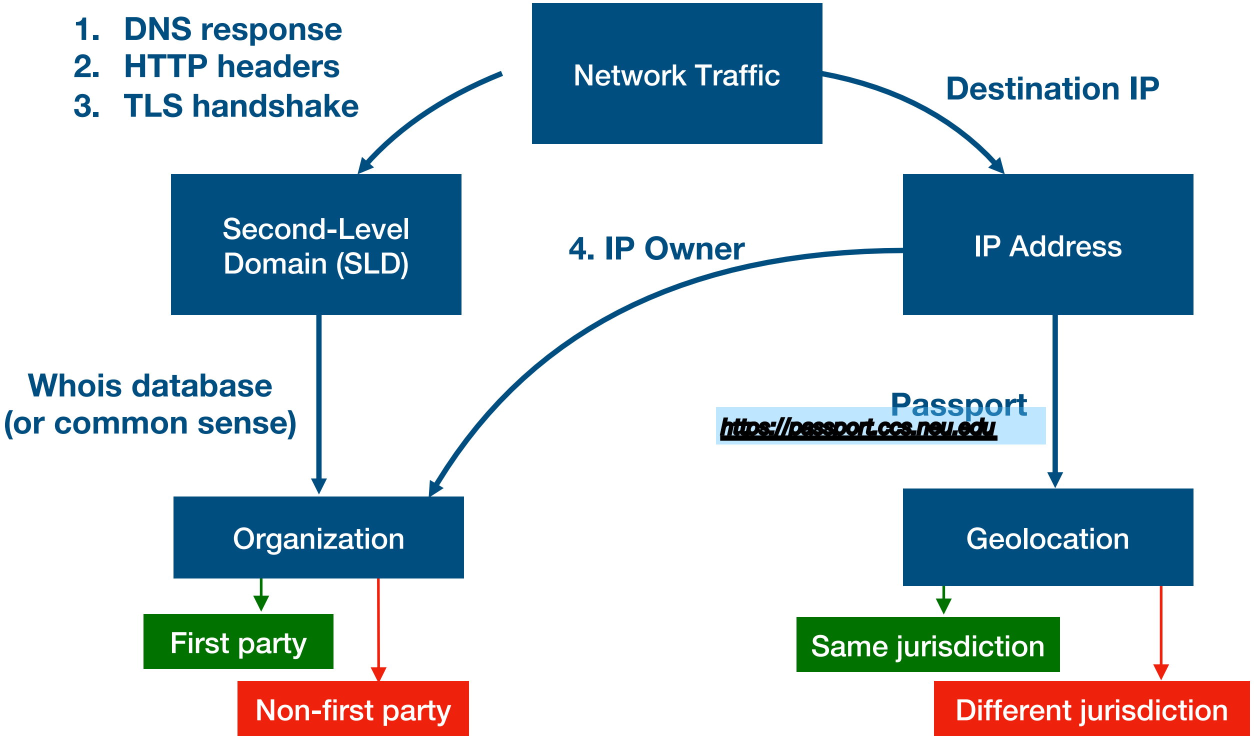

What Non-First Parties Are Contacted?

- Number of devices contacting non-first party organizations

High reliance on cloud and CDN providers

Nearly all TVs contact Netflix w/o being logged in

Chinese cloud pro

Organization	US 46	UK 35

Regional differences

Most traffic goes beyond Europe

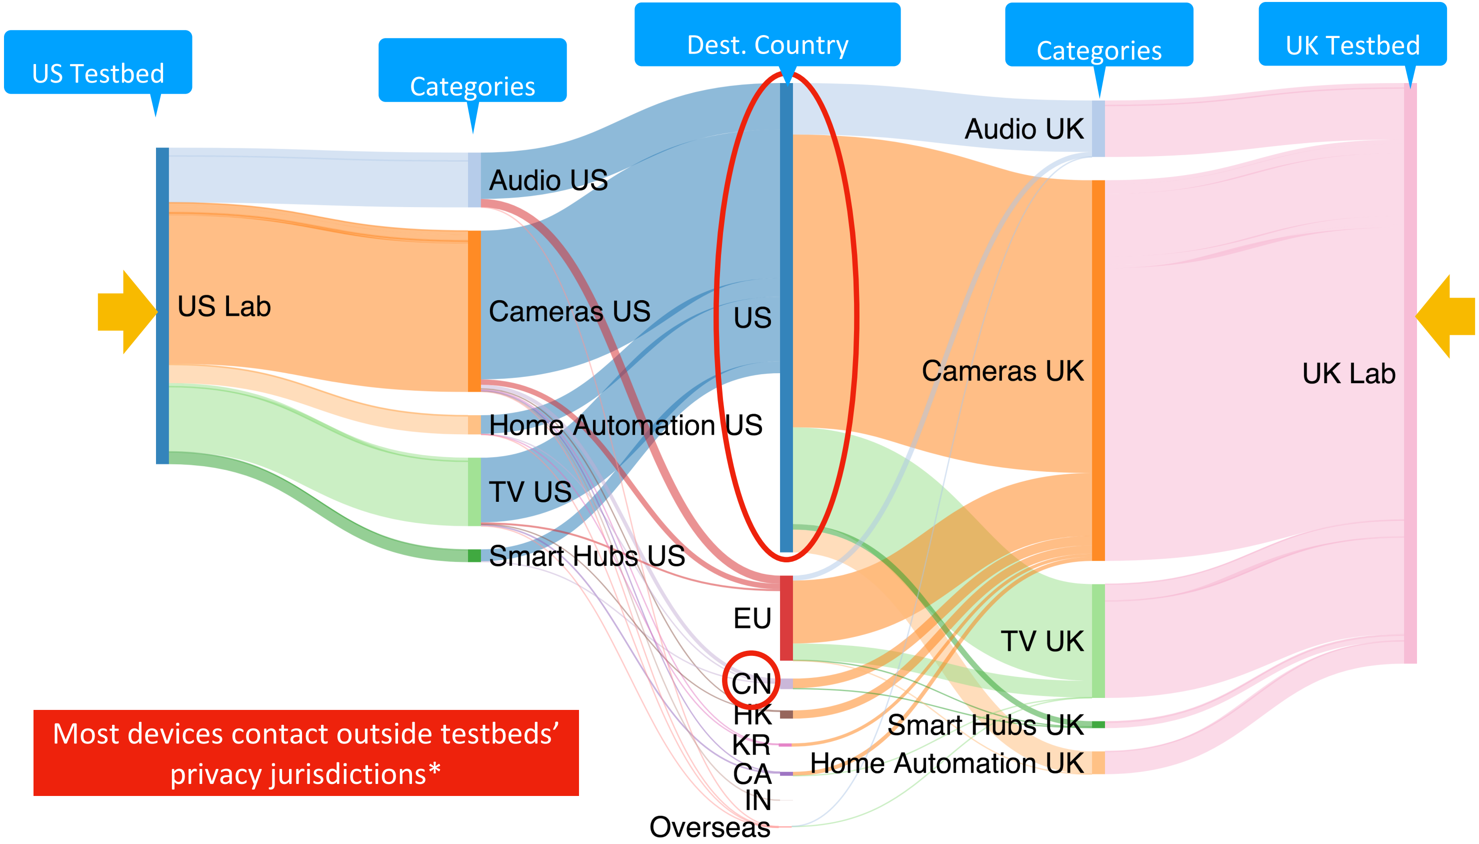

Most devices contact outside testbeds' privacy jurisdictions*

Cases of Unexpected Behavior

Popular doorbells

Video recording on detected motion (cannot be disabled)

Popular smart TVs

Contact **Netflix**, **Google**, and **Facebook** unexpectedly

Alexa-enabled devices

Frequently falsely triggered (e.g. "**I like S**tar Trek")

/// Other notable cases of activities detected when <u>idle</u>

// Cameras reporting **motion** in absence of movement

// Devices spontaneously **restarting** or reconnecting

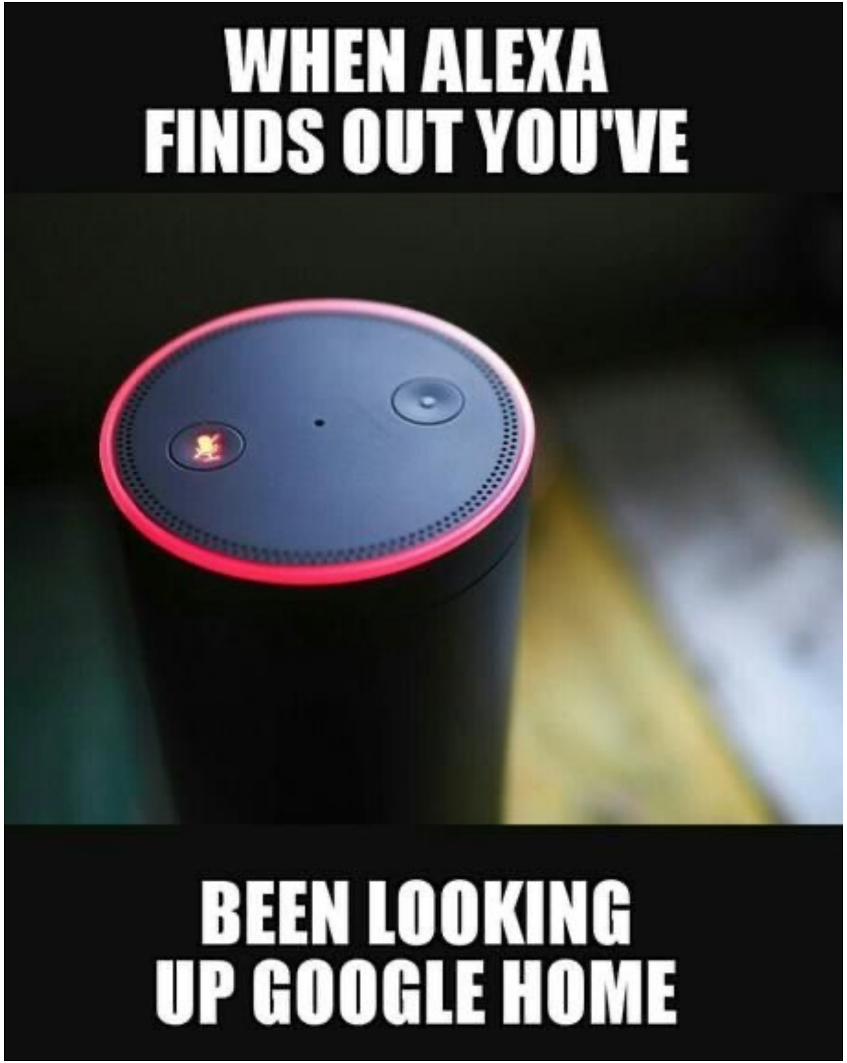
WHEN ALEXA
FINDS OUT YOU'VE
BEEN LOOKING
UP GOOGLE HOME

Are Smart Speakers Listening to Us?

What happens when the wake word is misunderstood?

- Smart speakers signal **activation** (wake word detection) by **lighting up**
- They **send the recording** to the voice assistant cloud service
- The cloud service may **store the recording** and produce an answer

Bloomberg

Technology

Amazon Workers Are Listening to What You Tell Alexa

A global team reviews audio clips in an effort to help the voice-activated assistant respond to commands.

Google updates Home Mini to address major privacy bug

Some units of the smart speaker are found to activate at random times and transmit the audio to Google's servers.

Consumer Reports

Electronics & Computers / Audio & Video / Smart Speakers / Smart Speakers That Listen When They Shouldn't

Smart Speakers That Listen When They Shouldn't

Goals

Understanding when smart speakers mistakenly record conversations

MISACTIVATIONS

How frequent?

Signaled?

Do they adapt?

Regionality?

For how long?

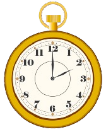

Biases?

Which words?

Measurement challenges and solutions

How to expose smart speakers to content?

1. PLAY AUDIO FROM POPULAR SHOWS

NETFLIX

12 TV shows
134 hours of video
played two times
in two regions

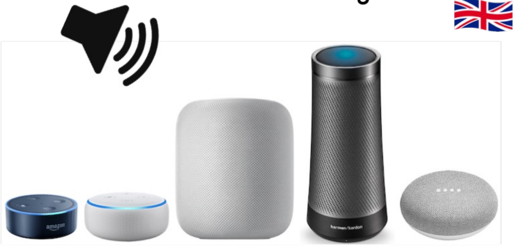

Smart speakers are exposed to audio content

How to detect activations?

2. DETECT ACTIVATIONS

Camera recordings

Detects when a smart speaker lights up and for how long

Cloud data

Detects smart speaker recordings (only for Amazon and Google)

Network traffic

Detects traffic patterns resembling voice transmissions

How to distinguish unwanted activations?

3. RECOGNIZE MISACTIVATIONS

CC

Analyze closed captions of each activation

Do they contain the wake word?

Yes – Legitimate activation

No – **Misactivation!**

Test environments

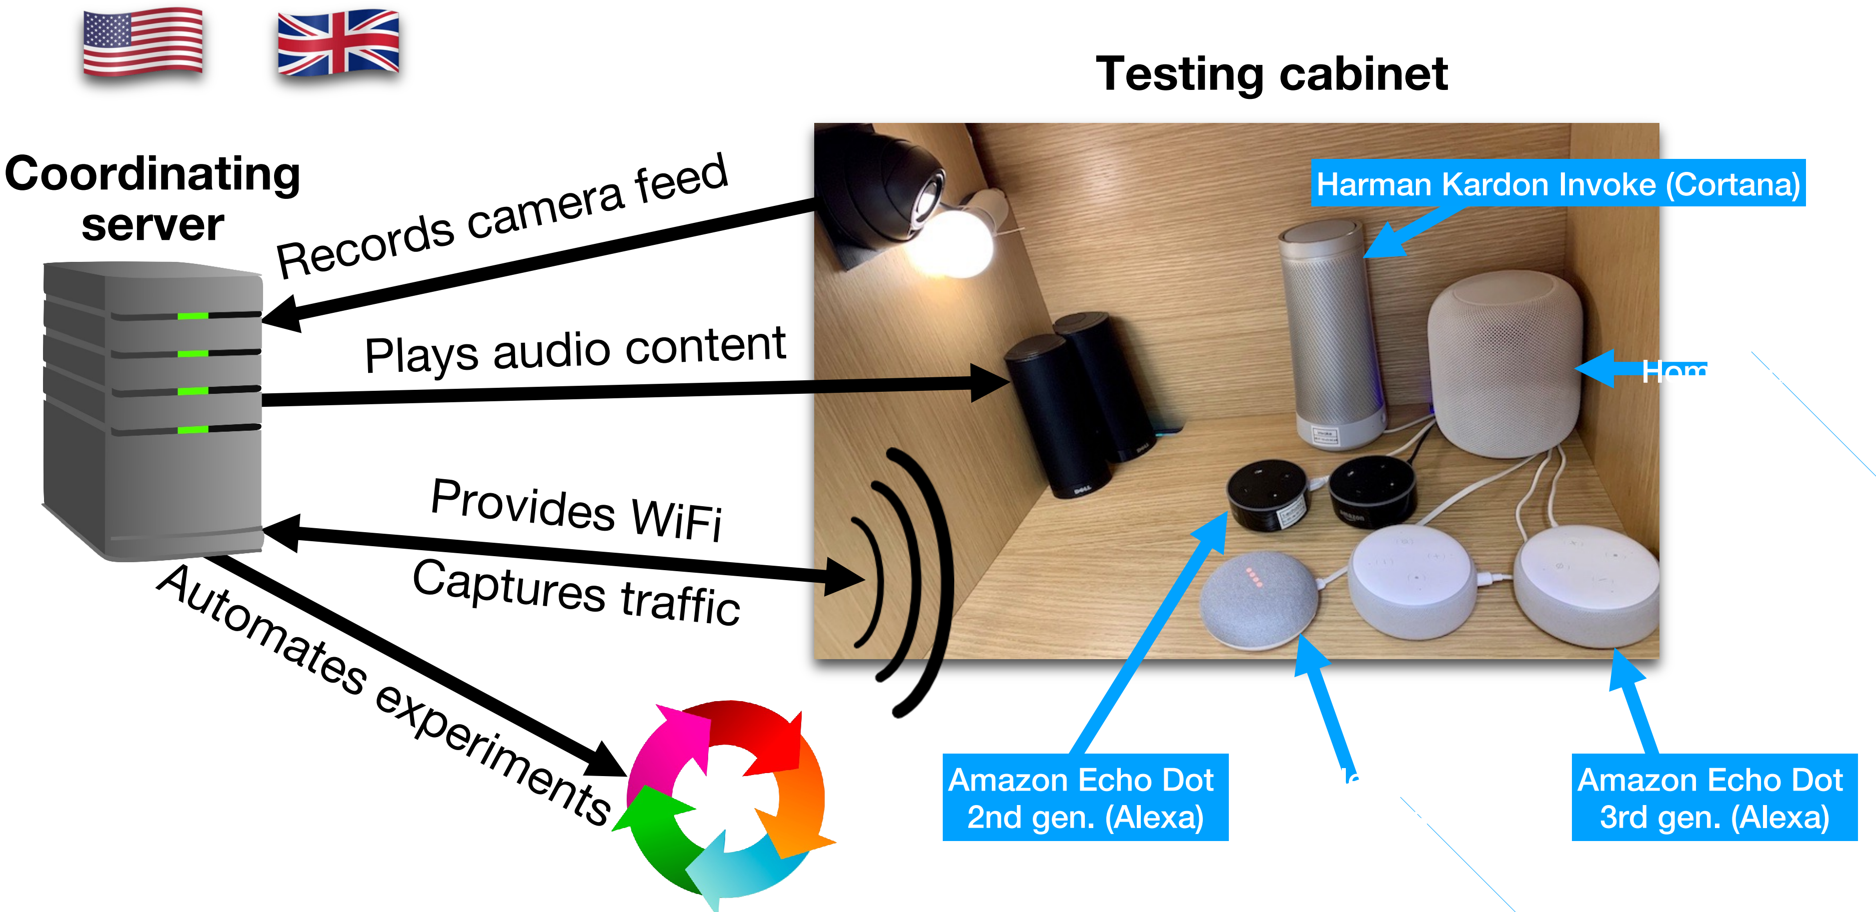

Activation detection methods

Camera activation

- search the video stream for frame changes

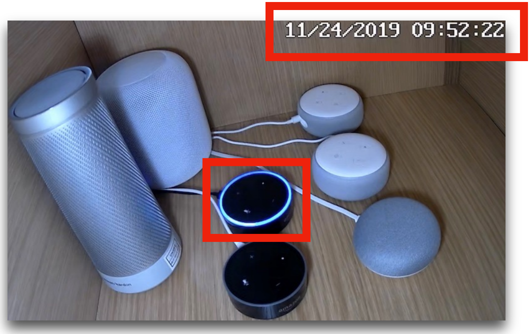

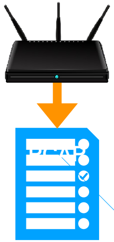

Traffic activation

- look for traffic spikes to certain destinations exceeding a certain threshold

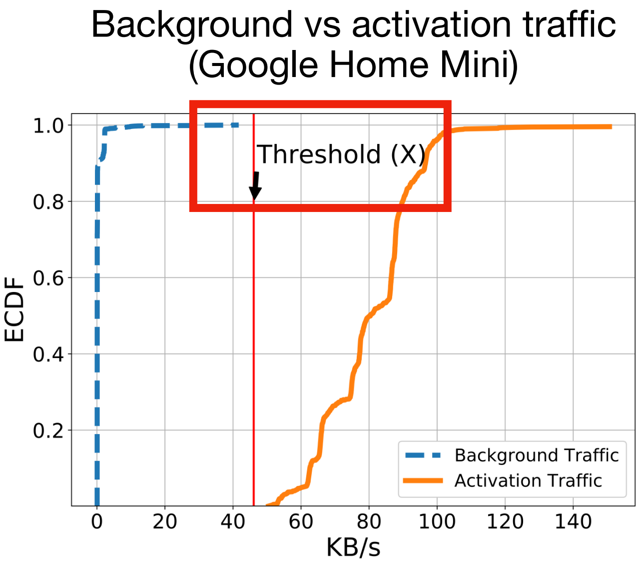

Cloud activation: download information from the voice assistant cloud

- Only for Google Assistant and amazon alexa

How frequently do smart speakers misactivate?

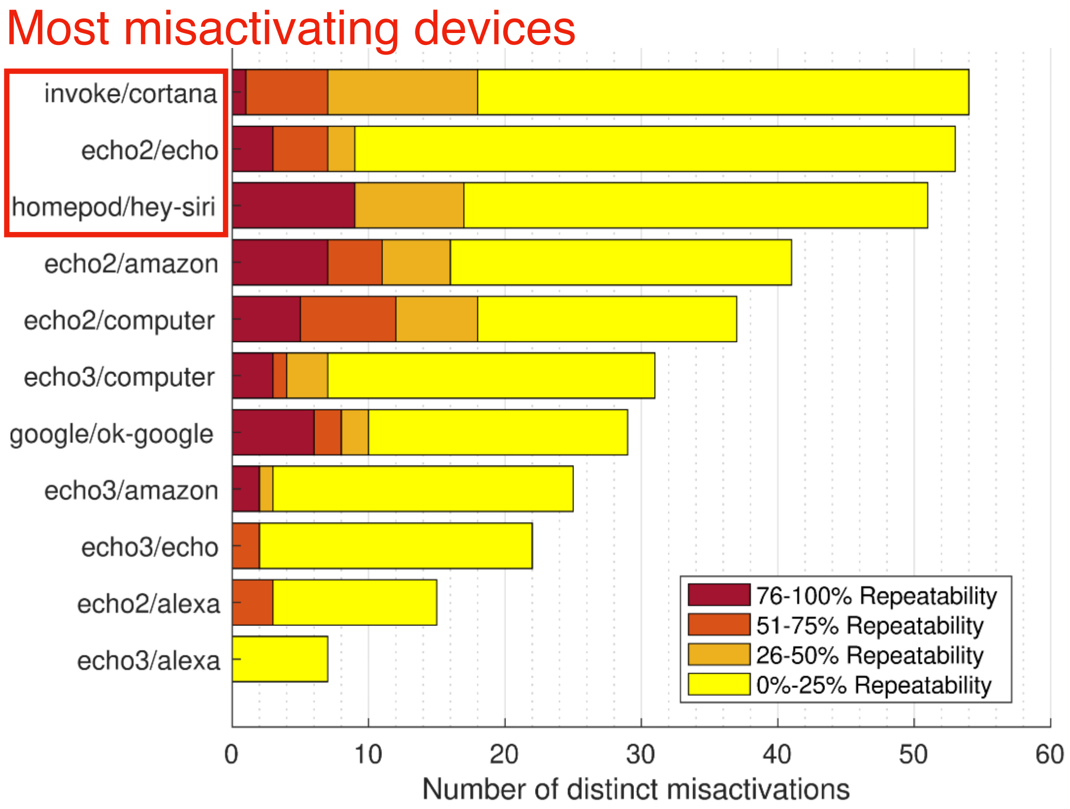

Repeatability

- Consistency of misactivations across experiments

Takeaways

- Devices with the most recordings (**Invoke**, **Echo2**, **Homepod**) expose user privacy more often
- Prevalence of low repeatability suggests low determinism

How long do smart speakers record?

Misactivation duration: amount of time the smart speaker is lit up after a misactivation

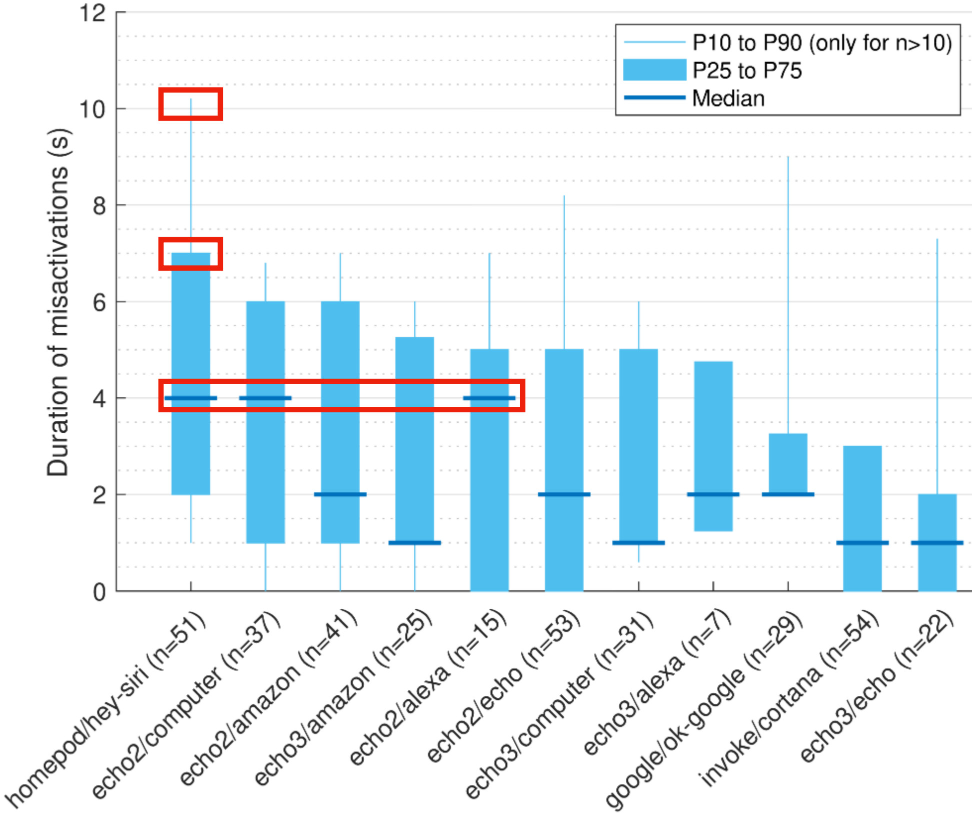

Most common case (median)

- up to **4s** (Homepod, Echo Dot 2G)

Less common case (top 25%)

- up to **7s** (Homepod)

Rare case (top 10%)

- up to **10s** (Homepod)

Enough to grasp a conversation?

What words cause most misactivations?

Words	Some patterns	Some examples from the closed captions of highly repeatable misactivations
OK/Hey Google	Words rhyming with "hey"/"hi" followed by "ol"/"g"/"w"	"**Okay, where were we?**", "hey ... you told", "A-P girl"
Hey Siri	Words rhyming with "hey" or "hi" followed by voiceless "s", "f", "th" sound and "i"/"ee" vowel	"yeah. I was thinking", "Hi. Mrs. Kim", "they ... secretly"
Alexa	"i" followed by "k" sound or a voiceless "s"	"I care about", "I messed up", "I got something"
Echo	"e"/"ee"/"i" vowel followed by hard "k" or "g" sounds	"head coach", "I got", "that cool", "pickle"
Computer	Words starting with "comp" or that rhyme with "here"	"Comparisons", "come here", "nuclear accident"
Amazon	Combinations of "was"/"as"/"goes"/"some"/"I'm" followed by "s"/"z"/"on"/"om"	"it was a", "life goes on", "want some water?", "I was in"
Cortana	"k" sound closely followed by "r" or "t"	"lecture on", "quarter", "courtesy", "according to"

- Most are wake word variations, no evidence of secret wake words
- Potential for some patterns to be used by an attacker to forge commands


WHEN YOU REALIZE
YOUR SMART HOME ISN'T SMART
imgflip.com

Course Overview

- ❑ Benchmarking privacy in IoT devices
- ❑ **IoT devices identification**
- ❑ Benchmarking security in IoT devices
- ❑ Benchmarking security solutions for IoT devices
- ❑ Privacy solutions for IoT devices at the edge
- ❑ Security solutions for IoT devices at the edge
- ❑ IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security

Providers need to “**identify**” and “**locate**” IoT devices in the network

Detecting IoT Devices at the Provider is Challenging

Traffic patterns across IoT devices are diverse

Deploying an agent inside at each ISP customers is not scalable

Active measurements do not work with devices behind NATs

Deep packet inspection raises privacy concerns

Our contribution: a methodology for ***detecting*** and monitoring IoT devices with ***limited, passive,* and *sparsely sampled*** flow data in the ***wild.***(Detection rules available at *https://moniotrlab.ccis.neu.edu/imc20/*)

Key Insights

- Devices have repeating patterns of communication that appear even in sparsely sampled data
- Detection rules can be generated using limited packet fields
- Detected devices from 77% of studied IoT manufacturers in an ISP and IXP within minutes to hours

Methodology

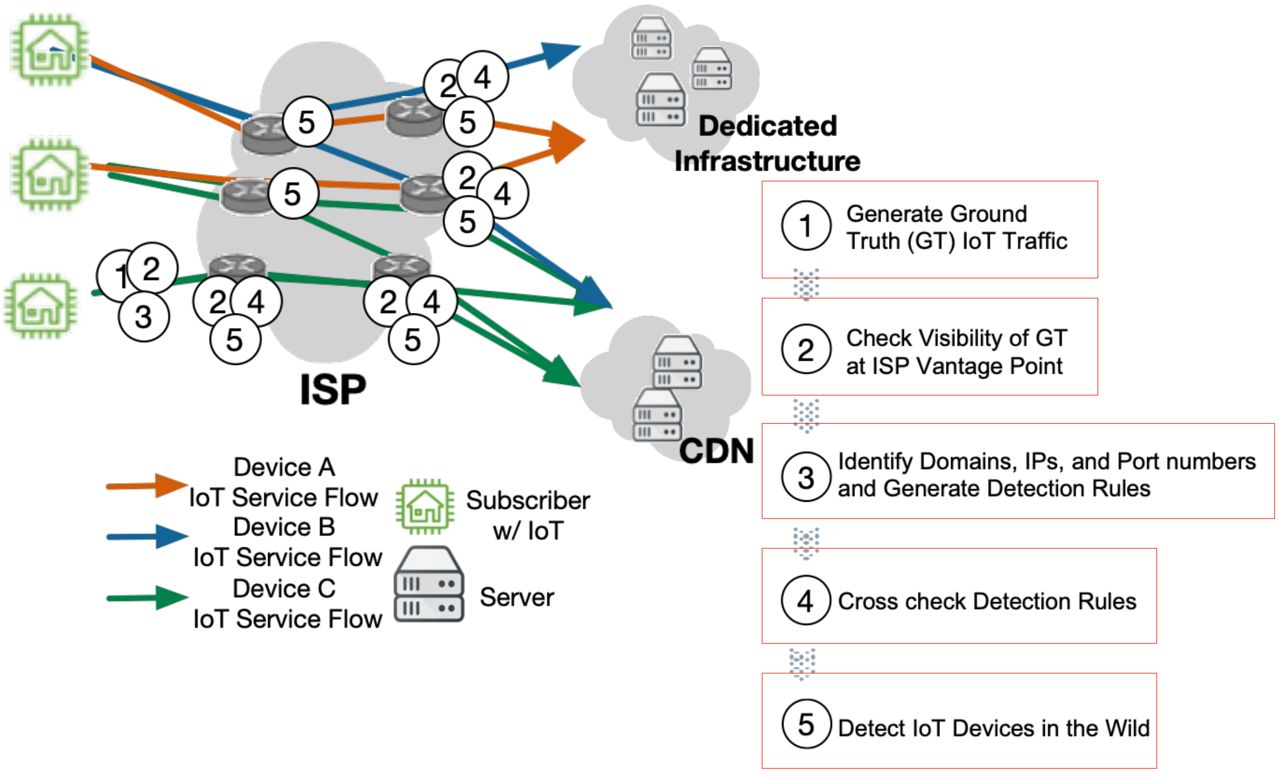

Generate Ground Truth IoT traffic

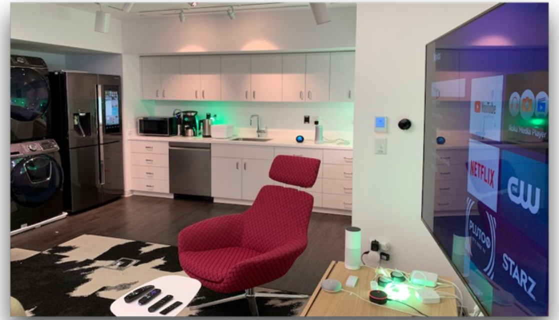

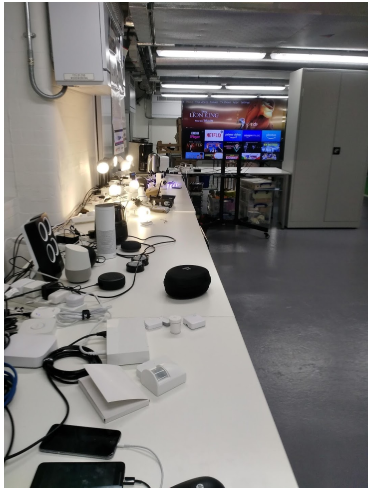

Activity	Description
Power	power on/off the device
Voice	voice commands for speakers
Video	record/watch video
On/Off	turn on/off bulbs/plugs
Motion	move in front of device
Others	change volume, browse menu

- Idle Experiments
- Active Experiments

56 different IoT products

ISP Setup

15M broadband subscriber lines linesExperiments

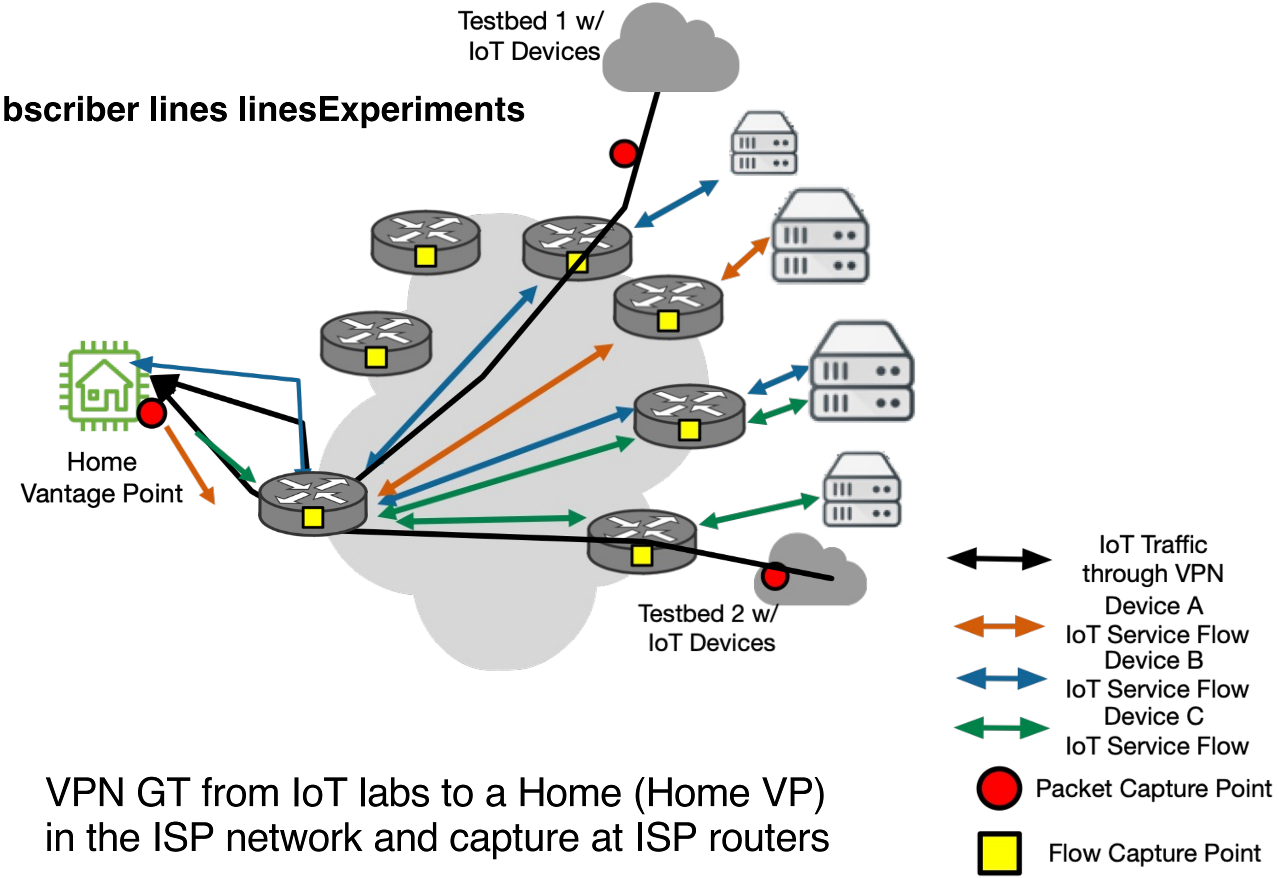

VPN GT from IoT labs to a Home (Home VP)
in the ISP network and capture at ISP routers

Generating Detection Rules

Detection Levels:

Product-level: Amazon Echo

Manufacturer-level: A Samsung Device

Platform-level: an IoT device

Detection Rules:

5 IoT Platforms

20 Manufacturers

11 Products

77% of the manufacturers in the testbeds

Cross Check Detection Rules

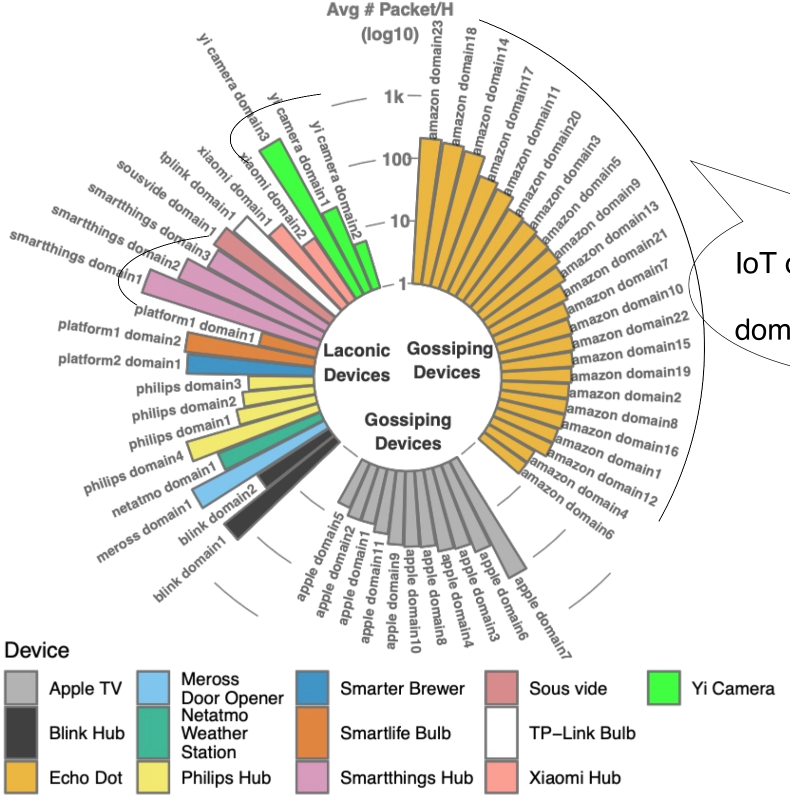

IoT devices talk to different domains at different rates

Number of ISP Subscribers with IoT Devices (Per hour/24h)

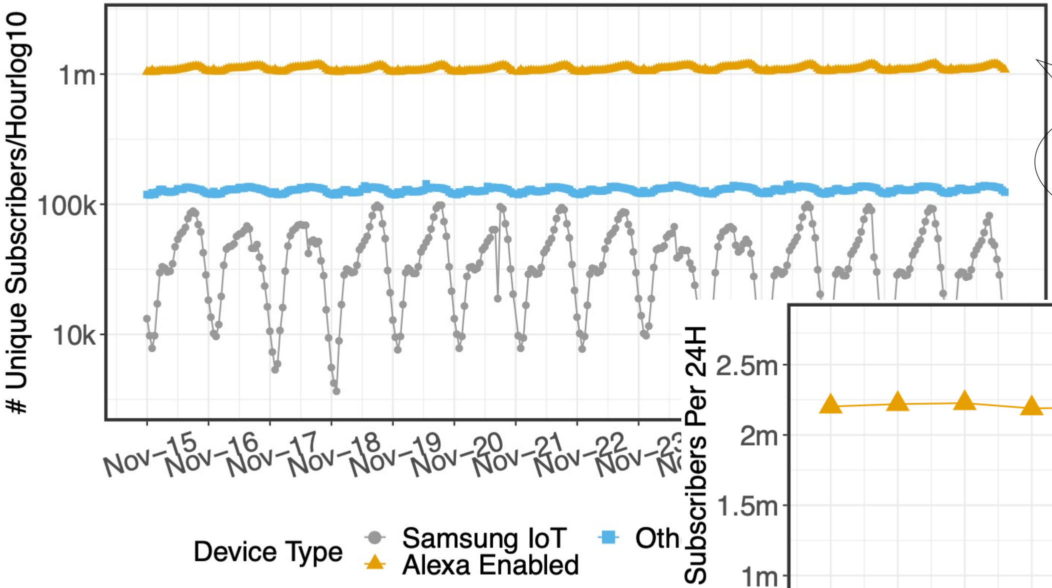

1m+ subscribers with Alexa-enabled devices

Increasing observation period helped detecting more devices

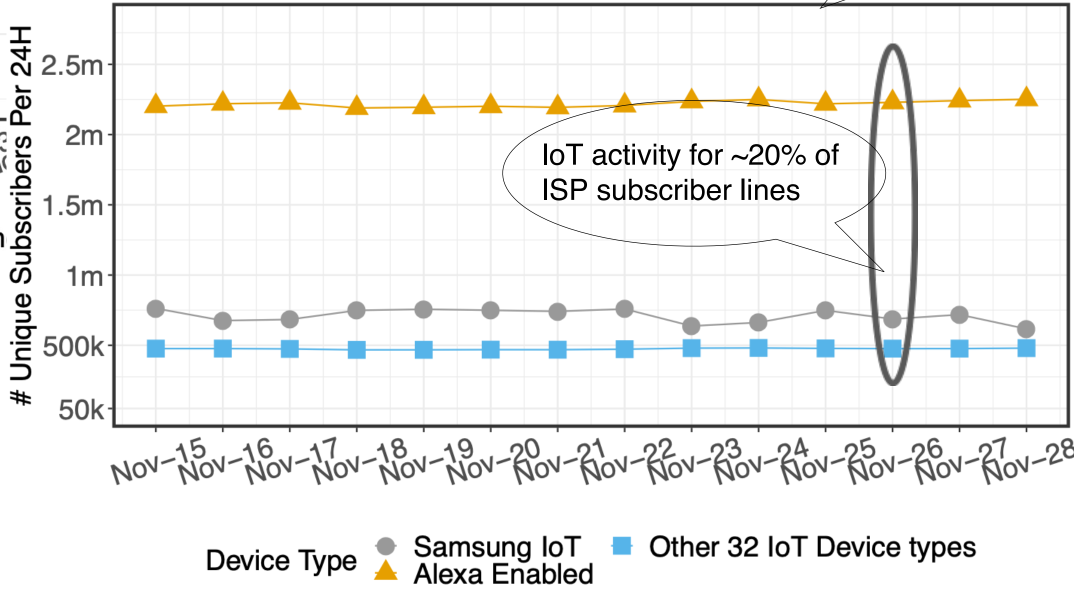

Some diurnal patterns for Alexa and Samsung IoT devices

Detecting IoT Devices Activity in the Wild

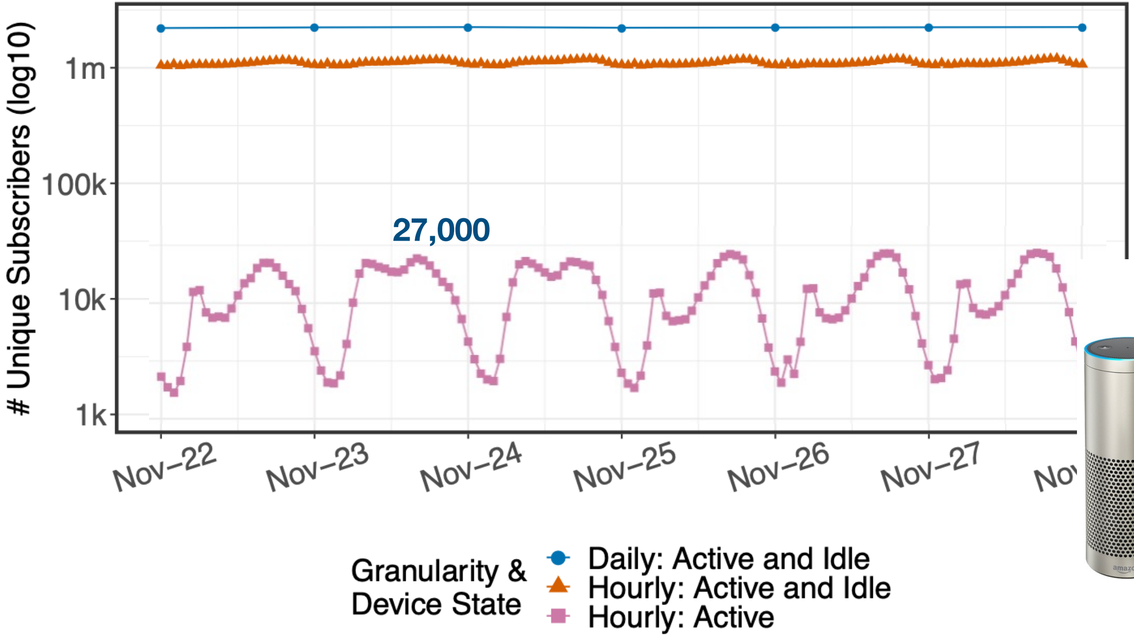

27,000

THRESHOLD

For some devices we can infer activity

Course Overview

- ❑ Benchmarking privacy in IoT devices
- ❑ IoT devices identification
- ❑ Benchmarking security in IoT devices
- ❑ Benchmarking security solutions for IoT devices
- ❑ Privacy solutions for IoT devices at the edge
- ❑ Security solutions for IoT devices at the edge
- ❑ IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security


I DON'T ALWAYS
THINK ABOUT CYBER SECURITY
BUT WHEN I
DO, IT'S TOO LATE
imgflip.com

Contributions

- We develop an automated methodology for evaluating security vulnerabilities in common consumer IoT devices using large-scale, diverse experiments and sets of attacks
- We assess the security vulnerabilities of popular IoT devices against existing network and device attacks and identify privacy risks

Assumptions

- Threat modelling
  - *Adversary*: Any party that can access the IoT device's network
  - *Victim*: The victim is anyone who enters the service area of the IoT device
  - *Threat*: We assume the presence of malicious or compromised IoT devices in a smart home. Adversaries may be incentivised to compromise other devices in the network to infer user activities or deny their usage of them.
- **Goals**
  - **Are consumer IoT devices vulnerable to common security attacks?**
  - **Do IoT devices detect threats?**
- Non-goals
  - We have no control over how an IoT device works internally.
  - We do not test all threats.
  - We only focus on consumer IoT devices.

Testbed

Testbed

Category	Device
Smart speaker	Bose Smart Speaker 500
	Sonos One (Gen2)
	Echo Dot 5
Smart doorbell	Ring Chime Pro
	Ring Video Doorbell (2nd Gen)
Smart camera	Google Nest Cam
	SimpliSafe Security Camera Indoor
	Furbo 360° Dog Camera
Appliances	WeeKett Smart Wi-Fi Kettle
	Govee Alexa LED Strip Lights
	Sensibo Sky Smart AC

Testbed

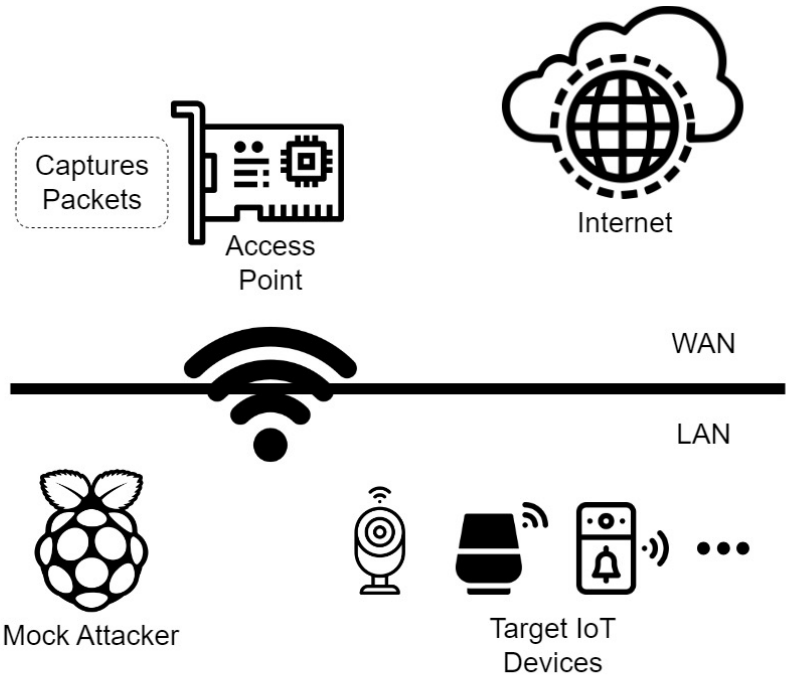

Category	Attacks
Flooding	SYN (port 80) flooding
	UDP flooding
	DNS flooding
	Fragmented IP flooding
Scanning	Port scanning
	OS scanning

- Within the same LAN
- Packets are captured on the access point
- Tshark for filtering responses
- Assess device reaction
  - Counter-measures detected – attack *unsuccessful*
  - No counter-measures detected – attack *successful*

Software

- We write and use configurable and automated scripts for simulating attacks and analysing the replies
- We setup tcpdump to continuously capture network traffic on the network access point
- Dedicated network traffic capturing for active experiments
- Devices are activated with their companion applications remotely and automatically using ADB
- We verify the attacks using two RPis

Software – testing usecase

- Activate the device with ADB
- Start running a simulated attack on the device’s IP address
- Wait until the attack stops
- Download the captured traffic
- Analyse the traffic using tshark

Results - flooding

Devices	SYN	UDP	DNS	Frag. IP
Bose Speaker	✓	✗	✗	✓
Sonos One (Gen2)	✗	✗	✗	✓
Echo Dot 5	✗	✗	✗	✓
Ring Chime Pro	✗	✗	✗	✓
Ring Doorbell	✗	✗	✓	✓
Google Nest Cam	✗	✗	✗	✓
SimpliSafe Cam	✗	✗	✗	✓
Furbo Camera	✗	✗	✗	✓
WeeKett Kettle	✗	✓	✓	✓
Govee Lights	✗	✗	✓	✗
Sensibo Sky	✗	✓	✓	✓

- Most of the devices are vulnerable to Frag. IP flooding, as opposed to SYN flooding, which is only successful on the Bose Speaker.

Results – port scanning

- Open ports can be detected on 7 devices out of 11.

Devices	Identified Open Ports
Bose Speaker	80/7000/8082/8083/8085/8091/8200/30030/40002/40031/40035
Sonos One (Gen2)	1400/1410/1443/1843/7000
Echo Dot 5	1080/4070/8888/55442/55443
Ring Chime Pro	847/1003/1020/1393/3736/7240/8173/12302/15986/16891/17704/17944/17993/18682/20307/21257/23825/24669/25781/25958/25997/26757/27234/28363/29161/32466/33377/33544/33616/33862/35470/38657/44100/46108/46194/47199/50852/51212/52663/54739/55524/55530/56621/65488
Ring Doorbell	Blocking ping probes & none found
Google Nest Cam	8012/10101/11095
SimpliSafe Cam	19531
Furbo Camera	None found
WeeKett Kettle	6668
Govee Lights	None found
Sensibo Sky	None found

Results – OS scanning

Devices	Operating System
Bose Speaker	Linux 3.2 - 4.9
Sonos One (Gen2)	Linux 3.2 - 4.9
Echo Dot 5	No exact match, can be Linux
Ring Chime Pro	Too many fingerprints match
Ring Doorbell	2N Helios IP VoIP doorbell (95%)
Google Nest Cam	Too many fingerprints match
SimpliSafe Cam	Too many fingerprints match
Furbo Camera	Too many fingerprints match
WeeKett Kettle	No exact OS matches
Govee Lights	Espressif esp8266 firmware (lwIP stack), NodeMCU firmware (lwIP stack)
Sensibo Sky	Philips Hue Bridge (lwIP 1.4.1), Philips Hue Bridge (lwIP stack)

- OS can be identified on 5 devices out of 11.

Discussion

- Potentially consequential user implications can be identified (e.g. a successful DoS attack on the LED light)
- Open ports and identified OS could be exploited for obtaining private info (e.g. camera feed)
- Limitations
  - We consider devices as black-boxes
  - We only tested 11 devices
- Ethical considerations
  - We follow the ethical guidelines of our affiliated organisation
  - We conduct our experiments locally


Install IoT security camera
camera has security flaw
strangers can watch the video feed
strangers can watch the video feed

Course Overview

- ❑Benchmarking privacy in IoT devices
- ❑IoT devices identification
- ❑Benchmarking security in IoT devices
- ❑**Benchmarking security solutions for IoT devices**
- ❑Privacy solutions for IoT devices at the edge
- ❑Security solutions for IoT devices at the edge
- ❑IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security

Problem: IoT Devices Expose Information Over the Internet

They “sense” a lot

Microphones
Cameras
User activities

…

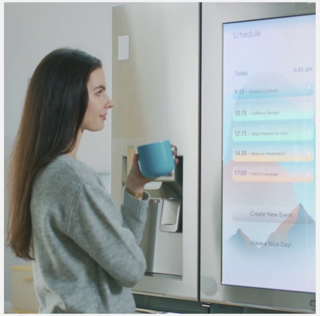

Privacy Threats

IoT devices collect user information

They share user information

Security Threats

Malware can affect IoT devices

An attacker can control them

User Frustration

IoT devices privacy/security is hard to control

Hard to protect users from IoT threats

IOT PROTECTION SYSTEMS: SAFEGUARDS

Why Were We Interested in This?

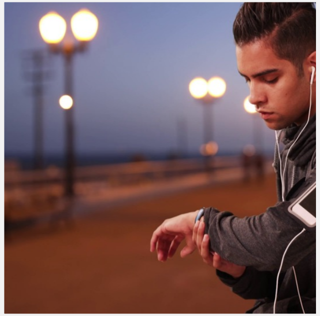

Control

Device detection

Intelligent profiles

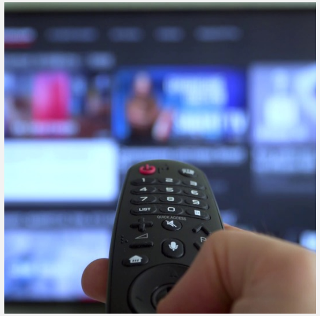

Security

Vulnerability Assessment

Brute Force Protection

Anomaly Detection

Privacy

Content filtering

Network Intrusion Prevention

- These safeguards may currently be ineffective in preventing risks.
- Their cloud interactions and data collection operations may introduce privacy risks.

Research Questions

- ❑ **Goal 1:** What are the privacy and security implications on how a safeguard works?
- ❑ **Goal 2:** Do the safeguards detect threats?
- ❑ **Goal 3:** What are the side effects of the safeguards?

IoT Safeguards

Challenges for Measuring IoT Safeguards

Difficult to automate the testing of commercial IoT safeguards

- Closed systems
- Blackbox approach

Difficult to perform IoT experiments and generalize

- Lack of automation and emulation tools
- Lack of standard testbed

Our contribution: a large IoT testbed used to test IoT safeguards in real-world scenarios (software and data available online).

Selecting IoT Safeguards

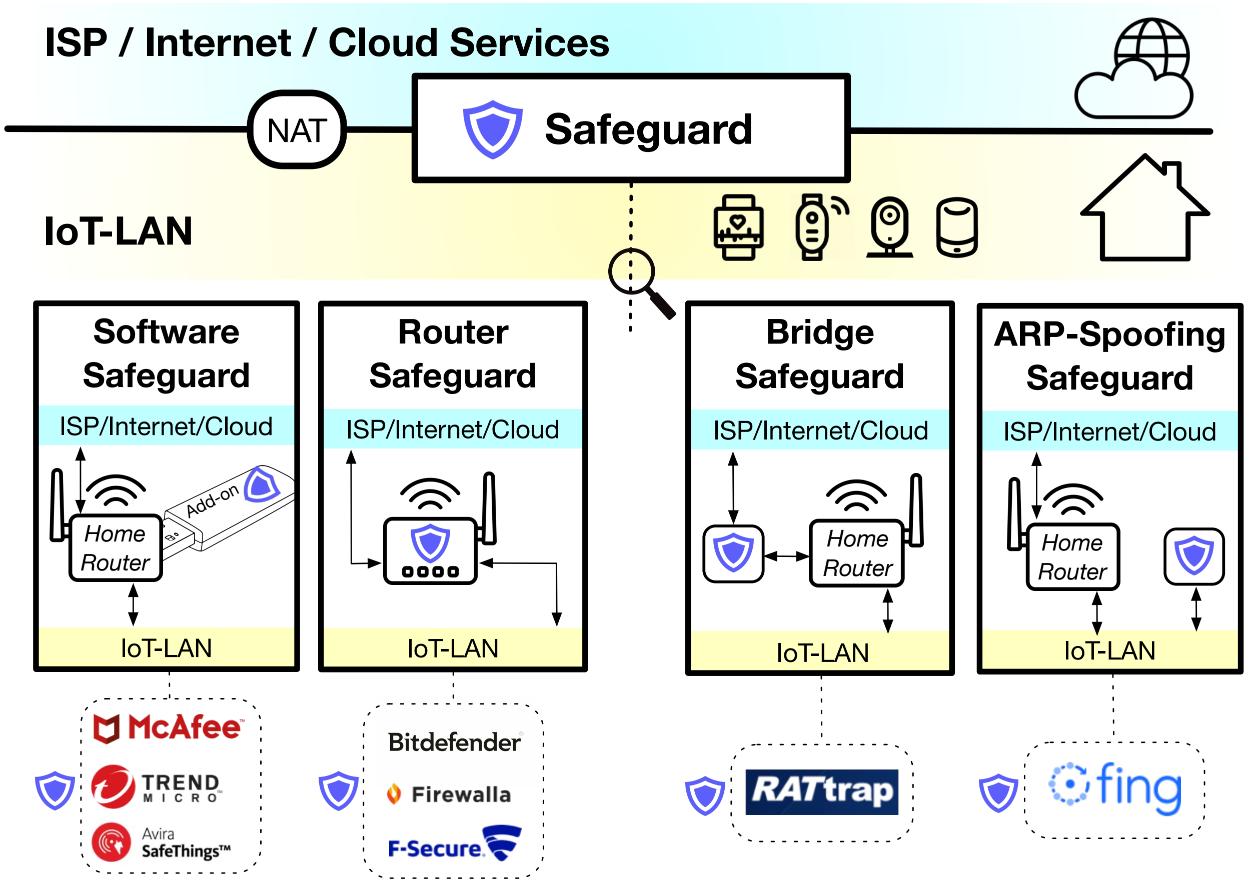

Testbed

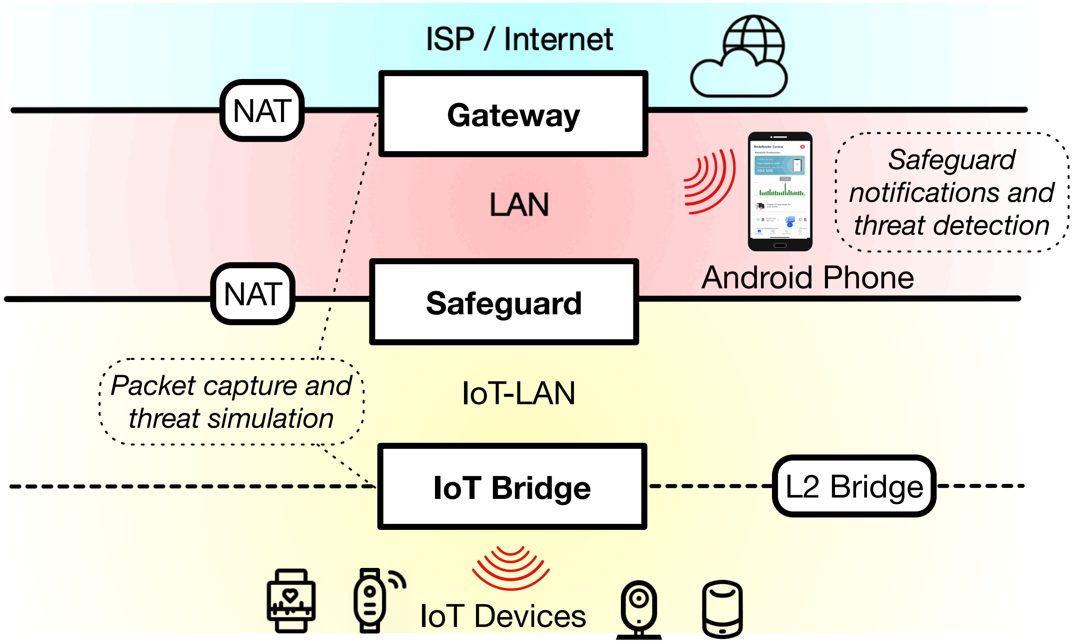

ISP / Internet
NAT
Gateway
LAN
Safeguard notifications and threat detection
Android Phone
NAT
Safeguard
Packet capture and threat simulation
IoT-LAN
IoT Bridge
L2 Bridge
IoT Devices

Research Questions

- **Goal 1:** What are the privacy and security implications on how a safeguard works?
  - **Identify locality**: cloud vs local operation
  - **Operation**: usage third-party services to operate

IoT Safeguards

Processing Locality & Party Characterization

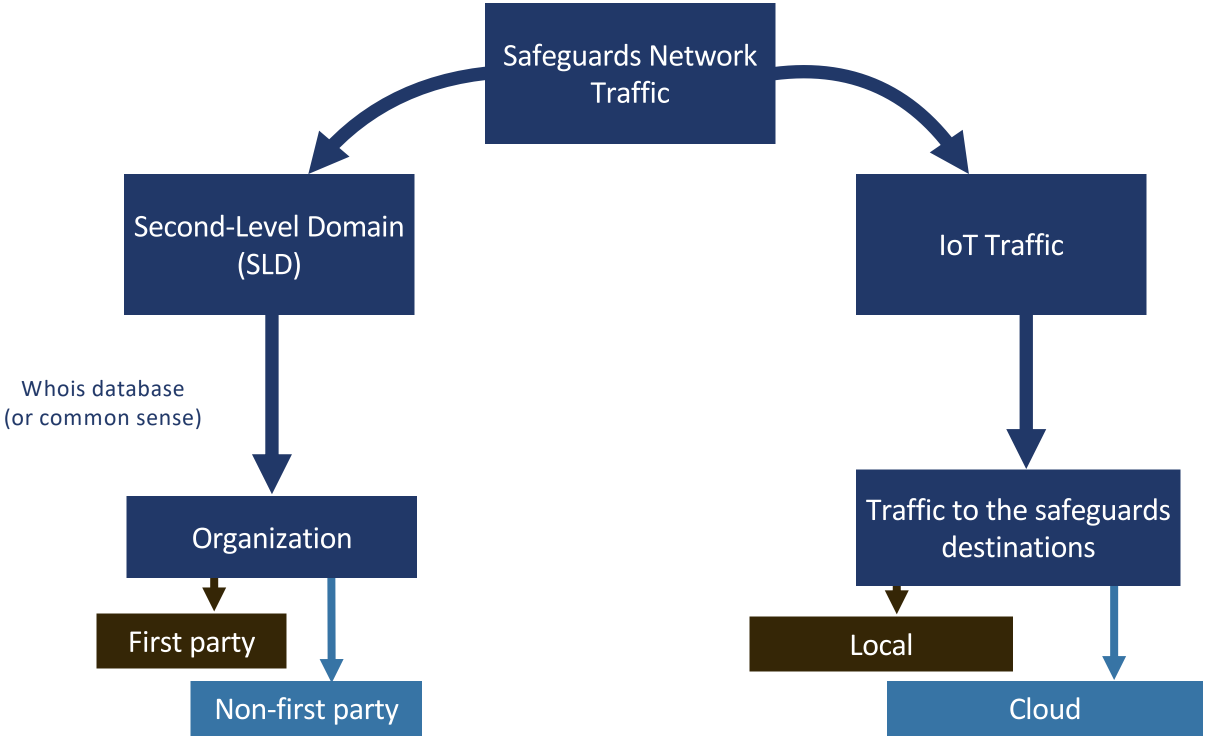

Processing Locality & Party Analysis

Safeguard	Destinations #	Cloud	# and list of Support/3rd Parties
Avira	10	Yes	(1) api.mixpanel.com
Bitdefender	5	Yes	-
F-secure	1	Yes	-
FingBox	5	Yes	(2) api.snapcraft.io, mlab-ns.appspot.com
Firewalla	4	No	(1) api.github.com
McAfee	22	Yes	(3) app-measurement.com, commscope.com, avast.com
RatTrap	1	Yes	-
TrendMicro	3	Yes	(1) policy.ccs.mcafee.com

Take away: - Usage of the cloud for performing analysis, potentially leaving the user vulnerable in the event of a data breach.
- Destinations contacted that are not first parties.

IoT Device Identification

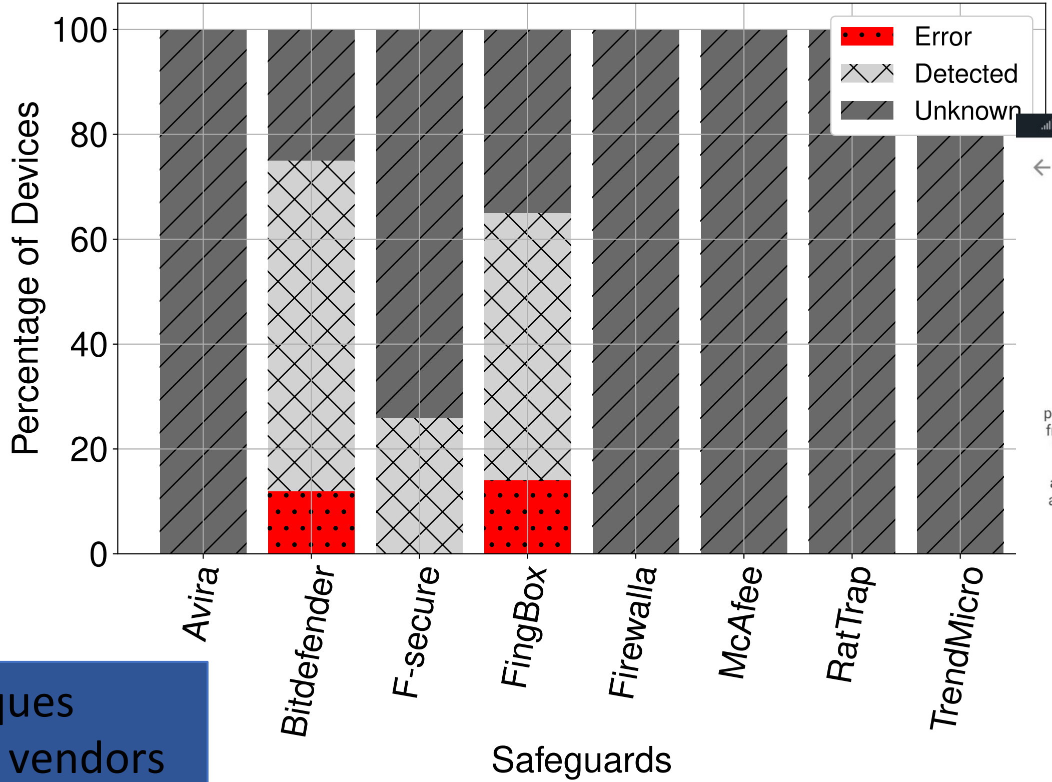

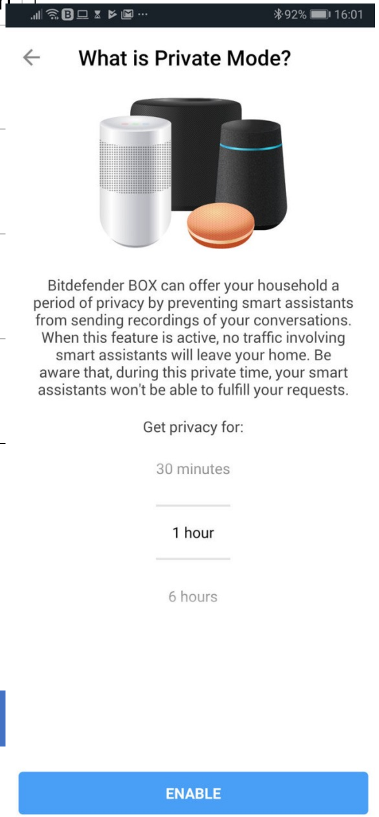

Protection techniques applied to specific vendors

percentage of IoT devices is correctly identified.

Research Questions

❑ **Goal 2:** Do the safeguards detect threats?

- Safeguards **notify** the user when detecting privacy or security threats

IoT Safeguards

Testing Threat Detection Capability

- Security

Threats
Anomalous behavior
Open Port
Weak Password
Device Quarantine
DoS attacks
Port/OS Scanning
MaliciousDestinations

- Privacy

Threats
PII Exposure
Unencrypted Traffic
DNS over HTTPS

Threat Detection Experiments

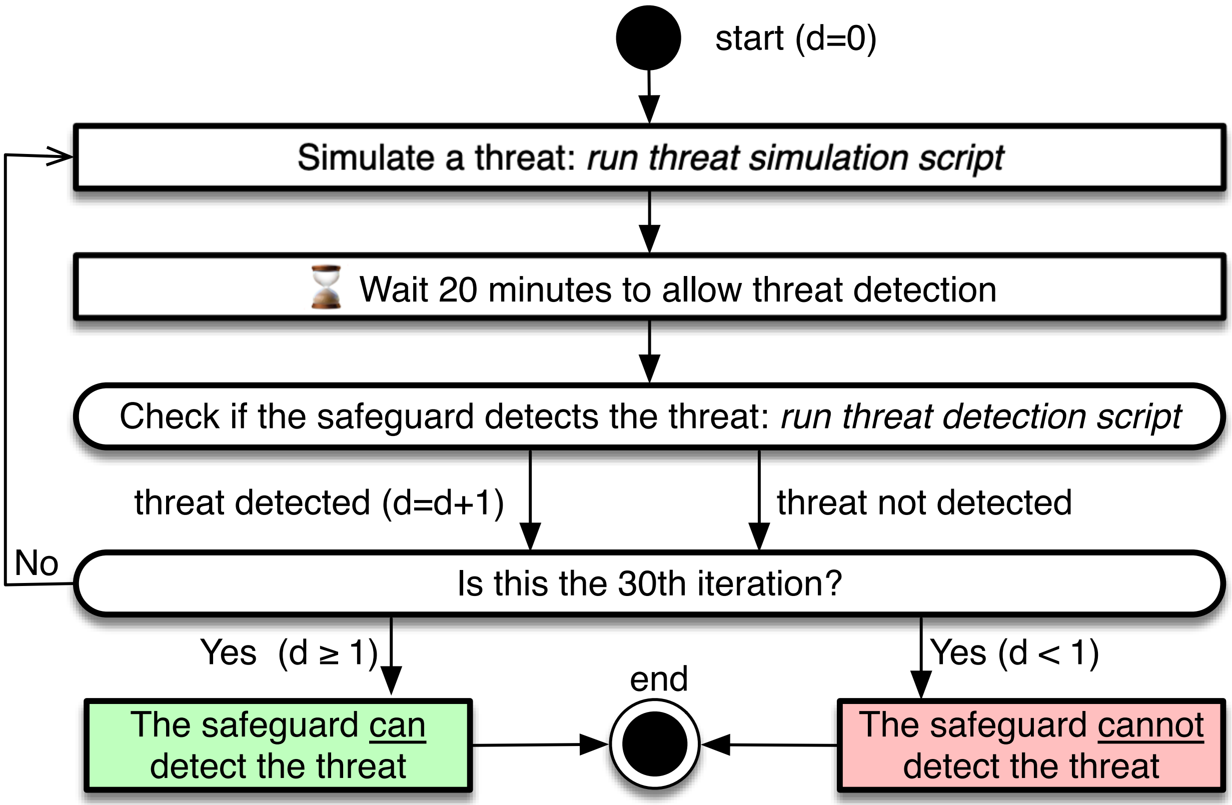

Evaluation of Threat Detection Capability

	Threat	Avira	Bitdefender	F-Secure	Fingbox	Firewalla	McAfee	RaTtrap	TrendMicro
Security	Anomaly ON/OFF	-	✗	✗	-	✗	✗	✗	-
	Anomaly Traffic Pattern	-	✗	✗	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
	Abnormal Upload	-	✗	✗	[illegible]	[illegible]	[illegible]	[illegible]	-
	Open Port	✗	✓(30s)	-	✗	✓(30s)	✗	-	✗
	Weak Password	✗	✗	-	-	-	✗	-	✗
	Device Quarantine	-	✓	-	✓	✓	-	✗	-
	SYN Flooding	✗	✓(30s)	✗	-	✓(40s)	✗	✗	✗
	UDP Flooding	✗	✗	✗	-	✗	✗	✗	✗
	DNS Flooding	✗	✗	✗	-	✗	✗	✗	✗
	HTTP Flooding	✗	✓(3m)	✗	-	✓(2m)	✗	✗	✗
	IP Fragmented Flood	✗	✗	✗	-	✗	✗	✗	✗
	Port Scanning	✓(45s)	✗	✗	-	✗	-	✗	✓(30s)
	OS Scanning	✓(45s)	✗	✗	-	✗	-	✗	✗
	Malicious Destinations	✓	✓	✗	-	✓	✗	✗	✓
Privacy	PII Exposure	✗	✗	-	-	✗	-	-	-
	Unencrypted Traffic	✗	✗	-	-	✗	-	-	-
	DNS over HTTPS	✗	✓	-	-	✓	-	-	-

Time consistency

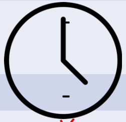

Take away: - only 3 out of 14 threats are detected by the safeguards. 3 out of 8 safeguards do not detect any threats at all, despite they claiming to do so in their specifications

- Some of safeguards take between 45 seconds and 3 minutes to detect a security threat.

Research Questions

❑ **Goal 3:** What are the side effects of the safeguards?

- **Traffic overhead, overprotection, privacy implications**

IoT Safeguards

Safeguard Side Effects

Overprotection

CONNECT 12 IOT DEVICES TO THE SAFEGUARDS AND CAPTURE THE TRAFFIC FOR ONE MONTH

Network traffic overhead

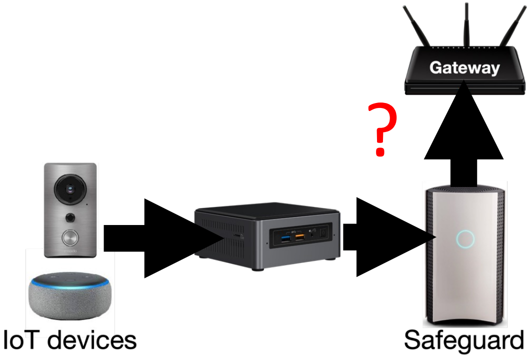

Privacy Policy

MANUALLY INSPECTING THE PRIVACY POLICY

Overprotection

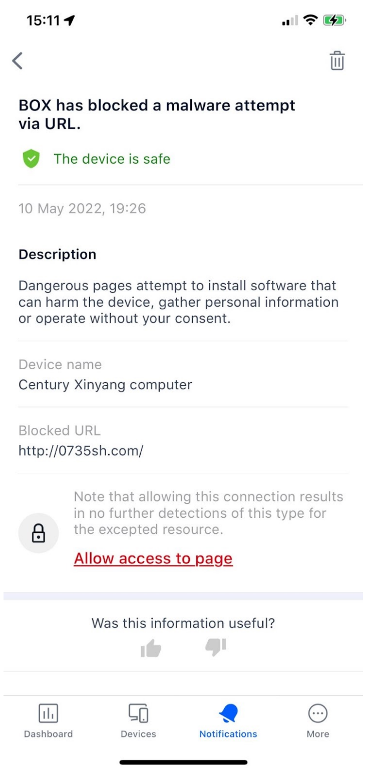

Take away: Most safeguards do not overprotect (i.e., they do not report threats that do not occur).

Traffic Overhead

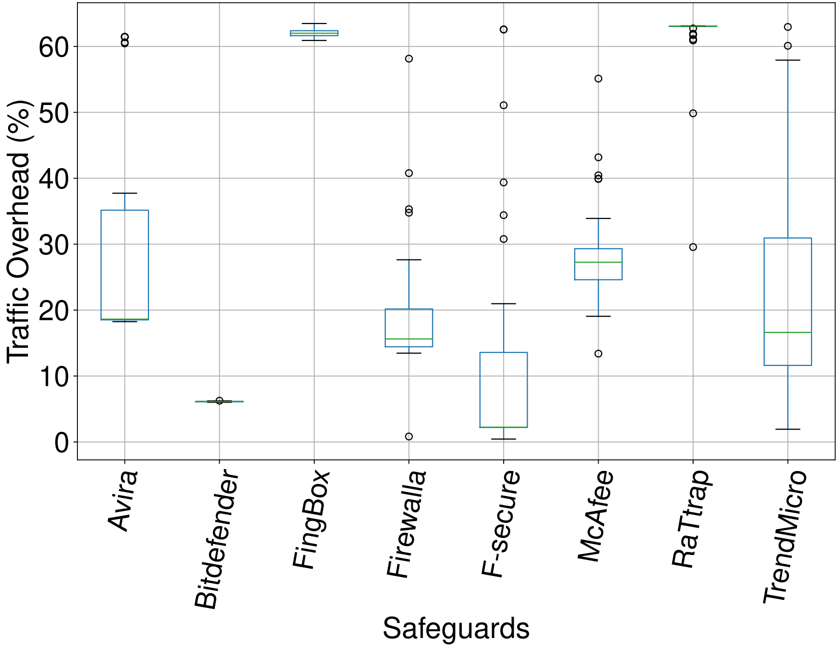

Take away: Some of the safeguards introduce significant traffic overhead. In general the overhead is never less than 10% of the traffic of the IoT devices.

Privacy Policy

Privacy Policy	Avira	Bitdefender	F-Secure	Fingbox	Firewalla	McAfee	RaTtrap	TrendMicro
Anonymization	✓	✓ [pseudonymize]	✗ [ceasing subscription]	✓	✗	✗	✗	✗
Usage of Personal Data	✓	✓	✓	✓	✓	✓	✓	✓
Retention Period	In accordance with legal requirements	10 years	6 months	As long as necessary	Indefinitely	Subscription period	Subscription period	Ongoing legitimate business need
Third Party	SaaS vendor, Akamai. Mixpanel, Ivanti	Partners	Partners	Partners	✗	Partners	Partners	Partners

Take away: Most user information is shared with third-party entities, sometimes without anonymization. Sharing data outside user's privacy jurisdiction.

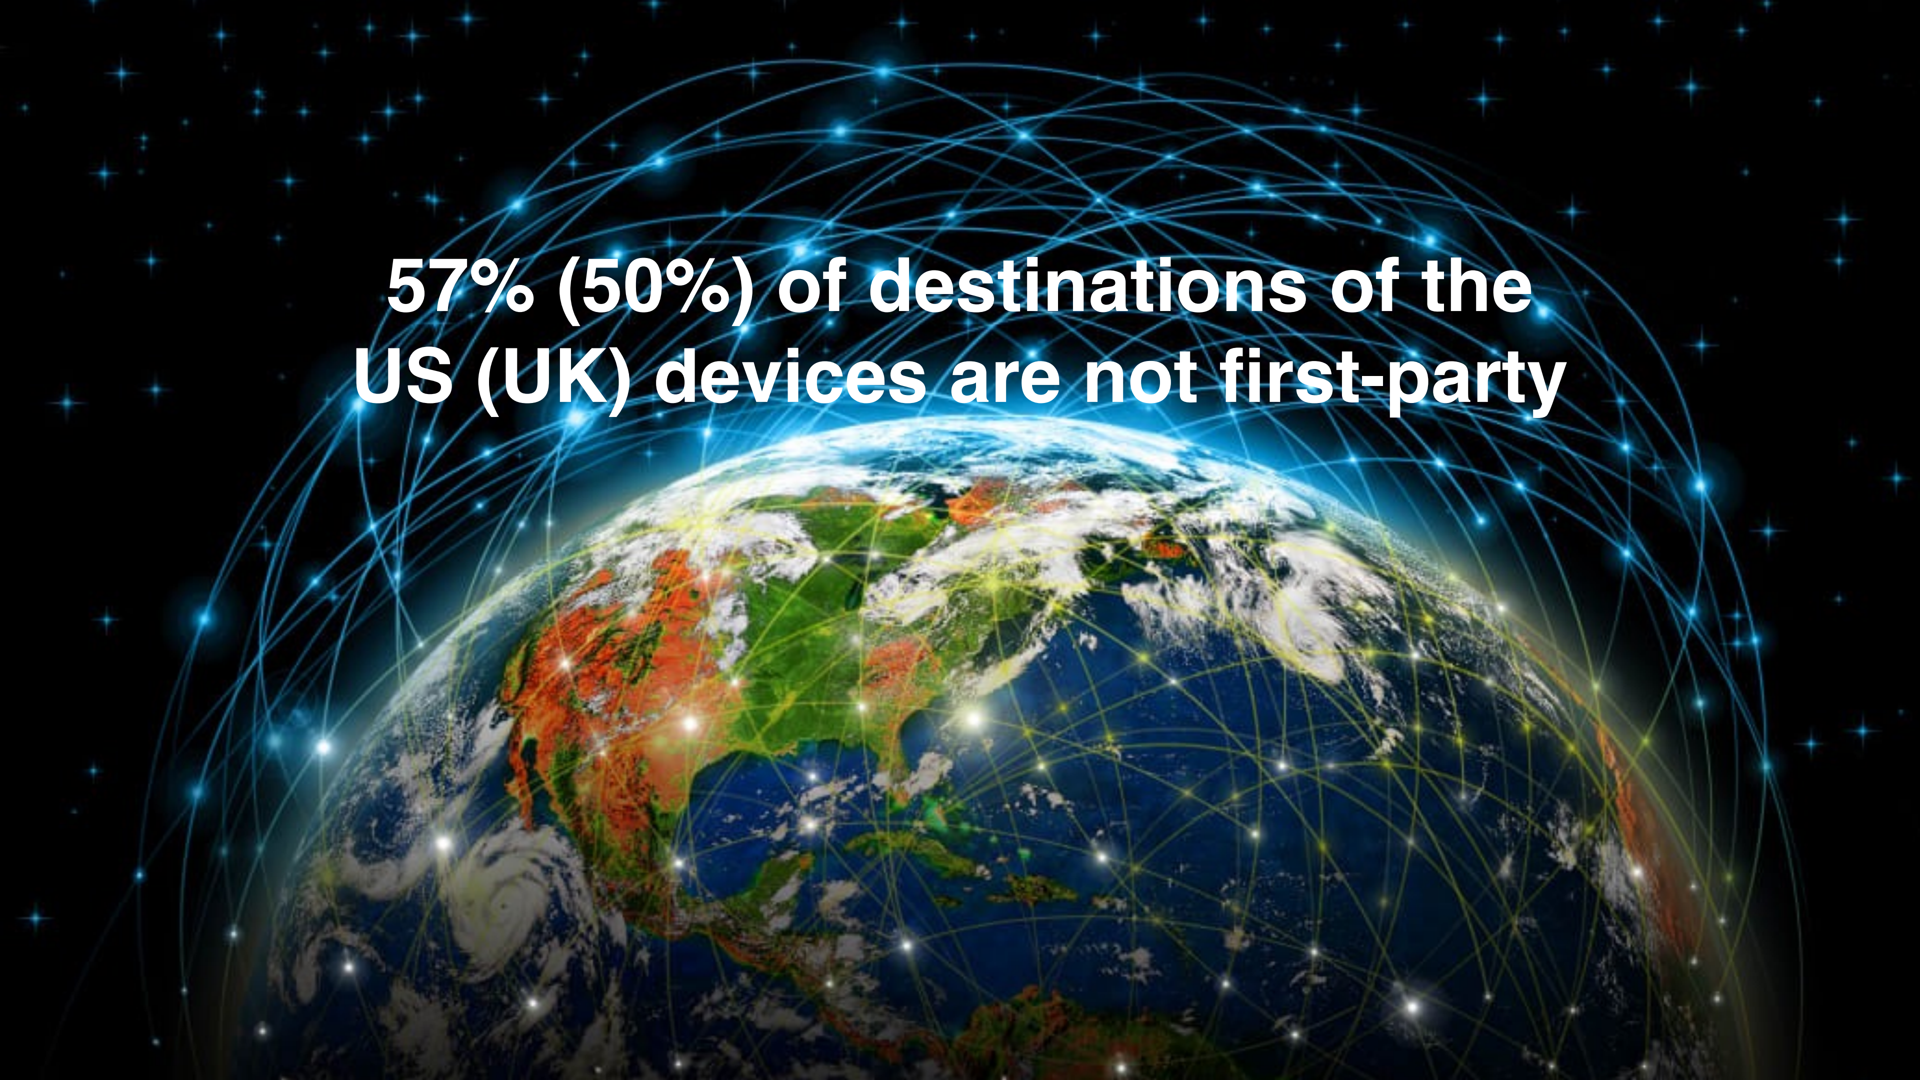
57% (50%) of destinations of the
US (UK) devices are not first-party

Why is this a problem?

Profiling

Mass-influence

User emotion


TODAY SECURITY ENGINEER SITUATION
DEAR HACKERS PLS DON'T DO ANYTHING TODAY!

What might this mean for the future?

Control

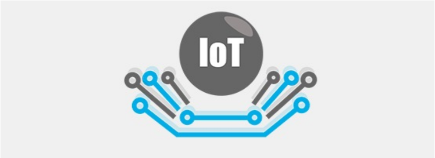

MUD profile

Certificate

Course Overview

- ❑ Benchmarking privacy in IoT devices
- ❑ IoT devices identification
- ❑ Benchmarking security in IoT devices
- ❑ Benchmarking security solutions for IoT devices
- ❑ **Privacy solutions for IoT devices at the edge**
- ❑ Security solutions for IoT devices at the edge
- ❑ IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security

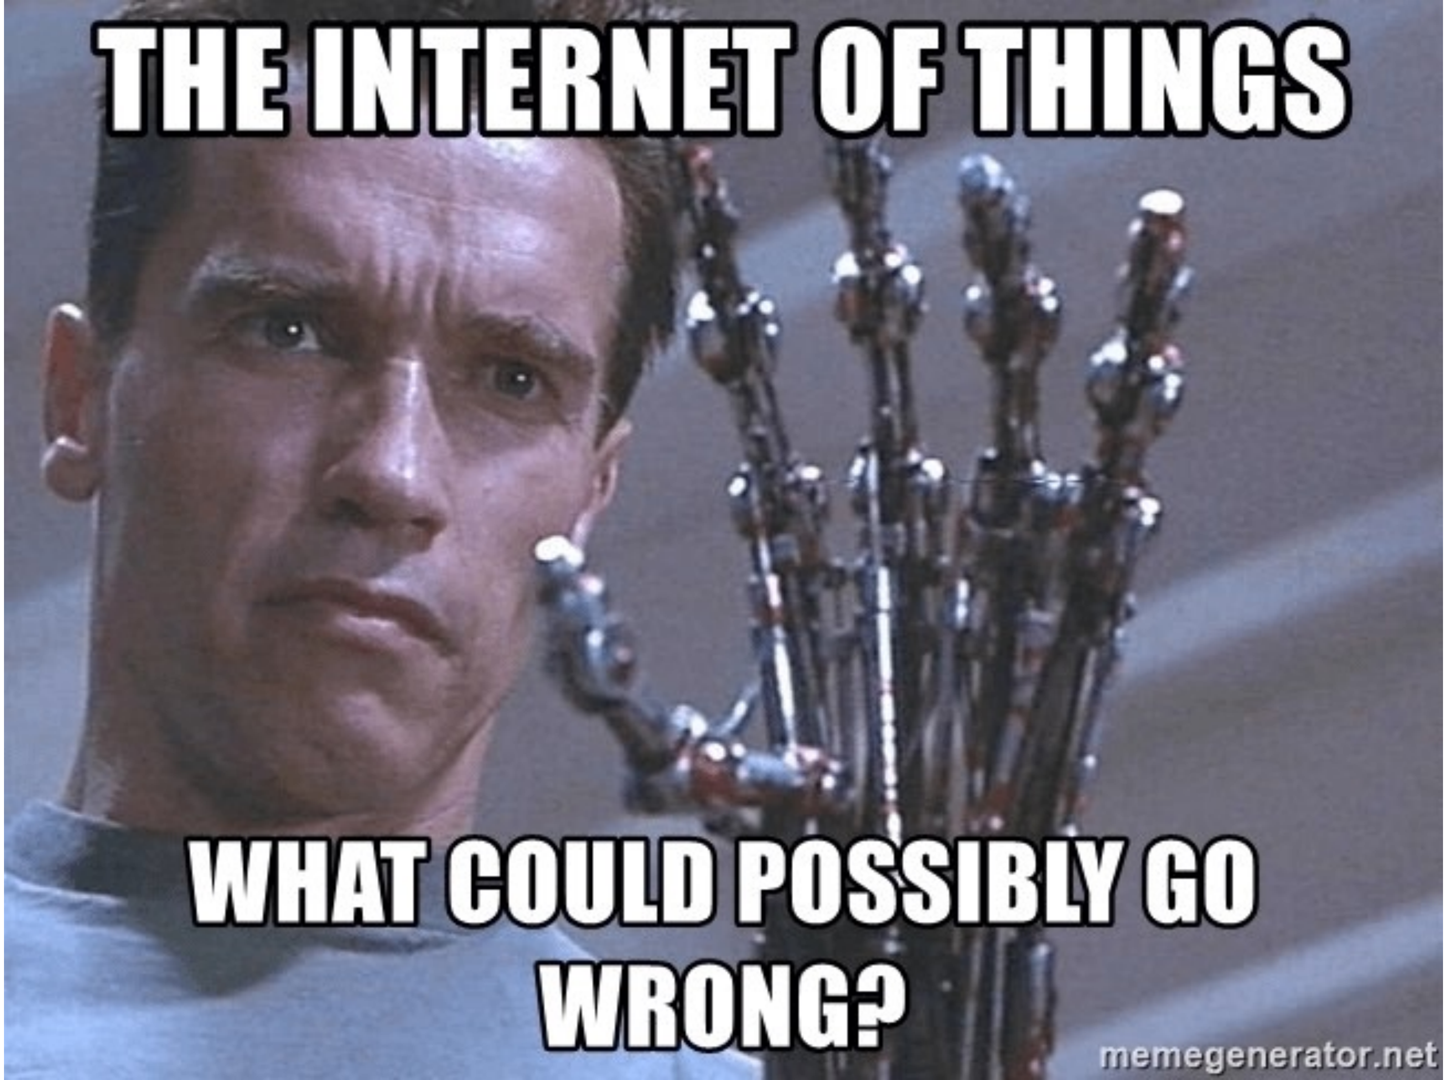
THE INTERNET OF THINGS
WHAT COULD POSSIBLY GO WRONG?
memegenerator.net

Solution at the Edge

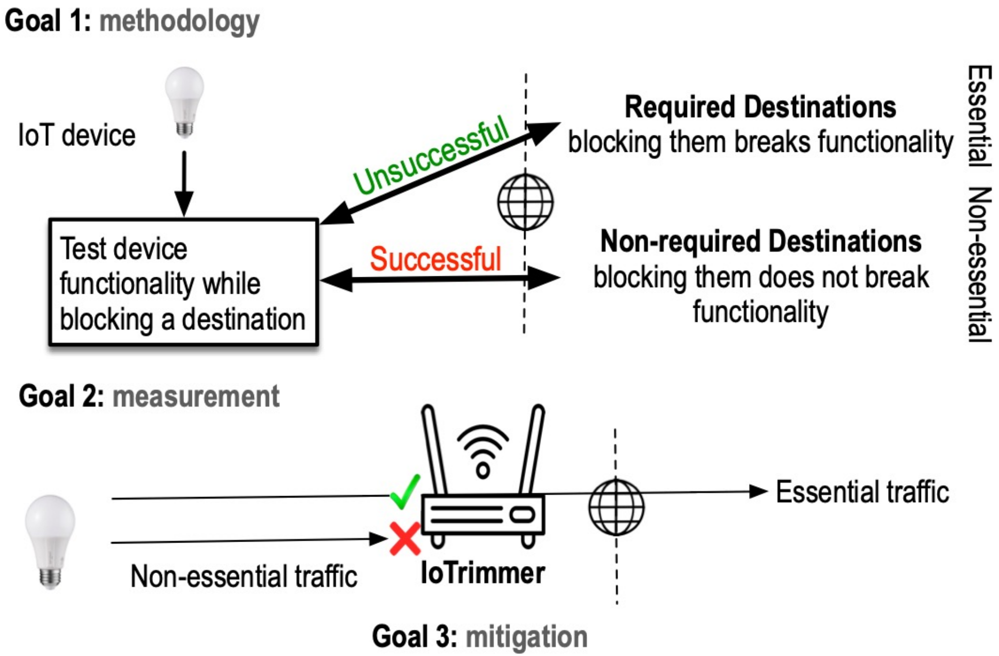

/ Generalizable

/ Self adaptive

/ Accurate IoT blocker

Idea

- What we learn: some IoT traffic is **essential** and some of it is **non-essential**
- Can we (partially) "silence" IoT devices and still be able to enjoy them?

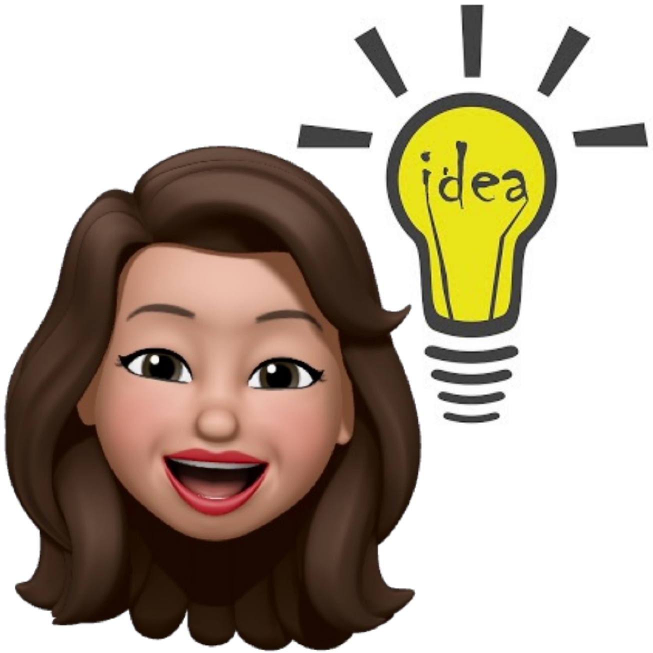

Goals

- *Measurement Methodology*:
  How to **automatically** separate **essential traffic** from **non-essential** traffic?
- *Identification*:
  How **prevalent** is non-essential traffic in our **testbed** of 31 IoT devices?
- *Generalizations*:
  Are there any **common patterns** in non-essential traffic?
- *Mitigation*:
  How to build a **system for filtering** non-essential traffic?

Challenges

- IoT devices are **hard to test** automatically
    - They offer very different functionalities
    - They suffer (in our experience) from frequent service outages that must be detected
    - They typically require user interaction (i.e., they are not directly programmable)
    - Hard to verify if a functionality was actually executed or not
- Ideas:
    - use **companion devices** (phones and voice assistants)
    - use **network traffic patterns** to classify IoT devices responses

Measurement Methodology

Hardware and Software of our IoT testbed

- **IoT devices**
    - 31 in total: 6 cameras, 15 home automation, 5 smart hubs, 3 smart speakers, 2 video
- **Router** with IP filtering and DNS filtering capabilities
- **Power plugs** and scripts to power cycle the devices
- **Trigger scripts** to invoke IoT devices functionality
    - *Companion app interaction* and *voice assistant interaction*
- **Probe scripts** to detect success or failure in functionality execution
    - Compare companion app *screenshots* and identification of *traffic peaks*

Functionality Experiment

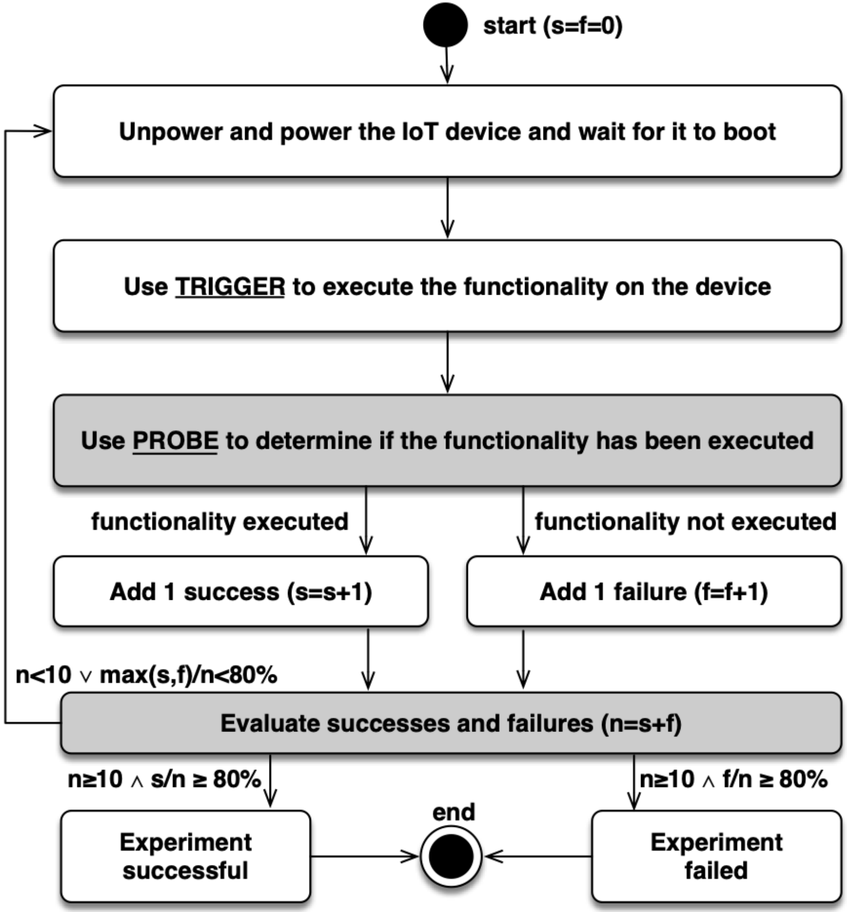

- **Goal: determine if a functionality works**
- Test the functionality at least **10 times**
- Terminate if **80% consensus** is reached
- When tested **30 times** against **ground truth**, probes have been **80% correct**
- If probes are 80% correct, the chance of an incorrect functionality experiment result is **less than 0.01%**

Identifying Non-essential Traffic

Distinguishing Required from Non-Required Destinations

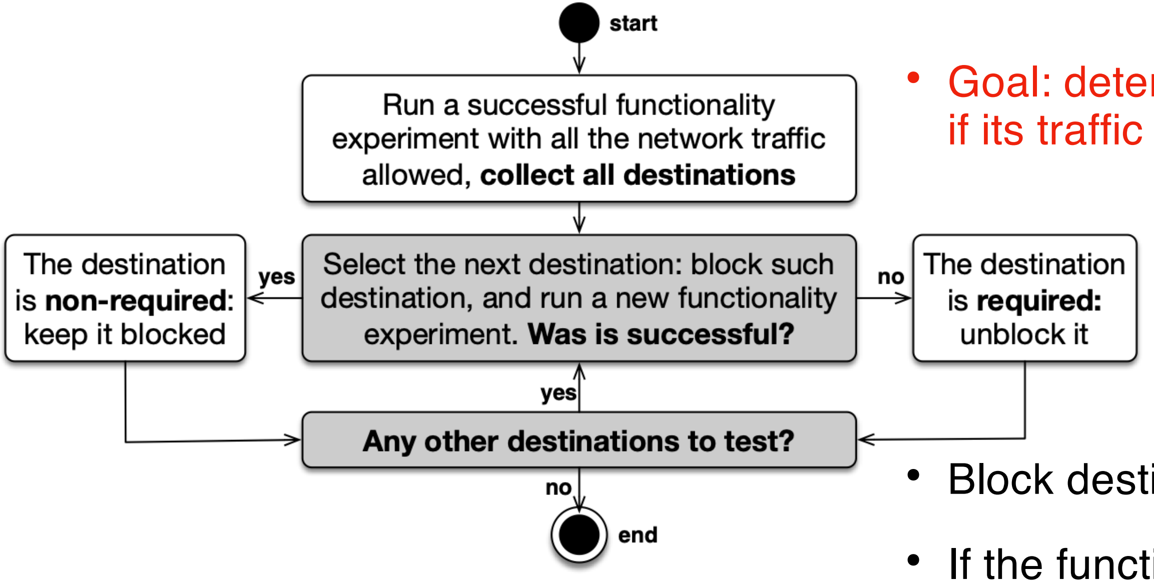

- Goal: determine if a destination is required (i.e., if its traffic is essential)

- Block destinations one by one
- If the functionality succeeds when a destination is blocked, such destination is **non-required**
- Otherwise it is **required**

Overall Results

Devices with at least one non-required destination

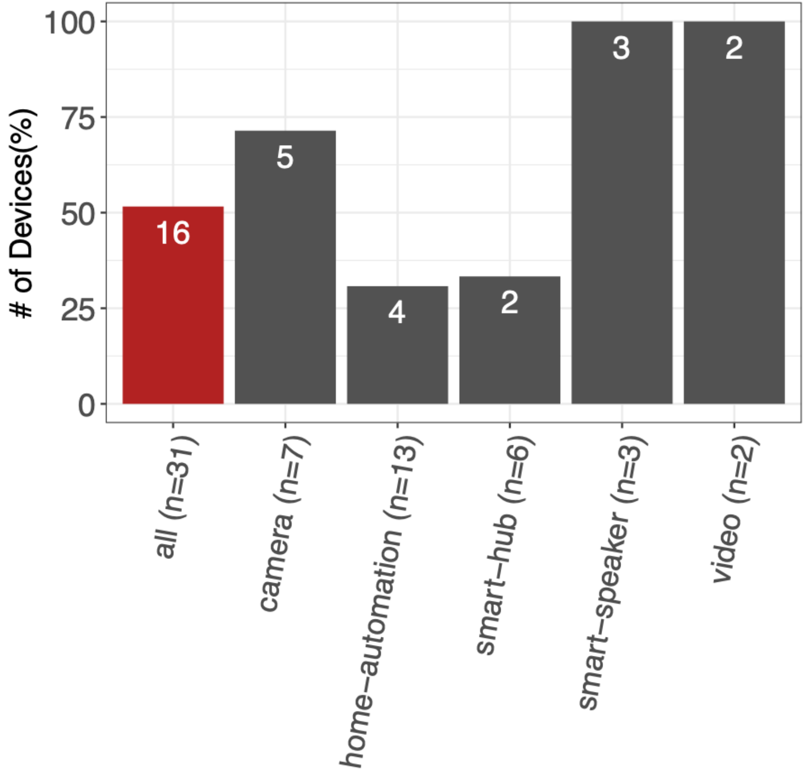

- 16/31 devices have non-essential traffic
- Mostly cameras, smart speakers, and video
- Possible explanations:
    - complexity (skills and apps)
    - uncommon vendors / rebranding (for cameras)

Required *vs.* Non-required Destinations

Device	# of Destinations	Required	Non-Required
Camera			
Blink-camera	2	2	0
Bosiwo-camera	4	2	2
Icsee-doorbell	6	2	4
Reolink-cam	2	1	1
Wansview-cam	9	3	6
Yi-camera	5	3	2
Home-automation			
App-kettle	2	2	0
Honeywell-thermostat	2	2	0
Magichome-strip	1	1	0
Meross-dooropener	1	1	0
Nest-tstat	3	2	1
Netatmo-weather-station	1	1	0
Smarter-coffee-mach	1	1	0
Smartlife-bulb	1	1	0
Smartlife-remote	1	1	0
Sousvide	1	1	0
Switchbot	1	1	0
Tplink-bulb	4	1	3
Tplink-plug	4	1	3
Wemo Plug	2	2	0
Xiaomi-ricecooker	7	3	4

Device	# of Destinations	Required	Non-Required
Smart-hub			
Insteon-hub	1	1	0
Lightify-hub	3	3	0
Philips-hub	4	2	2
Samsung Hub	3	2	1
Sengled-hub	2	2	0
Smart-Speaker			
Allure-speaker	3	1	2
Echodot	10	3	7
Google-home	9	4	5
Video			
FireTV	14	3	11
Roku TV	10	2	8
Total	*119*	*57*	*62*

- Non-required destinations are contacted the most by cameras, speakers, and video devices
- But it also happens on simpler devices such as the TP-Link smart plug and smart bulb

Amount of Data Sent During One Experiment

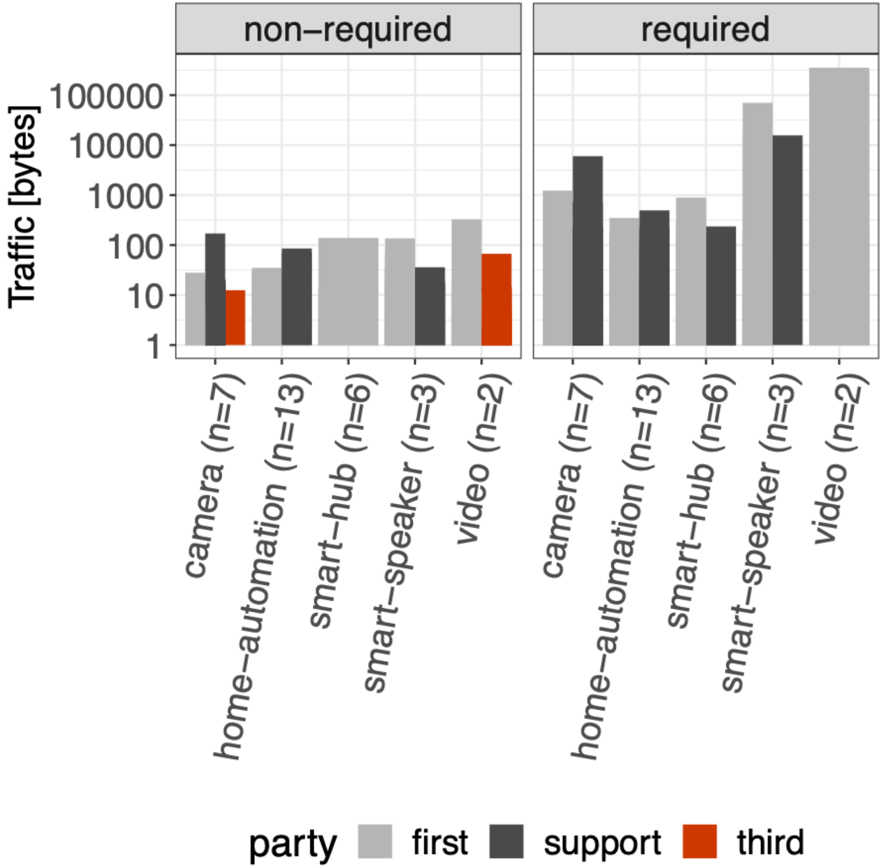

- **Good news**: non-essential traffic is relatively small (less than 1KB/device)
- However, it is still possible to transmit:
  - Presence of the device
  - Its status
  - Basic data from the sensors (e.g., open/close, motion/still, alarm/no alarm)

Similarities with Existing Blocklists

- We consider Pi-hole, Firebog, MoAB, StopAD lists
- No required destinations on such lists
- Up to 6 out of 62 non-required destinations present in existing blocklists
- Public blocklists are of limited help in blocking IoT non-essential traffic

Number of non-required destinations present in public blocklists

Device	Non-req Dest.	Pi-hole	Firebog	MoAB	StopAd
Allure Speaker	2	0	0	0	0
Bosiwo Camera	2	0	0	0	0
Echo Dot	7	1	1	0	0
Fire TV	11	2	3	1	0
Google Home	5	0	0	0	0
Icsee Doorbell	4	0	0	0	0
Nest Thermostat	1	0	0	0	0
Philips Hub	2	0	0	0	0
Reolink Camera	1	0	0	0	0
Roku TV	8	1	2	1	0
Samsung Hub	1	0	0	0	0
TP-Link Bulb	3	0	0	0	0
TP-Link Plug	3	0	0	0	0
Wansview Camera	6	0	0	0	0
Xiaomi Ricecooker	4	0	0	0	0
YI Camera	2	0	0	0	0

Mitigating Non-essential IoT Traffic

- A blocking system: **IoTrim**
    - Filtering router between the IoT devices and the Internet
    - Block/allow lists based on (non-)required destinations → crowdsourced
    - Software to declare device types and manage the lists / blocking rules
    - **A proof-of-concept prototype is available for download**

Course Overview

- ❑ Benchmarking privacy in IoT devices
- ❑ IoT devices identification
- ❑ Benchmarking security in IoT devices
- ❑ Benchmarking security solutions for IoT devices
- ❑ Privacy solutions for IoT devices at the edge
- ❑ **Security solutions for IoT devices at the edge**
- ❑ IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security

Motivation

- Inefficiency of existing IoT solutions
- Most of them are cloud-based: might share users' personal/sensitive data

Research Questions

- Can we replace cloud-based IoT protection systems by a local IDS/IPS running on a home router?
- If so, what is the performance overhead?

Benefits

- **Security improvement:** cover wider spectrum of IoT threats in a home network
- **Privacy improvement:** All users' data processed locally and not shared with cloud

SunBlock Architecture

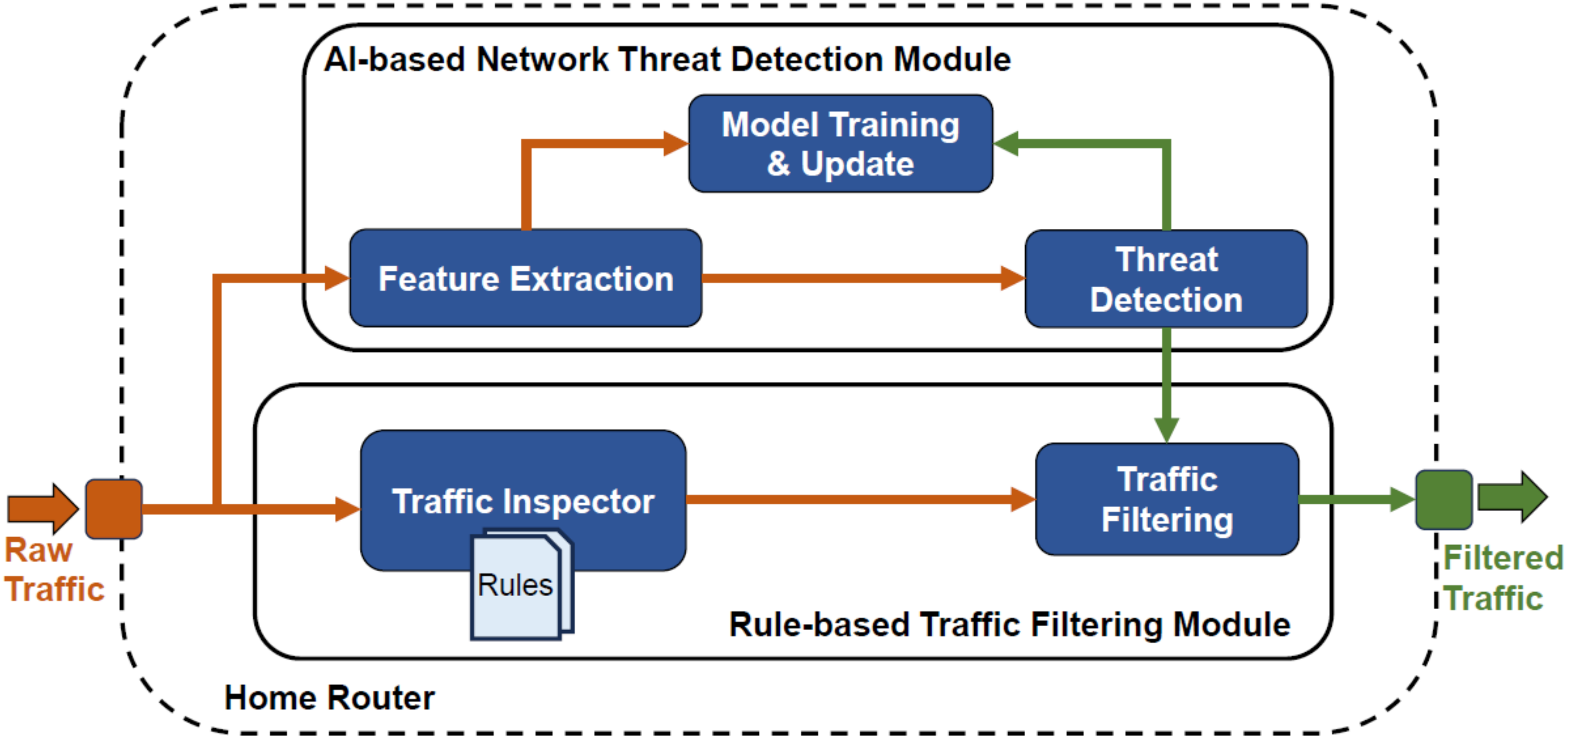

Implementation: home router with IoT protection

- LinkSys WRT3200ACM, OpenWRT Linux-based OS
- ~4GB flash, 512MB swap (for ML training only), 512 MB RAM
- Snort3 for rule-based filtering, netml with OCSVM for AI-based module

Testbed

- 10 most popular IoT device types (according to IoT Inspector paper)
- Smart speakers (Echo spot, Google Home), Video (FireTV), Camera (Yi, Blink), Home automation (Nest thermostat, TP-Link/Wemo plugs, Gosund/TP-Link bulbs)
- Devices were triggered daily using the methodology similar to the S&P paper

Evaluation: threat coverage and prevention time

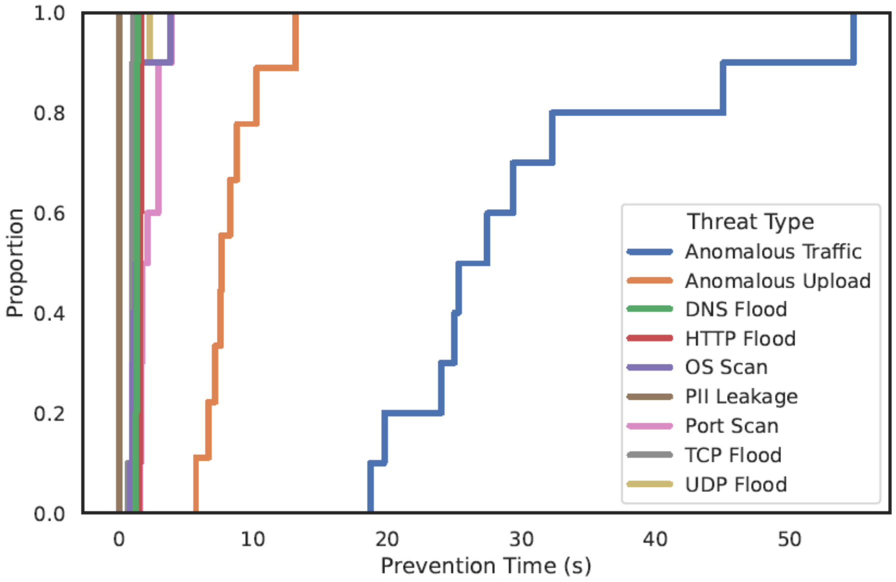

Threat	IoT Protection Systems	SunBlock
Anomalous Traffic	✗	✓
Anomalous Upload	✗	✓
SYN Flooding	✓	✓
UDP Flooding	✗	✓
DNS Flooding	✗	✓
HTTP Flooding	✓	✓
Port Scanning	✓	✓
OS Scanning	✓	✓
PII Leakage	✗	✓

Evaluation: performance overhead

Model training

Protection Level	CPU (%)	RAM (MB)	swap (MB)	Training Time (s)
Rule-based & AI-based	18 ±3	444 ±4	296 ±21	924 ±253
AI-based only	26 ±2	442 ±6	197 ±28	429 ±171
Rule-based only	32 ±4	423 ±9	132 ±20	180 ±22
Unprotected	39 ±2	410 ±3	55 ±1	113 ±10

HTTP flood

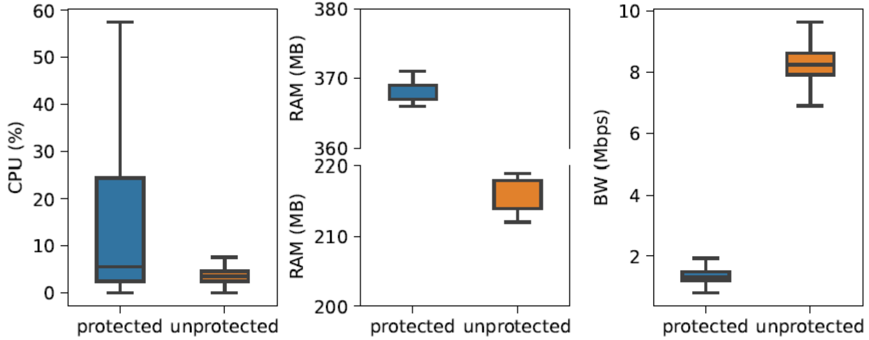

Regular IoT traffic

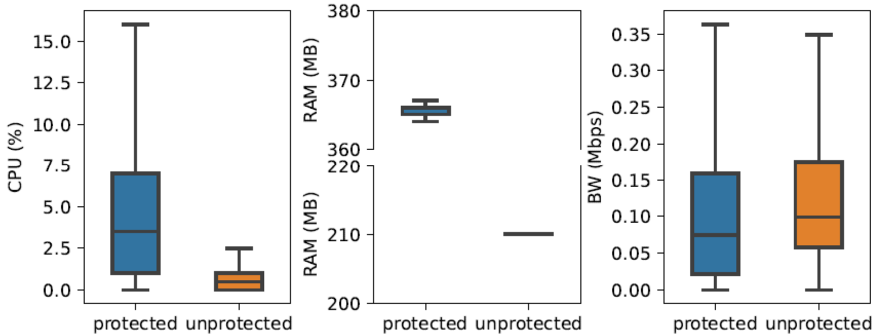

UDP flood

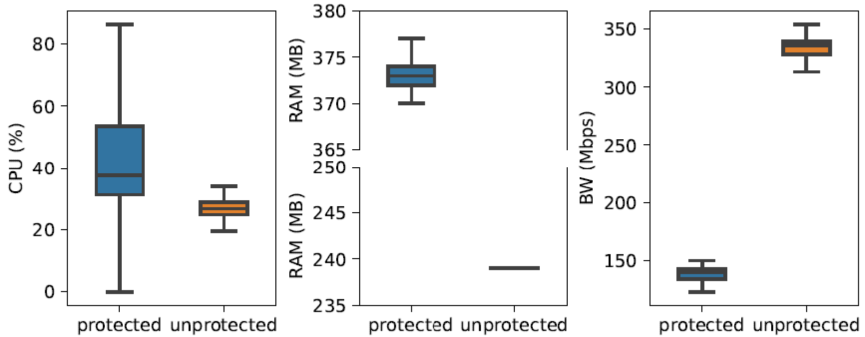

Takeaways

- IoT threats can be rapidly detected on a home router with Rule&AI-based filtering algorithms

- No need in cloud-based solutions and in sharing your personal data

- Increase in CPU and RAM doesn’t affect main router functions leaving plenty of free resources: >50% free CPU and ~30% free RAM

- Further plans: beta testing and precise performance benchmarking against existing IoT solutions


OUR DEVICES ARE NOW 100% SECURE.
I TURNED THEM ALL OFF.
HOW DID YOU DO THAT?
© D.Fletcher for CloudTweaks.com

Strengthening the IoT Ecosystem

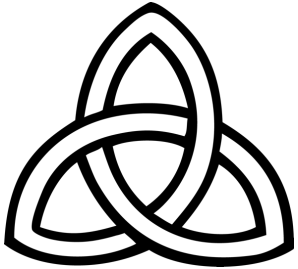

Trust

- Endpoints’ practices
- Trusted platform modules
- Domain-specific guidelines and frameworks
- Access networking system & machine learning

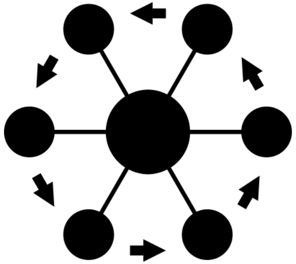

Interconnectivity

- Understand threats in real world scenario
- Inferences on crowdsourced IoT data
- New secure IoT (wireless) networking protocols & systems
- Privacy preserving technologies at the edge

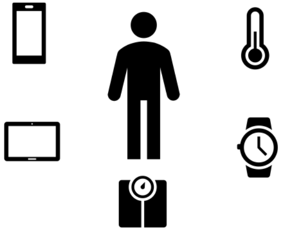

Awareness, Authentication & Management

- Usable monitors for IoT
- Context-aware privacy
- Personalised privacy

Is Your Kettle Smarter Than a Hacker?

- **Assessing Replay Attack Vulnerabilities on Consumer IoT Devices using AI**
    - Automated methodology for large-scale testing replay attack vulnerabilities on IoT devices
    - Using AI for detecting the success of the attack

Methodology

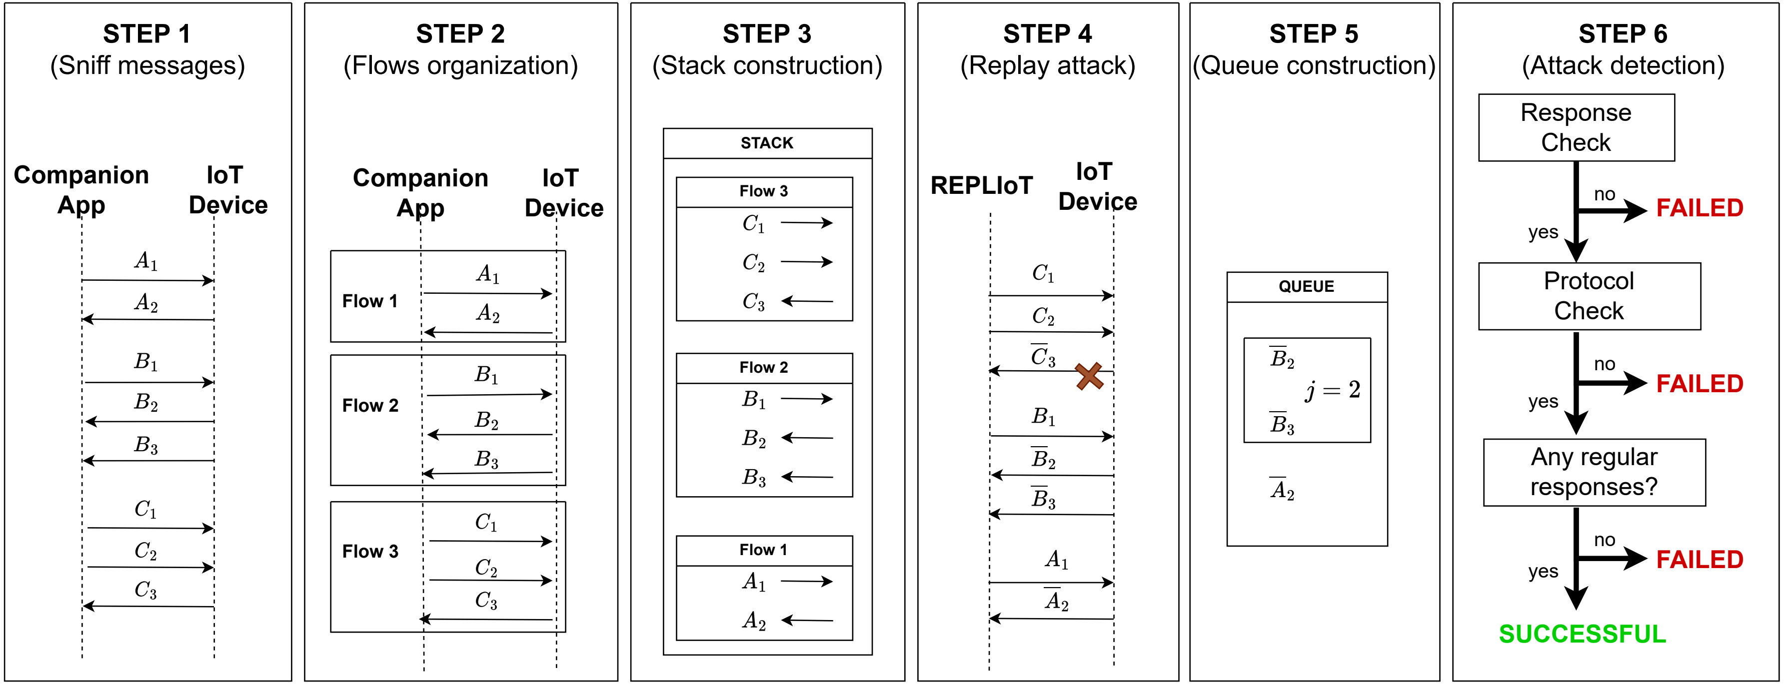

Methodology

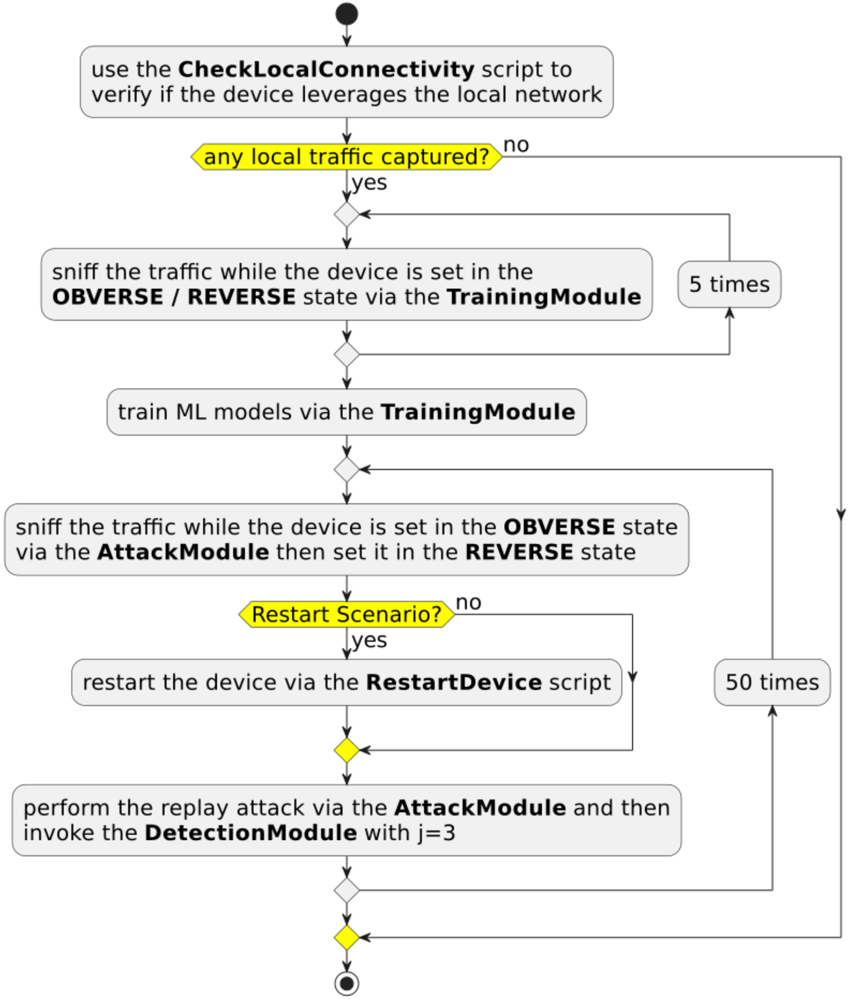

Results

REPLAY ATTACK RESULTS. ✓ INDICATES WHETHER THE REPLAY ATTACK IS SUCCESSFUL OR NOT (*X*).

Device (*Tested via APIs)	**Non-Restart Scenario**	**Restart Scenario**
Yeeligth lightstrip	✓	✓
Yeelight bulb	✓	✓
Wiz ligthbulb	✓	✓
Lifx bulb	✓	✓
Lepro bulb	✓	✓
Govee lightstrip *	✓	✓
Nanoleaf triangle *	✓	✓
Tapo smartplug	✓	*X*
Meross smartplug	✓	✓
WeeKett Kettle	✓	✓
Eufy robovac 30C	✓	✓
OKP vacuum	✓	✓
iRobot roomba i7	*X*	*X*
Sonos Speaker *	✓	✓
Bose Speaker *	✓	✓
Wyze cam pan	*X*	*X*
Vtech baby monitor	*X*	*X*
Boyfun Baby monitor	*X*	*X*
Furbo camera	*X*	*X*
Meross Garage Opener	✓	✓

Responsible Disclosure

[TP-Link Support]-[TKID231119229] Responsible disclosure

To: vincenzo.deangelis@dimes.unical.it, Cc: Mandalari, Anna Maria, Francesco Buccafurri & 2 more

Details

⚠ Caution: External sender

Many thanks for your valued reply.

After confirming with our security team, the vulnerability has been resolved in the latest firmware of P110.

Since this firmware is currently in gray release, we are not sure whether your P110 could receive the firmware right now, you could check it in Tapo App.

If there is no firmware update for your P110, please provide the MAC address with us, we will release the firmware for your P110 and hope you could help verify the remediation work in the latest firmware.

If you have additional concerns or information, please feel free to let us know. If you have any subsequent plan or processing of the vulnerability, we also hope that you can further synchronize to us.

Thank you for your cooperation and patience.

Ian.xu

TP-Link Technical Support

Website: https://www.tp-link.com/support/

Feedback: Report a suggestion/complaint on this email service by clicking here

Why Were We Interested in This?

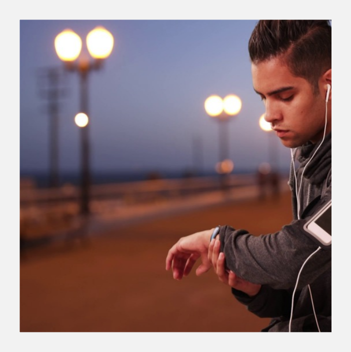

Control

Device detection

Intelligent profiles

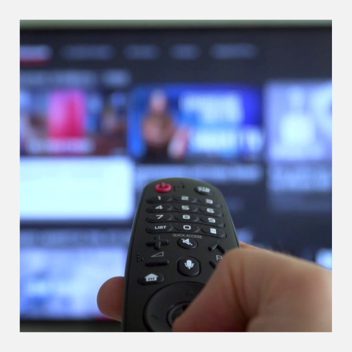

Security

Vulnerability Assessment

Brute Force Protection

Anomaly Detection

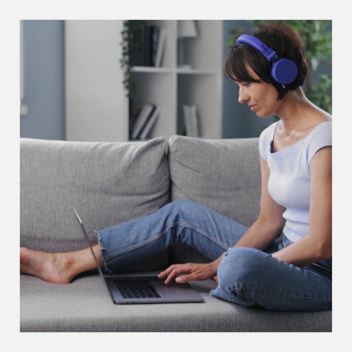

Privacy

Content filtering

Network Intrusion Prevention

- These devices may introduce privacy and security risks.
- Their cloud interactions and data collection operations may introduce privacy risks.

Aim and Contribution

- ❑ **Goal 1:** Develop a system abled to inject realistic anomalies for healthcare IoT devices.

- ❑ **Goal 2:** Explore how the time window used for training affects the accuracy of the anomaly detection, for three different types of anomalies.

- ❑ **Goal 3:** Demonstrated that training the model at the edge of the network on a representative edge device (Raspberry Pi) is feasible.

PRISM

Challenges for Measuring IoT Devices

Difficult to automate the testing of commercial IoT safeguards

- Closed systems
- Blackbox approach

Difficult to perform IoT experiments and generalize

Our contribution: a system for injecting and detecting IoT anomalous behavior in real-world scenarios (software and anomaly data available online).

- Lack of standard testbed

Dataset

- Collected by the UK Dementia Research Institute and Technology Centre (UK DRI).

- In-home activity of **people living with dementia** (PLWD), from motion sensors, wearable devices and physiological measurements.

- **44 different households**, each fitted with **22 IoT devices**.

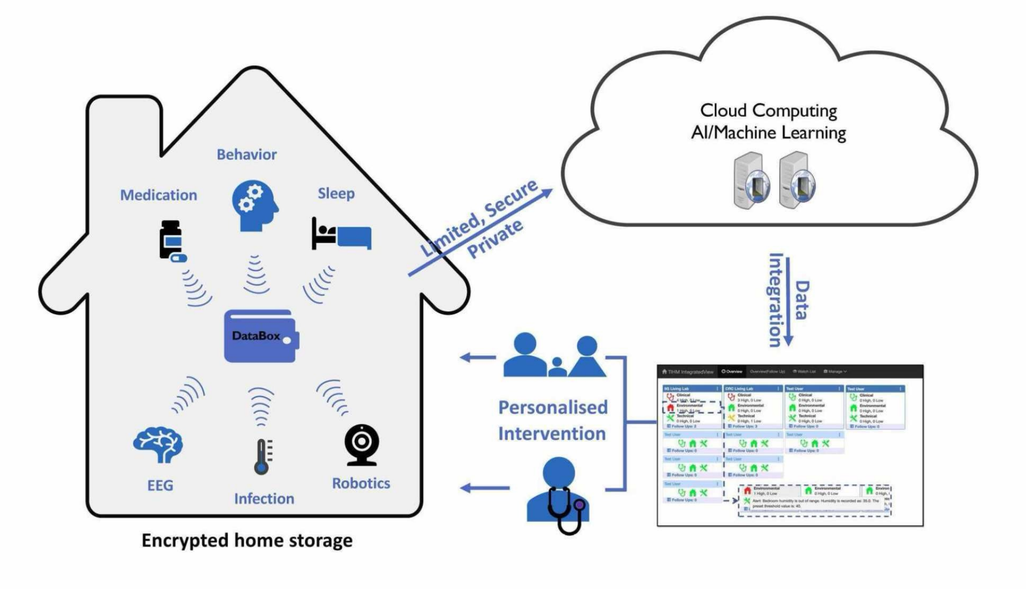

Dataset

Function	Format	IoT Device	Continuous
Location	**Binary**	WC, bathroom, bedroom, corridor dining room, hallway - kitchen, living room, lounge office, study	-
Door	**Binary**	back door, conservatory fridge door, front door garage, main door secondary, utility	✓
Appliances	**Binary**	iron use, kettle use, microwave use - socket use, toaster use	-
Temperature	**Float**	temperature, body temperature skin temperature	✓
Health Related	**Float**	blood pressure, body mass index body muscle mass, body weight - heart rate, body fat body water, bone mass	-
Light	**Integer**	light level	✓
Sleep Event	**Binary** **Float** **Integer**	sleep event, sleep mat snoring sleep mat heart rate ✓ sleep mat respiratory rate sleep mat state, agitation	✓ -

Threat Model

Victim: A person that uses a healthcare IoT device.

Adversary: Any party with access to the IoT device Traffic.

Threat:

- Adversaries may be incentivized to share privacy-sensitive information of patients.
- Malicious attacks hijack the communication channel, modifying the data sent by the IoT device.

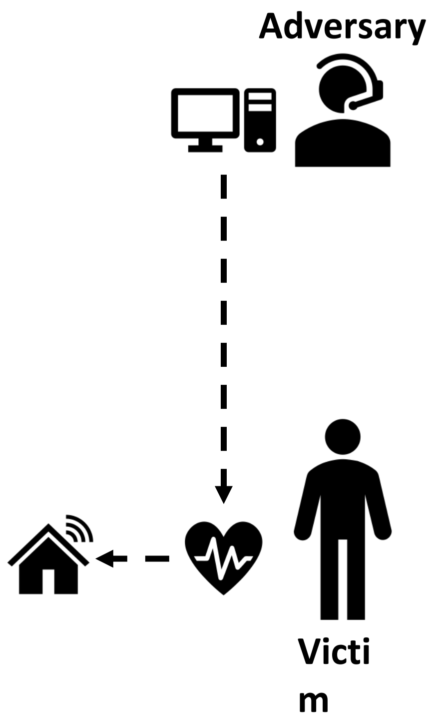

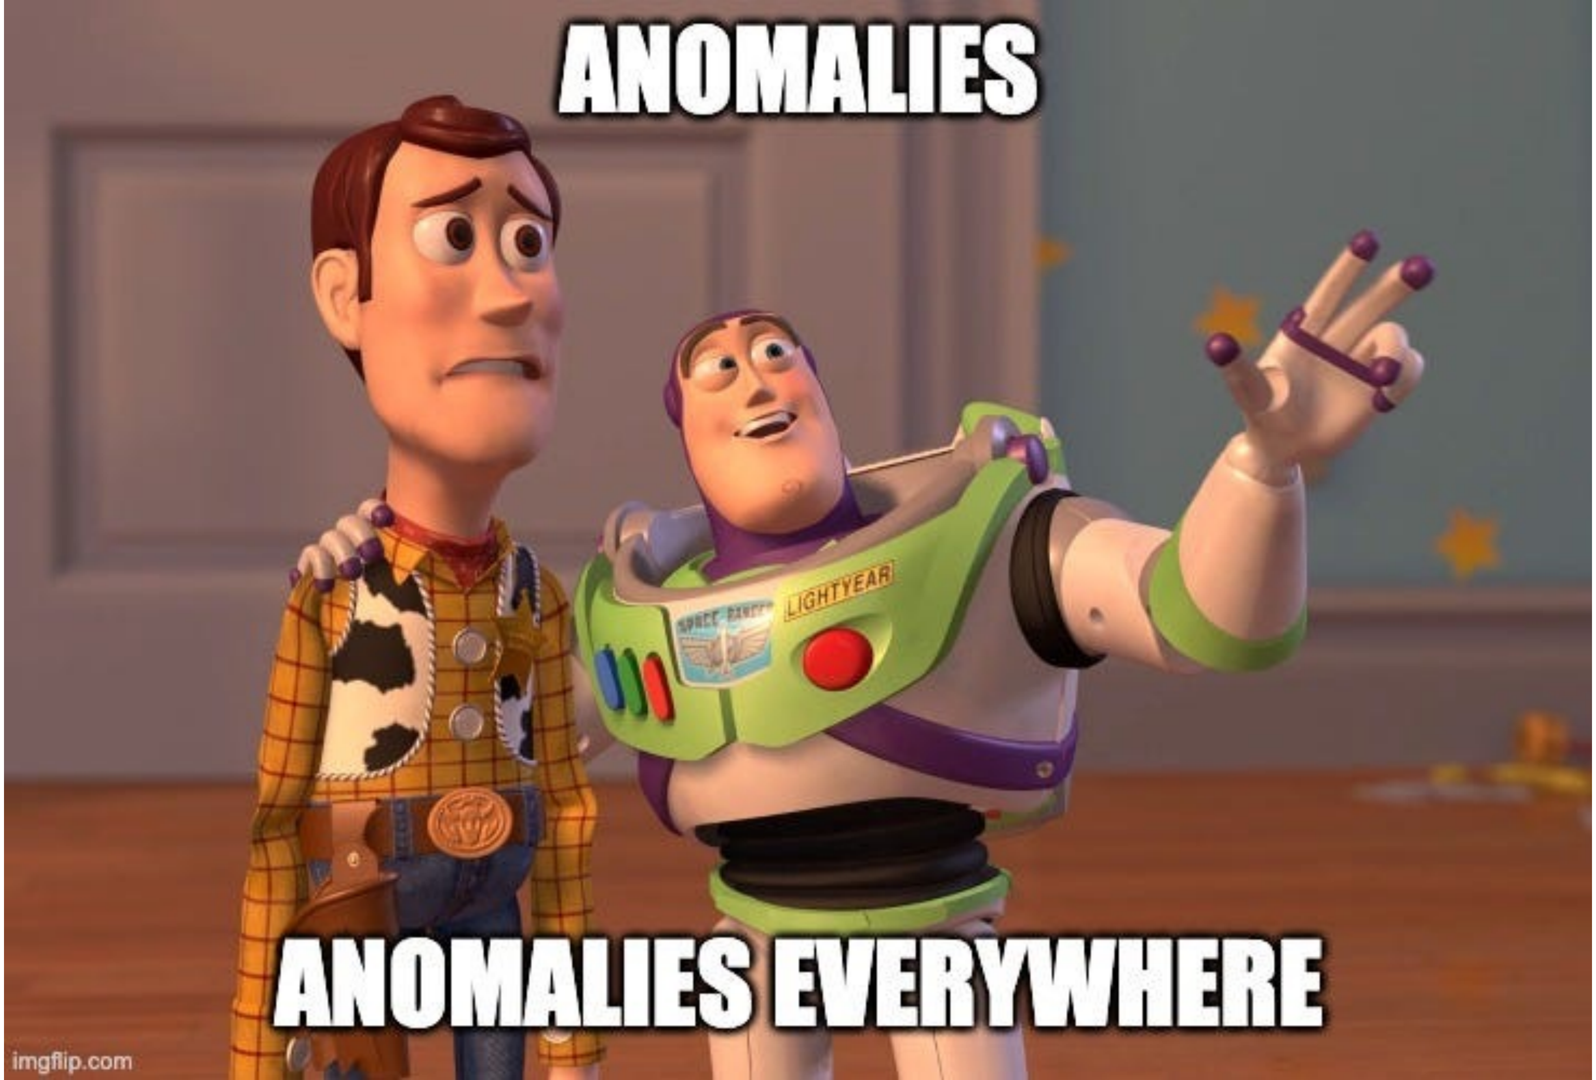
ANOMALIES
LIGHTYEAR
ANOMALIES EVERYWHERE
imgflip.com

Types of anomalies

On-off:

- For Binary sensors (i.e switches, doors)
- Recreates a sensor which repeatedly switches on and off.

Variance:

- For sensors which record floats or integers (i.e thermometer, blood pressure)
- Recreates noise or randomized readings.

Spike:

- For sensors which record floats or integers
- Recreates a random abnormal spike in the readings

Anomaly	IoT device
On-Off	Room Location, Appliance Use, Sleep Event
Variance	Ambient Temp, Body Temp, Light
Spike	Sleep Respiratory, Hearth Rate, Sleep Hearth Rate

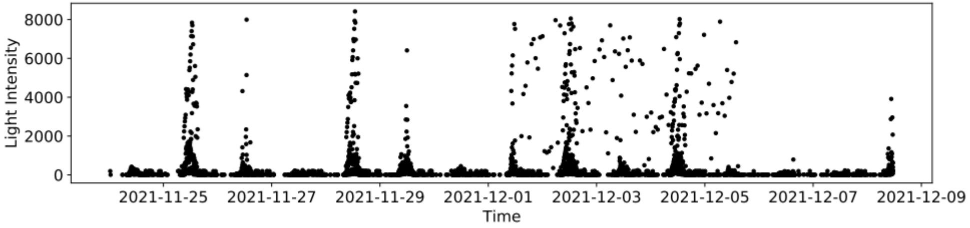

“Variance anomaly injected in the Light intensity sensor for 4 days ”

System Design

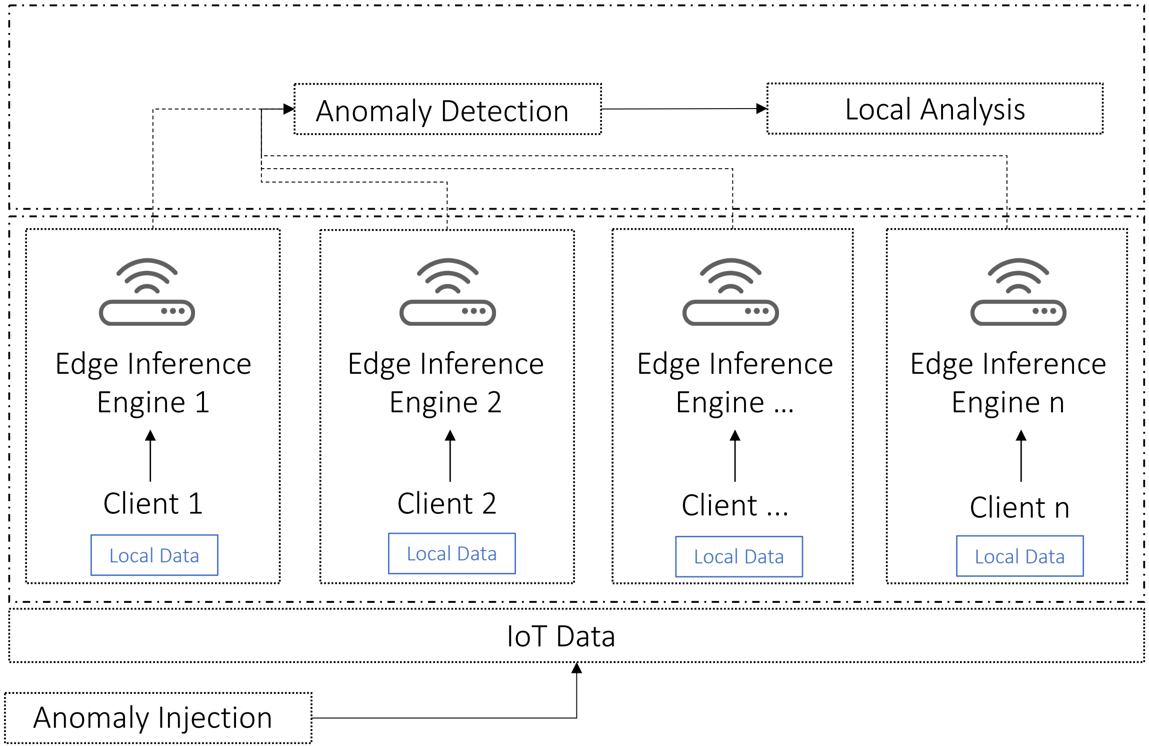

Overview of Methodology

Data Loader

Options:

- Select specific patients.
- Select one or multiple patients.

Anomaly Injection

Options:

- Type of anomaly.
- Number of anomalies.
- Time window length that anomalies will be injected in.

Data Pre-Processing

- Sliding time window.
- Train, Valid. split based on time window.
- Features: Sensor readings, Time interval between readings (δt).

Model Inference

- Architectures tested: DNN, CNN, KNN.
- Library used: PyTorch.
- Early stopping for training.
- Unsupervised Learning.

Anomaly Detection

- The average training loss of the final epoch is calculated.
- Multiplied by a coefficient, it acts as the threshold to detect anomalies.

Anomaly Detection Accuracy

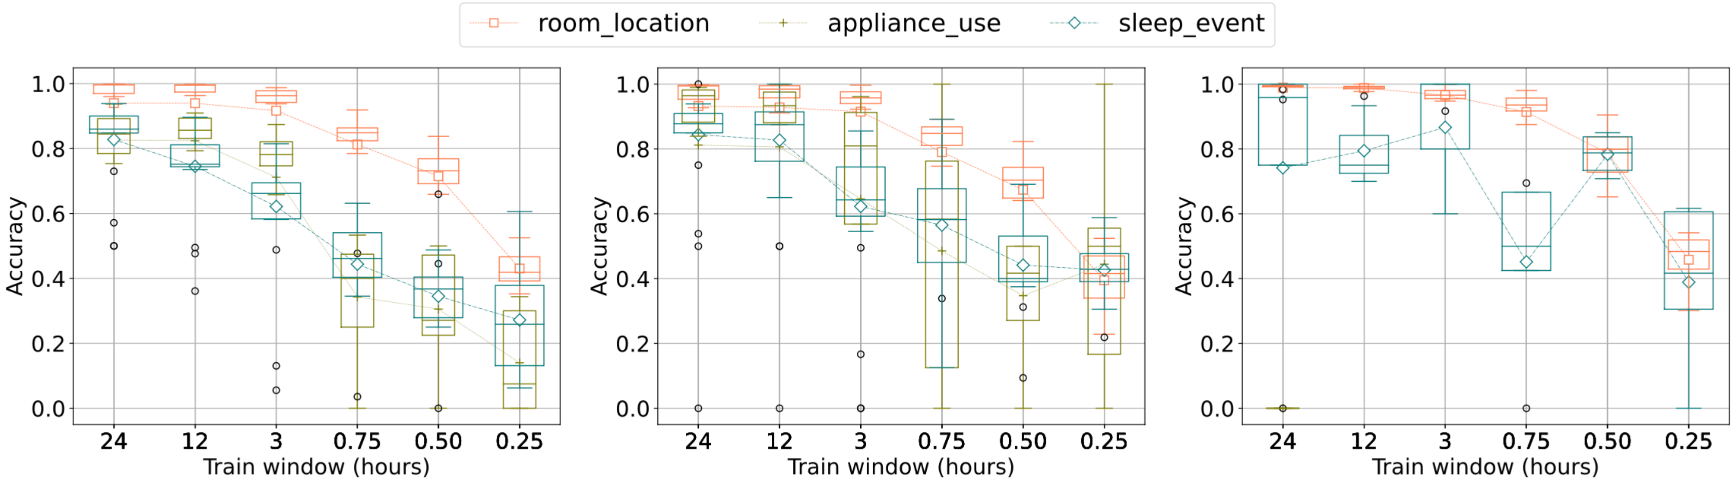

(a) 24h (b) 3h (c) 15min

Take away: On-Off Anomaly. The anomaly detection accuracy changes with training window size and different validation window sizes.

Personalized Models

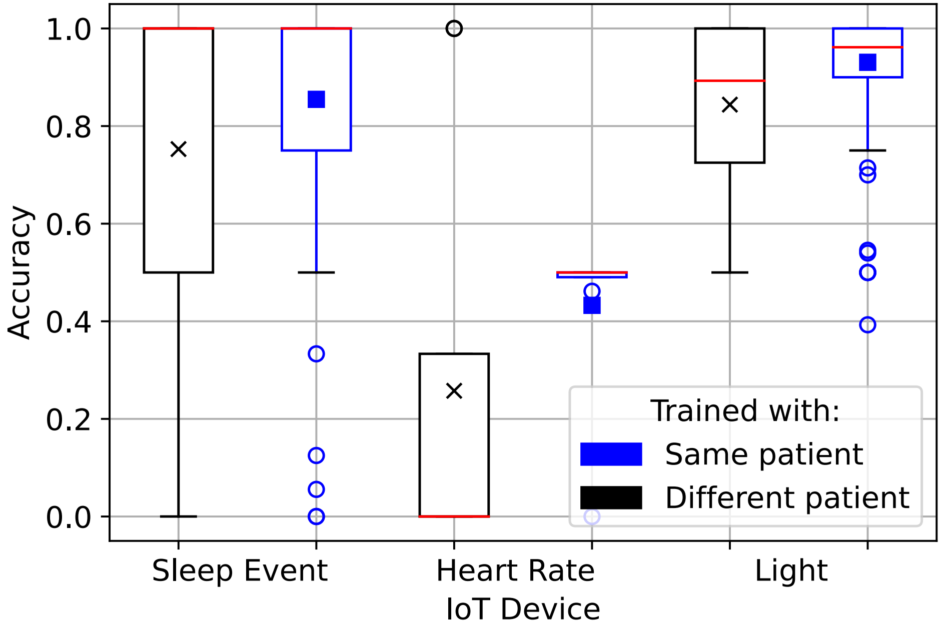

Average accuracy across all patients while training and validating with the same and different patients.

Take away: A model updated using data from one patient does not perform well on another patient and vice versa.

Personalized Models

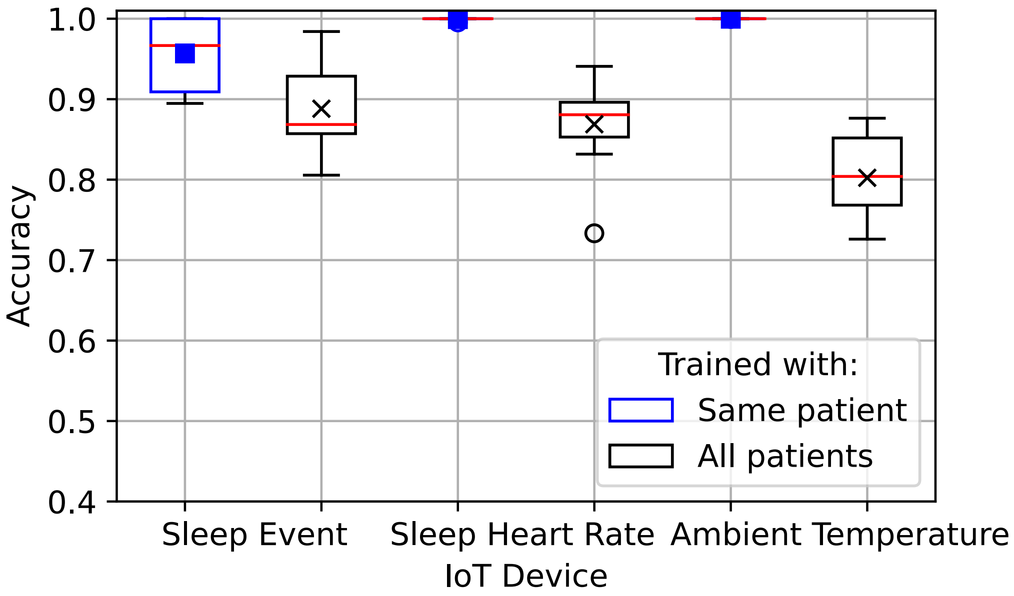

Average accuracy across all patients while training with all patients and validating with one patient, compared to training with all and validating with one patient.

Take away: The accuracy decreases when training the model with all patients. This shows that a model updated with data specific to each patient will achieve better performance.

Course Overview

- ❑Benchmarking privacy in IoT devices
- ❑IoT devices identification
- ❑Benchmarking security in IoT devices
- ❑Benchmarking security solutions for IoT devices
- ❑Privacy solutions for IoT devices at the edge
- ❑Security solutions for IoT devices at the edge
- ❑IoT devices certification scheme

The Problem

- 21.5 billion IoT devices in the world
- They have access to user private information
- They are a threat for user privacy and security

Mitigation

- Regularly train the ML models at the edge to keep up with the changes in device usage trends
- Approaches that rely on local traffic analysis: edge-based solutions running on the home gateway

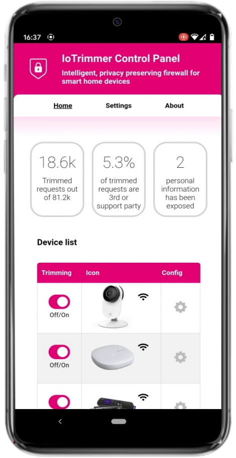

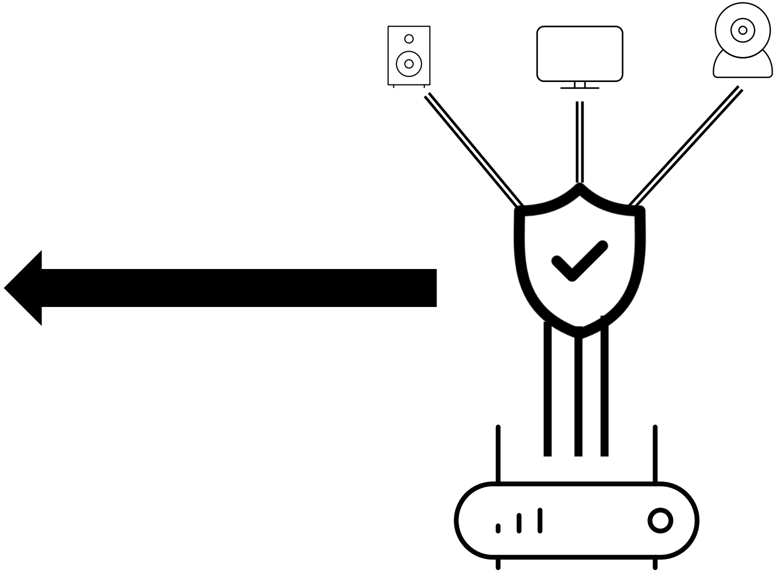


WHEN YOU SEE HOW SOPHISTICATED CYBERTHREATS HAVE BECOME
imgflip.com

COPSEC: Compliance-Oriented IoT Security and Privacy Evaluation Framework

Cybersecurity guidelines* such as ENISA, NIST, *IoT Regulation Policy (UAE)* have been released for improving IoT design practice

Privacy regulations** such as GDPR (in EU) and CCPA (in California)

There is a lack of understanding whether IoT devices comply with them

***NOT mandatory**
****Mandatory**

Motivation

- In 2023 the Cyber Resilience Act (in EU) and the US Cyber Trust Mark (in US) make further step towards a certification program of smart devices
- For consumer IoT devices, the certification process is thought as a <u>self-assesement</u> performed by the vendors themselves
- Should we trust vendors?

Methodology

Select security guidelines and privacy regulations → Turn them into metrics → Define experiments to test the extracted metrics on IoT devices → Produce a certification label for the tested device

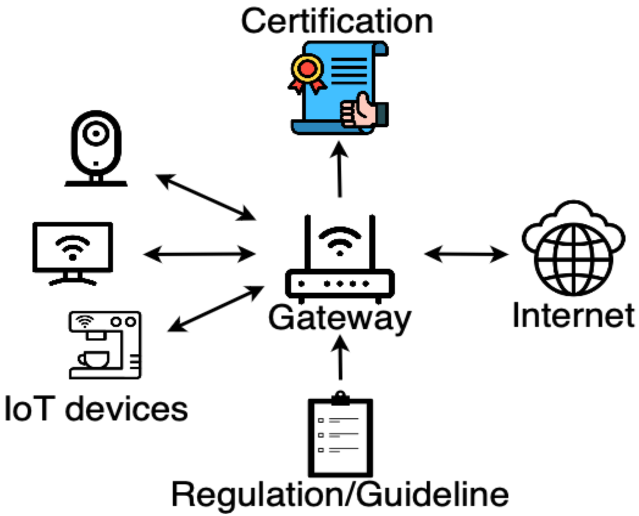

Results

Device	# of Unused Open Ports	# of Unrecognized Protocols	Compliant with GDPR Art. 32 (a)
Bose Speaker	✗ (11 ports)	✓ (0 protocols)	✓
Echo Dot 5	✗ (5 ports)	✗ (3 protocols)	✓
Furbo Dog Camera	✓ (0 ports)	✗ (1 protocol)	✓
Google Nest Cam	✗ (3 ports)	✗ (1 protocol)	✓
Govee lights	✓ (0 ports)	✓ (0 protocols)	✓
Ring Video Doorbell	✓ (0 ports)	✗ (2 protocols)	✓
Sensibo Sky Sensor	✓ (0 ports)	✓ (0 protocols)	✓
SimpliSafe Cam	✗ (1 ports)	✓ (0 protocols)	✓
Sonos One	✗ (5 ports)	✗ (1 protocol)	✗ (mac in the clear)
WeeKett Kettle	✗ (1 ports)	✗ (2 protocols)	✓

IoTrim

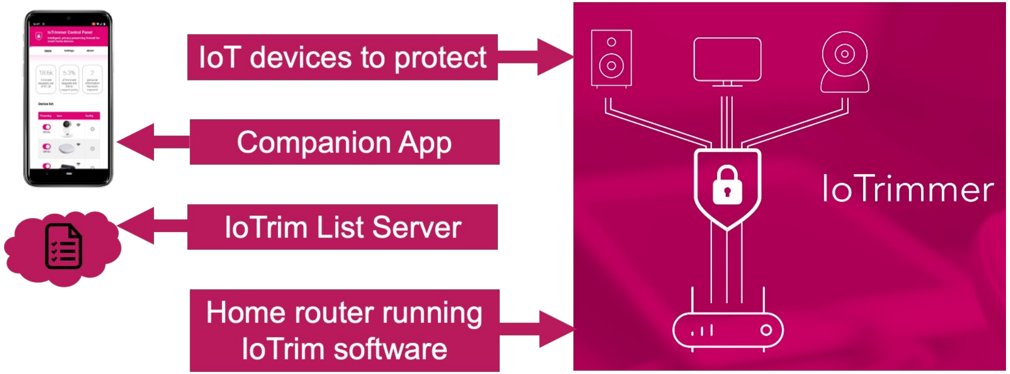

IoT devices to protect
Companion App
IoTrim List Server
Home router running
IoTrim software
IoTrimmer

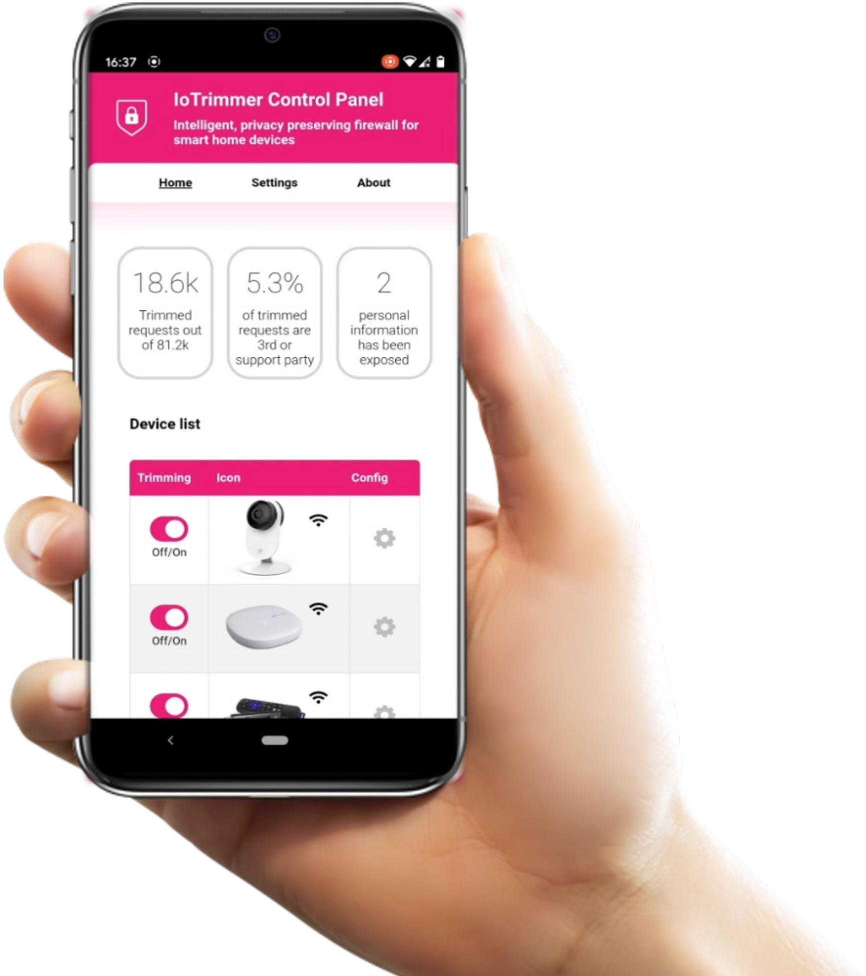

16:37
IoTrimmer Control Panel
Intelligent, privacy preserving firewall for smart home devices
Home
Settings
About
18.6k
Trimmed requests out of 81.2k
5.3%
of trimmed requests are 3rd or support party
2
personal information has been exposed
Device list
Trimming
Icon
Config
Off/On
Off/On

What’s Next?

Privacy Preserving IoT Security Management

- Real industrial gateway
- Medical IoT Devices
- Real-world trial

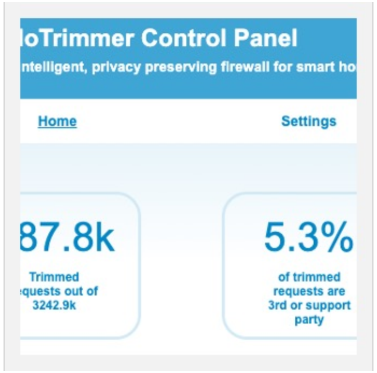

Mitigation

- Real deployment and evaluation
- Third party certification

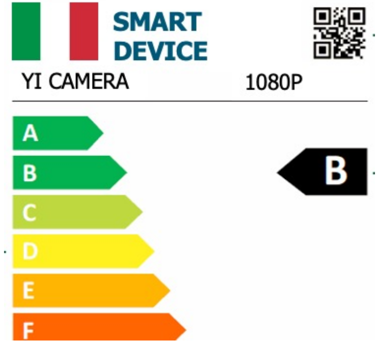

Privacy and Security Label/Certification

- Privacy and security by default


GET INTO IOT THEY SAID
38
IT WILL BE FUN THEY SAID
makeameme.org

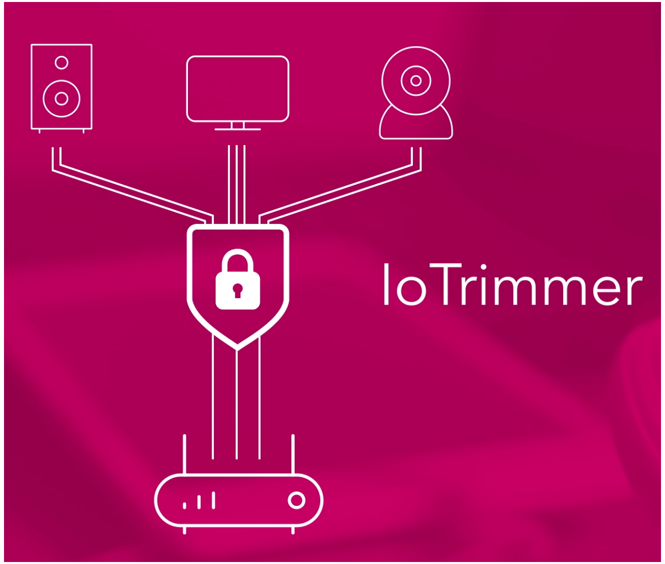

IoT Hacking Lab

Follow us

Twitter: @iotrim @ammandalari

https://youtu.be/mMAH5UhEfxQ

https://youtu.be/P9AyJsMnX88

annamandalari.com