Consumer IoT Devices: Privacy Implications and Mitigation

Anna Maria Mandalari

Imperial College London

Based on joint works with:
H. Haddadi, Roman Kolcun, D. Dubois, D. Choffnes

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)

y may know what
atch (e.g., smart TVs)

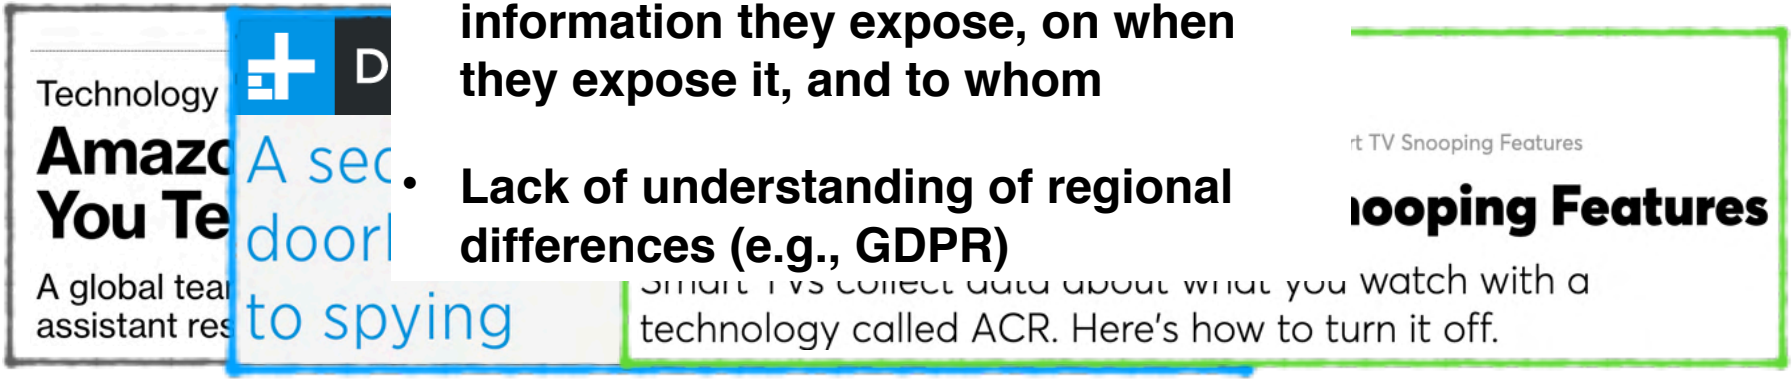

123 devices in two different countries

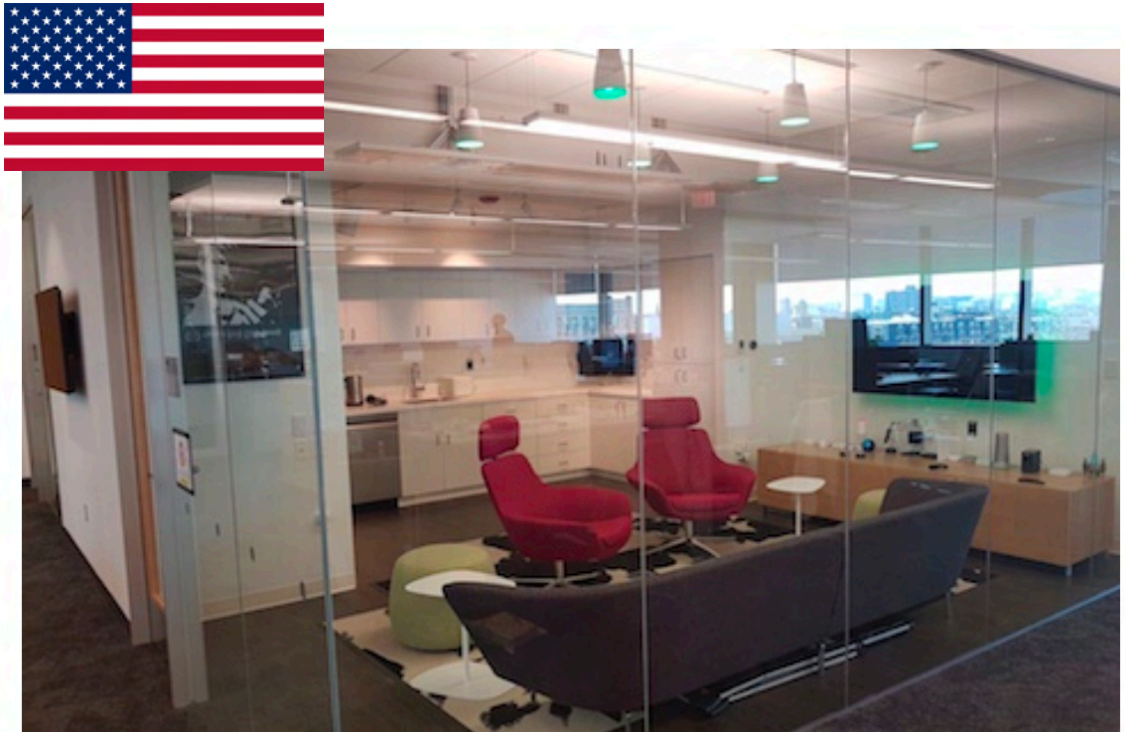

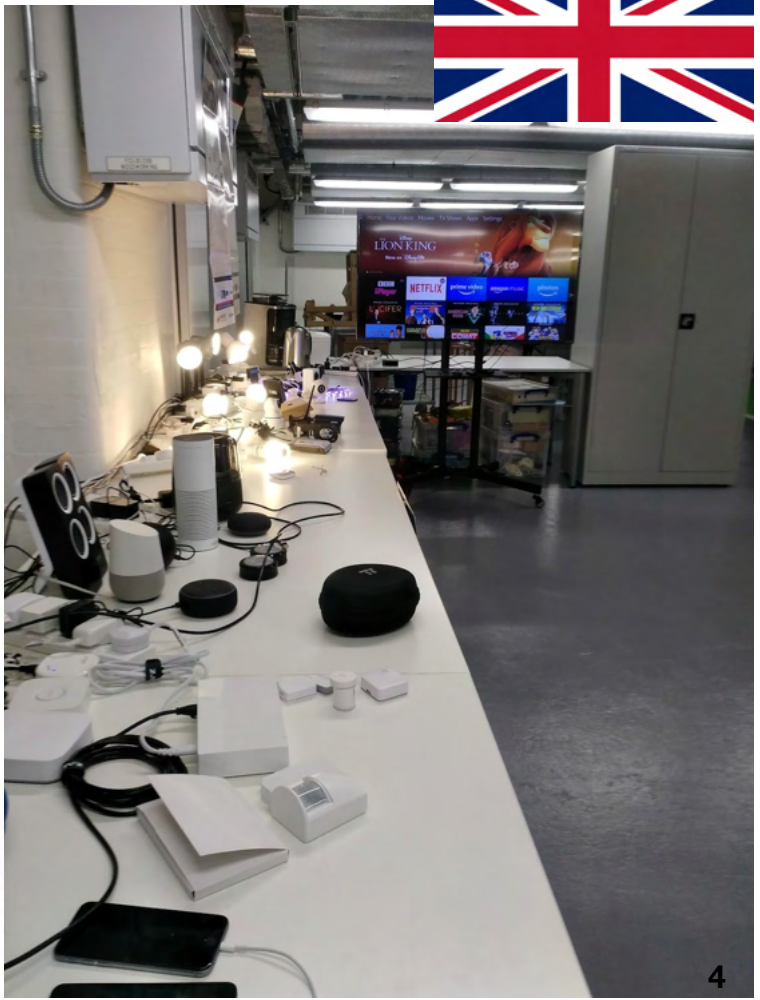

Blocking without Breaking

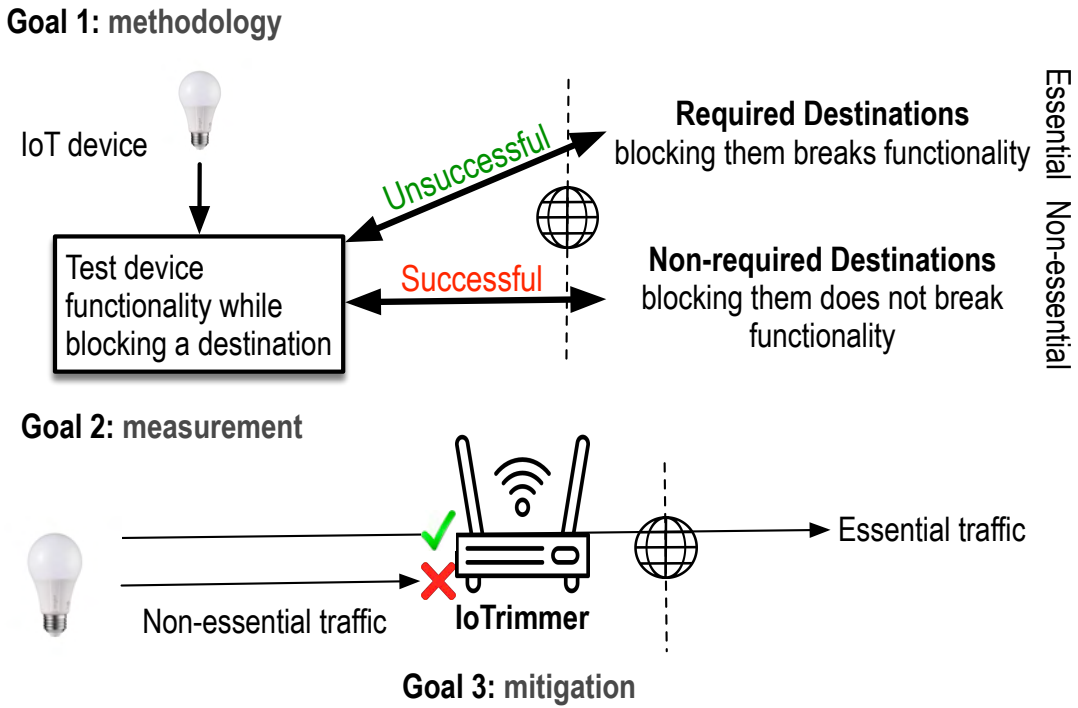

IoTrimmer

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?

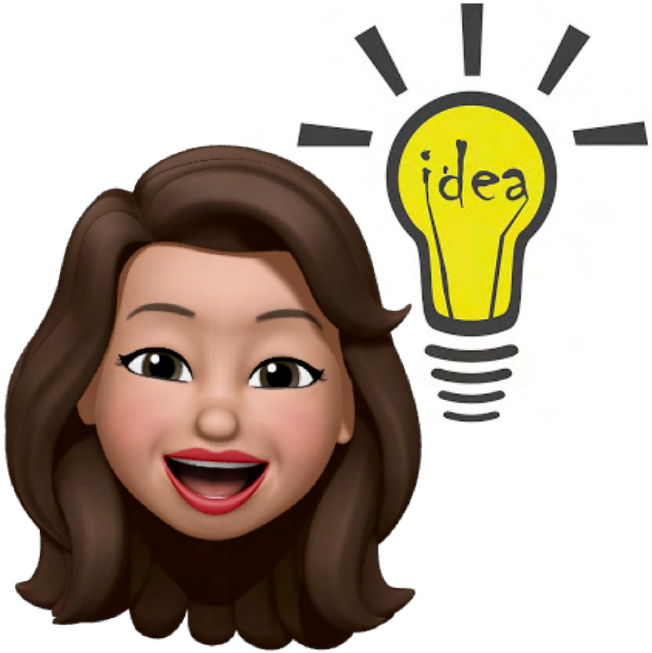

Goals

- *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*

Design of Experiments

- **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%**

Activity	Description
Power	power on/off the device
Voice	voice commands for speakers
Video	record/watch video
On/Off	turn on/off bulbs/plugs

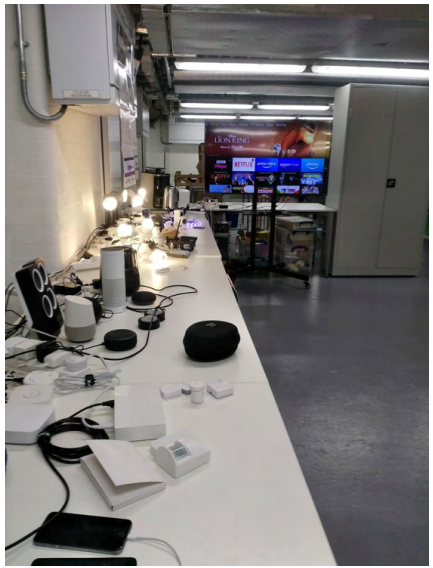

Identifying Non-essential Traffic

Distinguishing Required from Non-Required Destinations

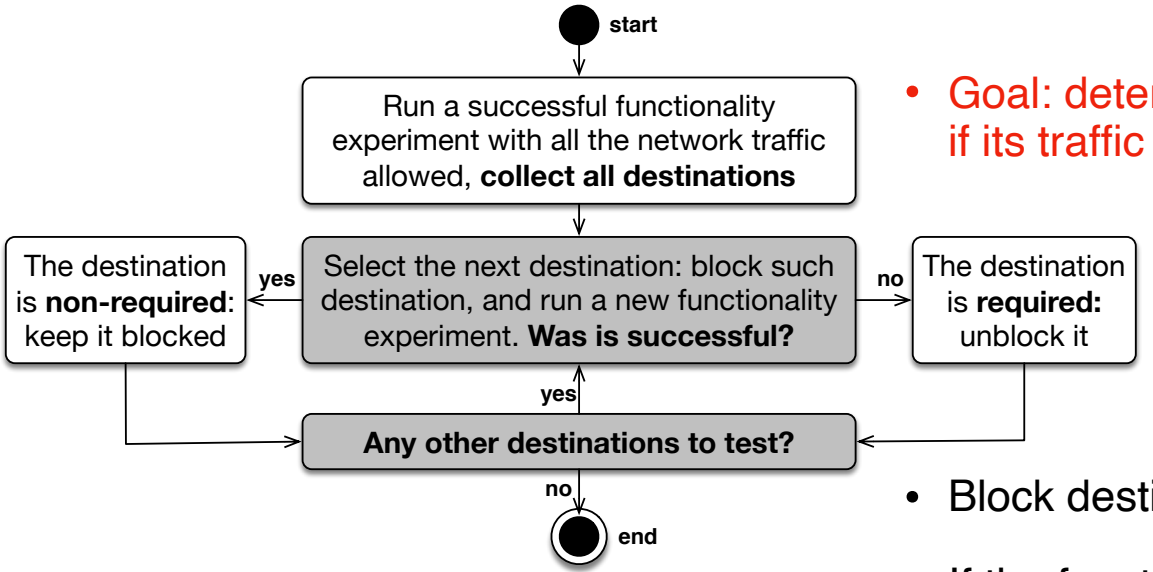

- Goal: determine if a destination is required (i.e., if its traffic is essential)

- 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

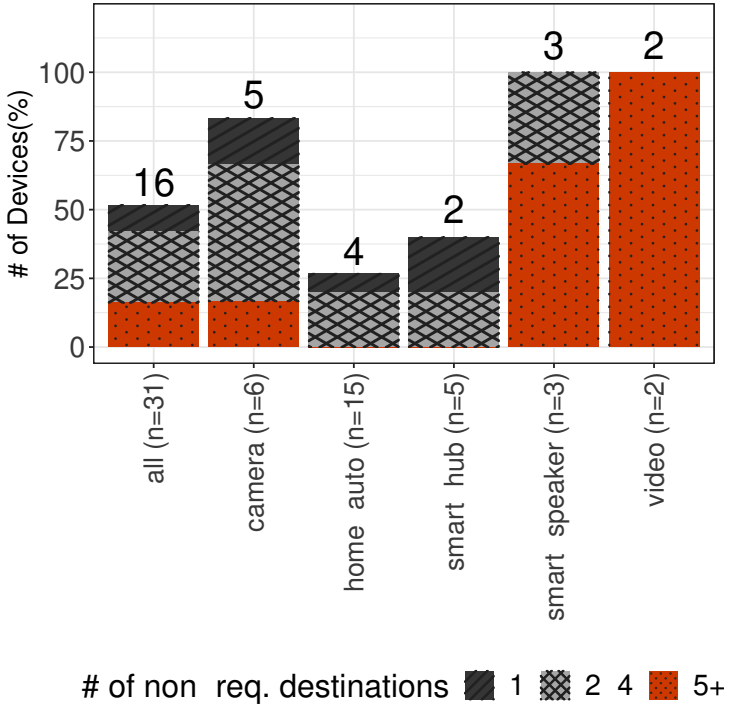

- 16/31 devices have non-essential traffic
- Mostly cameras, smart speakers, and video
- Possible explanations:
  - complexity (skills and apps)
  - uncommon vendors / rebranding (for cameras)

Impact of Destination Party

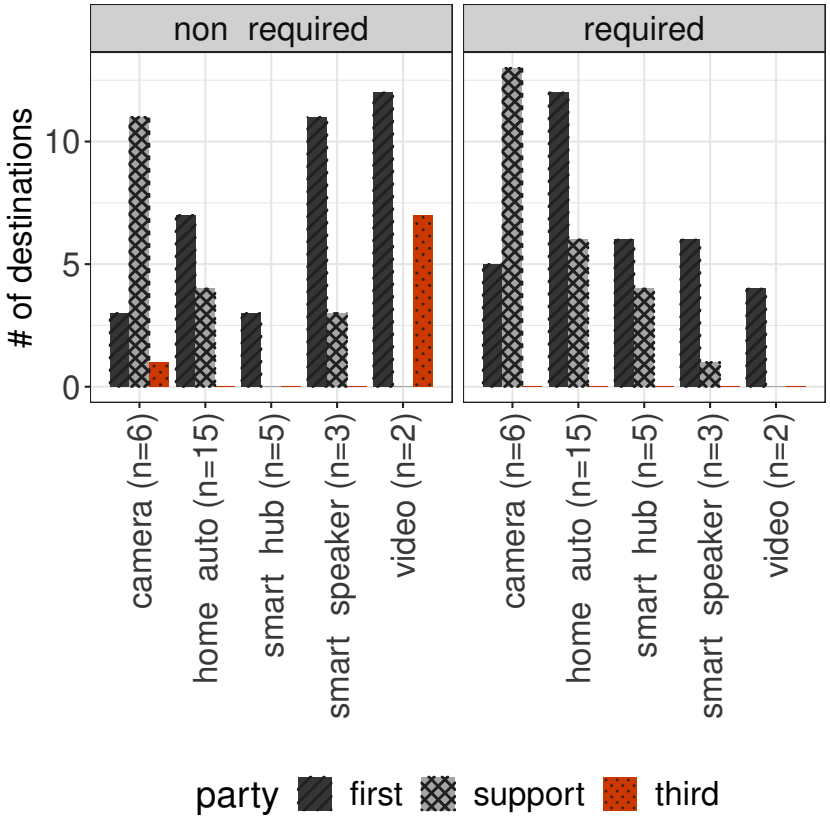

- Third parties are always non-required
  - Probably because of background app activity (Netflix)
- Some first/support parties also non-required
  - Best guess: firmware upgrade
  - Worst guess: data collection

Amount of Data Sent During One Experiment

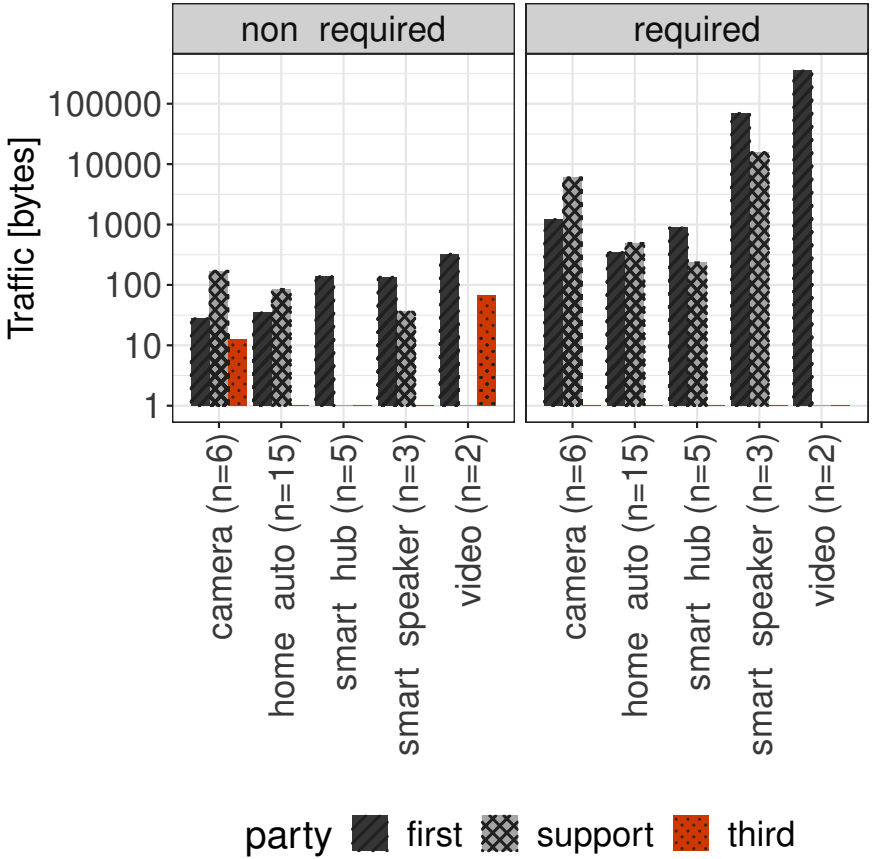

- **Good news**: non-essential traffic is relatively small (less than 1KB/device)
- However, it is still possible to transmit:
  - Presence of the device
  - Its status
  - Basic data from the sensors (e.g., open/close, motion/still, alarm/no alarm)

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

Open Challenges

- **Testing more devices**
- **Do protocol and ports help in detecting non-essential traffic?**
- **Do required and non-required destinations change over time?**

Mitigating Non-essential IoT Traffic

- A blocking system: **IoTrimmer**
  - 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**

IoTrimmer Control Panel