DO WE NEED MYRIADS OF CLOUD-BASED SAFEGUARDS?

Using a home router for AI-powered IoT threats detection

Anna Maria Mandalari, Vadim Safronov

Hamed Haddadi, Daniel J. Dubois, David Choffnes

UCL

Imperial College London

Northeastern University

PETRAS

EPSRC Pioneering research and skills

NSF

Problem: IoT Devices Expose Information Over the Internet

They "sense" a lot

Microphones
Cameras
User activities

…

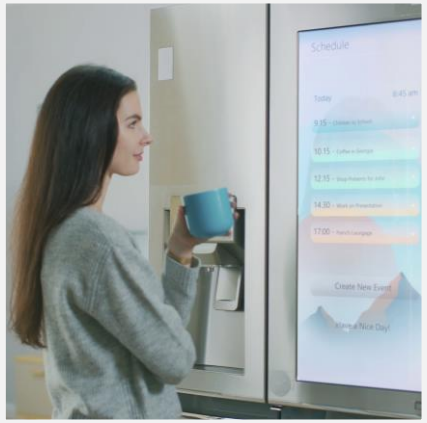

Privacy Threats

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?

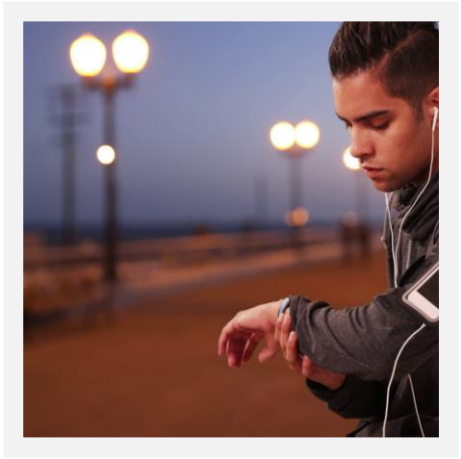

Control

Device detection

Intelligent profiles

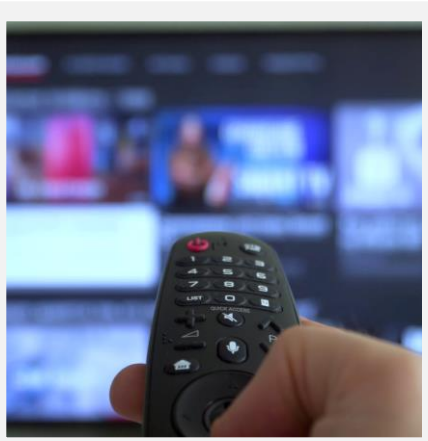

Security

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

Testbed

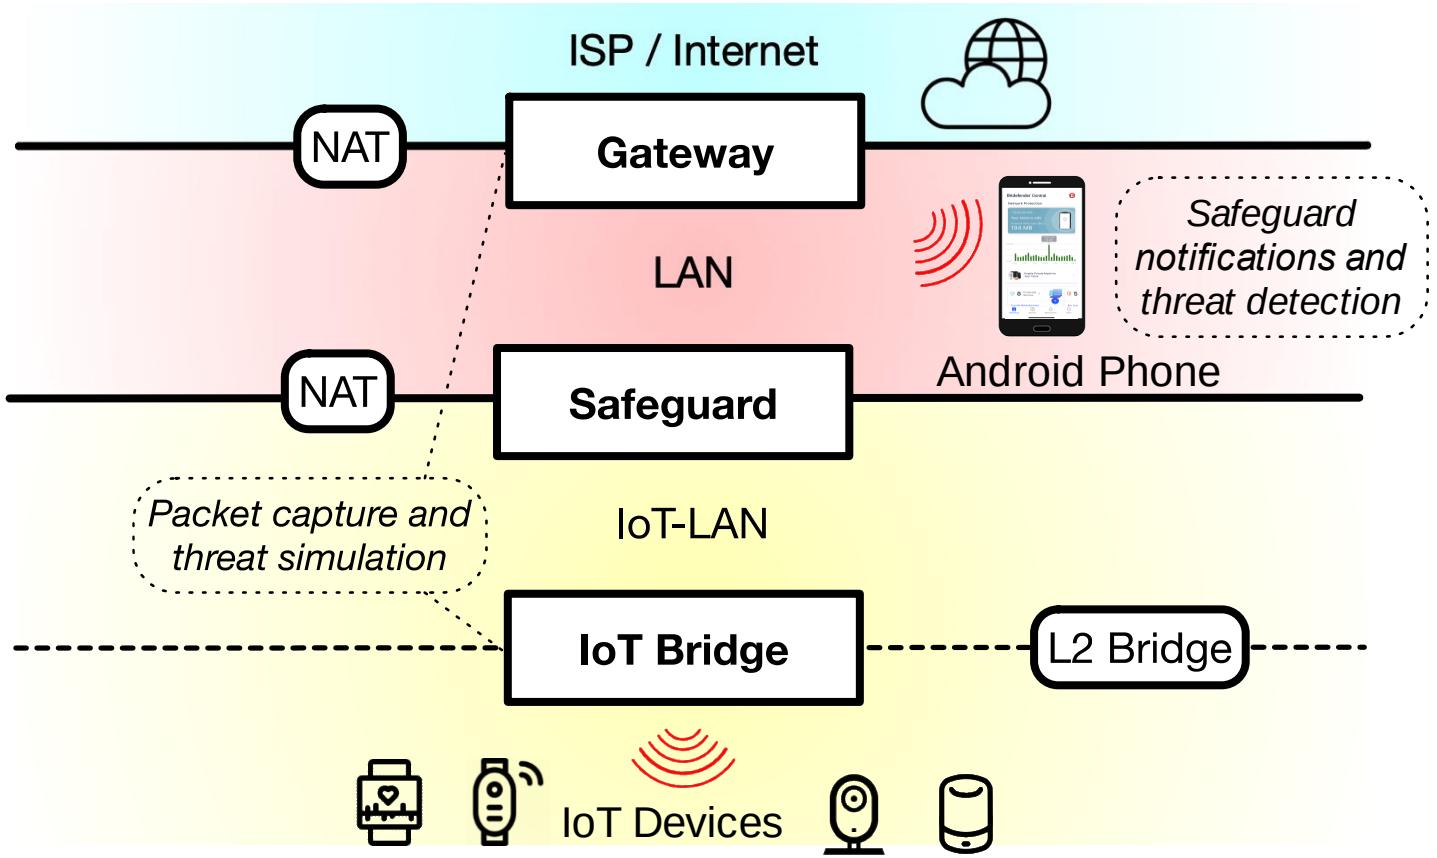
ISP / Internet
NAT
Gateway
LAN
Safeguard notifications and threat detection
Android Phone
NAT
Safeguard
Packet capture and threat simulation
IoT-LAN
IoT Bridge
L2 Bridge
IoT Devices

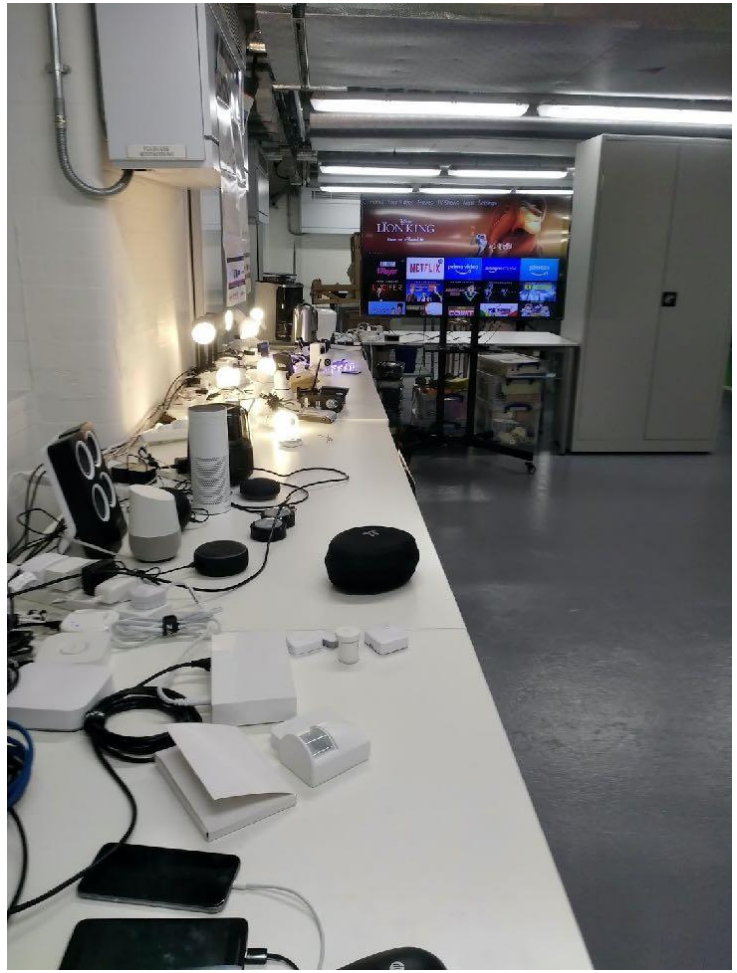

Research Questions

- **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

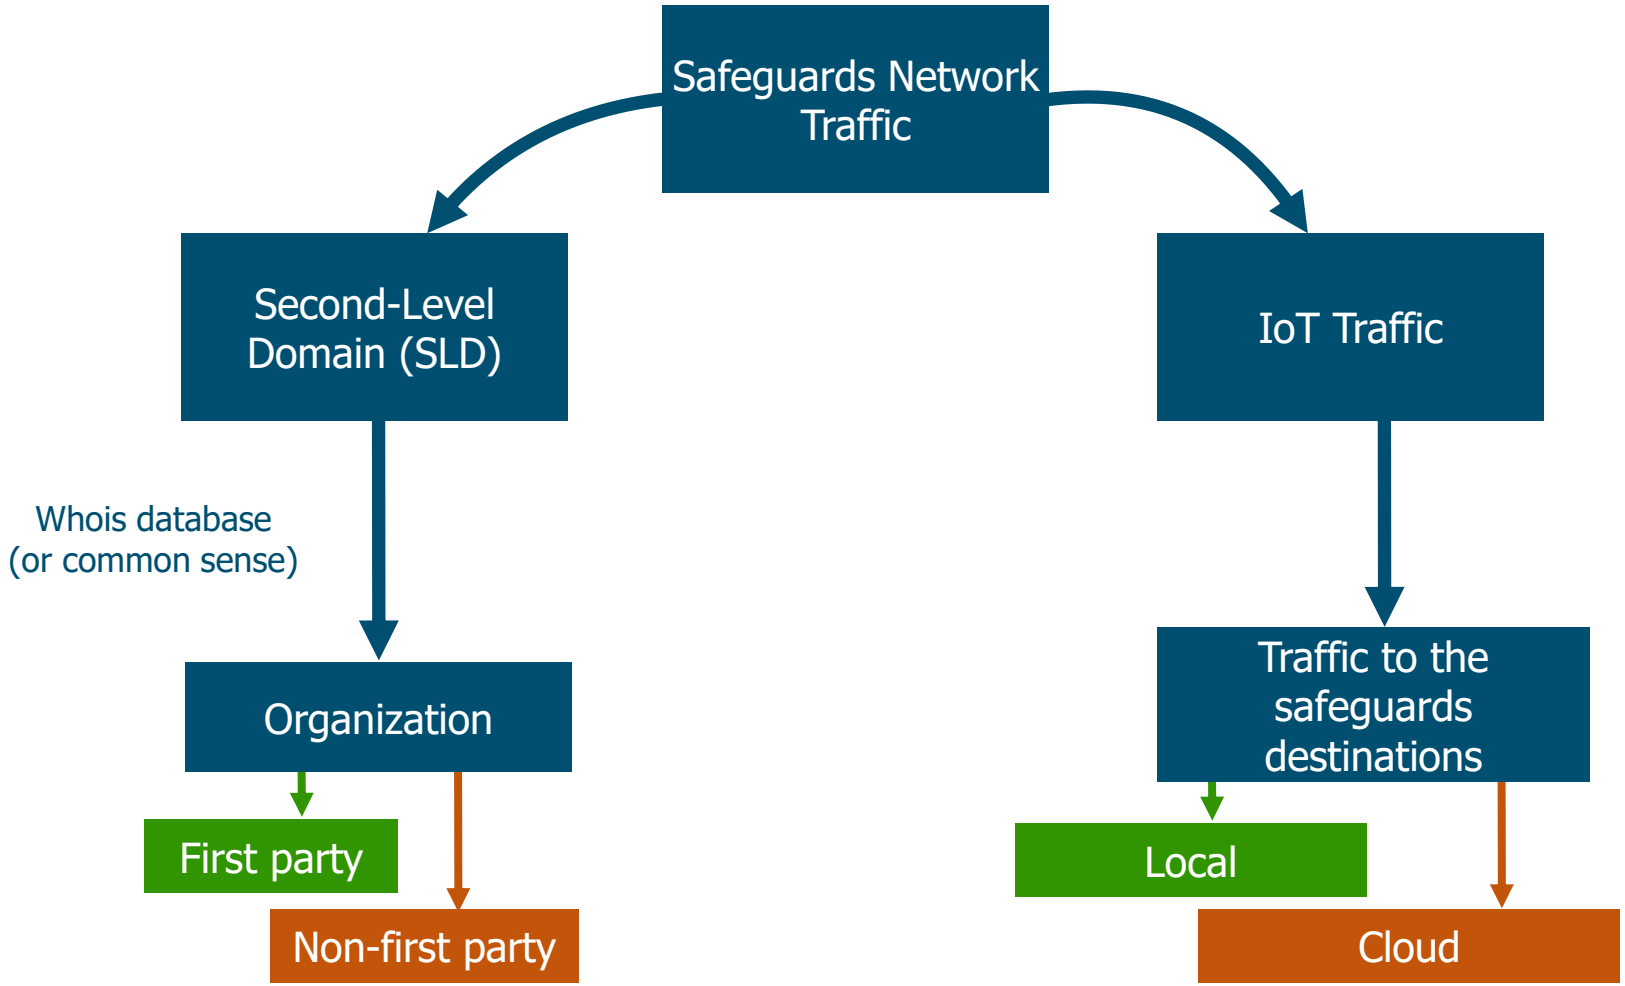

Processing Locality & Party Analysis

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

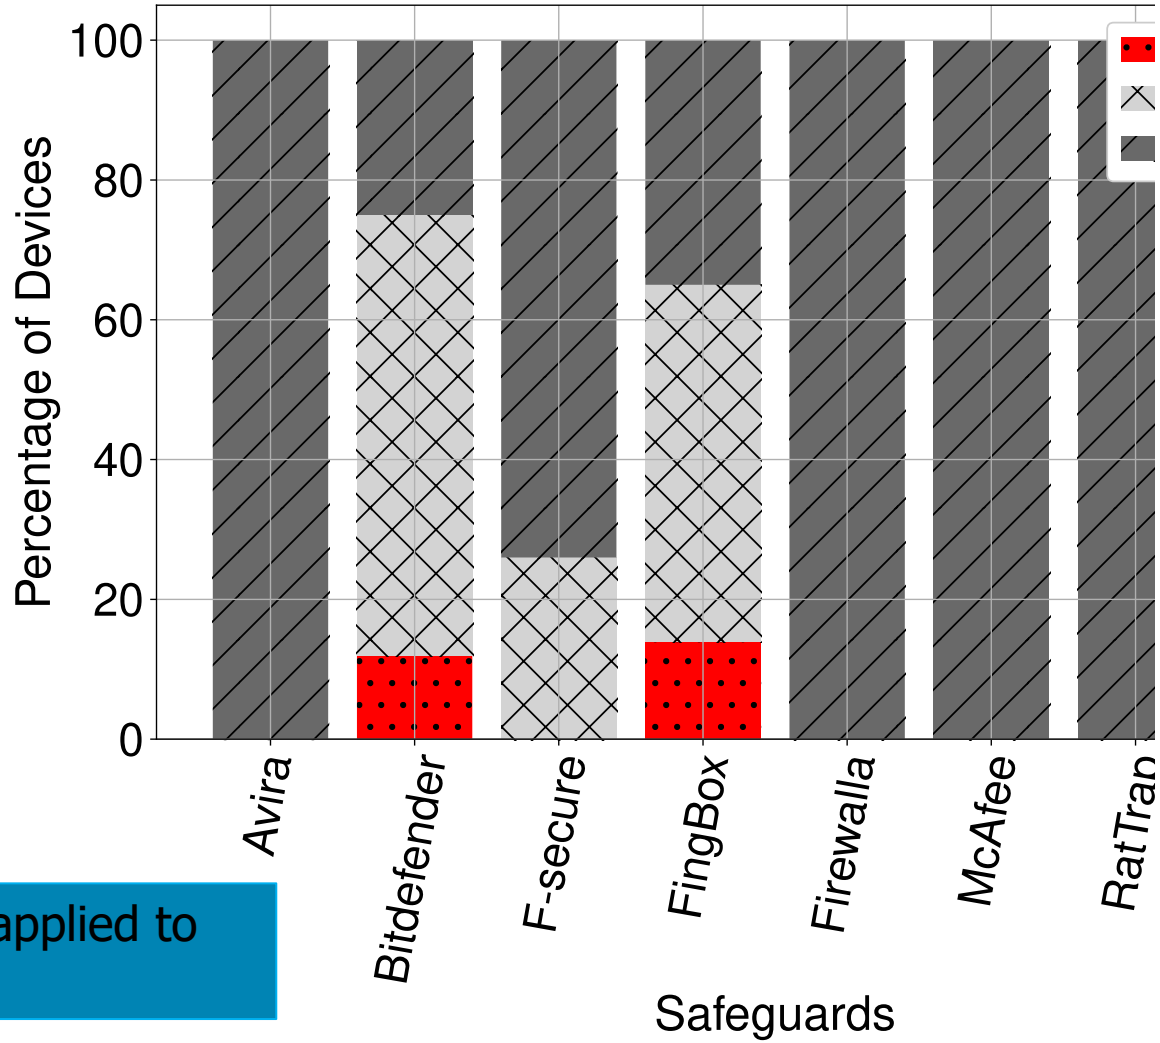

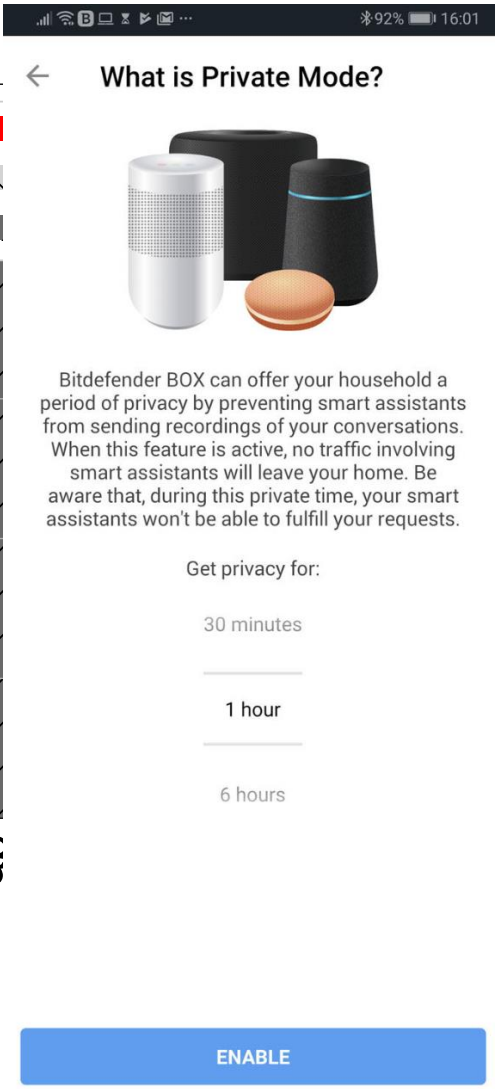

Protection techniques applied to specific vendors

Take away: only a small 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

Threat Detection Experiments

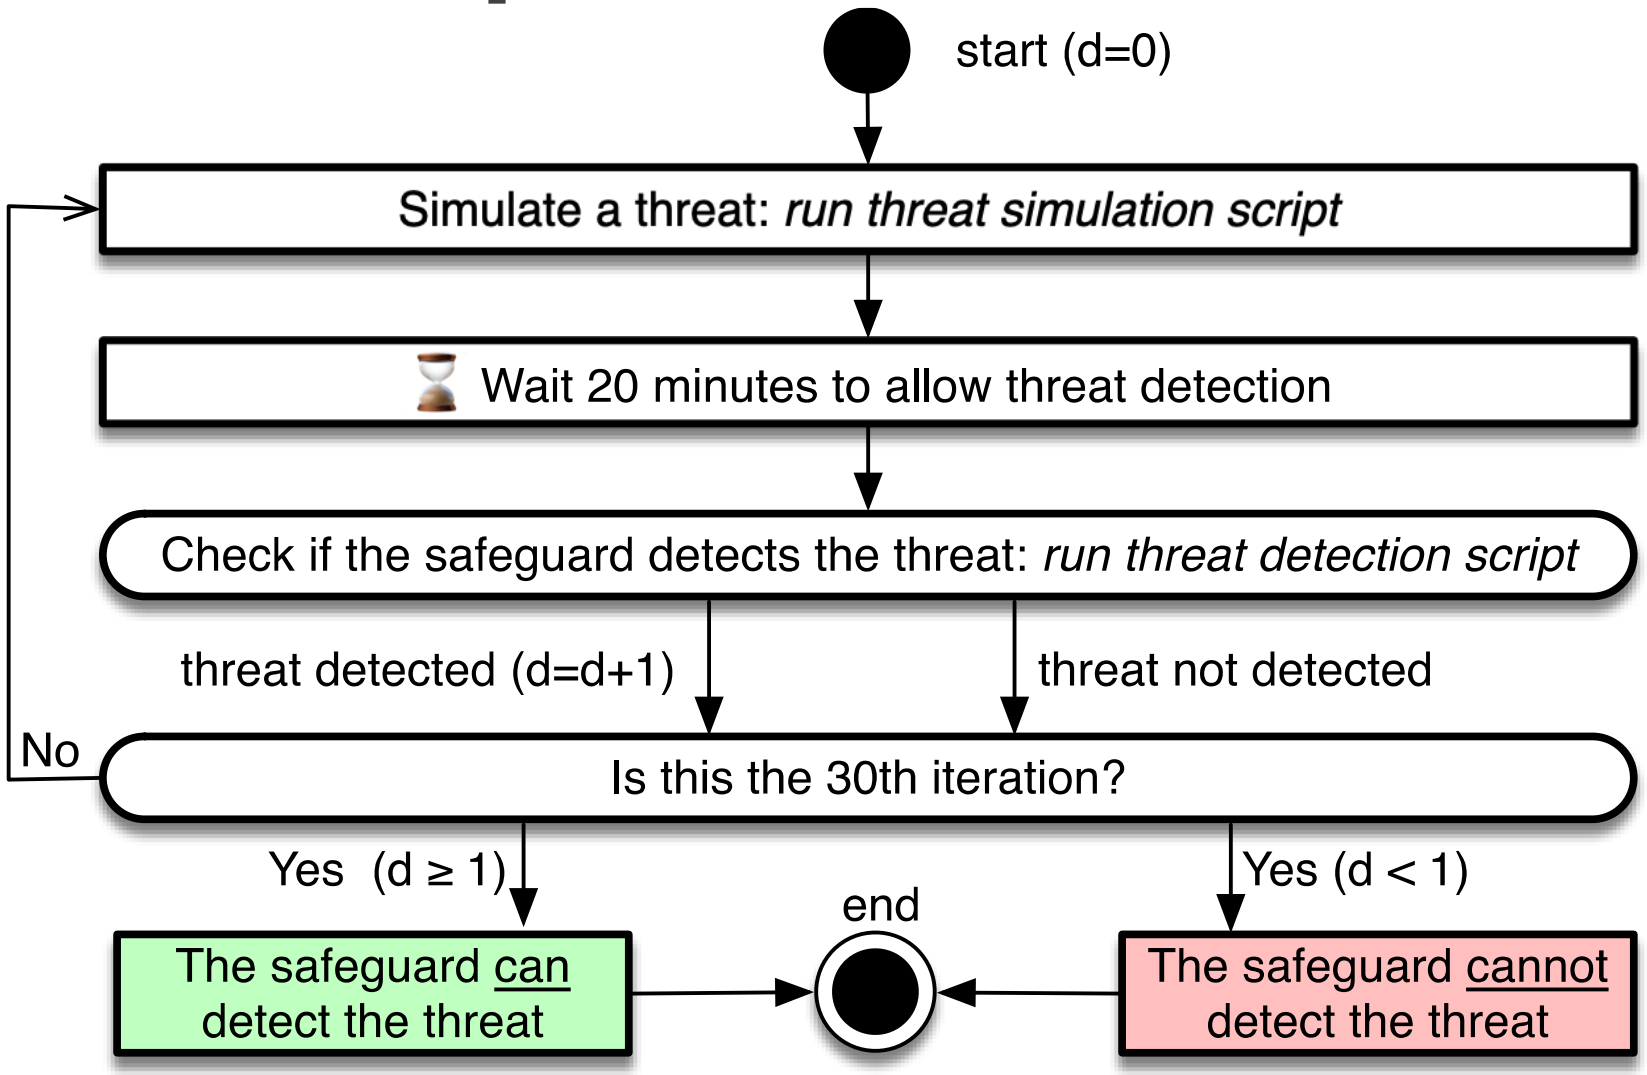

Evaluation of Threat Detection Capability

	Threat	Avira	Bitdefender	F-Secure	Fingbox	Firewalla	McAfee	RaTtrap	TrendMicro
Security	Anomaly ON/OFF	-	✗	✗	-	✗	✗	✗	-
	Anomaly Traffic Pattern	-	✗	✗					
	Abnormal Upload	-	✗	✗					-
	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

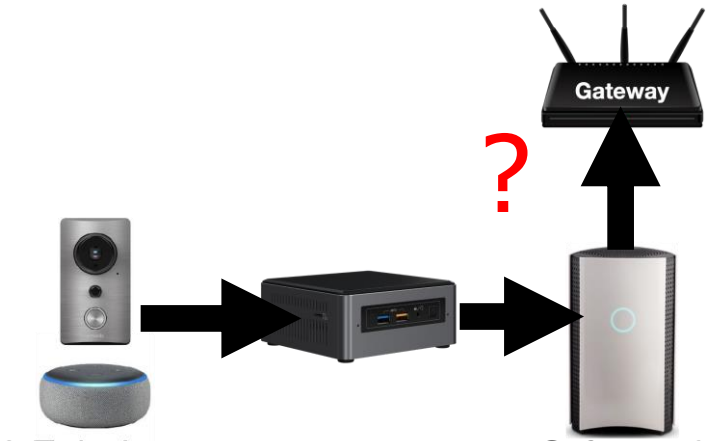

Privacy Policy

MANUALLY INSPECTING THE PRIVACY POLICY

Traffic Overhead

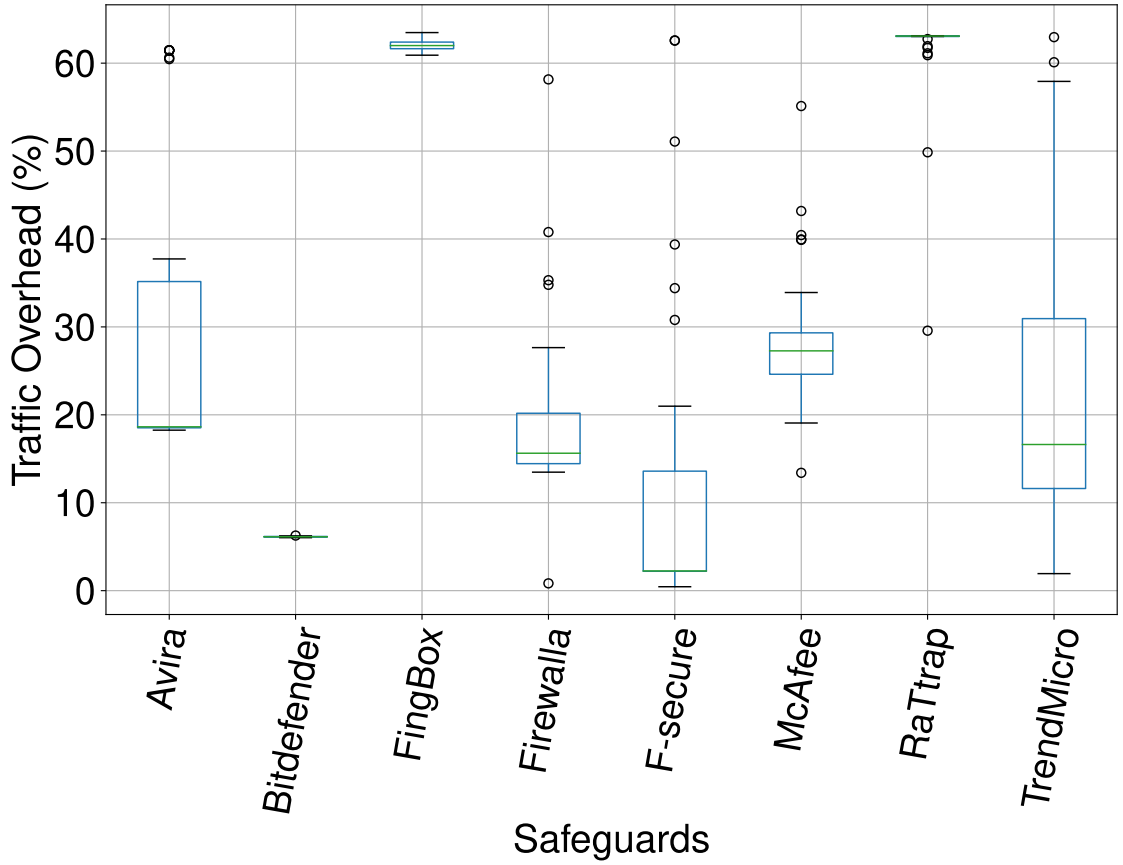

Take away: Some of the safeguards introduce significant traffic overhead. In general the overhead is never less than 10% of the traffic of the IoT devices.

Privacy Policy

Privacy Policy	Avira	Bitdefender	F-Secure	Fingbox	Firewalla	McAfee	RaTtrap	TrendMicro
Anonymization	✓	✓ [pseudonymize]	X [ceasing subscription]	✓	X	X	X	X
Usageof 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	X	Partners	Partners	Partners

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

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

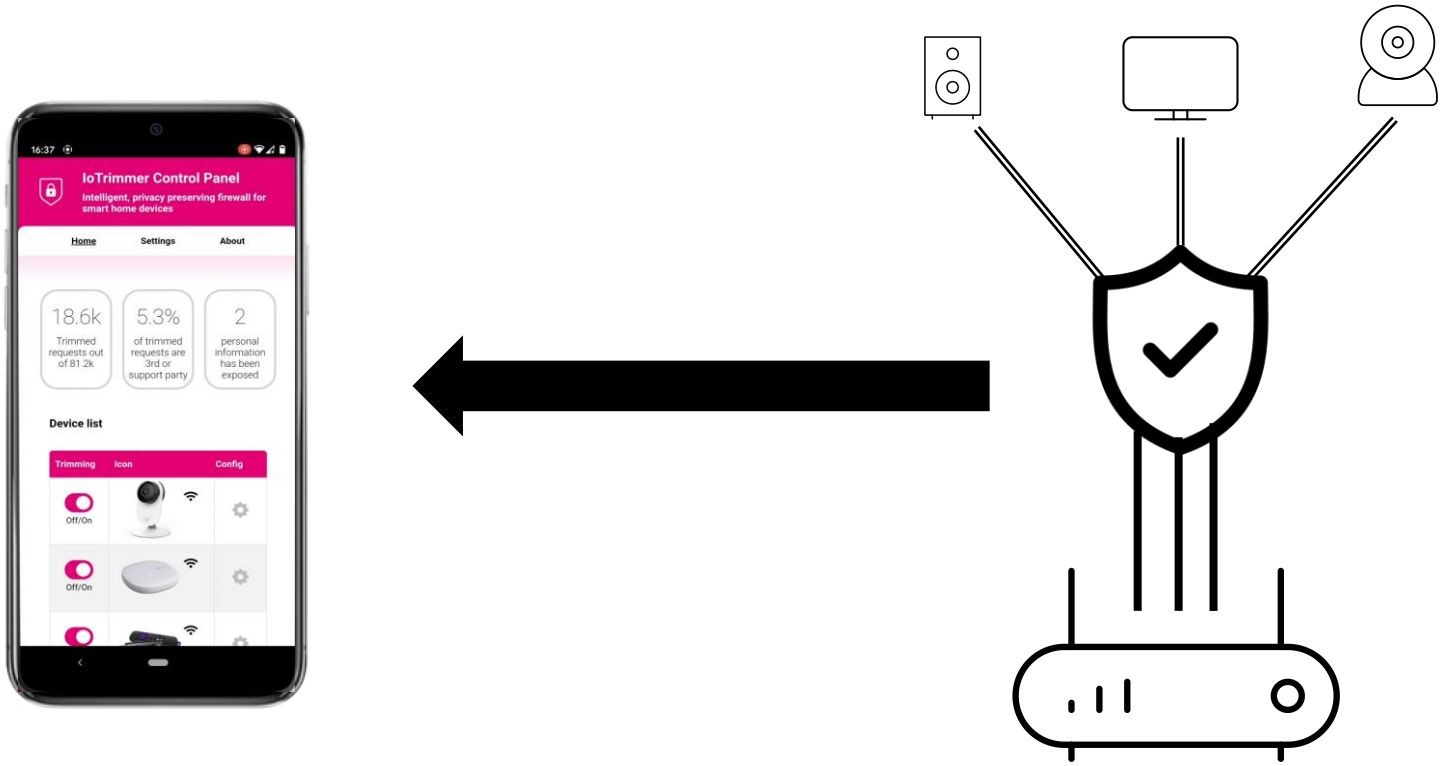

Motivation

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

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?

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

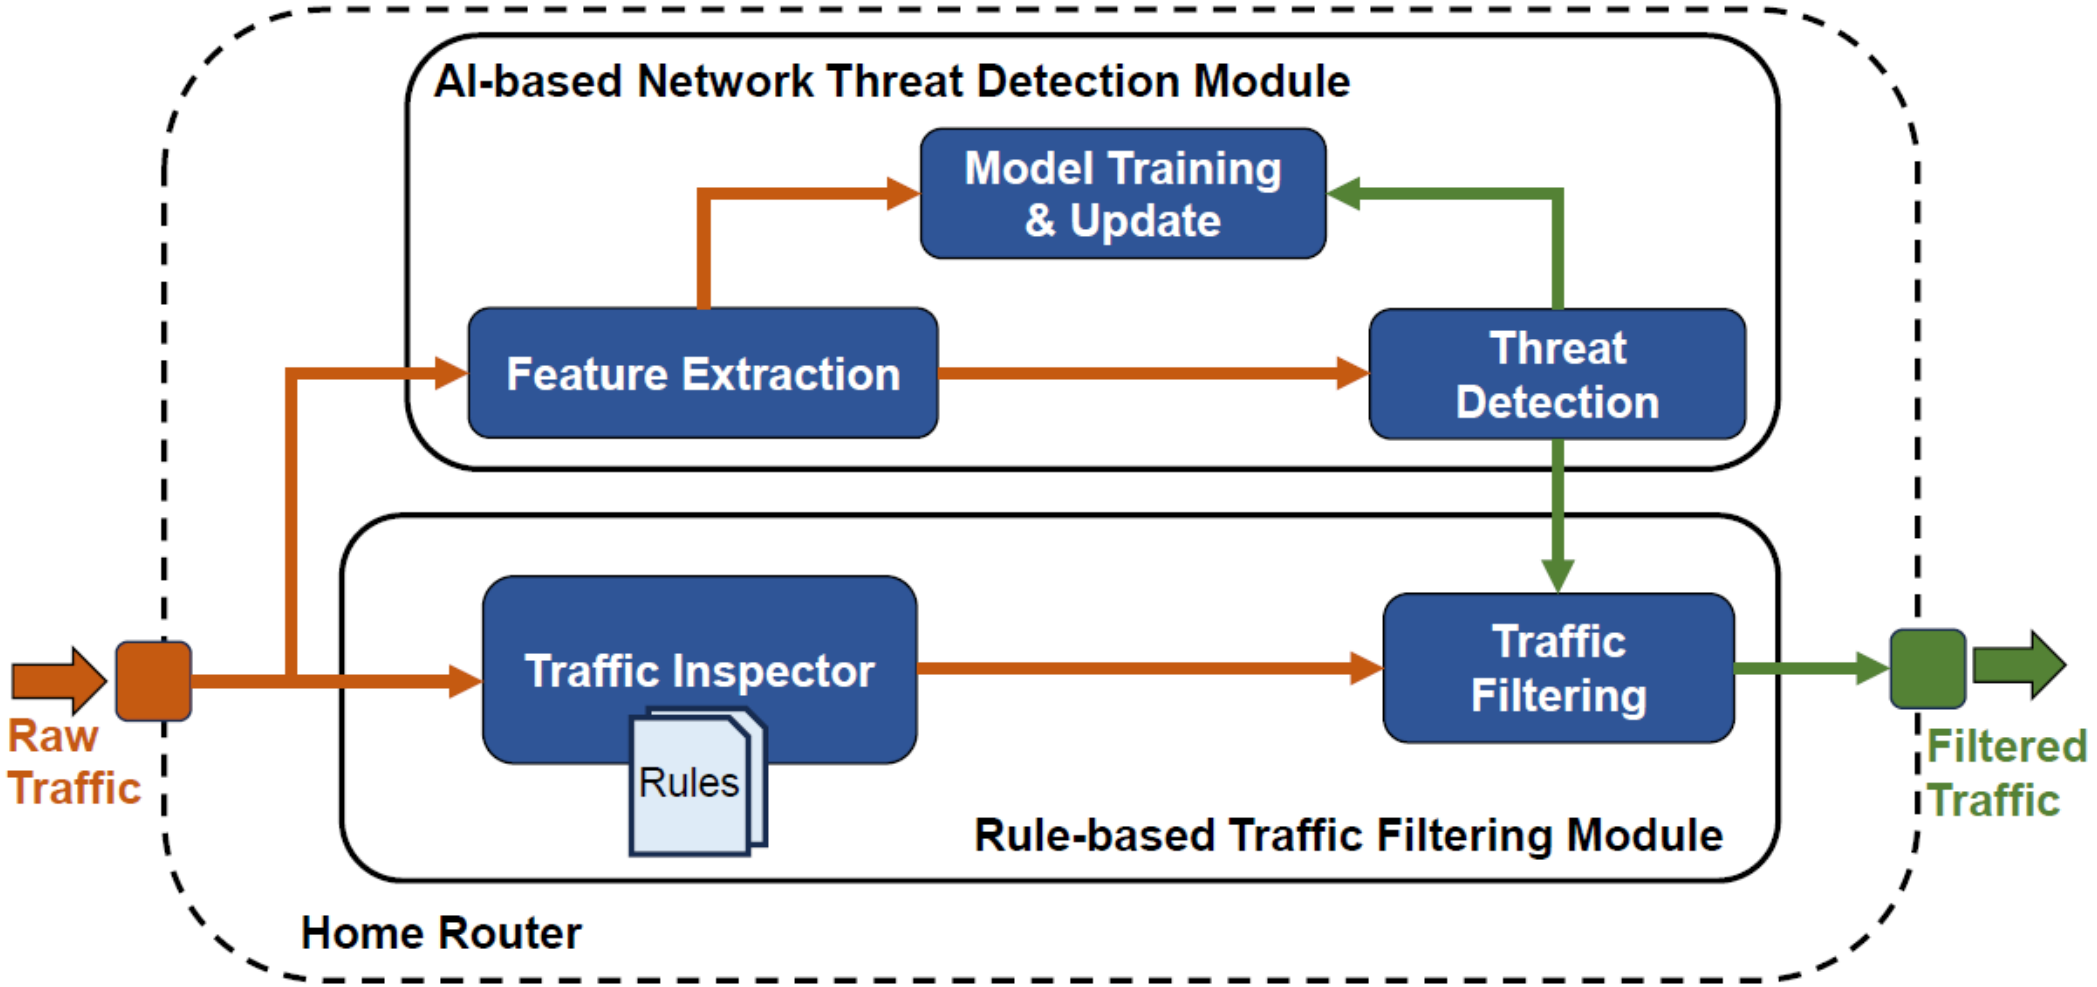

Implementation: home router with IoT protection

- 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

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

snort3: https://www.snort.org/snort3 ; netml lib: https://github.com/noise-lab/netml

Evaluation: threat coverage and prevention time

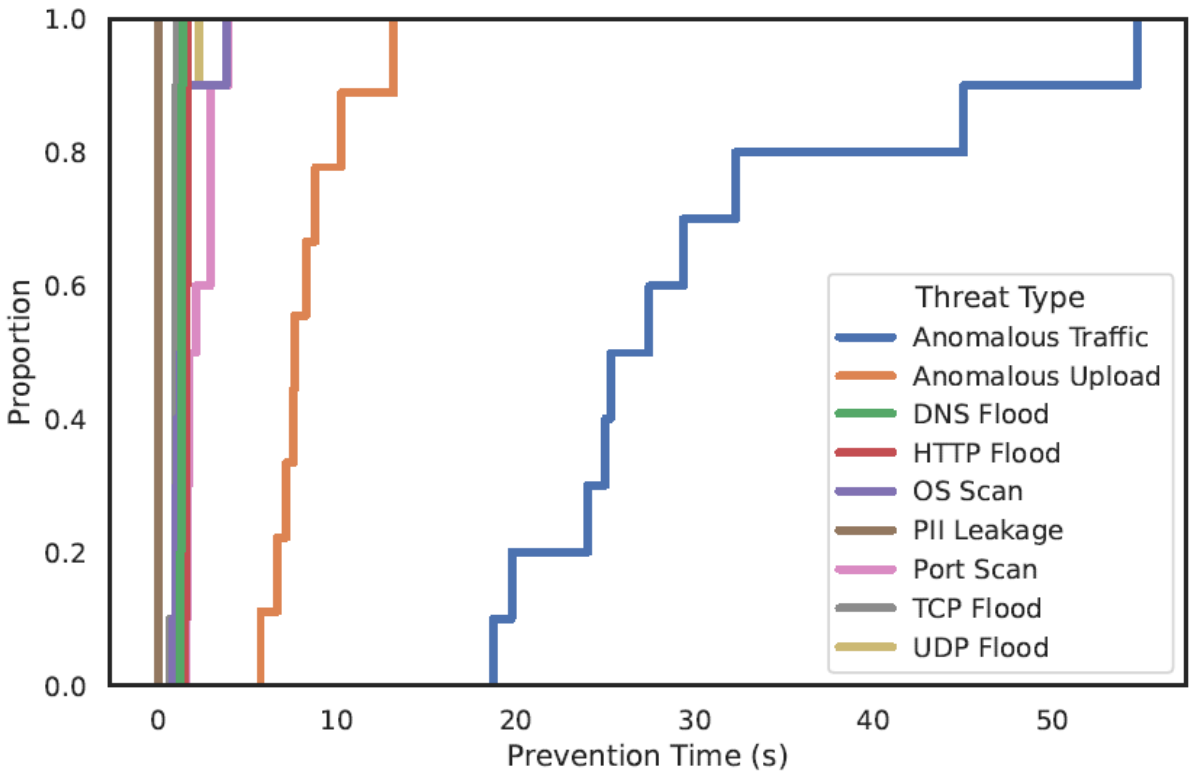

Threat	IoT Protection Systems	SunBlock
Anomalous Traffic	✗	✓
Anomalous Upload	✗	✓
SYN Flooding	✓	✓
UDP Flooding	✗	✓
DNS Flooding	✗	✓
HTTP Flooding	✓	✓
Port Scanning	✓	✓
OS Scanning	✓	✓
PII Leakage	✗	✓

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

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

Regular IoT traffic

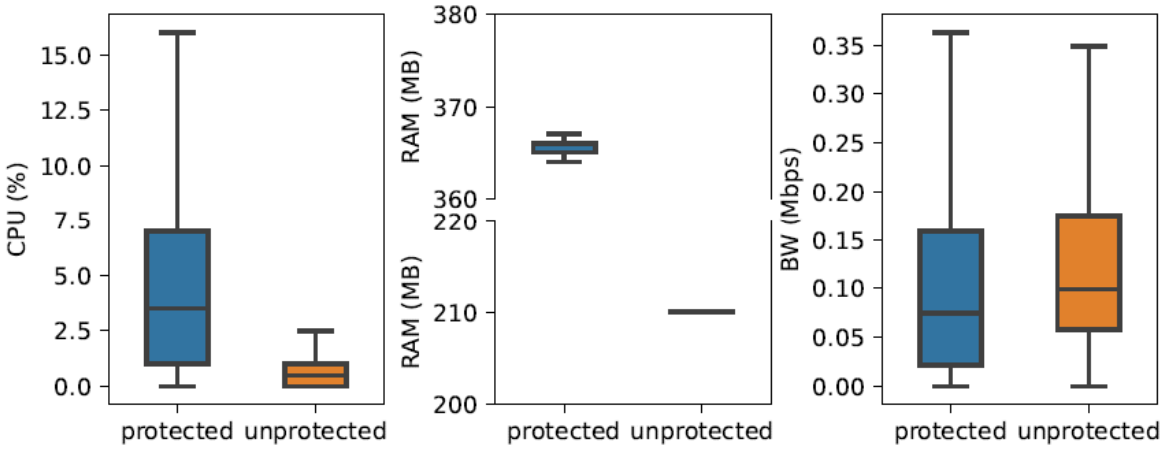

Evaluation: performance overhead

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

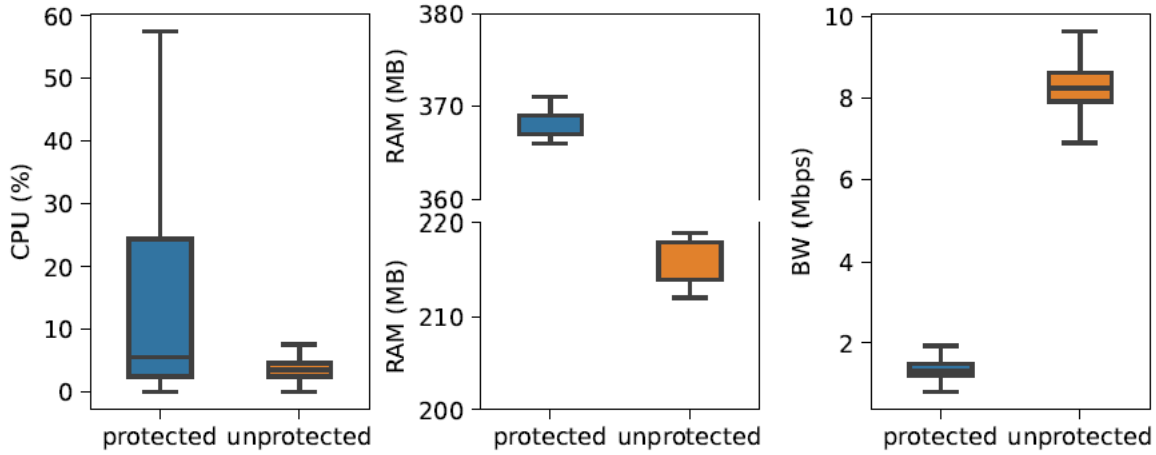

Regular IoT traffic

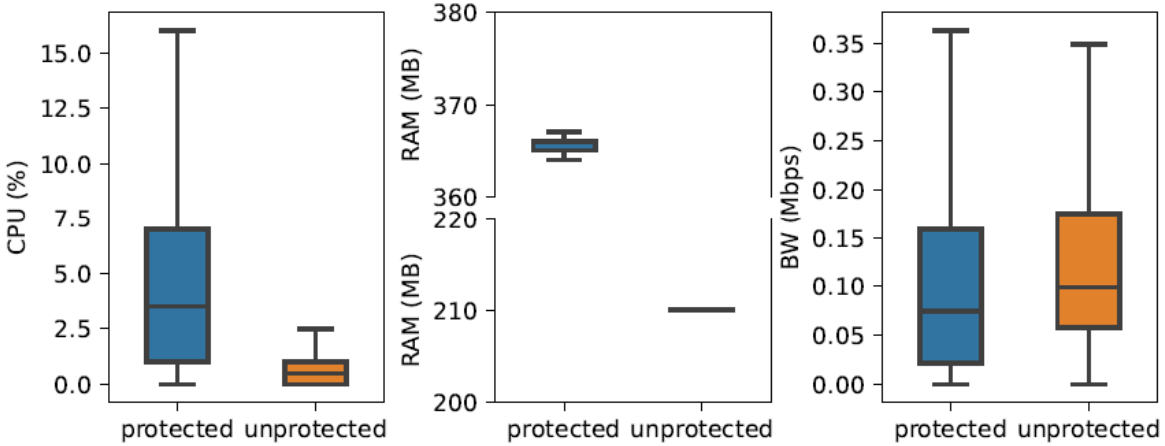

UDP flood

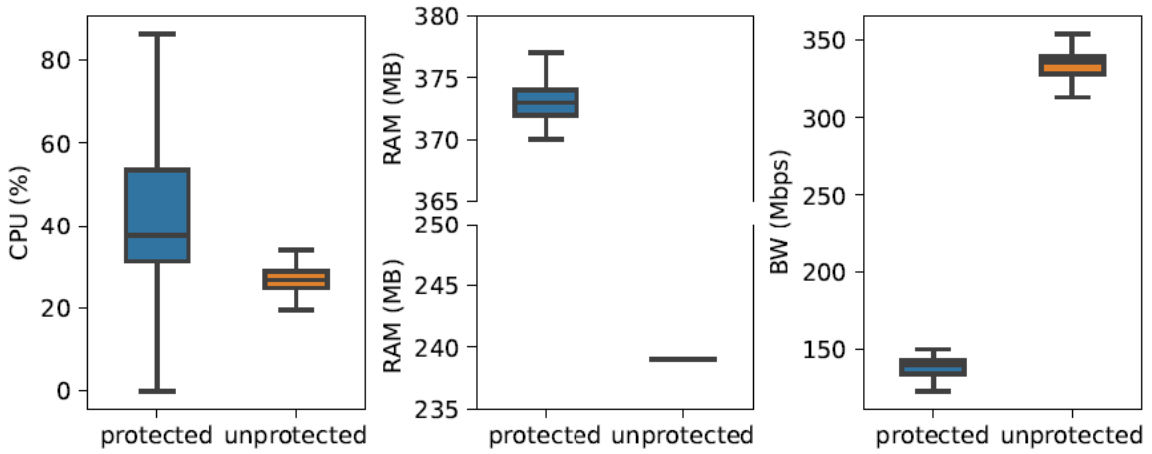

Takeaways

- 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

Anna Maria Mandalari

University College London

Vadim Safronov

Cambridge University

Hamed Haddadi

Imperial College London

Daniel Dubois

Northeastern University

David Choffnes

Northeastern University

DO WE NEED MYRIADS OF CLOUD-BASED SAFEGUARDS?

Using a home router for AI-powered IoT threats detection

IEEE S&P 2023