

MADEA: A Malware Detection Architecture for IoT blending Network Monitoring and Device Attestation

Renascence Tarafder Prapty
University of California Irvine
rprapty@uci.edu

Rahmadi Trimananda
University of California Irvine
rtrimana@uci.edu

Sashidhar Jakkamsetti
Bosch Research
sashidhar.jakkamsetti@us.bosch.com

Gene Tsudik
University of California Irvine
gts@ics.uci.edu

Athina Markopoulou
University of California Irvine
athina@uci.edu

Abstract—Internet-of-Things (IoT) devices are vulnerable to malware and require new mitigation techniques due to their limited resources. To that end, previous research has used periodic Remote Attestation (\mathcal{RA}) or Traffic Analysis (\mathcal{TA}) to detect malware in IoT devices. However, \mathcal{RA} is expensive, and \mathcal{TA} only raises suspicion without confirming malware presence. To solve this, we design MADEA, the first system that blends \mathcal{RA} and \mathcal{TA} to offer a comprehensive approach to malware detection for the IoT ecosystem. \mathcal{TA} builds profiles of expected packet traces during benign operations of each device and then uses them to detect malware from network traffic in real-time. \mathcal{RA} confirms the presence or absence of malware on the device. MADEA achieves 100% true positive rate. It also outperforms other approaches with 160 \times faster detection time. Finally, without MADEA, effective periodic \mathcal{RA} can consume at least $\sim 14\times$ the amount of energy that a device needs in one hour.

I. INTRODUCTION

Internet-of-Things (IoT), namely smart devices are becoming increasingly widespread. They are encountered in many settings, including home, office, factory, public venues, and transportation of all sorts. Such devices are specialized sensors, actuators, or hybrids thereof; and their computational and communication resources are usually meager due to size, cost, and power constraints. Unsurprisingly, they represent an attractive set of targets for malware and other exploits. For example, the most infamous IoT malware incident was the Mirai botnet attack that infected hundreds of thousands of DVRs, IP cameras, routers, and printers in 2016 [1]. There are many other types of IoT-focused malware exploiting diverse and often new attack vectors. Malware’s job is made easier in terms of infection scalability by the general consumer tendency towards monocultures, as witnessed by the immense popularity of certain “smart” voice assistants, doorbells, cameras, and appliances. This motivates the need for new malware mitigation techniques geared for the IoT ecosystem, distinct from those used to combat more traditional malware [2].

Limitations of Prior Work. The research community has been keenly aware of the IoT security issues and many prior works proposed various mitigation techniques. A large body of such work is focused on Remote Attestation (\mathcal{RA}) – a means of measuring the internal software state of a remote device to determine whether it is infected [3], [4], [5], [6]. Most of prior work proposed and evaluated *new* \mathcal{RA} techniques but often side-stepped an important practical aspect: how to schedule the \mathcal{RA} ? Doing it infrequently makes instantaneous detection of malware impossible. In the best case, it increases the detection latency and, in the worst case, allows the malware to come and go between successive measurements and thus escape undetected. Meanwhile, doing it too often is impractical and wasteful since each \mathcal{RA} instance consumes resources by taking the attesting device away from its main task(s), as well as expending its power due to computation and communication.

Another major direction of prior work is detection of malicious activity via Traffic Analysis (\mathcal{TA}) [7], [8], [9], [10], [11]. Most \mathcal{TA} approaches rely on Machine Learning (ML) techniques, which are typically not lightweight and use many traffic features. Packets sent and received by IoT devices are captured and their relevant features are extracted and used to build data samples for training ML models. Then, during the monitoring phase, classifiers label each new data sample as either benign or suspicious. Nevertheless, \mathcal{TA} , by itself, only flags suspicious activity and does nothing to confirm whether a device is indeed infected. Note that some \mathcal{TA} techniques blacklist suspected devices without confirming infection and discard all of their traffic, which is overkill.

MADEA Approach. To address the above issues, this paper constructs MADEA, a system that blends \mathcal{RA} and \mathcal{TA} , resulting in a more comprehensive IoT malware detection than either of the two alone. To the best of our knowledge, MADEA is the first to leverage the benefits of both \mathcal{TA} and \mathcal{RA} . We provide a comparison of MADEA with prominent existing approaches of \mathcal{TA} [12], [9] and \mathcal{RA} [6] in Table I.

TABLE I
COMPARISON OF MADEA WITH PRIOR APPROACHES.

Approach	Search Vector	Instantaneous Detection	Malware Confirmation
MADEA	Network Packet & Device Memory	✓	✓
N-BaIoT	Network Packet	✓	✗
D ² IoT	Network Packet	✓	✗
VRASED	Device Memory	✗	✓

MADEA is aimed at a typical smart home/office IoT device deployment. Conveniently, traffic patterns of such devices tend to be limited and predictable. They usually communicate with only a handful of external endpoints and, as shown in prior work [13], [14], variations in packet sizes are minor. Infected devices show diverging traffic patterns by communicating with different IP addresses and using different packet sizes, typically talking to their Command-and-Control Center (CCC) or other infected devices.

MADEA consists of three main components: *Profiler*, *Monitor*, and *Attester*. Profiler trains one $\mathcal{T}A$ profile per device by collecting all packets generated by the device during a typical (benign) operation, and extracting and storing three features per packet: endpoint addresses, packet length, and direction. Next, Monitor detects potentially malicious activity in real time through *whitelisting*: a packet is considered as suspicious if it does not match the features of any packets stored in the corresponding device profile. Whenever that occurs, MADEA invokes Attester on the suspected device to determine whether it is infected. Infection can be detected by changes in the program memory of the device. If an infection is confirmed, MADEA reports this incident; otherwise, it includes the triggered traffic pattern in the benign patterns database. Our experiments demonstrate that MADEA achieves 100% True Positive Rate (TPR) with at most 1.2% false positive rate (FPR), on average.

Contributions. First, MADEA performs *lightweight TA* of IoT devices via *per-packet* detection, using only three features, to match each packet in real-time against device profiles while maintaining high detection accuracy. Notably, MADEA outperforms machine learning based state of the art approach, [12] with a $1.05\times$ better TPR and $160\times$ faster detection time.

Second, MADEA benefits from triggering $\mathcal{R}A$ only when a suspicious traffic pattern is detected. This also allows periodic/scheduled $\mathcal{R}A$ to be done infrequently, which avoids unnecessary $\mathcal{R}A$ instances that would otherwise waste device resources and take devices away from their actual primary tasks that might be safety-critical. In our evaluation, without the $\mathcal{R}A$ - $\mathcal{T}A$ combination, a camera with 1.84 W power consumption would need to consume *an additional ~ 26 Wh of power in a year for periodic attestation without any guarantee of detecting the malware.*

Furthermore, the entire MADEA system is made available as an open-source implementation at [15]. Its proof-of-concept

includes two components: (1) a prototype IoT device based on Raspberry Pi 4, in which the Attester module is implemented by extending the Raspbian OS to measure the currently running processes and (2) an emulated router, also based on Raspberry Pi 4, which profiles and monitors device traffic.

Organization. Section II provides background and related work in $\mathcal{R}A$ and $\mathcal{T}A$ for IoT malware. Next, Section III describes MADEA’s system and adversary models and Section IV lays out the high-level design and implementation details. Then, Section V presents the evaluation results, followed by Section VI, which discusses various aspects and limitations of MADEA. Section VII concludes the paper.

II. BACKGROUND & RELATED WORK

This section provides some background on targeted devices, $\mathcal{R}A$, and $\mathcal{T}A$, and on their interaction.

A. Targeted Devices

This work focuses on resource-constrained IoT devices specifically designed for sensing and actuation tasks, some of which perform safety-critical functions (e.g., fire/smoke detectors, smart locks, glass vibration alarms, and surveillance/motion cameras). They are equipped with microcontroller units (MCUs) such as ARM Cortex-M series or mid-range MCUs like ARM Cortex-A series. To support security features, some MCUs provide extended hardware features like memory protection units (MPUs) or memory management units (MMUs).

This work is applicable to a wide range of IoT devices equipped with some form of root of trust. However, very simple devices (so-called “bare metal”) that lack security features are considered out-of-scope.

Network Connectivity. IoT devices are often connected to the Internet and other peer devices, either directly or via an intermediary, such as a controller hub or a router. This paper focuses on devices that use routers for connectivity. Such devices are equipped with at least one network interface, such as WiFi, Bluetooth, cellular, wired Ethernet, or Zigbee. We focus on WiFi traffic since WiFi is the most common network interface present on consumer IoT devices [16]. Albeit, any other network media (wired or wireless) can be easily supported by MADEA.

B. Remote Attestation: $\mathcal{R}A$

$\mathcal{R}A$ is a security service that enables a trusted party ($\mathcal{V}rf$) to measure current software state of an untrusted remote device ($\mathcal{P}rv$). $\mathcal{R}A$ protocol is usually realized as follows:

- $\mathcal{V}rf$ sends an $\mathcal{R}A$ request to $\mathcal{P}rv$ that contains a challenge.
- $\mathcal{P}rv$ receives the request, computes a secure measurement of its current software state and sends it to $\mathcal{V}rf$.
- $\mathcal{V}rf$ verifies the result and decides whether $\mathcal{P}rv$ is in a valid/healthy state.

The measurement is obtained by: (1) hashing (using a suitable cryptographic hash function, such as SHA-256) the software, and then (2) computing either a Message Authentication Code (MAC) or a digital signature over the hash digest and the

challenge. Computing a MAC requires $\mathcal{P}rv$ to share a key with $\mathcal{V}rf$, while computing a signature requires $\mathcal{P}rv$ to have a private key with the corresponding public key known to $\mathcal{V}rf$. Both approaches require $\mathcal{P}rv$ to have some form of secure storage for the (shared or private) key. The measurement is typically computed within a Trusted Computing Base (TCB) that relies on trusted hardware and/or software, depending on the type of $\mathcal{R}A$: software-based, hardware-based, or hybrid. Hardware-based $\mathcal{R}A$ [17], [3], [18], [19] relies on dedicated secure hardware components, e.g., TPM [17] or Intel SGX [20]). However, such hardware is normally too costly (in multiple aspects) for most IoT devices. Software-based $\mathcal{R}A$ [21], [22], [5], [4], [23] requires no secure hardware features. However, it assumes an ideal environment with reliable communication even in the presence of a remote attacker, which is difficult to achieve in real world scenarios.

Hybrid $\mathcal{R}A$ [24], [6], [25], [26] claims security equivalent to hardware-based $\mathcal{R}A$, at a lower cost. Such methods require minimal hardware support and shift the bulk of complexity to the (trusted) software component. However, hybrid $\mathcal{R}A$ still requires hardware modification and, therefore, is not readily applicable to off-the-shelf IoT devices.

For our proof of concept implementation, we use a software-based $\mathcal{R}A$ method suitable for off-the-shelf consumer IoT devices. However, we expect that MADEA could be used in tandem with many $\mathcal{R}A$ methods. Section VI discusses the compatibility of MADEA with other $\mathcal{R}A$ techniques.

C. Traffic Analysis: $\mathcal{T}A$

In $\mathcal{T}A$, a specific software/hardware combination (device x) is profiled in a controlled setting to identify its traffic characteristics under normal operating conditions. These characteristics (also referred to as x 's traffic fingerprint or profile) are then used to identify x on any network by matching live traffic patterns to x 's fingerprint. Network fingerprints can include bandwidth, timing of TCP connections, port numbers, IP addresses, packet sizes, protocol identifiers, directions, etc. For example, [13] develops packet-level event signatures of IoT devices using only sequence numbers and the size and direction of TCP packets.

Some early $\mathcal{T}A$ techniques relied on deep packet inspection (DPI) [27], [28]. Surprisingly, DPI also initially proved effective for smartphones [29] and IoT [30], perhaps due to lax security practices in the early days of these technologies. However, gradually traffic encryption became more common [31], [32], [33]. Consequently, new fingerprinting techniques were proposed that identify individual functionality on IoT devices even if the application-layer payload is encrypted, e.g., [34], [13], [35], [36]. Network fingerprinting has various useful applications such as intrusion detection, traffic prioritization, and user activity profiling. However, attackers also use it to reconnoiter a network for an effective attack.

Although they vary greatly in terms of complexity and collected traffic features, these techniques focus only on a subset of IoT traffic that uniquely identifies a specific device or a certain event occurring on it. In contrast, MADEA monitor

component operates as a whitelist. Therefore, it needs to learn the complete (normal) traffic profile of an IoT device, which it achieves by blending a simple passive $\mathcal{T}A$ technique (to flag suspicious patterns) with $\mathcal{R}A$ (to perform infection checking), as described in Section IV).

Another body of prior work focused on detecting malicious activity via $\mathcal{T}A$ based on Machine Learning (ML) techniques. [7], [8], [12], [10], [11]. These techniques tend to use numerous traffic features, resulting in heavy-weight ML models. Typically, packets to/from IoT devices are captured and relevant features are extracted in order to build data samples. During the training phase, ML models are trained with labeled benign and malicious data samples. During the monitoring phase, classifiers label each data sample as either benign or suspicious. A related approach involves ML models trained with only benign data [9], [37]. In contrast, MADEA uses very few traffic features which results in a very lightweight, heuristic-based detection model.

D. $\mathcal{R}A$ - $\mathcal{T}A$ Synergy

One important practical aspect of $\mathcal{R}A$ mostly avoided by prior work is: *exactly when to perform it?* Two intuitive choices are to perform $\mathcal{R}A$ at: (1) scheduled, e.g., at fixed and/or random intervals, and/or (2) on-demand, e.g., before using $\mathcal{P}rv$. Both approaches have pros and cons. Though easy to configure, the former has two disadvantages: First, it leaves potentially long time gaps between successive $\mathcal{R}A$ instances which allows malware to come and go in the interim and remain undetected. Second, it is generally quite wasteful due to $\mathcal{R}A$'s resource consumption on $\mathcal{P}rv$. If the device performs safety-critical tasks, performing $\mathcal{R}A$ can divert it from such tasks or at least slow it down.

Meanwhile, most $\mathcal{T}A$ techniques do not go beyond detection of suspicious activity, i.e., they do not inspect the device to confirm that it is actually infected, though some involve suspected device quarantine. This is clearly insufficient for comprehensive IoT malware mitigation. Even quarantine represents potential over-reaction since the suspected device might in fact not be infected. The only way to make sure is to inspect the suspected device, e.g., via $\mathcal{R}A$. MADEA uses lightweight $\mathcal{T}A$ to flag suspicious IoT traffic patterns and triggers $\mathcal{R}A$ for all suspected devices. If an $\mathcal{R}A$ result reflects malware presence, the user (i.e., administrator/owner) is immediately notified. Otherwise, in case of a false positive (no malware detected based on the $\mathcal{R}A$ result), MADEA "learns" by whitelisting the observed traffic pattern.

III. SYSTEM & ADVERSARY MODELS

This section describes MADEA system model and adversarial assumptions.

A. System Model

MADEA considers four entities: an IoT device, a gateway, a user device, and a remote server.

IoT device is a typical smart device (as in Section II-A) that is the main target of the adversary. It is assumed to be

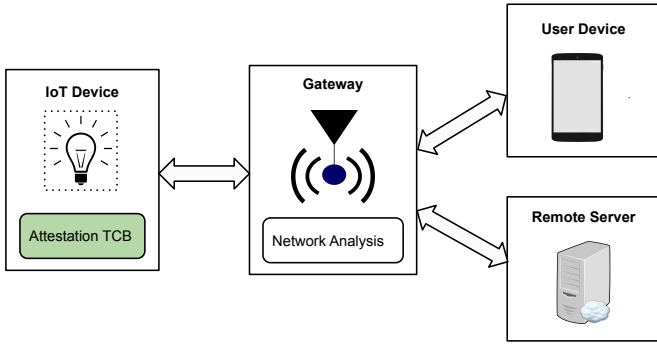


Fig. 1. MADEA system model.

installed in a private space (such as a home/office setting) or a public place (such as a grocery store or a hotel), connected to a gateway. A gateway can be a local router, an IoT controller hub, or any other gateway that connects a set of IoT devices directly to (or towards) the Internet. We assume that an IoT device is equipped with a Trusted Computing Base (TCB), that performs device attestation, as in Section II. A remote server is a remote back-end server (generally hosted in the cloud) maintained by the IoT service provider. The server is responsible for registering, updating, and assisting the IoT device. A user device is a personal item (such as a smartphone, smartwatch, laptop, or tablet) that controls the IoT device via a dedicated app or a web interface.

MADEA works in an IoT deployment with at least one IoT device, one or more user devices, and one gateway, as shown in Figure 1. User devices communicate with IoT devices either directly or via the server. All packets from or to IoT devices pass through the gateway: we assume that peer-to-peer communication is not allowed. The gateway captures packets and makes decisions based on their features, without keeping the state of on-going traffic flows/connections. In the context of \mathcal{RA} , IoT devices and gateway play the roles of \mathcal{P}_{rv} and \mathcal{V}_{rf} , respectively.

B. Adversary Model

We make the following assumptions about adversarial capabilities:

- *Remote adversary*: infects IoT devices with malware over the network without having physical access to them, for example, by exploiting a software vulnerability in an application running on IoT devices or similar to Mirai, through telnet by exploiting default/weak telnet credentials.
- *TCB*: IoT device's TCB is trusted and not subject to compromise. Each TCB has a distinct pair of public and secret keys for computing signatures as part of \mathcal{RA} . Each public key is known to the gateway.
- *Malware-generated traffic*: We assume that malware must generate some amount of traffic, in order to communicate with its Command-and-Control Center (CCC). In most cases, malware activities are coordinated by CCC – a

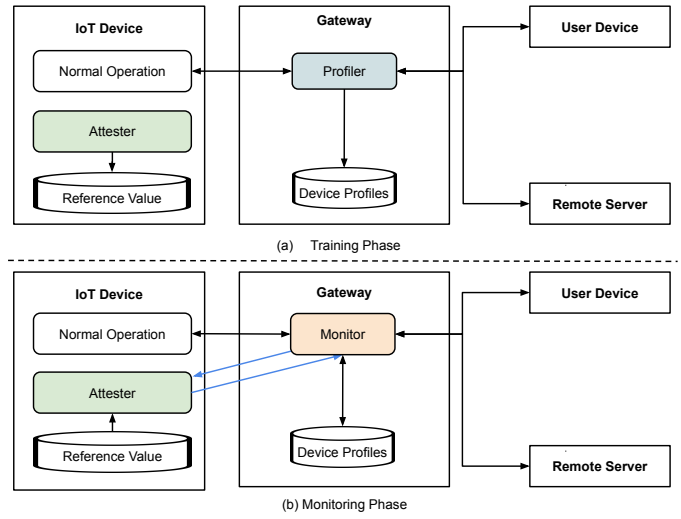


Fig. 2. MADEA overview: blue arrows represent the feedback loop between Attester and Monitor.

remote server maintained by the adversary which controls a network of infected devices. This network is usually called a botnet and infected devices are called zombies/bots. CCC sends various commands to be executed by IoT devices, e.g., to download code, infect other devices, or launch attacks.

- *Trusted entities*: We assume that the gateway, server, and user devices are sufficiently powerful computing platforms that can protect themselves from malware and remain trusted. Moreover, since the service provider is trusted, each newly deployed IoT device is assumed to be initially healthy. DNS resolver and other routing protocols are considered to be uncompromised: an adversary cannot spoof IP address of CCC to match the IP address of the remote server.

1) *DoS and Other Attacks*: An adversary can launch a DoS attack on an IoT device by flooding it with a high volume of traffic from the outside, i.e., from beyond the gateway. Alternatively, the adversary can use infected IoT devices to launch a DoS attack on other peer (healthy) devices. MADEA detects both DoS attacks. Since we assume that the gateway runs on a trusted high-end platform, DoS attacks on it are considered out-of-scope. DoS attacks on IoT devices are further discussed in Section VI.

We consider all physical attacks (i.e., those requiring adversary's physical presence) to be out of the scope of MADEA. The same holds for malware that does not involve any of its own incoming or outgoing traffic. (Malware that targets actuation capabilities and acts autonomously would fall into this category.) Section VI discusses this in more detail.

IV. DESIGN & IMPLEMENTATION

This section discusses the main components of MADEA: Profiler, Monitor, and Attester (see Figure 2).

A. Profiler

The first step of MADEA is building a device profile for each IoT device during the training period when each device is newly deployed and assumed to be malware-free. This task is performed by Profiler running on the gateway. All normal events of an IoT device (e.g., turning on/off as well as brightness and color control in case of a smart bulb) are triggered manually using the corresponding controller app in the user device. User manuals and options in controller apps are consulted to develop the list of all possible events for the IoT device. This list is easy to build since MADEA targets simple to moderately complex devices such as smart bulb, smart plug, robot vacuum cleaner, camera, etc. Profiler can build profiles for multiple IoT devices simultaneously, and the training is performed online.

Profiler refreshes its state after every 20,000 packets to avoid stalling caused by large PCAP files due to limited RAM and processing power. It stores IoT device profiles in a database shared with Monitor, using a hash map for fast access. The MAC addresses of devices to be trained are also stored in the database and read by Profiler at the beginning of each iteration, allowing for dynamic changes in the list in each iteration.

Device Profile. A device profile refers to the profile of the expected network traffic to/from the IoT device. Figure 3 shows example profiles for two commodity IoT devices: Lumiman bulb [38], and Sensi thermostat [39]. Each profile can be viewed as a table, where each row represents a unique group of packets called profile entry. Thus, each profile consists of a list of profile entries. Each profile entry consists of the device MAC address, external address, packet direction and length.

External address is the IP address belonging to the entity that exchanges packets with the IoT device. It refers to either its corresponding server or user device. This is the most important feature since malware-generated packets usually have different IP addresses. This holds as long as malware does not compromise DNS and/or the routing protocol.

Legitimate packets sent to/received from a specific external address by an IoT device usually have slight length variations. It is worth mentioning that MADEA can effectively detect malware packets without this feature. However, we include it because packet length provides additional information while being very lightweight.

Packet direction reflects whether the IoT device is the source or destination of the packet, i.e.,

- *Server-to-Device*: sent by the server to the IoT device.
- *Device-to-Server*: sent by the IoT device to the server.
- *Device-to-User*: sent from the IoT device to the user device in the same network.
- *User-to-Device*: sent from the user device to the IoT device within the same network.

Note that we choose four different values for packet direction to make profile entry more intuitive and informative for human users. Just two directions (to/from the IoT device) would suffice for MADEA to work properly. Moreover, MADEA does not require users to be in the same LAN as the IoT device.

A remote user device communicates with the IoT device through the server and is treated the same as communications between the IoT device and server.

Packet direction plays an important role since suspicious packets generated by the IoT device might cause more concern than those sent to the IoT device. For example, packets sent to the IoT device may indicate malware payload being downloaded or malicious command. In contrast, packets sent from the IoT device might indicate exfiltration of sensed data or malware being sent to another IoT device. Moreover, used together, packet direction and length make the packet matching criteria stricter. Oftentimes, packets of a specific length are exchanged only in one direction and not the other.

Other features that could be taken into account include the packet sequence, frequency, protocol, or packet volume per unit of time. We do not include them in MADEA because, as we found out, the set of three aforementioned features suffices to detect malware with high accuracy (see Section V). Furthermore, this makes the packet matching algorithm in Monitor lightweight and efficient.

Packets with the same external address, direction, and length are grouped together. This grouping is independent of packet sequencing and packets in the same group can be captured at any time during the training phase. We assume that each profile contains entries for all possible groups of packets that can be transmitted to and from the IoT device.

Building Profile. Figure 2 shows how Profiler builds a profile.

PyShark [40] library captures and processes packets from the IoT device via its LiveCapture interface. A BPF filter, configured with the MAC address of the IoT device, filters captured traffic.

To construct a profile entry, five fields are extracted from each packet: source IP address, destination IP address, source MAC address, destination MAC address, and packet length. Further, MADEA assumes that the IoT device is connected to the LAN defined by the gateway. Thus, either the source or destination MAC address must be the actual MAC address of the IoT device, whereas the other MAC address is the MAC address of the gateway.

Configuration Packets. Configuration packets of standard protocols, such as NTP, DHCP, DNS, MDNS, BOOTP, ARP, EAPOL, are not included in the profile. Such packets are typically transient and do not provide meaningful insight into the recurrent nature of IoT device traffic. However, MADEA considers both TCP and UDP packets for the profile. Many prior network fingerprinting mechanisms [13], [14] consider only TCP packets. The fact that many modern IoT devices use UDP makes MADEA more inclusive and robust.

Determining Direction & External Address. Packets are associated with a specific IoT device using its MAC address. This way, if the IoT device changes its IP address, its profile is not affected. MAC addresses are used to determine directions and IP addresses. If the list of MAC addresses in the database contains the source MAC address of the packet, then the destination IP address is considered the external address. In

Dir.	External Address	Len.	Dir.	External Address	Len.
S-D	m2.tuyaus.com	235	S-D	icda.sensicomfort.com	155
D-S	m2.tuyaus.com	235	D-S	icda.sensicomfort.com	539
D-U	192.168.4.111	196	U-D	192.168.4.1	86
S-D	ec2-35-164-195-39.us-west-2.compute.amazonaws.com	123	S-D	icda.sensicomfort.com	58
D-S	a3.tuyaus.com	251	D-S	icda.sensicomfort.com	129
			D-S	224.0.0.251	46

(a)

(b)

Fig. 3. Partial profiles of: (a) Lumiman Bulb, (b) Sensi Thermostat (S-D=Server-to-Device, D-S=Device-to-Server, D-U=Device-to-User, U-D=User-to-Device). Note: See Appendix A for full profile of Sensi Thermostat. For more examples, see [41].

that case, the direction is *Device-to-Server*. On the other hand, the source IP address becomes the external address if the destination MAC address is in the list of MAC addresses and the direction is *Server-to-Device*. Finally, if both source and destination IP addresses are local addresses, the direction is either *Device-to-User* or *User-to-Device*.

Mapping IP Addresses to Hostnames. In our experiments with numerous devices, we observed that IP addresses of servers can change frequently, leading to unstable profiles. Hostnames, on the other hand, remain unchanged or change rarely. Therefore, using hostnames as external addresses is preferable. However, this requires an IP-to-hostname mapping in MADEA, as hostnames are resolved to IP addresses through DNS packets, and the associated IP addresses are used in subsequent communication.

MADEA builds IP-address-to-hostname mappings using DNS response packets instead of creating profile entries for captured DNS packets. It extracts hostnames from the *Query* field and resolved IP addresses from the *Answer* field to create a mapping from each IP address to its corresponding hostname.

If source/destination IP address is not a local address, MADEA tries to find the corresponding hostname from the IP-address-to-hostname mapping. If this fails, MADEA performs reverse DNS lookup using the IP address. If a hostname is returned, MADEA uses it; otherwise, it uses the IP address.

Example 1. In this example, we use profiles of Lumiman Bulb and Sensi Thermostat from Figure 3.

Lumiman Bulb. In the first profile entry, the bulb is the destination of the packet. Thus, the direction is *Server-to-Device* and the source address is the external address. The source IP address is resolved to a hostname (i.e., *m2.tuyaus.com*) using IP-address-to-hostname mapping and that hostname denotes the external address. During the development of MADEA, the hostname *m2.tuyaus.com* resolves to 21 different IP addresses, e.g., *54.212.163.173*. Interestingly, apart from Lumiman Bulb, another IoT device (i.e., HBN plug) also exchanges packets with the same external address. Thus, these two devices are likely to be associated with the same manufacturer and talk with the same servers. The second profile entry has the

same external address and length as the first entry, but the opposite direction. In the third profile entry, both source and destination IP addresses are local and the bulb is the source. Thus, the direction is *Device-to-User* and the IP address of the user device becomes the external address. In the fourth profile entry, the external address is a hostname resolved through reverse DNS lookup (i.e., the source IP address cannot be resolved using IP-address-to-hostname mapping) and the direction is *Server-to-Device*. The original source IP address is 35.164.195.39 and the hostname implies that the server is hosted on AWS cloud.

Sensi Thermostat. In the Sensi Thermostat profile in Figure 3(b), the external address of the last profile entry is an IP address and the direction is *Device-to-Server*. It means that the destination IP address cannot be resolved even through reverse DNS lookup.

B. Monitor

Monitor attempts to match each packet to/from the IoT device to some profile entry. We assume that all normal events have been reflected in the profile during the training phase. Hence, if no match is found for a packet, MADEA assumes it to be suspicious (i.e., potentially generated by malware). Appendix B provides more insight into the differences between observed malware-generated and normal packets. One concern about this approach is that updates in either the IoT device or server can change the benign traffic pattern. Please refer to Section VI for details about how MADEA handles updates. Any device can be added to the list of IoT devices to be monitored as Monitor runs constantly. Similar to Profiler, it refreshes its state after every 20,000 packets.

Matching a Packet. During the monitoring phase, Monitor constructs an entry for every packet following the same approach that Profiler performs (see Section IV-A). Then, Monitor checks for a matching profile entry based on MAC address, external address, packet direction and length. If no exact match is found, it performs partial string matching for external address with at least 80% similarity while ensuring the top-level domain name is the same.

For instance, *ec2-35-164-195-39.us-west-2.compute.amazonaws.com* will match with *ec2-13-235-158-83.ap-south-1.compute.amazonaws.com* since both partially match and the top-level domain is *amazonaws.com*. 80% similarity is chosen based on empirical observation. This percentage gives a negligible number of false positive packets for hostnames found through reverse DNS lookup while keeping the level of similarity significant. We define false positive as a benign packet being detected as suspicious by MADEA. With 85% and 90% similarity matching, we observe more false positives. Partial string matching is considered because the external address changes slightly periodically in some IoT devices, e.g., Ring doorbell and Blink camera. MAC address, and packet direction and length are still matched exactly.

Example 2. In this example, let us consider a packet captured by Monitor. The following five properties are extracted from it:

- 1) source MAC address: *64:16:66:49:3e:cb*.
- 2) destination MAC address: *dc:a6:32:ce:31:63*
- 3) source IP address: *192.168.4.230*
- 4) destination IP address: *35.186.98.64*
- 5) packet length: *284*

MADEA maintains a list of MAC addresses of the monitored devices in the database. It also maintains a hash map of all profile entries. Monitor observes that the list contains source MAC address. Thus, Monitor determines that the device MAC address is *64:16:66:49:3e:cb* and the packet direction is *Device-to-Server*. Furthermore, the destination IP address is translated to the hostname *oculus9353-us1.dropcam.com* that becomes the external address. Monitor forms an entry as a tuple (MAC address: *64:16:66:49:3e:cb*, external address: *oculus9353-us1.dropcam.com*, packet direction: *Device-to-Server*, packet length: *284*) and finds a matching profile entry from the hashmap. If a matching profile entry is found for Monitor's entry, it means that such a packet has been seen by Profiler during the training phase and, thus, the packet is considered legitimate. The Monitor's entry may fail to match any profile entry because of two reasons: (1) there is no profile entry that can be matched with Monitor's entry when Monitor inspects the MAC address, external address, and packet direction; or (2) there are profile entries that can be matched with Monitor's entry when Monitor inspects MAC address, external address, and packet direction; however, the packet length in profile entry does not match the packet length in Monitor's entry.

We have observed that mismatched packet lengths contribute the most to FPR. However, the security impact of mismatched packet lengths is relatively low, as long as external address is matched. Thus, MADEA offers relaxed matching, in which Monitor excludes packet length when matching Monitor's entry to a profile entry.

Finally, in the case that Monitor detects a suspicious packet, i.e., a packet whose entry does not match with any profile entry, it then calls a program called Attester to verify the state of the IoT device.

C. Attester

Attester runs locally on the IoT device and it is responsible for attesting the software state of the IoT device when requested. It works on the assumption that the software state of the IoT device is deterministic, measurable, and unchanged unless the IoT device is infected by malware. It is to be noted that the software state can also change due to legitimate software updates. Section VI discusses how this case is handled by MADEA.

During the training phase, the software state of a healthy IoT device is measured as the *reference value* that indicates that the IoT device is not compromised. During the monitoring phase Attester measures the software state when called by Monitor. Monitor's call includes a new challenge every time. The

measured value is compared with the reference value locally by Attester as sending the full measurement result to Monitor can be expensive in terms of network bandwidth. If the values match, Attester reports an acknowledgment indicating that the IoT device is healthy. Otherwise, it reports that the IoT device has been infected by malware. In this case Attester's report includes the divergent measurement result for further analysis. The Attester's report is always signed using the IoT device's secret key, and the signature includes the challenge so that the gateway can authenticate that the report is fresh and indeed generated by the IoT device.

In MADEA proof of concept implementation, Attester is a Python program running on a Raspberry Pi emulated smart bulb. Traditional device attestation that involves measuring the full program memory is not suitable for RPi's virtual memory which is managed by OS. We can't determine the physical addresses of processes in memory, and, thus, cannot determine which memory region to attest. Additionally, measuring the large RAM (in the scale of gigabytes) of a RPi would be slow and impractical. Therefore, Attester attests the binaries of the currently running processes on the RPi instead.

Attester works on the assumption (backed by observation 4, see V-C) that the processes running on the IoT device, at any given time, are deterministic and limited in number. Hence, it is possible to maintain a list of expected processes and hash values of their corresponding binaries. If any IoT device is infected by malware, a new malware process would be running on it or the malware would embed itself in an existing process which would change the hash value of the corresponding binary file. Attester inspects this hash value change to detect malware on the IoT device. This attestation is supported by extensive experiments conducted on the emulated bulb for evaluation. Section V provides more details on these experiments.

During the training period, Attester builds a list of expected processes and corresponding hash values for the IoT device. When called by Monitor, Attester obtains a list of all currently running processes by calling the Linux command `ps aux | less > [output file]` which directs the output to a text file. Next, the output file is parsed to extract the PID of each process. PID is used to obtain the location of the binary file for each process. This location is in the format of `/proc/[PID]/exe`. Finally, MADEA calculates the hash value of the binary file of each process using the command `sha256sum /proc/[PID]/exe`.

This list is then compared against the list of expected processes. If there is any new process in the current process list or an expected process with a wrong hash value, the IoT device will be considered compromised (i.e., infected by malware) and Attester reports this to Monitor: Attester will also send the list of the new processes and/or the processes with the wrong hash values. Otherwise, if there is no new process and all current processes have the same hash values as calculated during training, Attester will send the attestation report back to Monitor indicating that the IoT device is healthy.

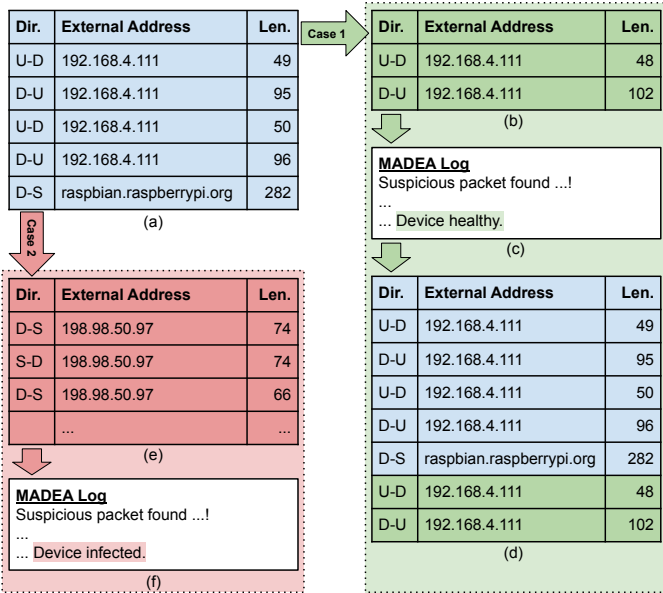


Fig. 4. **Case 1**—Malware not detected by Attester: (a) the RPi smart bulb’s profile, (b) packets for a new command (i.e., status message), (c) Monitor’s log, and (d) the RPi smart bulb’s profile updated by Monitor. **Case 2**—Malware detected by Attester: (a) the RPi smart bulb’s profile, (b) packets generated by malware, and (c) Monitor’s log.

D. Feedback Loop

Monitor and Attester complement each other and create a feedback loop, which is a key contribution of MADEA. During monitoring, if no matching profile entry is found for Monitor’s entry of a particular packet, Monitor calls Attester to attest the IoT device. If Attester finds no malware (Case 1), it reports to Monitor that the IoT device is healthy and Monitor considers the triggering packet as a false positive. Next, it updates the IoT device’s profile by creating a corresponding profile entry and appending it to profile to avoid future false positives prompted by similar packets. If Attester finds malware on the IoT device (Case 2), it reports to Monitor that the IoT device has been infected.

Section V provides an example of the feedback loop operation; Figure 4 illustrates the process. Currently, Monitor logs the report message for further analysis by device owner. Ideally, Monitor would send the report directly to the owner as an alert for Case 2.

V. EVALUATION

We next present the evaluation for MADEA. Section V-A outlines the experimental setup for MADEA. Section V-B presents the evaluation of MADEA in monitoring devices in three datasets. Section V-C presents a case study to illustrate the effectiveness of MADEA’s feedback loop, which is a main contribution in this paper. Section V-D characterizes MADEA’s performance overheads. Section V-E compares the performance of MADEA with three prior works. Finally, Section V-F illustrates the energy savings over periodic attestation due to the feedback loop.

A. Experimental Setup

Our experimental setup follows the MADEA system model and architecture shown in Figures 1 and 2 respectively. We use a Raspberry Pi 4 (RPi4) [42] machine as the gateway, configured as a WiFi access point. It is based on quad-core Cortex-A72 (i.e., ARMv8-M architecture) 64-bit Systems-on-Chip (SoC). The board runs at 1.5GHz with 8GB LPDDR4-2400 SDRAM. Additionally, it is equipped with 2.4GHz and 5.0GHz IEEE 802.11B/g/n/ac WLAN, Bluetooth 5.0, double-true Gigabit Ethernet, and a 32GB SD card as storage. Thus, with these specifications, RPi4 is comparable with commodity home routers. Profiler and Monitor are implemented as Python programs which run on the gateway.

Sandbox for Malware. The gateway is connected to a NETGEAR 4G LTE Broadband Modem (LM1200) [43] with Ethernet cable to access the Internet. NETGEAR Broadband Modem uses a T-Mobile SIM card for network connectivity. With this, we create a sandbox for safety reasons since we experiment with malware. Connecting the gateway directly to our institutional LAN would potentially make other LAN devices vulnerable to malware infection.

Smart Home Setup. We use a smart home setup to evaluate MADEA. However, MADEA can be implemented at scale to suit the deployment in a commercial setting. Our setup consists of 13 off-the-shelf devices and one RPi4 emulated smart bulb as the IoT device. Table II provides details about the devices.

Off-the-shelf devices are chosen based on their popularity and ratings on Amazon [44]. We also consider a wide range of functionality: from plugs to cameras and vacuum-cleaning robots (i.e., iRobot Roomba). The RPi smart bulb is implemented on a RPi4 with 2GB RAM. The RPi bulb is controlled by a Python program. This program can accept requests and send back responses over UDP connections. It has two main parts: bulb controller and attester. The bulb controller receives commands from the user device over the internet and turns on/off a LED based on the received command. Attester receives attestation request from the gateway, performs attestation, and sends back attestation result.

Additionally, we implement three malware CCC (Mirai, Bashlite, Lizkebab) on a Linux laptop with an Intel(R) Core(TM) i7-8550U CPU running at 1.80GHz and equipped with 8GB RAM. This Linux laptop is also connected to the gateway. Alternatively, malware CCC can be hosted outside the network, but we choose to host it inside our sandbox network due to safety issues.

We first collect network traffic (in the training phase) from the 13 devices to build MADEA dataset and, in turn, the device profiles without introducing any malware infection. Table III shows the numbers of packets from the traffic collection and profile entries for each device profile.

B. Monitoring Devices

Monitoring Healthy Devices. First, we use the MADEA dataset to evaluate the performance of MADEA. In particular, we evaluate the performance of Monitor in monitoring

TABLE II
SMART HOME DEVICES USED IN MADEA EVALUATION.

Device Name	Details
Amazon Smart Plug	Amazon Smart Plug
HBN Smart Plug Mini	HBN Smart Plug Mini 15A
Ring Doorbell	Ring Doorbell with camera (2020 release)
Blink Mini Camera	Blink Mini compact indoor plug-in smart security camera
Nest Camera	Nest indoor security camera
Lumiman Smart Bulb	Lumiman multi-colored and warm-to-cool white LED light bulb
Kasa Smart Bulb	Kasa Smart Bulb KL110 white light bulb
LIFX Smart Bulb	LIFX A19 multi-colored light bulb
ULTRALOQ U-Bolt Pro	ULTRALOQ U-Bolt Pro smart lock with Bluetooth and keypad
Sensi Thermostat	Emerson Sensi smart thermostat
Nest Protect Smoke Alarm	Nest Protect smoke alarm, smoke detector, and carbon monoxide detector
Rachio Sprinkler	Rachio Sprinkler third generation with 8-zone sprinkler controller
iRobot Roomba	iRobot Roomba 690
RPi Smart Bulb	Emulated smart bulb on Raspberry Pi 4 board

TABLE III
EXPERIMENTAL RESULTS OF MONITORING UNINFECTED DEVICES IN SMART HOME SETUP. THE LAST ROW GIVES THE **TOTAL** FOR # TRAINING PACKETS (TRNG. PKTS.), # MONITORING PACKETS (MNT. PKTS.) AND # PROFILE ENTRIES, AND AVERAGE FOR FPR.

Device Name	# Trng. Pkts.	# Mnt. Pkts.	# Profile Entries	False Positive Rate (FPR)	
				End. & Len.	End.
Amazon Plug	8,317	11,017	119	0.87%	0.64%
HBN Plug Mini	22,096	24,688	81	0.01%	0.00%
Ring Doorbell	143,925	91,576	1590	0.39%	0.01%
Blink Mini Cam.	52,839	1,477,004	132	0.00%	0.00%
Nest Cam.	283,798	645,873	734	0.02%	0.02%
Lumiman Bulb	42,274	49,855	113	0.07%	0.00%
Kasa Bulb	8,667	6,434	146	1.09%	0.93%
LIFX Bulb	1,207	408	21	0.00%	0.00%
ULTRALOQ U-Bolt Pro	7,487	55541	137	0.05%	0.03%
Sensi Thermostat	1,492	796	52	1.13%	0.00%
Nest Protect Smoke Alarm	806	89	59	1.12%	1.12%
Rachio Sprinkler	3,508	4,614	112	4.76%	0.11%
iRobot Roomba	8,821	6,756	207	0.65%	0.25%
Total/Average	585,239	2,374,651	3,503	0.78%	0.24%

uninfected devices using the created device profiles. Table III provides FPR for each device. We measure FPR in two scenarios: (1) matching external address, packet direction, and length (see **End. & Len.**); and (2) matching external address and packet direction (see **End.**) in Monitor’s entry with the corresponding features in any profile entry in the device profile. MADEA incurs low FPR for all devices: 0.78% and 0.24% on average for both scenarios respectively.

Public Dataset 1: BehavIoT. We also evaluate MADEA using a public dataset called BehavIoT [45]. The BehavIoT dataset contains 32 uninfected IoT devices, among which three devices (i.e., Amazon Plug, Ring Doorbell, and Nest Thermostat) are present in our setup as well. This dataset contains both LAN and WAN traffic for all possible events of each device. We only use the LAN traffic in alignment with

TABLE IV
FALSE POSITIVE RATE OF UNINFECTED DEVICES FROM BEHAVIOT DATASET. THE LAST ROW GIVES THE **TOTAL** FOR # TRAINING PACKETS (TRNG. PKTS.), # MONITORING PACKETS (MNT. PKTS.) AND # PROFILE ENTRIES (PROF. ENTS.), AND AVERAGE FOR FPR.

Device Name	# Trng. Pkts.	# Mnt. Pkts.	# Prof. Ents.	False Positive Rate (FPR)	
				End. & Len.	End.
Tp-Link Bulb	14,928	15,026	94	0.07%	0.00%
Tp-Link Plug	3,467	3,474	51	0.14%	0.00%
Magic Home Strip Light	2,592	2,597	19	0.00%	0.00%
Smart Life Bulb	5,599	5,511	23	0.18%	0.00%
Nest Thermostat	8,197	7662	245	1.59%	0.00%
Bulb1	1,860	1,830	16	0.05%	0.00%
Govee LED	3,600	3,600	12	0.36%	0.36%
Gosund Bulb	3,587	3,627	21	0.03%	0.00%
Meross Door Opener	621	647	21	2.94%	0.00%
Amazon Plug	1,297	1,263	41	0.39%	0.00%
Echo Spot	17,822	15,827	618	3.28%	1.24%
Echo Show	17,127	16,700	575	3.03%	0.25%
Echo Dot	3,803	4,247	131	2.33%	0.47%
Philips Hub	17,936	17,793	254	0.35%	0.00%
SmartThings Hub	1,568	1,498	81	0.47%	0.00%
Aqara Hub	862	801	40	0.49%	0.00%
IKEA Hub	2,170	2,172	51	0.00%	0.00%
Wemo Plug	11,390	10,795	87	0.17%	0.00%
Wansview Cam.	290,372	293,693	1104	0.05%	0.00%
Lefun Cam.	105,723	105,362	1,717	0.36%	0.00%
Dlink Cam.	229,960	236,918	1,576	0.04%	0.00%
Microseven Cam.	290,639	298,642	1,668	0.13%	0.00%
Ring Door Bell	238,173	234,583	3241	3.10%	0.00%
Ring Cam.	311,528	301,520	12387	3.82%	0.00%
Ubell Doorbell	237,400	230,272	2581	0.06%	0.00%
Wyze Cam.	478,060	506,694	1,146	0.39%	0.00%
Yi Cam.	189,899	189,395	1,365	0.19%	0.00%
Tuya Cam.	477,039	497,072	293	0.01%	0.00%
ICsee Doorbell	268,628	271,611	1635	0.06%	0.00%
SwitchBot Hub	728	722	11	0.00%	0.00%
iKettle	18	33	6	6.06%	0.00%
Google Home Mini	8,089	7,900	861	9.29%	0.51%
Total/Average	3,244,673	3,289,487	31,971	1.23%	0.09%

our setup. We divide the LAN traffic per device per event into two equal parts. One part is used by Profiler to build device profiles (i.e., training phase). The other part is used to evaluate the performance of Monitor (i.e., monitoring phase).

Table IV provides the number of packets, profile entries, and FPR of each healthy device from this dataset. MADEA shows relatively high FPR for a few devices from this dataset when the external address, packet direction, and length are matched. However, FPR is zero or very low when only the external address and packet direction are matched. Most of these devices are comparatively complex devices (e.g., cameras and voice assistants). We observe that exact length matching is less effective for complex devices that may generate a more diverse traffic pattern compared to simpler devices (e.g., plugs and bulbs). Instead, if we relax the matching criteria (e.g., allowing a slight variation of 2 bytes in packet length for Google Home Mini), FPR is almost totally eliminated.

Monitoring Infected Devices. We use the RPi smart bulb for

this purpose. In the training phase, we build the device profile for the RPi smart bulb without infecting it with malware. Next, we infect the RPi smart bulb with Mirai, Lizkebab, Bashlite, and 100 randomly selected ARM malware binaries from [2]. Monitor is able to detect suspicious packets with **100% TPR** for all malware in the following aspects, namely: (1) the network payload sent from the malware CCC during bot installation, (2) the periodic heart-beat communication between the bot and the CCC, and (3) the attack commands sent from the CCC to the bot.

Public Dataset 2: Cryptojacking. Finally, we evaluate MADEA using another public dataset that we call the Cryptojacking dataset [7]. The Cryptojacking dataset consists of both benign and IoT cryptojacking malware traffic collected from a Raspberry Pi, a laptop, a LG Smart TV, and a server. We exclude server data from our evaluation since it is not an IoT device. For each of these devices, MADEA successfully detects malware traffic with **100% TPR**.

Thus, our evaluation of MADEA on the three datasets and two scenarios (i.e., healthy and infected devices) indicates that MADEA is effective in profiling and monitoring IoT devices.

C. Effectiveness of Feedback Loop

We use the RPi smart bulb to evaluate the feedback loop’s effectiveness. During the training phase, the bulb receives *turn on*, *turn off* commands and sends responses. Entries for the corresponding packets generated by the two commands are included in the device profile of the RPi smart bulb. In the monitoring phase, we create two scenarios: (1) *healthy bulb (Case 1)*: the user device sends a new command to the RPi smart bulb. The RPi smart bulb is uninfected at this stage, (2) *infected bulb (Case 2)*: the RPi smart bulb becomes infected with malware and malware tries to communicate with its CCC.

Healthy Bulb (Case 1). During the training phase the commands and responses exchanged between the user device and RPi smart bulb, when turning on and off the RPi smart bulb are:

- **Command:** `turn_on` (packet length=49)
response: `{“status”:“200”,“message”:“Light bulb turned on.”}` (packet length=95)
- **Command:** `turn_off` (packet length=50)
response: `{“status”:“200”,“message”:“Light bulb turned off.”}` (packet length=96)

From Figure 4 we can see that the profile of the RPi smart bulb includes profile entries for these packets. However, during the monitoring phase, another previously unseen command (i.e., `status` message) is exchanged:

- **Command:** `status` (packet length=48)
response: `{“status”:“200”,“message”:“Light bulb is currently off.”}`
(packet length=102)}

Monitor detects the command packet as suspicious and requests attestation from Attester inside the RPi smart bulb. Attester confirms that the RPi bulb is healthy since it does not find any new process or any expected process with a

wrong hash value. As a result, Monitor adds the command packet to the bulb’s profile. The same process is repeated for the response packet. Afterward, subsequent `status` command and response packets are treated as benign and not subject to further attestation.

Infected Bulb (Case 2). The RPi smart bulb is intentionally infected with a malware binary. This binary is randomly selected from the malware repository of [2]. Figure 4 shows some packets exchanged between the malware CCC and the zombified RPi smart bulb. All of those packets are confirmed as suspicious through the feedback loop. However, we illustrate the process for the first packet only.

Figure 4 shows that external address of the packet is `198.98.50.97`, the packet length is 74 and direction is *Device-to-Server*. However, no profile entry in the RPi smart bulb’s profile matches those values. Thus, Monitor detects the packet as suspicious and calls Attester. Attester reports that new processes are running on the RPi smart bulb, which confirms that the RPi smart bulb is infected by malware, as shown in the warning message in the log file.

D. Performance Overheads

The *TA* technique deployed by MADEA to detect anomaly in the network traffic is very lightweight. The fourth column in Tables III and IV denotes the number of profile entries, i.e., maximum search space for matching packets of a specific device. We observe that FPR is consistently low and TPR is consistently high in experimental settings and public datasets.

Runtime Overhead. In addition, the detection is fast. The total time for MADEA to extract endpoints and length properties from a packet, classify the packet as benign or suspicious, and verify the detection through the feedback loop is on average 1.6 ms. Note that this is an upper bound since our proof of concept implementation is in Python. An implementation in low-level languages such as C/C++ would result in better timing performance.

Storage Overhead. We can also determine the absolute upper bound of the total number of device profile entries stored in the memory of the router during the monitoring phase: $n = d * e * l$, where n represents the number of profile entries of all devices being monitored, d represents the number of devices being monitored, e represents the average number of packet direction-external address pairs for each device, and l represents the average number of unique packet lengths per packet direction-external address pair.

The average value of l is 23 for the MADEA dataset and 58 for the BehavIoT dataset. The average value of e is 13 for the MADEA dataset and 18 for the BehavIoT dataset.

E. Comparison with Prior Works

In this section, we compare the performance of MADEA with two *TA*-based [9], [12] and one *RA*-based [6] systems from prior work. Table V presents the details. MADEA’s TPR is 100%, which means that MADEA can detect every malware-generated packet correctly. In other words, MADEA is more robust than DIoT [12] and the same as N-BaIoT [9],

TABLE V
PERFORMANCE COMPARISON OF MADEA WITH PRIOR APPROACHES.

Approach	# Required Features	TPR	FPR	Detection Time (ms)
MADEA	3	100.0%	0.78%	1.6
N-BaIoT [9]	114	100.0%	0.007%	174
DIoT [12]	7	95.6%	0%	256
VRASED [6]	N/A	100.0%	0%	attestation freq. + 450

VRASED [6]. The attack/malware infection detection time of MADEA is far better than all other approaches. This is due to the simple heuristic algorithm of comparing only three features from each packet. MADEA is able to detect the very first packet generated by malware, and it requires only 1.6 ms on average to do so. In comparison, N-BaIoT needs 174 ms and DIoT requires 256 ms for detection. The heavy machine learning models used in N-BaIoT (Autoencoder) and DIoT (Gated Recurrent Unit) cause these longer detection times. The detection time of VRASED depends on the frequency of remote attestation, i.e., how soon it is performed after infection, plus the time required to complete attestation. Additionally, the attestation time is proportional to memory size since VRASED calculates HMAC over the entire memory as part of the operation. It takes 450 ms to complete attestation for 4KB memory. This attestation time is significantly longer compared to MADEA’s. MADEA’s FPR is slightly higher compared to the other approaches. However, this FPR for off-the-shelf devices is overestimated since it is calculated without considering the feedback loop. The same packet is counted as false positive many times. For example, the packet with external address *www.google.com*, packet direction *Server-to-Device*, and length *409* is counted 13 times in the calculation of the FPR of the Google Home Mini device from BehavIoT dataset. With the feedback loop, the packet would only be considered suspicious once.

F. Energy Savings

In this section, we attempt to quantify the energy saved by the feedback loop of MADEA over plain periodic remote attestation. For this purpose, we consider two IoT devices: (1) DCS-930L Security Camera [46], a higher-end device that was compromised by Mirai [12]; and (2) TI MPS430, a bare metal IoT device, on top of which VRASED is implemented. Table VI shows the summary of the energy consumption rates for periodic \mathcal{RA} on these two devices in a year.

DCS-930L camera is not capable of \mathcal{RA} . However, if it was, we assume that its \mathcal{RA} time would be similar to the emulated RPi bulb since both have comparable operating systems and processing powers. The power consumption of the DCS-930L camera is 1.84 W [47], whereas the \mathcal{RA} time is considered to be 1.6 ms. This results in 0.49 mWh energy consumption per \mathcal{RA} . If \mathcal{RA} is performed once every 5 minutes to detect an attack within a reasonable time frame, the total energy consumption to perform \mathcal{RA} over a year becomes 5,157.9 mWh. Energy consumption for “somewhat instantaneous” detection (i.e., within 1 minute of infection) is even higher,

TABLE VI
COMPARISON ON TOTAL YEARLY ENERGY CONSUMPTION OF PERIODIC \mathcal{RA} (WITHOUT MADEA) FOR VARIOUS INTERVALS.

Device Name	Energy Consumption (mWh) for Attestation Intervals				
	1 hour	30 mins	10 mins	5 mins	1 min
DCS-930L Cam.	423.9	859.6	2,578.9	5,157.9	25,789.4
TI MPS430	77.8	157.7	473.0	946.1	4,730.4

namely 25,789.4 mWh or ~ 26 Wh. *In other words, the camera would need at least $\sim 14\times$ the amount energy that it needs to operate in one hour, just to perform \mathcal{RA} in a year without any guarantee of actually detecting the malware.*

TI MSP430 consumes 1.2 mWh power in active mode [48], and the attestation time of VRASED is 450 ms. This results in 0.009 mWh energy per attestation, 946.1 mWh over the year when the attestation period is 5 minutes and 4,730.4 mWh for an attestation period of 1 minute. This is a critical energy burden for an ultra-low power device that requires only 400 μA current during the active period.

Energy consumption of periodic attestation can be drastically reduced by adopting MADEA since MADEA triggers attestation *only* when there is a suspicious packet because of the rare events: actual malware infections or firmware updates.

VI. DISCUSSION AND LIMITATIONS

A. MADEA’s Scope

Closing the Loop. We acknowledge that MADEA does not quite “close the loop” on malware detection. The missing component is the feedback from scheduled \mathcal{RA} to \mathcal{TA} . Ideally, whenever scheduled/periodic (i.e., not triggered by \mathcal{TA}) \mathcal{RA} detects malware, \mathcal{TA} heuristics would be adjusted based on recent traffic logs that must contain infection-relevant traffic. Otherwise, the device must have been infected via direct/physical means. This is part of future work.

Malware without Network Traffic MADEA cannot detect malware that sits on a device silently without generating any network traffic. However, if such malware is transmitted to the device through the network, then MADEA catches it during transmission. Additionally, certain types of Mirai-variant malware have the capability to remain dormant for a specific duration before becoming active to initiate a DoS attack. In such scenarios, the MADEA can detect the malware when it becomes active and starts generating network traffic.

B. Attestation TCB

In MADEA Attester is implemented by extending the operating system (OS). This increases the TCB size by including the OS in the TCB and makes it vulnerable. To mitigate this, a TPM [17] can be integrated and a secure boot process can be configured to verify the TCB binary’s validity. This ensures that the OS always boots in the expected state. However, advanced attacks such as privilege escalation and runtime control-flow attacks can still compromise the TCB, which is out of scope for this paper.

C. Compatibility with Other Attestation Techniques

There are more secure ways of implementing attestation in IoT devices besides the OS-based Attester module. Notably, two options would be: (1) using off-the-shelf ARM TrustZone, or (2) using theoretical Root-of-Trusts (RoTs) such as SMART [24] and Sancus [3].

ARM TrustZone-M [49] provides secure isolation between normal and secure regions of memory for devices with ARM Cortex-M23/33/55 microcontroller units (MCUs). Installing the attestation TCB in the secure region prevents illegal access or modification by programs in the normal region. Further, the TCB can verify if the device is running expected user binary by accessing and measuring the normal region when an attestation request is received. We implemented such a technique on NXP LPC55S69-EVK [50] board with ARM Cortex-M33. However, we could not test and evaluate the prototype due to lack of real-world malware on low-end devices. Higher-end processors such as ARM Cortex-A32/72 are equipped with TrustZone-A [51], Trusted Execution Environment (TEE) extensions, which can access the physical memory of normal regions, but it is difficult to implement an attestation scheme that can measure user processes. This is because the normal region's separate Memory Management Unit (MMU) virtualizes user processes for the normal OS. Creating an attestation mechanism for TrustZone-A without trusting the normal OS is a non-trivial task. For ultra-low-power MCUs like TI MSP430 [52] and AVR ATmega8 [53], several RoTs [24], [3], [6], [26] propose attestation techniques based on custom hardware extensions. Since MADEA only relies on Attester as a black box, it can seamlessly integrate with these RoTs without modifications.

D. Physical Attacks

There are two types of physical attacks: non-invasive and invasive. Both are considered out of the scope of this paper.

Non-invasive attacks involve physical reprogramming of device software using direct/wired interfaces, such as USB/UART, SPI, or I2C. These attacks are challenging to detect, particularly when the attestation TCB becomes compromised or when the measurement values are maliciously updated to reflect the corrupted state of the applications being attested. However, well-known security mechanisms, such as secure boot, can mitigate these attacks.

Invasive attacks, on the other hand, involve attempting physical damage to the device, such as inducing hardware faults, tampering with hardware components, or exploiting physical side channels to extract secrets. Protection against invasive attacks can be achieved through standard tamper-resistant techniques [54].

E. Mitigation of DoS attacks

There are several ways an adversary can launch a DoS attack on an IoT device. Recall from Section III-B that an adversary can overwhelm an IoT device by sending a large volume of packets, or it can utilize the infected devices to flood the network with excessive traffic, affecting other devices in the same network. These attacks can be detected when the gateway

observes abnormal traffic to/from an IoT device, triggering attestation. Subsequently, since an IoT device is unresponsive during the attack, the attestation request will be disregarded, and after a timeout, MADEA will raise a failure.

Next, the adversary can also trick the gateway by sending abnormal packets to an IoT device and forcing the gateway to perform attestation repeatedly. Repeated attestation not only consumes power but can also potentially degrade the functionality of an IoT device. Although such attacks are out of the scope of MADEA, we note that they can be mitigated by implementing rate-limiting to perform attestation only a fixed number of times within a certain interval.

Furthermore, if the adversary has prior knowledge of the IoT device profile, it can launch a DoS attack by sending a burst of packets that conform to the profile. For instance, the adversary can construct malicious packets with the same packet length and forged source address to match one of the profile entries. In such cases, Monitor would allow the traffic to pass, resulting in a successful attack. Such attacks are also out-of-scope. Note that detecting such traffic would require stateful profile monitoring to check the timing interval between consecutive packets. However, implementing such monitoring is computational and storage intensive. Considering the general IoT traffic observed in our evaluation, we believe the benefits of deploying such resource-intensive measures are not substantial enough to justify their implementation.

F. Software Updates

There can be two different types of software updates: (1) the server's software changes resulting in different external address and/or packet lengths. However, the software state of the IoT device remains unchanged; or (2) the IoT device's software state changes. Similar to a new command that generates benign traffic (see Section V-C), software updates of the first type can be handled by the feedback loop. When a software update for the IoT device occurs, it triggers Monitor with a new packet. Monitor then calls Attester on the IoT device to perform attestation, which does not detect any change in the software state. Finally, this new packet is included as part of the IoT device's profile. It can also be handled manually by providing the gateway with the updated profile information. However, software updates of the second type cannot be incorporated automatically by MADEA. This is because when called by Monitor, Attester sees that the software state does not match the reference value. As a result, the legitimate new packet is considered malicious. To solve this problem, the server must provide the updated reference value to the IoT device and gateway in this case.

G. IoT Devices in Ad Hoc Network

IoT devices in ad hoc mode can exchange packets without going through the gateway, making it impossible for Profiler or Monitor to capture and process them. Such devices are out of the scope of MADEA. However, we assume that the owner of the IoT device is trusted. Thus, in other words, the IoT device

is configured to operate in infrastructure mode that allows all packets to go through the gateway.

VII. CONCLUSION

We present MADEA, a system that combines attestation and network traffic analysis to detect malware in IoT devices. MADEA stores benign traffic patterns from IoT devices that are usually limited and predictable. On the other hand, devices with malware infection exhibits network traffic patterns that are likely divergent from the usual traffic pattern. When unusual patterns are detected, MADEA performs \mathcal{RA} on the IoT device that produces suspicious traffic pattern. If malware infection is confirmed, MADEA will report this incident; otherwise, it will whitelist the traffic pattern as benign traffic. Our evaluation shows that MADEA is both effective and efficient in profiling and monitoring uninfected and infected devices. It achieves 100% TPR with low false positive rate (FPR) in most cases with low runtime (i.e., ~ 1.6 ms traffic classification time) and storage overheads. Finally, it outperforms the other approaches in terms of efficiency and effectiveness.

REFERENCES

- [1] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. Understanding the mirai botnet. In *26th {USENIX} security symposium ({USENIX} Security 17)*, pages 1093–1110, 2017.
- [2] Omar Alrawi, Charles Lever, Kevin Valakuzhy, Kevin Snow, Fabian Monrose, Manos Antonakakis, et al. The circle of life: A {Large-Scale} study of the {IoT} malware lifecycle. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 3505–3522, 2021.
- [3] Job Noorman, Jo Van Bulck, Jan Tobias Mühlberg, Frank Piessens, Pieter Maene, Bart Preneel, Ingrid Verbauwhede, Johannes Götzfried, Tilo Müller, and Felix C. Freiling. Sancus 2.0: A low-cost security architecture for iot devices. *ACM Trans. Priv. Secur.*, 20(3):7:1–7:33, 2017.
- [4] Michele Grisafi, Mahmoud Ammar, Marco Roveri, and Bruno Crispo. {PISTIS}: Trusted computing architecture for low-end embedded systems. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 3843–3860, 2022.
- [5] Mahmoud Ammar, Bruno Crispo, and Gene Tsudik. Simple: A remote attestation approach for resource-constrained iot devices. In *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPs)*, pages 247–258. IEEE, 2020.
- [6] Ivan De Oliveira Nunes, Karim Eldefrawy, Norrathep Rattanavipanon, Michael Steiner, and Gene Tsudik. VRASED: A verified hardware/software co-design for remote attestation. In *USENIX Security*, 2019.
- [7] Ege Tekiner, Abbas Acar, and A Selcuk Uluagac. A lightweight iot cryptojacking detection mechanism in heterogeneous smart home networks. In *Proc. of the ISOC Network and Distributed System Security Symposium (NDSS)*, 2022.
- [8] Gonzalo Marín, Pedro Casas, and Germán Capdehourat. Deep in the dark-deep learning-based malware traffic detection without expert knowledge. In *2019 IEEE Security and Privacy Workshops (SPW)*, pages 36–42. IEEE, 2019.
- [9] Yair Meidan, Michael Bohadana, Yael Mathov, Yisroel Mirsky, Asaf Shabtai, Dominik Breitenbacher, and Yuval Elovici. N-baiot—network-based detection of iot botnet attacks using deep autoencoders. *IEEE Pervasive Computing*, 17(3):12–22, 2018.
- [10] Ibrahim Alrashdi, Ali Alqazzaz, Esam Aloufi, Raed Alharthi, Mohamed Zohdy, and Hua Ming. Ad-iot: Anomaly detection of iot cyberattacks in smart city using machine learning. In *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*, pages 0305–0310. IEEE, 2019.
- [11] Marcin Woźniak, Jakub Silka, Michał Wiczorek, and Mubarak Al-rashoud. Recurrent neural network model for iot and networking malware threat detection. *IEEE Transactions on Industrial Informatics*, 17(8):5583–5594, 2020.
- [12] Thien Duc Nguyen, Samuel Marchal, Markus Miettinen, Hossein Feridooni, N Asokan, and Ahmad-Reza Sadeghi. Diot: A federated self-learning anomaly detection system for iot. In *2019 IEEE 39th International conference on distributed computing systems (ICDCS)*, pages 756–767. IEEE, 2019.
- [13] Rahmadi Trimananda, Janus Varmarken, Athina Markopoulou, and Brian Demsky. Packet-level signatures for smart home devices. In *Network and Distributed Systems Security (NDSS) Symposium*, volume 2020, 2020.
- [14] TJ OConnor, Reham Mohamed, Markus Miettinen, William Enck, Bradley Reaves, and Ahmad-Reza Sadeghi. Homesnitch: Behavior transparency and control for smart home iot devices. In *Proceedings of the 12th conference on security and privacy in wireless and mobile networks*, pages 128–138, 2019.
- [15] Madea source code. <https://github.com/sprout-uci/madea>.
- [16] Very. Bluetooth vs. wi-fi for iot: Which is better? <https://www.verytechnology.com/iot-insights/bluetooth-vs.-wi-fi-for-iot-which-is-better>, 2019.
- [17] Trusted Computing Group. Trusted platform module (tpm), 2017.
- [18] Chongkyung Kil, Emre C Sezer, Ahmed M Azab, Peng Ning, and Xiaolan Zhang. Remote attestation to dynamic system properties: Towards providing complete system integrity evidence. In *2009 IEEE/IFIP International Conference on Dependable Systems & Networks*, pages 115–124. IEEE, 2009.
- [19] Jonathan M McCune, Yanlin Li, Ning Qu, Zongwei Zhou, Anupam Datta, Virgil Gligor, and Adrian Perrig. TrustVisor: Efficient TCB reduction and attestation. In *IEEE S&P '10*, 2010.
- [20] Intel. Intel Software Guard Extensions (Intel SGX). <https://software.intel.com/en-us/sgx>.
- [21] Rick Kennell and Leah H Jamieson. Establishing the genuinity of remote computer systems. In *USENIX Security Symposium*, 2003.
- [22] Arvind Seshadri, Adrian Perrig, Leendert Van Doorn, and Pradeep Khosla. SWATT: Software-based attestation for embedded devices. In *IEEE Symposium on Research in Security and Privacy (S&P)*, pages 272–282, Oakland, California, USA, 2004. IEEE.
- [23] Virgil D Gligor and Shan Leung Maverick Woo. Establishing software root of trust unconditionally. In *NDSS*, 2019.
- [24] Karim Eldefrawy, Gene Tsudik, Aurélien Francillon, and Daniele Perito. SMART: Secure and minimal architecture for (establishing dynamic) root of trust. In *NDSS*, 2012.
- [25] Ferdinand Brasser, Brahim El Mahjoub, Ahmad-Reza Sadeghi, Christian Wachsmann, and Patrick Koeberl. Tytan: tiny trust anchor for tiny devices. In *Proceedings of the 52nd Annual Design Automation Conference, San Francisco, CA, USA, June 7-11, 2015*, pages 34:1–34:6. ACM, 2015.
- [26] Patrick Koeberl, Steffen Schulz, Ahmad-Reza Sadeghi, and Vijay Varadarajan. TrustLite: A security architecture for tiny embedded devices. In *EuroSys*, 2014.
- [27] Andrew W. Moore and Konstantina Papagiannaki. Toward the accurate identification of network applications. In Constantinos Dovrolis, editor, *Passive and Active Network Measurement*, pages 41–54, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [28] Justin Ma, Kirill Levchenko, Christian Kreibich, Stefan Savage, and Geoffrey M. Voelker. Unexpected means of protocol inference. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, IMC '06*, page 313–326, New York, NY, USA, 2006. Association for Computing Machinery.
- [29] Stanislav Miskovic, Gene Moo Lee, Yong Liao, and Mario Baldi. Appprint: Automatic fingerprinting of mobile applications in network traffic. In Jelena Mirkovic and Yong Liu, editors, *Passive and Active Measurement*, pages 57–69, Cham, 2015. Springer International Publishing.
- [30] Xuan Feng, Qiang Li, Haining Wang, and Limin Sun. Acquisitional rule-based engine for discovering internet-of-things devices. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 327–341, Baltimore, MD, August 2018. USENIX Association.
- [31] Abbas Razaghpanah, Arian Akhavan Niaki, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Johanna Amann, and Phillipa Gill. Studying TLS Usage in Android Apps. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*,

CoNEXT '17, page 350–362, New York, NY, USA, 2017. Association for Computing Machinery.

[32] Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monrose. SoK: Security Evaluation of Home-Based IoT Deployments. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1362–1380, 2019.

[33] Danny Yuxing Huang, Noah Aporthe, Frank Li, Gunes Acar, and Nick Feamster. IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 4(2), June 2020.

[34] TJ OConnor, Reham Mohamed, Markus Miettinen, William Enck, Bradley Reaves, and Ahmad-Reza Sadeghi. HomeSnitch: Behavior Transparency and Control for Smart Home IoT Devices. In *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '19, page 128–138, New York, NY, USA, 2019. Association for Computing Machinery.

[35] Abbas Acar, Hossein Fereidooni, Tigist Abera, Amit Kumar Sikder, Markus Miettinen, Hidayet Aksu, Mauro Conti, Ahmad-Reza Sadeghi, and Selcuk Uluagac. Peek-a-boo: I see your smart home activities, even encrypted! In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '20, page 207–218, New York, NY, USA, 2020. Association for Computing Machinery.

[36] Janus Varmarken, Jad Al Aaraj, Rahmadi Trimananda, and Athina Markopoulou. FingerprinTV: Fingerprinting Smart TV Apps. *Proceedings on Privacy Enhancing Technologies*, 2022(3), 2022.

[37] Valerian Rey, Pedro Miguel Sánchez Sánchez, Alberto Huertas Celdrán, and Jérôme Bovet. Federated learning for malware detection in iot devices. *Computer Networks*, 204:108693, 2022.

[38] Lumiman smart bulb. <https://www.amazon.com/Lights-Multi-Colored-Google-Assistant-Required/dp/B07DLSNNDS/>.

[39] Sensi thermostat. <https://www.amazon.com/Emerson-Thermostat-Version-Energy-Certified/dp/B01NB1OB0I/>.

[40] PyShark. <https://github.com/KimiNewt/pyshark/>.

[41] Madea/dataset. <https://github.com/sprout-uci/madea>.

[42] Raspberry pi 4. <https://www.amazon.com/Raspberry-Pi-Computer-Suitable-Workstation/dp/B0899VXM8F>.

[43] Netgear 4g lte broadband modem. <https://www.amazon.com/NETGEAR-LTE-Broadband-Modem-LM1200/dp/B08R813HLW/>.

[44] Amazon. <https://www.amazon.com/>.

[45] Tianrui Hu, Daniel J Dubois, and David Choffnes. Behavior: Measuring smart home iot behavior using network-inferred behavior models. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, pages 421–436, 2023.

[46] DCS-930L Wireless N Network Camera. <https://eu.dlink.com/uk/en/products/dcs-930l-wireless-n-network-camera#Specs>.

[47] DCS-930L Datasheet. https://eu.dlink.com/uk/en/-/media/consumer_products/dcs/dcs-930l/datasheet/dcs_930l_b1_datasheet_en_eu.pdf.

[48] MSP430 Ultra-low-power Microcontrollers. <https://www.ti.com/sc/docs/products/micro/msp430/39113.pdf>.

[49] ARM Limited. ARM TrustZone for Cortex-M. <https://www.arm.com/technologies/trustzone-for-cortex-m>, 2009.

[50] NXP LPC55S69-EVK. <https://www.nxp.com/design/software/development-software/mcuxpresso-software-and-tools-/lpcxpresso-boards/lpcxpresso55s69-development-board:LPC55S69-EVK>.

[51] Arm Ltd. Arm TrustZone. <https://www.arm.com/products/security-on-arm/trustzone>, 2018.

[52] Msp430 flash memory characteristics. <https://www.ti.com/lit/an/slaa334b/slaa334b.pdf?ts=1638460551489>, 2018.

[53] Avr atmega 1284p 8-bit microcontroller. <http://www1.microchip.com/downloads/en/DeviceDoc/doc8059.pdf>, 2009.

[54] Srivaths Ravi, Anand Raghunathan, and Srimat Chakradhar. Tamper resistance mechanisms for secure embedded systems. In *VLSI Design*, 2004.

APPENDIX A: COMPLETE DEVICE PROFILE

This appendix provides the complete device profile of an off-the-shelf device, Sensi Thermostat in Table A1. The partial profile of this device is shown in Figure 3 of Section IV.

APPENDIX B: DIFFERENCE BETWEEN MALWARE PACKETS & IOT DEVICE PACKETS

In this appendix we briefly describe the difference between malware packets and IoT device packets that we observed in the course of developing MADEA.

External Addresses. Packets generated by malware binaries have completely different external addresses from external addresses of packets generated by the observed IoT devices. Most malware programs try to connect with their CCC using static IP addresses. MADEA is sometimes able to map these IP addresses to hostnames through reverse DNS lookup. Interestingly, we observe that none of those hostnames indicates a service hosted in AWS cloud, whereas most hostnames found through reverse DNS lookup for IoT devices belong to AWS cloud. However, few malware programs use hostnames resolved through DNS lookup. Figure B1 shows a partial list of external addresses found for malware packets.

EXTERNAL_ADDRESS
web.sandvine.cloud
wirejunkie.com
124.220.8.193
147.12.185.31
147.141.65.32
146.135.30.107
ai126181062177.70.access-internet.ne.jp
047-036-232-224.res.spectrum.com
agg2.wttwnycf02m.nyroc.rr.com
072-239-141-038.res.spectrum.com
035-139-043-241.dhcp.bhn.net
124-140-55-197.rev.home.ne.jp
adsl-074-173-233-133.sip.bct.bellsouth.net
016-191-171-102.res.spectrum.com
129.235.172.34.bc.googleusercontent.com
zaq7d04b4dd.rev.zaq.ne.jp
vz17.securehostserver.com
016-203-236-107.res.spectrum.com
075-129-211-247.res.spectrum.com
146.125.224.206
072-185-008-125.res.spectrum.com

Fig. B1. Some of the external addresses used by malware.

Packet Lengths. The length distribution of malware packets is also drastically different from the length distribution of IoT device packets. This is because observed malware packets are mostly TCP SYN, ACK, and RST packets. Figure B2 and Figure B3 show the length distributions for malware packets and device packets respectively.

TABLE A1
COMPLETE DEVICE PROFILE OF SENSI THERMOSTAT.

DEVICE MAC	PACKET DIRECTION	EXTERNAL ADDRESS	PACKET LENGTH	Device Name
34:6f:92:1d:ba:50	USER_TO_DEVICE	192.168.4.1	62	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	224.0.0.251	46	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	58	Sensi Thermostat
34:6f:92:1d:ba:50	USER_TO_DEVICE	192.168.4.1	86	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	58	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	54	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	104	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	54	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	144	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	255	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	129	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	129	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	139	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	139	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	539	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	155	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	427	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	619	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	299	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	491	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	299	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	267	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	443	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	219	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	107	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	107	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	635	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	619	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	283	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	459	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	155	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	1450	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	172	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	864	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	512	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	917	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	61	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	192.168.47.2	54	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	330	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	61	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	475	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	411	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	171	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	571	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	118	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	224	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	384	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	192	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	315	Sensi Thermostat
34:6f:92:1d:ba:50	SERVER_TO_DEVICE	icda.sensicomfort.com	789	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	315	Sensi Thermostat
34:6f:92:1d:ba:50	DEVICE_TO_SERVER	icda.sensicomfort.com	347	Sensi Thermostat




Fig. B2. Length Distribution of Malware Packets.




Fig. B3. Length Distribution of IoT Device Packets.