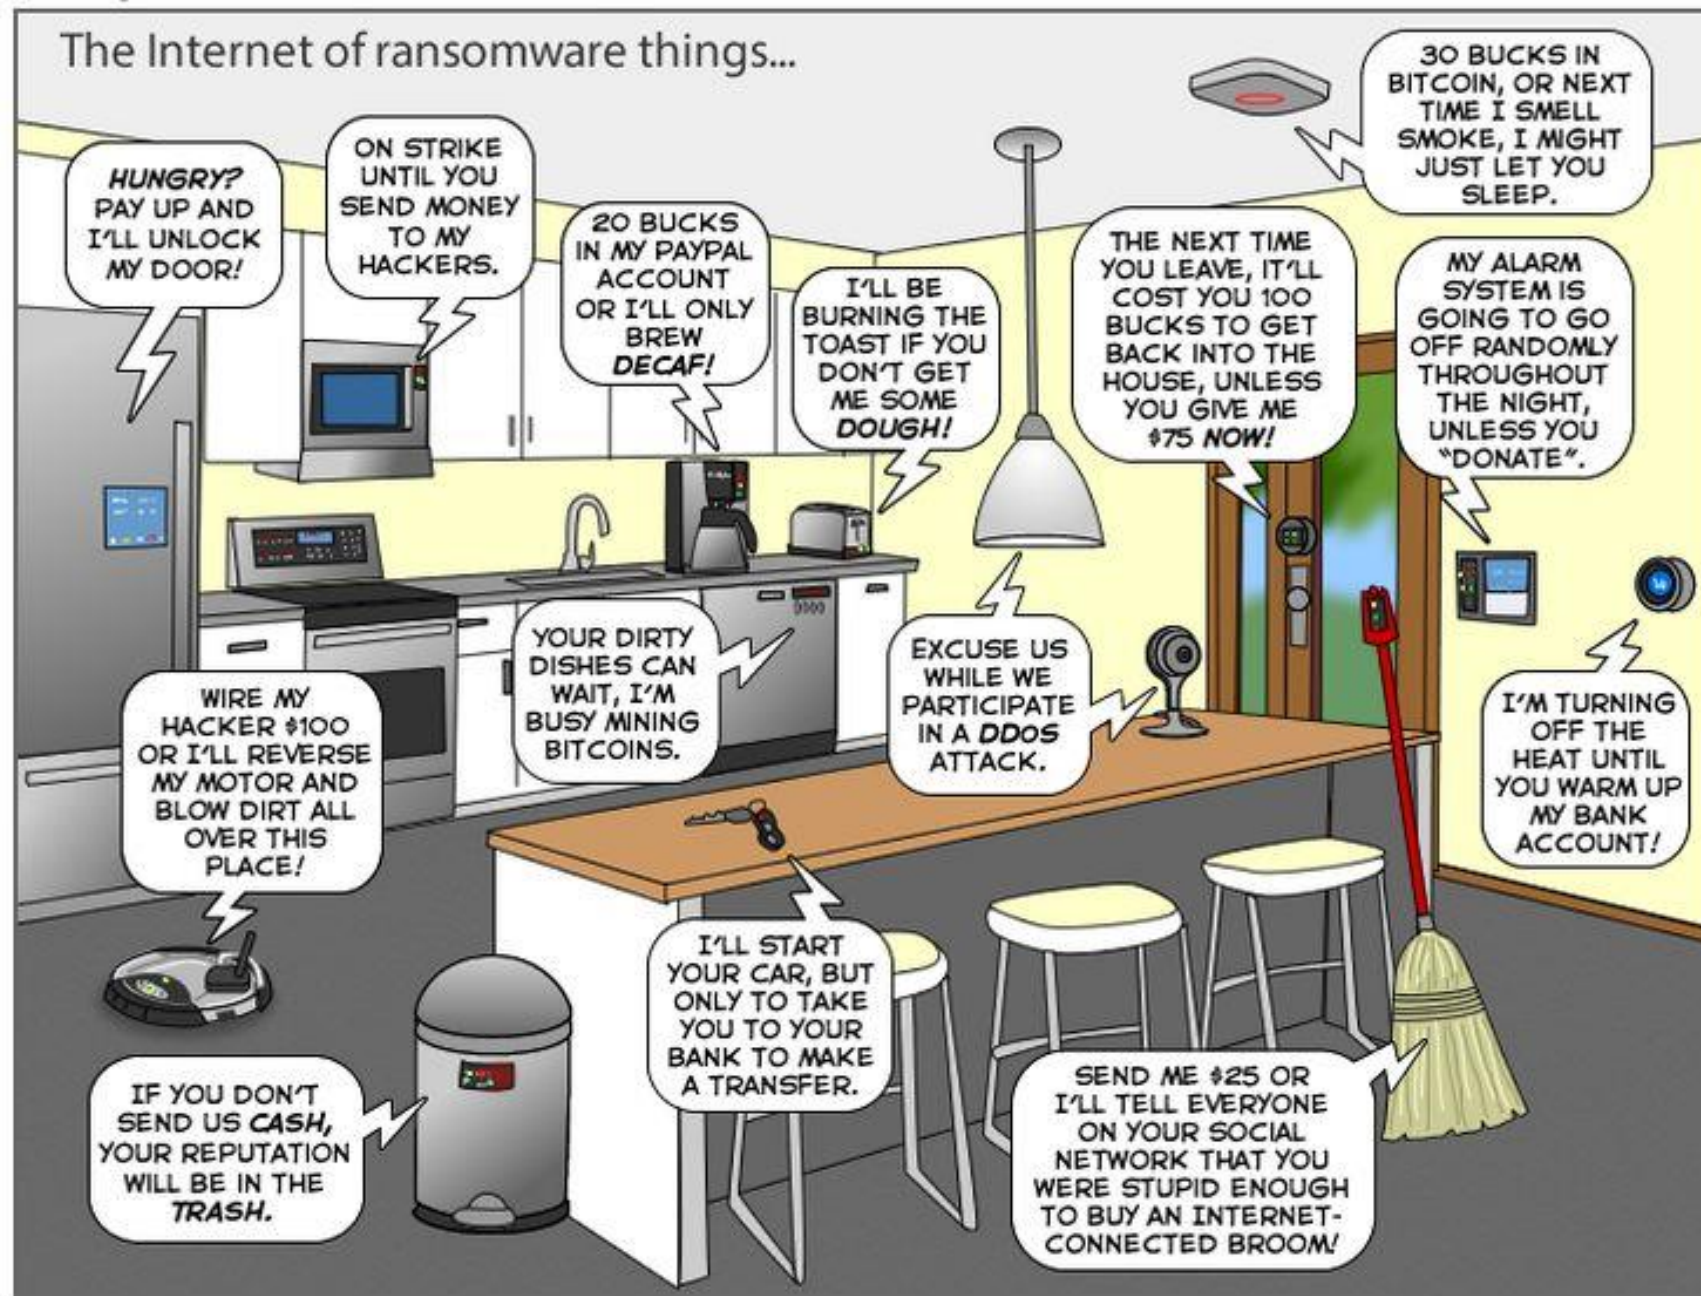


Lecture #7: IoT forensics

Antonia Affinito, Etienne Khan, Ting-Han Chen,
Savvas Kastanakis, and Cristian Hesselman

University of Twente | June 18, 2025



Today's agenda

- Admin
- Introduction to today's lecture
- Paper #1: analysis of IoT malware behavior with RIOTMAN
- Paper #2: IoT honeypot framework based on Honware
- Feedback (last 10 mins)

Admin

Schedule

Lecture	Date	Contents
R1	Apr 25	Course Introduction
G1	Apr 30	How the core of the Internet works.
R2	May 9	Principles of IoT security
R3	May 16	Internet Core Protocols
R4	May 23	IoT Botnet Measurements
R5	May 27	IoT TLS and Q&A lab assignment
G2	Jun 6	IoT and post-quantum crypto
R6	Jun 13	IoT Security Vulnerabilities
R7	Jun 18	IoT Forensics

Important dates



- All summaries due: **Fri Jun 20**
- Written exam: **Mon Jun 23, 08:45-10:45**
- Slides (PDF), PCAP, MUD, README files due: **Wed Jun 25, 9AM CEST**
- Presentations:
 - **Fri June 27**, from 8:45 to 12:30 in NH 115 and NH 124
 - **Mon June 30**, from 8:45 to 12:30 in NH 115 and NH 124

Official feedback forms

- Survey by EEMCS Quality Assurance folks, will be sent out on in the next week or so
- Please fill it out, your feedback is **crucial** for us to further improve the course!
- Next year's students will thank you for it 😊
- We'll let you know how we handled your feedback

EvaSys EEMCS Master Student Experience Questionnaire Corona Electric Paper

University of Twente Quality Assurance EEMCS UNIVERSITEIT TWENTE.

Faculty of EEMCS ()

Mark as shown: ☐ ☒ ☐ ☐ Please use a ball-point pen or a thin felt tip. This form will be processed automatically.
Correction: ☐ ☒ ☐ ☐ Please follow the examples shown on the left hand side to help optimize the reading results.

1. Administrative

1.1 Which Master programme do you attend? ☐ Applied Mathematics ☐ Business Information Technology ☐ Computer Science
☐ Electrical Engineering ☐ Embedded Systems ☐ Interaction Technology
☐ Internet Science and Technology ☐ Systems & Control ☐ Other

1.2 Which other Master programme do you attend?
☐ Applied Physics ☐ Biomedical Engineering ☐ Business Administration
☐ Chemical Engineering ☐ Civil Engineering & Management ☐ Communication Science
☐ Construction Management & Engineering ☐ Educational Science & Technology ☐ Environmental & Energy Management
☐ European Studies ☐ Geo-information Science and Earth Observation ☐ Geographical Information Management and Applications
☐ Health Sciences ☐ Industrial Design Engineering ☐ Industrial Engineering & Management
☐ Mechanical Engineering ☐ Methodology & Statistics for the Behavioural, Biomedical & Social Sciences ☐ Nanotechnology
☐ Philosophy of Science, Technology & Society ☐ Psychology ☐ Public Administration
☐ Science Education and Communication ☐ Social Sciences and Humanities Education ☐ Spatial Engineering
☐ Sustainable Energy Technology ☐ Technical Medicine ☐ Water Technology

1.3 At which university are you primary enrolled in (hoofdinschrijving)? ☐ University of Twente ☐ Delft University of Technology ☐ Eindhoven University of Technology
☐ Other

2. Online/hybrid education

2.1 How did you experience the online/hybrid education as offered in this course? Insufficient ☐ ☐ ☐ ☐ Excellent ☐ N/A

2.2 Which teaching activities helped you the best?

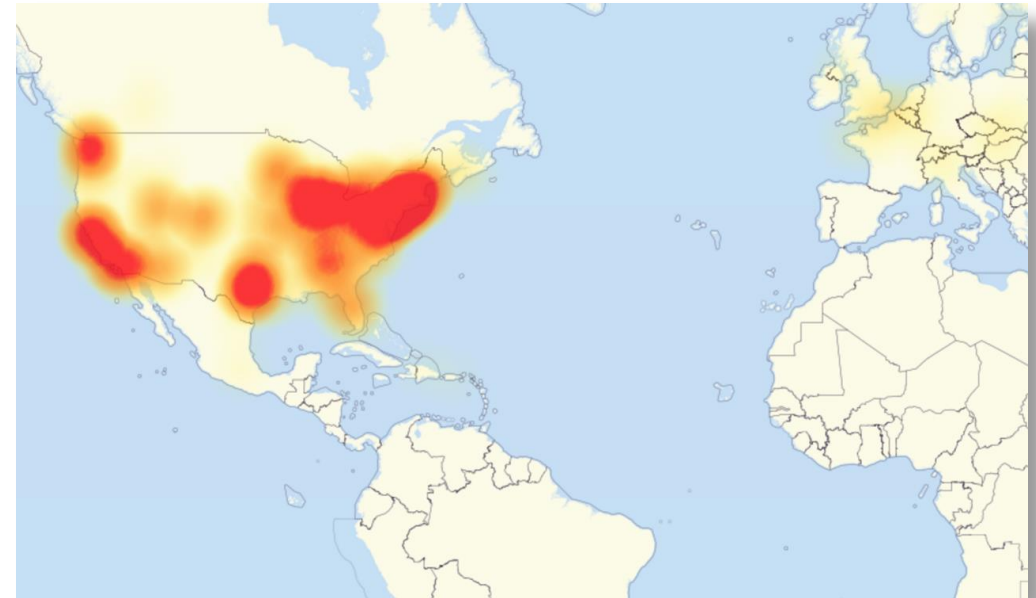
2.3 Which teaching activities worked counterproductive for you?

F5261U0P1PL0V0 31.05.2021, Page 1/2

Introduction to today's lecture

Malware we have seen in the course: Mirai and Hajime

- Requires scalable mechanisms to understand the characteristics and behavior of IoT bots
- Challenging because of wide variety of IoT devices and their increasing number and distribution across multiple network operators
- One approach is emulation: use real-world artifacts (malware, bots) in a controlled environment

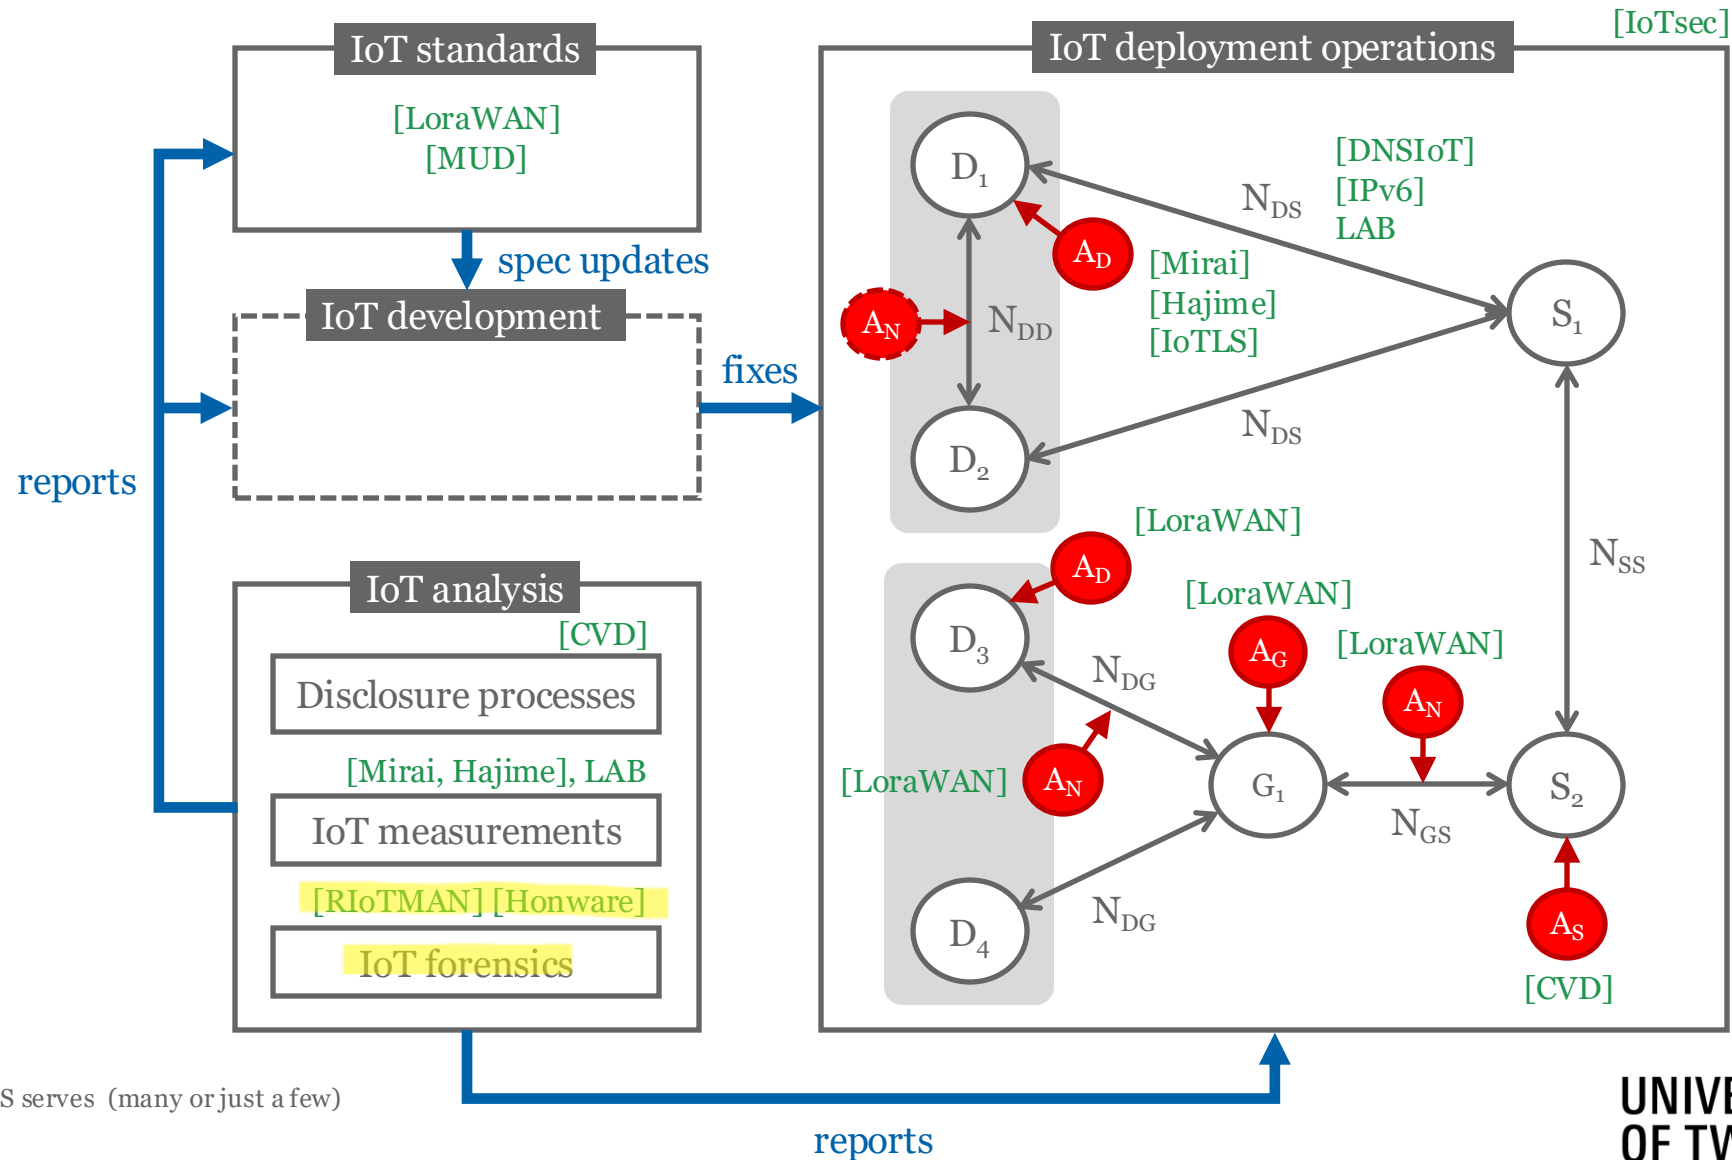


So that's why we selected today's papers for you

[RIoTMAN] A. Darki, and M. Faloutsos, “RIoTMAN: a systematic analysis of IoT malware behavior”, CoNEXT '20: Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies, November 2020

[Honware] A. Vetterl, and R. Clayton. “Honware: A virtual honeypot framework for capturing CPE and IoT zero days”, Symposium on Electronic Crime Research (eCrime), IEEE, 2019

SSI covers different parts of the IoT ecosystem



Today's learning objective

- After the lecture, you will be able to discuss mechanisms to analyze the behavior of IoT devices that have been infected with a bot/malware
- Contributes to SSI learning goal #1: “Understand IoT concepts and applications, security threats, technical solutions, and a few relevant standardization efforts in the IETF”

“RIoTMAN: a systematic analysis of IoT malware behavior”

16th International Conference on emerging Networking
EXperiments and Technologies (CoNEXT), November 2020

Get your phones ready!



1

Go to wooclap.com

2

Enter the event code in the top banner

Event code
GKZKMO



Enable answers by SMS



What struck you about the paper?


Challenge: profiling IoT malware

- What needs to be profiled?
- Why is profiling a challenge?
- Why do we need to solve it?


RIoTMAN: profiling IoT malware binaries

- What's their overall approach?
- What's the advantage of their approach?
- What malware states does RIoTMAN distinguish?

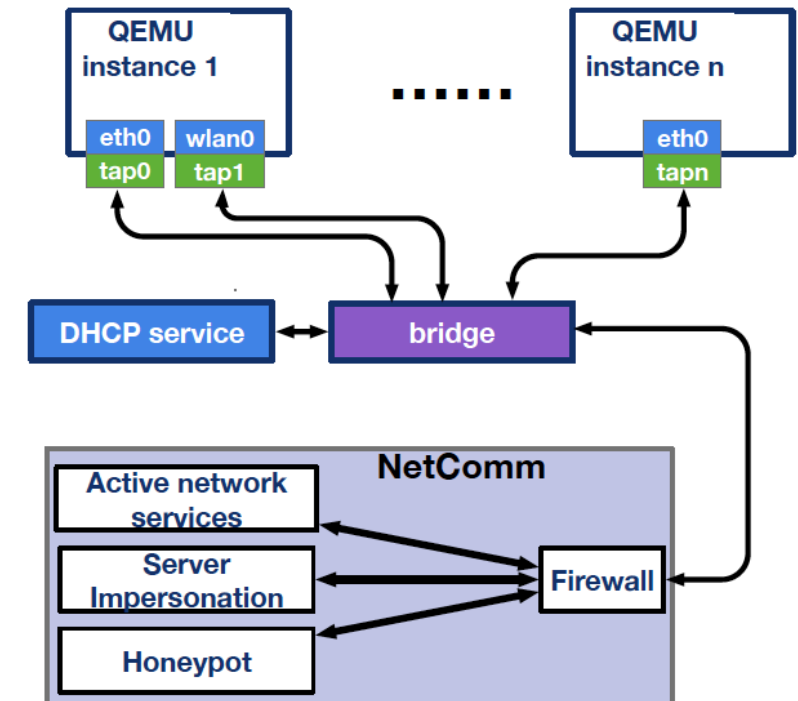
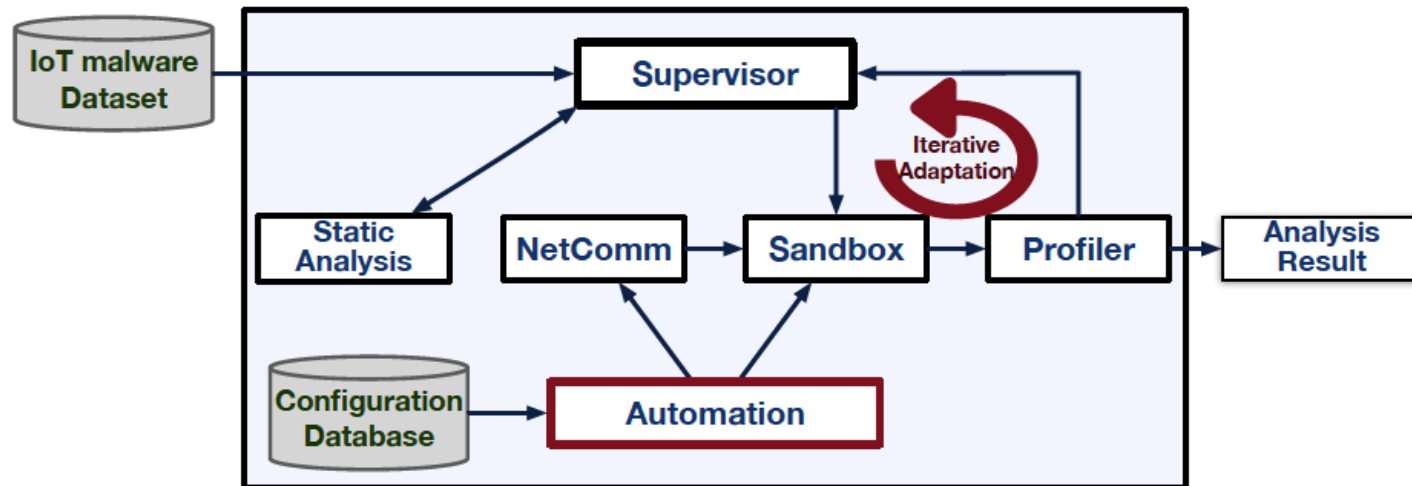
Example: Linux.Tsunami



Key measurement result – what are we looking at?



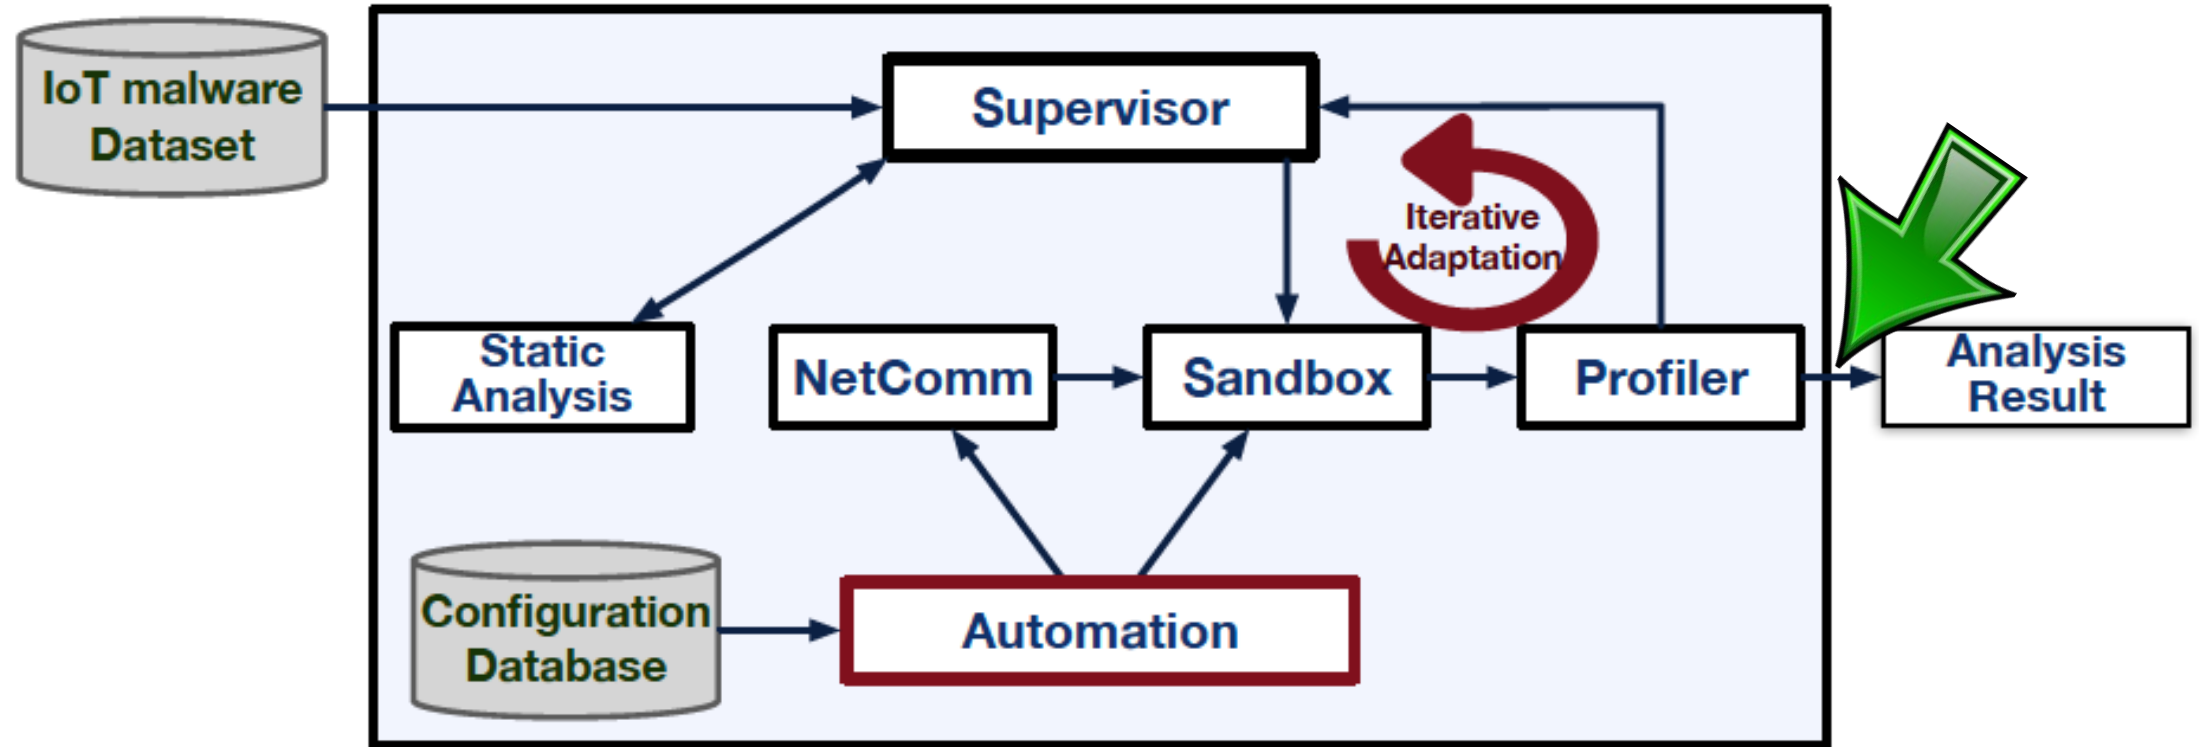
RIoTMAN emulation architecture



What are the responsibilities of the components?



RIoTMAN profiles





Emulation results

Total binaries	2885	
Activated	2688	93%
Engaged	2291	79%

Command Type	Malware	
Configuration or Report	1750	61%
Attack	2031	70%
Scanning	1842	64%
Termination	1684	58%



IoT malware behaviors – how can we leverage that?

C&C discovery

IP address	Single	2261
	Multiple	62
Domain	Fixed	257
	DGA	5

Malware Procedure	Most common techniques					
	Bin.	Technique 1	Bin.	Technique 2	Bin.	Technique 3
Infection	1676	Brute-force login	166	Exploit public facing apps	-	None observed
Persistence	375	Add routine in rc script	333	Add a job to cronjob	15	Specific to IoT device
Defense evasion	1494	Process masquerading	648	Malware binary removal	128	Software packing
Identifying device	1445	Use network config	843	Use config files	286	List processes in device
Impact on host	414	Block OS level access	413	Stop remote services	6	Bitcoin mining

Advanced behaviors



Limitations

- Linux-based IoT devices only
- They exclude botnets that use encryption, P2P botnets, and IPv6 communications

Key takeaways

- Dynamic analysis of IoT malware, limited manual effort
- Important to understand, detect, and mitigate IoT botnets at scale
- One piece of the “IoT botnet mitigation puzzle”
- Next challenge: how will RIoTMAN-like systems work in practice (higher TRLs)?



Coffee break

“Honware: a virtual honeypot framework
for capturing CPE and IoT zero days”
14th Symposium on Electronic Crime Research (eCrime), 2019

The Arms Race Between Hackers and Defenders

The Arms Race Between Hackers and Defenders

- IoT attacks evolve faster than defenses.
 - Attackers scan the internet for vulnerable devices 24/7.
 - Tools like Shodan and ZMap make it easy to find and target exposed systems.
 - New vulnerabilities are exploited within days often before anyone notices.



The Arms Race Between Hackers and Defenders

- IoT attacks evolve faster than defenses.
 - Attackers scan the internet for vulnerable devices 24/7.
 - Tools like Shodan and ZMap make it easy to find and target exposed systems.
 - New vulnerabilities are exploited within days often before anyone notices.
- Traditional detection relies on known signatures.
 - Prior solutions work only for previously seen malware.
 - Prior solutions are blind to novel exploits (i.e., zero-days).




The Arms Race Between Hackers and Defenders

- IoT attacks evolve faster than defenses.
 - Attackers scan the internet for vulnerable devices 24/7.
 - Tools like Shodan and ZMap make it easy to find and target exposed systems.
 - New vulnerabilities are exploited within days often before anyone notices.
- Traditional detection relies on known signatures.
 - Prior solutions work only for previously seen malware.
 - Prior solutions are blind to novel exploits (i.e., zero-days).
- We need better tools to shorten response time.
 - Faster detection means less time to cause damage.




Zero-Days: The Silent Killers in CPE and IoT



Zero-Days: The Silent Killers in CPE and IoT

- **Zero-day** = unknown vulnerability with no patch available. 
 - The defenders have had zero days to fix the vulnerability -> there are no known defences!

Zero-Days: The Silent Killers in CPE and IoT

- **Zero-day** = unknown vulnerability with no patch available. 
 - The defenders have had zero days to fix the vulnerability -> there are no known defences!
- IoT/CPE devices are rarely monitored for signs of compromise.
 - They aren't monitored in real-time. They don't log behaviors.
 - Hence, zero days can persist for months/years, turning millions of IoT into botnet nodes (Mirai).


Zero-Days: The Silent Killers in CPE and IoT

- **Zero-day** = unknown vulnerability with no patch available. 
 - The defenders have had zero days to fix the vulnerability -> there are no known defences!
- IoT/CPE devices are rarely monitored for signs of compromise.
 - They aren't monitored in real-time. They don't log behaviors.
 - Hence, zero days can persist for months/years, turning millions of IoT into botnet nodes (Mirai).
- **Honeypots** are key to discovering zero-days in the wild. 
 - They act as traps; they show early signs of compromise.

Honeypots: A Key Tool in the Defender's Arsenal


Honeypots: A Key Tool in the Defender's Arsenal

- Security traps to attract and observe attackers.
 - Appear like real systems but are isolated and monitored.
 - Allow attackers to engage without risking production environments.




Honeypots: A Key Tool in the Defender's Arsenal

- Security traps to attract and observe attackers.
 - Appear like real systems but are isolated and monitored.
 - Allow attackers to engage without risking production environments.
- Used for malware collection, attack pattern analysis.
 - Capture payloads, scripts, and exploit techniques in the wild.
 - Reveal CnC infrastructure and post-exploit behavior and methods.



Honeypots: A Key Tool in the Defender's Arsenal

- Security traps to attract and observe attackers.
 - Appear like real systems but are isolated and monitored
 - Allow attackers to engage without risking production environments
- Used for malware collection, attack pattern analysis.
 - Capture payloads, scripts, and exploit techniques in the wild.
 - Reveal CnC infrastructure and post-exploit behavior and methods.
- Especially useful when defenders lack prior knowledge.
 - Don't need to know the specific exploit or signature in advance.
 - Effective against zero-days and novel attack techniques.




But... Honeypots Have a Realism Problem

But... Honeypots Have a Realism Problem

- Many honeypots use fake or generic service responses.
 - They simulate basic behavior (e.g., fake login prompt, dummy web server).
 - Often don't run real firmware or OS-level services.


But... Honeypots Have a Realism Problem

- Many honeypots use fake or generic service responses.
 - They simulate basic behavior (e.g., fake login prompt, dummy web server).
 - Often don't run real firmware or OS-level services.
- Skilled attackers can fingerprint and avoid them.
 - Use timing attacks to tell if it is a trap.
 - Once detected, attackers may skip or behave differently.



But... Honeypots Have a Realism Problem

- Many honeypots use fake or generic service responses.
 - They simulate basic behavior (e.g., fake login prompt, dummy web server)
 - Often don't run real firmware or OS-level services.
- Skilled attackers can fingerprint and avoid them.
 - Use timing attacks to tell if it is a trap.
 - Once detected, attackers may skip or behave differently.
- Realism is crucial for detecting advanced threats.
 - Advanced malware often checks for real environments.
 - More believable honeypots attract deeper interactions.



Inside Honware: QEMU + Custom Kernel + Logging

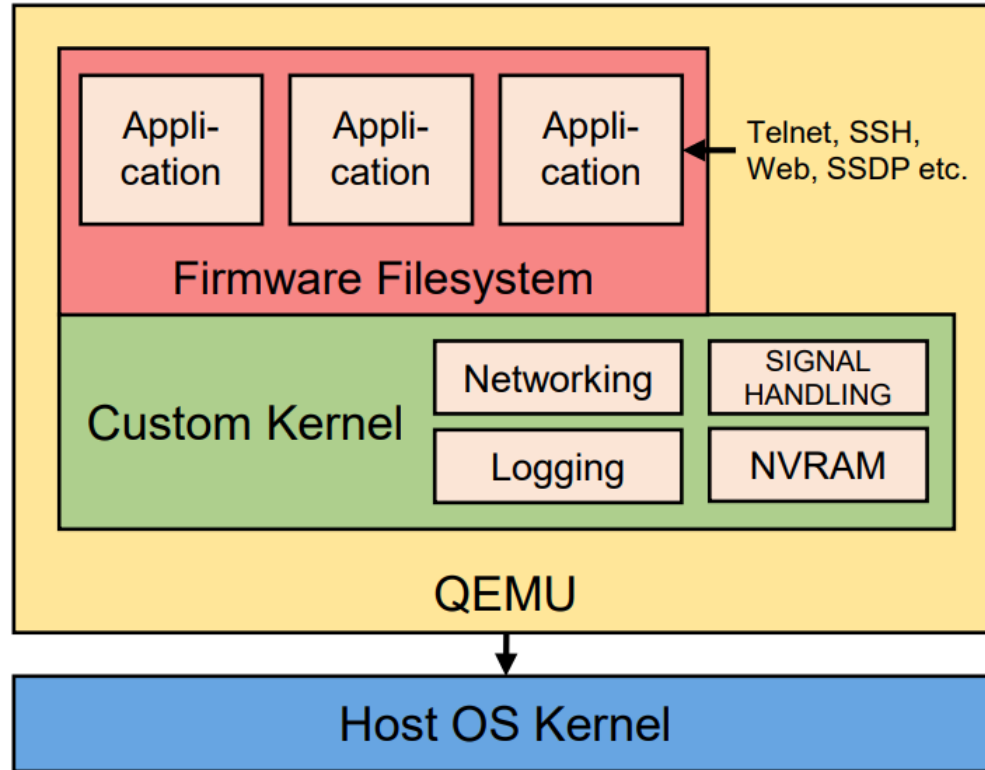


Fig. 1. Honware architecture overview: Honware consists of four main parts, a host operating system and kernel, Quick Emulator (QEMU), a custom kernel, and the firmware filesystem itself which contains specific applications such as telnet and web servers.

- Uses actual firmware images from manufacturers.
- Virtual setup, no need for physical hardware.
- Emulates actual services with deep interaction capabilities.
- QEMU supports multiple CPU architectures (MIPS, ARM).
- Custom kernel handles signal interception and module loading.
- Firmware filesystem is modified to run on the virtualized kernel
- Extensive logging: process IDs, commands, timestamps.

Emulated, Not Simulated: Why it matters?

- Simulation only imitates behavior (often poorly).
 - Many honeypots fake responses with hardcoded logic (i.e., *low-interaction* honeypots).
 - Attackers notice missing nuances, eg. protocol characteristics, file structure.
- Emulation produces realistic attack surfaces.
 - Emulation-based honeypots run actual services and software (i.e., *medium-interaction* or *high-interaction* honeypots).
 - Real vulnerabilities can still be exploited.
 - Malware behaves naturally (as on real devices).

Honware vs Firmadyne: Scalability and Reachability

- Honware extracted more firmware images (55% vs 35%).
- More devices were reachable over the network (40.9% vs 15.8%).

TABLE I
COMPARISON BETWEEN HONWARE AND FIRMADYNE: WE OBTAINED THE LIST OF FIRMWARE IMAGES (23 035) USED IN THE EVALUATION OF FIRMADYNE (2016-02) AND DOWNLOADED ALL THAT REMAINED ACCESSIBLE (8 387) IN 2019-03. WE USED FIRMADYNE AND HONWARE TO EXTRACT THESE AND TEST THEIR NETWORK REACHABILITY BY SENDING THEM ICMP ECHO REQUEST PACKETS.

# Brand	Available (2019-03/2016-02/ Δ)	Extracted (honw./firm./ Δ)	Network reach. (honw./firm./ Δ)
Total	8387/23035 14648↓	4650/2920 1730↑	1903/460 1443↑

More Real Services => Better Attacker Engagement

- More services detected by nmap:
Telnet, HTTP, UPnP, etc.
- Enables attack stages beyond login
brute-force.

TABLE II
COMPARING HONWARE AND FIRMADYNE: TOP 15 LISTENING SERV.

Prot.	Port/Service	Honware	Firmadyne	Δ
TCP	23/telnet	879	149	730↑
TCP	80/http	676	293	383↑
UDP	67/dhcp	316	160	156↑
UDP	1900/UPnP	239	128	111↑
UDP	53/various	239	174	65↑
TCP	3333/dec-notes	222	102	120↑
TCP	5555/freeciv	203	57	146↑
TCP	5431/UPnP	177	48	129↑
UDP	137/netbios	154	82	72↑
TCP	53/domain	139	73	66↑
TCP	443/https	107	105	2↑
UDP	5353/mdns	102	34	68↑
UDP	69/tftp	104	26	78↑
TCP	1900/UPnP	56	60	4↓
TCP	49152/UPnP	53	62	9↓


Case Study: Zero-day DNS Hijack on ipTIME Router

- What is DNS hijacking?
 - DNS = Phonebook of the Internet
 - DNS Hijacking -> Changing the "phonebook" entries
 - Eg, when you try to visit www.google.com, your router instead of translating it to 142.250.64.132 (legitimate address) it instead translates it to 1.2.3.4 (malicious address)
- Real-world zero-day attack captured by Honware
 - Attackers changed DNS settings to rogue servers
 - Traffic was redirected to malicious DNS resolvers


Honware emulated the real router firmware and captured the attacker's steps!

Can Honware be Fingerprinted?


- Attackers can probe for delays or odd behavior.
- Real devices have natural delays due to weak hardware.
- Cloud-hosted honeypots might respond too quickly.




(a) ASUS RT-AC52U FTP server: Time to welcome message




(b) ASUS RT-AC52U FTP server: Time between resource request (carriage return) and login message




(c) Zyxel VMG1312-B10A Telnet server: Time to telnet negotiation characters



(d) Zyxel VMG1312-B10A Telnet server: Time to Login message



(e) D-Link DIR-825 HTTPS server: Time to complete the TLS handshake



(f) D-Link DIR-825 HTTPS server: Time between ClientHello and resource received (web page)

Honware's Impact: Faster Detection, Safer Devices


- Core Technology:
 - Combines QEMU, custom Linux kernels, and real firmware filesystems.
- High Fidelity:
 - Emulates actual services (Telnet, HTTP, UPnP) with authentic behavior.
- Real-world impact:
 - Reduces attacker dwell time (the gap between breaking-in and getting detected).
 - Captures advanced malware and zero-days.

Could Honware be Deployed at Internet Scale?





Feedback?



Next up:

Written exam: Mon Jun 23, 08:45-10:45

Lab presentations: Fri June 27 and Mon
Jun 30, 8:45-12:30