

SIPHON: Towards Scalable High-Interaction Physical Honey pots

Juan Guarnizo ^{1 *}

Amit Tambe ¹

Suman Sankar Bhunia ¹

Martín Ochoa ¹

Nils Ole Tippenhauer ¹

Asaf Shabtai ¹

Yuval Elovici ^{1,2}

¹Singapore University of Technology and Design

²Ben Gurion University, Israel

ABSTRACT

In recent years, the emerging Internet-of-Things (IoT) has led to rising concerns about the security of networked embedded devices. In this work, we propose the SIPHON architecture—a Scalable high-Interaction Honey pot platform for IoT devices. Our architecture leverages IoT devices that are physically at one location and are connected to the Internet through so-called *wormholes* distributed around the world. The resulting architecture allows exposing few physical devices over a large number of geographically distributed IP addresses. We demonstrate the proposed architecture in a large scale experiment with 39 wormhole instances in 16 cities in 9 countries. Based on this setup, five physical IP cameras, one NVR and one IP printer are presented as 85 real IoT devices on the Internet, attracting a daily traffic of 700MB for a period of two months. A preliminary analysis of the collected traffic indicates that devices in some cities attracted significantly more traffic than others (ranging from 600 000 incoming TCP connections for the most popular destination to less than 50 000 for the least popular). We recorded over 400 brute-force login attempts to the web-interface of our devices using a total of 1826 distinct credentials, from which 11 attempts were successful. Moreover, we noted login attempts to Telnet and SSH ports some of which used credentials found in the recently disclosed Mirai malware.

Keywords

Low-Interaction Honey pot, High-Interaction Honey pot, Internet of Things, Scalability

1. INTRODUCTION

The Internet of Things (IoT) has gained immense popularity, creating a bridge between the physical world and

the Internet. According to the International Telecommunications Union (ITU) specifications, IoT can be defined as Internet connected physical objects or sensors which may be called as *things* [3]. This implies that, unlike conventional network nodes in the Internet such as servers or PCs, IoT devices have a stronger link to physical reality by means of their sensors, and their behavior relies on that interaction with the analog world. During the last decade, rapid advancement in sensor fabrication and miniaturization has fueled the growth of the *things* tremendously. It is expected that by the end of 2016, 6.4 billion *things* will be connected to the Internet, growing to 11.8 billion by 2018 [5].

IoT has propelled the usage of applications such as smart wearable devices, intelligent transportation, and industrial automation (among others). As a large number of *things* are being used or deployed commercially, it is becoming challenging to provide suitable security mechanisms. Incorrectly configured *things* will provide opportunities for attackers to perform malicious activities, ranging from sabotage of critical infrastructure to privacy violations. It is estimated that by the year 2017, 90% of organizations that install IoT devices in their premises will suffer from attacks to their back-end IT systems, due to vulnerabilities in IoT devices [6].

In fact, recent Distributed Denial of Service attacks on core internet services such as DYN's Managed DNS infrastructure were performed partially using IoT devices (IP cameras) that had default or weak hard-coded passwords on Telnet and SSH interfaces [12].

Therefore, there is a strong need to develop suitable and cost efficient methods to find vulnerabilities in IoT devices, in order to address them before attackers take advantage of them. In traditional IT security, Honey pots [7, 29, 30] are commonly used to better understand the dynamic threat landscape without exposing critical assets. Usually honey-pots attempt to mimic a certain interaction in a realistic way (such as a login shell), encouraging unsolicited connections and potentially encouraging attempts to perform (possibly unknown) attacks.

Within the context of IoT, it is challenging to realistically mimic the interaction of IoT devices. For instance, consider the case of IP cameras. In order to virtualize or simulate their behaviour in a realistic way, one would need not only to broadcast some video to an attacker, but also react faithfully to commands such as tilting the camera or zooming in. Although not impossible, this requires a significant amount of technical work that cannot be easily reused for mimicking

*Contact: {juan_guarnizo, amit_tambe}@sutd.edu.sg.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CPSS'17, April 02 2017, Abu Dhabi, United Arab Emirates

© 2017 ACM. ISBN 978-1-4503-4956-7/17/04...\$15.00

DOI: <http://dx.doi.org/10.1145/3055186.3055192>

different types of IoT devices from different vendors, due to the heterogeneity of such devices.

On the other hand, since IoT devices are sensing and actuating on the physical world, they might have a different perceived value to attackers depending on their location (e.g., different organizations or countries). This is similar to the trend seen on the botnet underground market where compromised servers in different regions are rented at different prices [18]. Therefore a Honeypot exposing IoT devices should ideally appear geographically distributed to attackers interacting with it.

Based on the requirements mentioned above, a framework for deploying IoT honeypots should: (1) be able to provide *high-interaction* in order to motivate skilled attackers to perform their activity and expose the vulnerabilities they are exploiting; (2) make it possible to easily “place” the IoT devices in a wide range of geographic locations; and (3) easily scale to devices of multiple vendors and of different kinds in order to keep track of a dynamic threat landscape.

In this paper, we propose an abstract architecture and insights on a prototype implementation of such a Honeypot, that we call SIPHON. Our implementation allows us to deploy over 80 high-interactive devices with a diverse set of IPs located in different regions of the world by using only seven real IoT devices in our lab. We ran SIPHON for a period of two months gathering 20GB of raw traffic data. On doing a preliminary analysis of the data, we were able to see that instances in different cities received significantly different amounts of attention (in terms of connections and traffic). Moreover, curious users from the Internet attempted to login by brute-force, gaining access to the admin interface and interacting with devices in some cases.

Problem Statement. IoT devices offer a rich set of potential attack vectors to the attacker. Our goal is to learn about existing and novel attack vectors on IoT devices. In particular, in this work we focus on gathering unsolicited traffic to IoT devices exposed to the Internet. Conventional honeypots use virtualization or simulation approaches to replicate the device under attack, and attract attackers. In the context of IoT, replicating the devices with virtual machines will be challenging. Only high interaction honeypots (e.g., allowing a user to move a camera) will convince skilled attackers to use advanced attack methods while exploiting zero-day vulnerabilities. The main question that motivates this work is *How can we construct a large scale honeypot consisting of high-interaction IoT devices?*

Approach. We propose the SIPHON architecture for scalable high-interaction IoT Honeypots. Our architecture uses physical IoT devices that are connected to various geographical locations through so-called “wormholes”. The proposed architecture was implemented in a large scale distribution with 85 instances in 16 cities, 9 countries while using only five physical IP cameras, one NVR and one IP printer (i.e., seven IoT devices). In our implementation, network traffic was forwarded leveraging an infrastructure of cloud service providers consisting of Amazon EC2 instances, DigitalOcean and Linode. We preliminary analyzed the collected traffic, in particular the amount of attackers’ interest in various geographical locations and types of interaction with the devices.

Contributions. In this work:

1. We explain the expected benefits from a high-interaction IoT honeypot and define the design challenges.

2. We propose an architecture for a high-interactive IoT honeypot using limited number of physical IoT devices.
3. We present an implementation of that architecture based on seven devices presented as 85 distinct services on the Internet.
4. We analyze captured traffic, and show that geographic locations of our wormholes matter - attackers show more interest in some of the locations.

We focus on high interaction distributed honeypot consisting of real IoT devices. That setup allows us to capture and analyze traffic and attacks across different geo-locations. To the best of our knowledge, we are the first to propose and implement such an architecture.

This work is organized as follows. The background of this work is presented in Section 2. We propose the main architecture in Section 3, and its implementation in Section 4. We analyze and discuss the traffic collected by our implementation in Section 5. Related work is summarized in Section 6. Finally, the paper is concluded in Section 7.

2. BACKGROUND

A honeypot is an infrastructure commonly used to attract potential attackers to a controlled and monitored environment, without exposing critical assets, with the goal of understanding the threat landscape at a given point in time. Usually, it may consist of real or virtual systems [30] that mimic production environments [33]. Interactions of the attackers with the honeypot are typically monitored and recorded for further analysis. Researchers can then design appropriate mitigation techniques for attacks after analyzing the traces collected by the honeypot. From the point of intrusion analysts, “honeypots provide another indication of wave of network attacks” [34].

Honeypots may be set up in such a manner that they will provide convincing fake information which may be desirable to potential attackers, in order to attract their interest. Also, honeypots can help to redirect attackers to a decoy system and thus indirectly protect critical systems or infrastructure from a possible compromise: if attackers are misled to think they are interacting with a real system, they are wasting time and resources to cause damage to real infrastructures.

Security researchers classify honeypots as either high-interaction or low-interaction in accordance with the privileges enjoyed by the attackers and real vulnerabilities being exposed to them. In the following, we briefly recap these two categories.

2.1 Low interaction honeypots

Low-interaction honeypots emulate well-known vulnerable network services by partially implementing the TCP and IP stacks, in order to attract attackers. They do not give an attacker access to a real system, but rather to some sort of emulated system [13]. Since usually they do not run a fully featured operating system, they are considered to be safer in terms of remote exploitation. However, due to their simplicity, the main disadvantage of low-interaction honeypots is the higher likelihood of being detected as artificial by the attacker, especially if the attacker is a human. The attacker may find that the real services are emulated and services

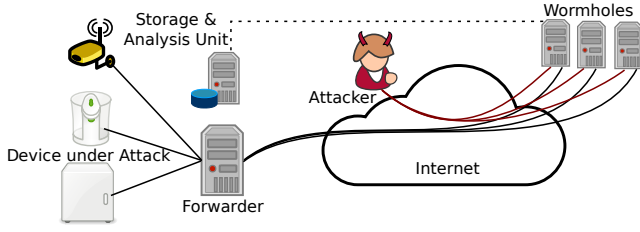


Figure 2: Abstract overview of distributed physical honeypot

environment, i.e. the cameras will show live images, the fridge will contain real food, etc. Devices will be configured with minimal effort (i.e. potentially with default passwords). No dedicated security solutions are present (such as firewalls or IDS). If the devices are reachable through a NAT service, then the exposed services are forwarded by the NAT.

3.3 Attacker model

We consider an attacker that searches for IoT devices using a search engine such as Shodan or an ad-hoc IP scanner to identify devices connected to the internet. Once a device is identified, an attacker interacts in an arbitrary way with it through the network. However, we assume the attacker will concentrate on the ports exposing open services of the IoT device, such as authentication to an administrator interface. These attacks can be automated or manual. The attacker goal is to compromise the device to steal sensitive privacy data (i.e. video, audio) or to install malicious software on it to attack other devices on his behalf. Also, attackers will perform actions to verify that the devices are highly-interactive, such as zooming in and out and tilting cameras.

3.4 Design Overview

Our design (see Figure 2) is based on real physical IoT devices, that are exposed to the Internet through a large number of *tunnels* that forward traffic from *wormholes* (remote public IP addresses) to the local physical IoT devices.

Wormhole. The wormhole device presents a number of open ports to the general Internet on a public IP address. Incoming traffic on those ports is either logged locally, or transparently forwarded to a specific port on a remote physical IoT device via the Forwarder device. There are a number of options on how this forwarding can be achieved, we discuss them in Section 3.6. More than one physical device can be associated to each wormhole, but each service forwarded will need a unique port on the wormhole.

Forwarder. The forwarder ensures that the traffic between the wormhole and the IoT device is re-written in real-time to hide the fact that devices are physically located somewhere else, or communicating with many other tunnels. If required, the forwarder can also act as a TLS man-in-the-middle because it will have access to the private keys of certificates used by the IoT devices.

Device under Attack (DUA). The local (IoT) device under attack in the SIPHON architecture can be a normal commercial off-the shelf device. It will expose one or more services via the network, most likely over TCP. The device is able to maintain a number of concurrent connections at any point in time. We note that the original design of these

off-the-shelf devices may limit the number of simultaneous connections.

Storage & Analysis Unit (SAU). The SAU obtains traffic records and general logs from the wormholes, and aggregates the data for offline analysis. For example, the recorded traffic can be analyzed by a suitable framework (e.g. Suricata or Snort [8]). In addition, dedicated analysis frameworks for other honeypot deployments could be used (similar to [27]).

3.5 Location of Wormholes

As IoT devices are embedded in the physical world and interact with it, our hypothesis is the physical location of devices under attack is correlated to its attractiveness to the attacker. While the direct physical location of a communication partner on the Internet is not easy to determine [9], many services on the web provide IP geolocation services (e.g., [17]).

For that reason, we propose to ensure that the wormholes have public IP addresses that are localized in a range of physical locations, ideally spanning multiple countries or even continents.

That goal can be achieved in a number of ways. However, the easiest way to obtain public IP addresses seems to be a) home users (who often get one IP per subscription line), and b) servers on the cloud (cloud providers own large IPv4 address spaces). In this work, we focus on the latter (see Section 4). To take advantage of public IP addresses of home users, we also considered preparing devices such as Raspberry PIs with a specific configuration as wormhole to forward traffic to the honeypot forwarder. The disadvantage of that approach is the fact that such home setups would likely require the home user to set up port forwarding on the home NAT device (often called *router* in that context). That configuration effort might be non-trivial for many users, and thus we have not followed this implementation path.

3.6 Forwarding Method

The tunnels between the forwarder and the wormholes can be realized using a range of different forwarding techniques. We now discuss a number of options that we considered.

Application-Layer Proxy. If the honeypot should represent only few selected services, it would be possible to use proxies on the wormholes for those services (E.g., an HTTP proxy like mitmproxy). There are a number of advantages of that approach, for example the option to cache static information and thus reduce the traffic load on the DUAs.

Encapsulation. On the transport and application layers, services such as SSH and TLS can be used to encapsulate selected traffic with low configuration overhead. In our implementation, we chose an SSH tunnel as method to forward the traffic.

Network-Layer VPN. A number of network-layer VPN services are available (e.g., IPsec), a detailed comparison of which is out of scope here (we refer to [23]). Such services might have an advantage over selective forwarding using SSH, as large numbers of ports could be forwarded with little configuration.

3.7 Compromise Detection and Restoration

As the SIPHON architecture exposes real devices to the attacker, it is possible for the attacker to fully compromise a

device as part of an attack. Once compromised, the devices could be used to attack other devices, the general neighboring infrastructure, or distribute malware to other attackers. There are several options to mitigate or prevent such scenarios.

Periodic Reset. A simplistic approach to mitigate the impact of compromised devices would be a setup that periodically resets the configuration and firmware of all IoT devices that are used in the lab. That approach has several disadvantages (i.e. timeliness, completeness of reset). On the other hand, it is easy to set up with some manual intervention.

IPS. Intrusion prevention systems like the open source Suricata [8] could be used together with custom rules tailored towards the IoT devices to prevent their ongoing compromise. In addition, an IDS/IPS could be used to detect the changed behavior of a compromised device to identify that such a compromise has happened.

Full instrumentation. The most reliable detection of a compromise could be achieved through low-level instrumentation of the IoT devices, e.g. through JTAG or similar connections. The setup and maintenance of such instrumentation is possible, but expected to be challenging and expensive [16].

3.8 Scalability

As we claim in the acronym SIPHON, the proposed architecture should be scalable. Our intuition for that is as follows: for n physical IoT devices and m wormholes, the number of devices presented to the attacker, without exposing the same device twice in the same wormhole, can be up to $n * m$. For instance, if $n = 100$ and $m = 1000$ up to 100,000 services could be exposed to the Internet. The number of services per wormhole is limited to around 65,000 due to port numbering restrictions. The number of wormholes per physical device is limited by the amount of parallel traffic that can be handled by each physical device. In our practical implementation, we chose lower values for both n and m due to budget and space constraints as we will discuss in the following.

4. A DISTRIBUTED IP-CAMERA HONEYPOT

In this section, we discuss an implementation of the proposed SIPHON architecture for a specific class of IoT devices: IP cameras used for home or public surveillance. We chose that class of devices as the number of IP cameras connected to the Internet is increasing daily, and recent research findings have shown that the IP cameras and Network Video Recorder (NVR) are vulnerable to attacks [12, 27], thus attracting more unsolicited traffic than other IoT devices in the market. In addition, IP cameras allow access to private video, and often allow complex interactions with the user (e.g. manual movement control).

In our implementation, locations of the cameras are spoofed using wormholes based on instances in cloud services. The instances are deployed in different cities around the world. From the perspective of the attacker, the cameras are located across cities (e.g., by using the location feature of search engines such as *Shodan*). All traffic is directed from the wormholes to our lab for in-depth analysis.

Device Model	Password	Difficulty
HP Pro Printer 6830	1234567890	Easy
D-Link DCS-9050	password123	Easy
D-Link DCS-930L	YAQvwrjy	Hard
D-Link DCS-942L	1234567890	Easy
Aztech WIPC409HD	admin	Default
Sineoji PT603V	9WgnTMxe	Hard
Trendnet Emulator	admin	Easy
HikVision NVR 7604NI-E1/4P	xDk2PKHU	Hard

Table 1: IoT Device Details

4.1 Implementation

Figure 3 shows our prototype infrastructure for the SIPHON deployment with IP cameras, a networked video recorder (NVR), and an IP printer. Table 1 shows the list of devices used.

By default, all the devices expose an administration web interface protected by a password. In order to attract complex interactions with the devices, we use default passwords for some of the cameras while we use weak passwords for others. A screenshot of one of the HTTP interfaces of the cameras after successful authentication is shown in Figure 4.

Using the administration web interfaces aligns well with our goal of capturing high interaction traffic. They provide apt opportunities to expose to the attacker, multiple device interaction scenarios, such as camera motion, wi-fi network scanning, etc. Even though more recent attacks, like the one using Mirai malware, rely completely on the discovery and exploitation of vulnerabilities via the Telnet and SSH ports [26], we do not actively seek information on attacks made on these ports. This is due to the following reasons: 1) Telnet/SSH would not provide a high-interaction environment as an environment that enables video transmission via HTTP and therefore would be less interesting 2) Telnet/SSH was already explored in this context by IoT POT [27] 3) None of the IoT devices used by us have open Telnet or SSH ports

The devices in our lab were connected to cloud servers from three different cloud service providers (viz. Amazon, LiNode and Digital Ocean) in various cities. During our experiment, a total of thirty nine servers were deployed and each of them had a distinct public static IP. Each of the cloud servers was designated as a wormhole with up to three (3) connected devices. The devices were connected on different ports of the cloud server. In total, 85 devices were visible to the attackers on the Internet, while only seven physical devices were used, thereby yielding in average about 10 possible virtual connections to each physical device. *Shodan* obtained the information about the devices using device fingerprinting heuristics such as parsing the HTTP response of the device and published the devices along with their geo-locations, accessible ports and other useful information as illustrated in Figure 1.

Eventually, all the traffic (i.e. device solicitations or attacks) was being diverted to the IP cameras or Network Video Recorders (NVR) in our lab. All traffic was captured using `tcpdump` on each of the wormhole servers and the resulting traces (in .pcap format) were stored locally for offline analysis.

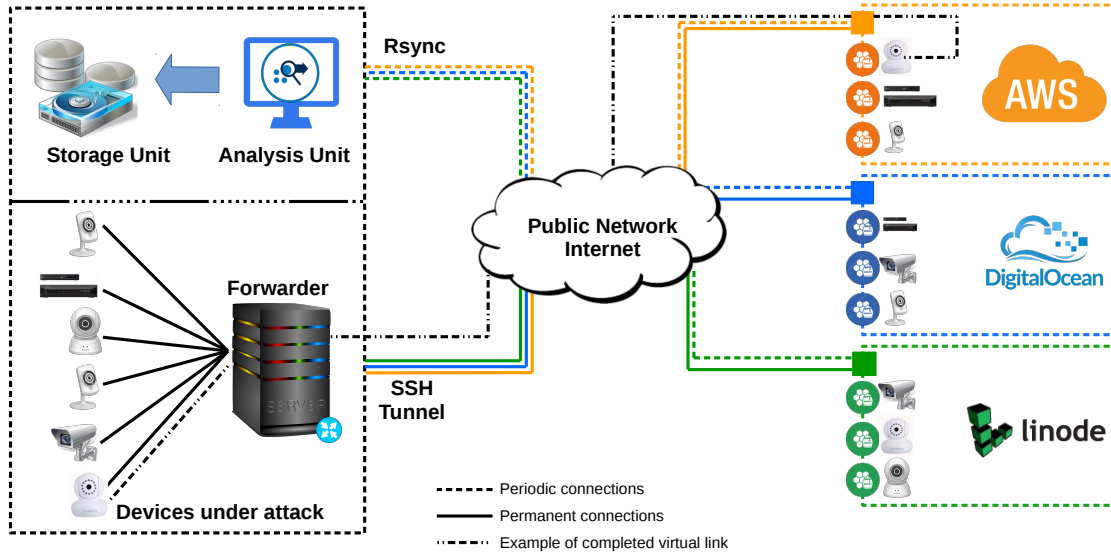


Figure 3: SIPHON prototype implementation in our lab

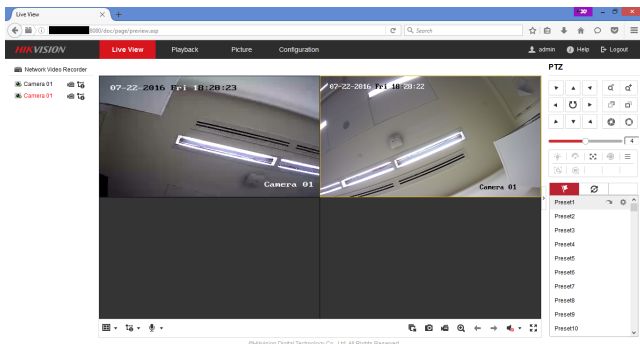


Figure 4: Example of device view through a worm-hole

4.2 Technical Setup for Traffic Forwarding

Our technical setup relies on a network gateway scheme, as described abstractly in Figure 2. To implement the *forwarder*, we used a local server in our lab which manages the TCP connections from the wormhole instances in the cloud. This server alone handles all m wormhole to n IoT devices communications. The server itself is a VM guest with 1 virtual core and 4GB of RAM, running Ubuntu 16.04.

The connection between the forwarder and the cloud instances is established with reverse `ssh` tunnels that redirect traffic of a specific port (i.e. port 80) on the wormhole to a port in the forwarder. Once the traffic has reached the forwarder, we complete the traffic redirection to the IoT devices by means of the `socat` linux command, which also forwards the device's responses to the cloud instance through the forwarder. This process is easy to automatize by using a different local port in the forwarder for each cloud replica of a physical device (wormhole).

The IoT devices are isolated from other devices in the lab through the use of 802.1Q VLANs. That ensures that even



Figure 5: Setup of Cameras and NVR in Lab.

compromised devices can at most interact with other IoT devices, and the server.

4.3 Locations of Wormholes

For the attackers, the IoT devices are placed in the cities where the actual wormhole is placed. Physically, the IoT devices can be located in a common area together with the forwarder.

Figure 6 shows the locations of honeypot deployment in various cities in different continents around the world. The SIPHON experimental deployment covered the following 16 cities in 9 different countries: a) USA (San Jose, Boardman, Ashburn, New York, San Francisco, Dallas, Fremont, Newark) b) Canada (Toronto) c) Europe (Frankfurt, London, Dublin, Amsterdam) d) Asia (Singapore, Bangalore) e) Australia (Sydney).

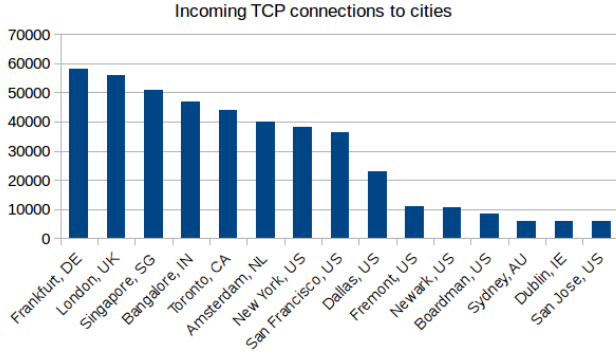


Figure 7: Distribution of TCP connections to wormholes according to their city

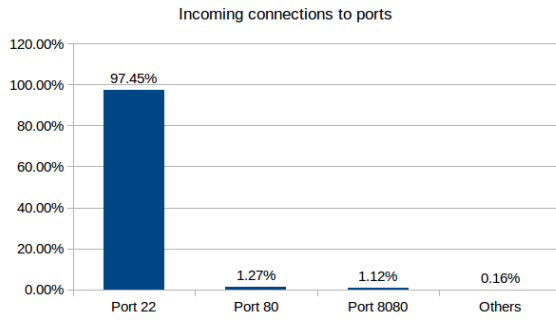


Figure 8: Connections to ports of wormholes

Distinct remote IPs. Upon analyzing the more than two months' worth of data, we observed incoming TCP connections from over 13000 distinct remote IP addresses.

Connections to wormhole per city. Based on the geographical distribution of our wormholes, we observed that wormholes in some cities received more attention than wormholes in other cities. Figure 7 depicts the distribution of incoming TCP connections received by wormholes per city. The wormhole(s) in Frankfurt received the most connections (almost 600 000), while San Jose in US received the least number of connections (about 50 000).

Wormhole ports. However, we observed that most of the incoming connections (about 97%) were on port 22 (SSH) of the wormholes, whereas HTTP ports like port 80 and port 8080 received just 1.27% and 1.12% of total connections respectively. Other HTTP ports which we had opened got a mere 0.25% of the total incoming TCP connections. Figure 8 shows the distribution of incoming TCP connections according to the wormhole ports.

Location of attackers. Although not a reliable indicator (since attackers can easily tunnel through IPs located anywhere in the world), we also gathered statistics on the location of remote IPs. We analyzed the origin of all the incoming connections by country. Figure 9 shows the distribution of incoming TCP connections by their countries of origin. We saw that more than 70% of the connections originated in China followed by 8% connections originating in USA. Netherlands, France and other countries make up the remaining list as shown in Figure 9.

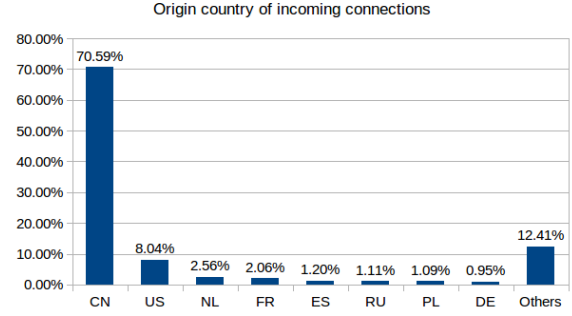


Figure 9: Incoming connections according to their country of origin

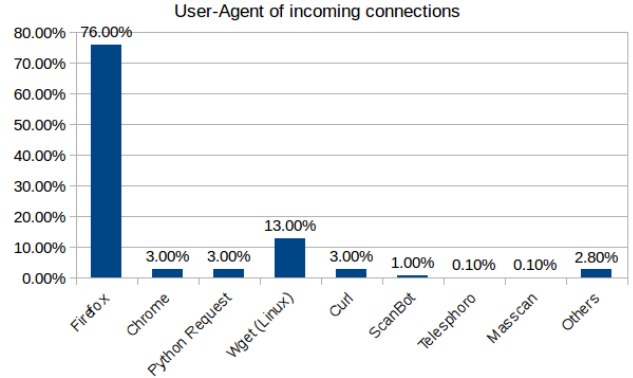


Figure 10: Incoming connections according to user agents

User-agents used by attackers. Next, we enumerated the agents being used in the connections to our devices (which might be arbitrarily set by the attacker). In Figure 10, usage of different user agents in incoming TCP connections has been shown. Most of the connections (76%) were made using Mozilla as the user agent. Apart from this, Chrome, Python Request, wget, Curl, Scanbot, Telesphoro and Masscan [2] agents were used. Other user agents were combined together into the 'Others' category that contributed 2.8% of the total connections and included IP Camera Viewer, Cam Finder, Morfeus F Scanner, Msqq agents, etc. Usage of Nmap was detected several times. We also noticed around 9000 Shell shock attempts on our devices; we note that our devices are not vulnerable to those attacks.

Attacked devices. In Figure 11, we depict the combined amount of interest generated (based on HTTP traffic only) by the different physical devices. We observe that the majority of traffic was directed towards the D-Link DCS-930L camera. We conjecture that this device attracted the most attention from attackers due to a recent report in the media about the presence of vulnerabilities in its firmware [1].

Brute force login attempts. We have setup the IoT devices with different levels of password difficulties (Table 1). We retained the default credentials for the web interface login on some, while we modified the credentials on others. We configured some cameras to have easy to guess passwords, whereas others were configured with hard to guess

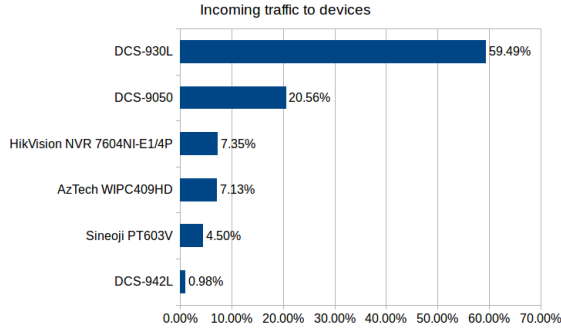


Figure 11: Traffic breakup for the devices

passwords. With this setup, we observed 404 brute force attempts on all the devices combined from 137 distinct remote IPs. We deem a login attempt as a brute force attempt if the same remote IP tries to perform more than three login attempts in an HTTP session. Of these 404 brute force attempts, we observed a total of 11 attempts succeed. All the successful login attempts were on devices with easy passwords, and no login was recorded on a device with a hard password. We note that these easy passwords are present in most of the dictionaries used by automated brute-force attacks. On the other hand, these passwords can be guessed or brute forced easily even when attacking manually. Moreover, none of the successful logins involved attackers performing high interaction functions like tilt or zoom-in.

Impact of device listing on Shodan

Another goal of this analysis was to investigate the impact of the listing of a wormhole on Shodan. We tried to understand this by computing the number of TCP connections received by wormholes before and after their Shodan listings. Shodan usually takes between one-two weeks to enlist an IP after it becomes live on the Internet. With respect to this, we took into consideration one week's worth of incoming connections before and after Shodan listing.

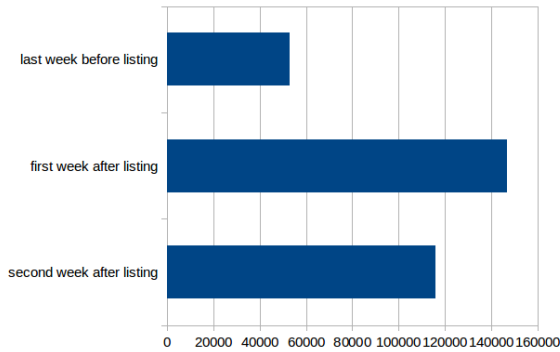


Figure 12: Impact of Shodan listing on no. of incoming connections

In Figure 12, we show the combined impact of Shodan listing for all wormholes. In particular, we show how many average connections the wormholes attracted per week, before and after the listing. In Figure 12, the top bar represents the average number of connections across all the wormholes

before they get listed on Shodan. The bar in the middle represents the average number of connections during the first week after getting listed on Shodan, while the bottom bar represents the average number of connections during the first two weeks after getting listed on Shodan. A surge in incoming connections is evident after listing on Shodan. We also observed that this surge in connections continues even in the second week after Shodan listing, but is reduced in increase. Thus, we can infer that there is an immediate impact of the Shodan listing, and attackers try to compromise the devices immediately after listing on Shodan. In the long term, the number of incoming connections decreases.

5.3 Other IoT devices

So far, we have used only IP-cameras and an NVR in the high interaction honeypot and gained knowledge about the attackers' behavior/interaction with them. But these devices do not represent the entire spectrum of existing IoT devices. We are interested in garnering knowledge about attackers' behavior towards other IoT devices as well. Therefore, we added another IoT device, the IP printer, to our honeypot, for about four weeks. This IoT device was connected to a cloud instance of Digital Ocean and located in London. Thus, similar to the IP cameras, even though the device was physically present in our lab, it appeared to be located in London.

After analyzing four weeks' worth of network traffic data from the IP printer, we present our results.

Wormhole ports. We observed that most of the connections (99%) were on port 22 (SSH) of the wormholes, whereas port 80 (HTTP) received a meager 1%. Even in the case of IP cameras, we observed that the SSH port received more number of connections compared to HTTP ports.

Location of attackers. Similar to the observation in the case of IP cameras, we noticed that the highest percentage of incoming connections to the IP printer (79%), originated in China. Taiwan was next with 11% of connections originating in Taiwan. France, India, Kazakhstan, Netherlands, Ukraine and USA made up the remaining portions.

Brute force login attempts. We had changed the credentials of the IP printer to a password that should be easy to guess (*admin:1234567890*). During our experiment, we observed only one brute force login attempt. Similar to the case of IP cameras, we deem a login attempt as a brute force attempt if the same remote IP tries to perform more than three login attempts in an HTTP session. A total of four successful login attempts were noted, from three distinct remote IPs, all from China.

User-agents used by attackers. Similar to the case of IP cameras, we tried to understand which different user-agents were used in the connections to IP printer. We observed that all the four successful login attempts were made using the Mozilla user-agent. Unlike in the case of IP cameras, we did not notice a few user-agents (viz. *cameraviewer*, *IPcamera finder*, *Scanbot*, *Morfeus F scanner*, *Msqq etc.*). In addition, we did not notice any *Shell shock* attempts on the IP printer.

5.4 Low interaction honeypot

One of our hypothesis for the low interaction honeypot is that most of the attacks are executed in an automated manner using bots. We assume that these bots are not capable to distinguish between real and virtual or fake devices.

So, these bots may conduct repeated attempts considering the device as a real one. In contrast, a human being will discover a fake device and will not make repeated attempts consequently.

We have implemented a low interaction honeypot using the same methodologies as of the high interaction one. However, instead of a real physical device, we used the Trendnet camera emulator in the low interaction honeypot. Hence, this camera is a fake camera and not a real one. The entry point (i.e. wormhole) to the low interaction honeypot has been created using an instance from cloud service provider Linode, located in Singapore. Thereafter, we gathered data from the low interaction honeypot for almost six weeks.

The results from six weeks of data are now presented.

Wormhole ports. We observed that most of the connections to the fake camera i.e. Trendnet camera emulator (87%) were on port 22 (SSH) whereas HTTP Ports 8080 and 80 received 12% and 1% of the total incoming connections respectively. In that respect, the behavior is consistent with the trends observed with high interaction honeypots.

User-agents used by attackers. Similar to the corresponding observation in the case of high interaction honeypot, we observed the maximum percentage of incoming connections to be having the user-agent as various versions of Mozilla. Apart from that, the attackers also used wget, python, masscan and Probenet Scanner. All these user agents were observed in the high interaction honeypot as well.

Login attempts. We have noted three successful login attempts with the default credentials `admin:admin` from three distinct remote IPs. One successful login attempt was made each from China, Russia and Iran.

We also observed that the attackers did not spend enough time to explore the device or its functionality. They spent on an average only around 30 seconds in the low interaction honeypot, compared to about one minute or even up to an hour in the case of real cameras. The time spent by the attacker exploring the device is calculated based on the timestamps of the first and last packets in the TCP session captured via tcpdump. It is understood that the duration of a session is correlated to human intervention. *session time* will be longer when the attacker will try to explore functionality of the device. Human beings can easily distinguish between a still image and a streaming video unlike automated bots. Thus, we can more easily determine manual attacks using this low interaction honeypot compared to our earlier high interaction honeypot.

5.5 Discussion and comments on validity

By analyzing the traffic, we see that our proposed architecture was successful in attracting a high number of connections, which confirms our overall design. In addition, we were able to observe interesting interactions in our preliminary analysis of the gathered data. We note that in our analysis, we assume that interaction between the attacker and IoT device over HTTP is more likely performed by humans than scripts, since interaction with unknown HTTP interfaces is harder to automate [27]. Our prototype is thus a promising first step towards effectively learning the threat landscape in IoT security.

We are aware of several issues with our current prototype. First, one might argue that an intelligent attacker can spot our honeypot in different ways. For example, the attacker

can detect that two wormholes represent the same physical device by simultaneously logging in to the same devices from two different wormholes. We acknowledge this shortcoming, but note that during our experiment, our data indicates that at least no single IP has discovered all the wormholes and the majority only interacted with a small subset (typically 5 or less). The issue can also be mitigated by having one different physical device per wormhole. Since we were interested in comparing different locations, we decided to scale the system in a 1 to n fashion instead of a bijective relation between devices and wormholes.

Compared to related work such as [27], we have so far observed very few attempts to upgrade the firmware, which is a very common objective of attackers who want to spread botnets or malware. We conjecture a possible explanation for this as follows. We have noticed that only the minority of traffic in the wormholes was directed to the ports exposing the web interface of our IoT devices. Attackers seemed much more interested in SSH connections, possibly because common automated attack tools are tailored to re-flash firmware using a shell, which is a general approach that works in many devices. We believe that the firmware upgrade attempts we received were performed manually, but we also think such attacks will increase in the future as attack tools automatize this process for different brands. It is also likely that in the future SSH access to many devices such as home cameras will be more restricted, whereas access to the web-interface might be still desirable for convenience of access.

Additionally, one might argue that an IoT device connected to the Internet through an IP that is owned by Amazon is suspicious to start with. However, given the amount of interest received in terms of traffic and the fact that there were some (potentially manual) logins, we think that using IPs belonging to cloud services might not be considered suspicious to a wide range of attackers. Most of our wormholes are not classified as Honeypots by Shodan’s Honeyscore (score of 0.0), which indicates that our setting appears to be realistic from a technical point of view.

Incidentally, we notice at least other 6 cloud instances exposing cameras to Shodan. All of those instances had a Honeypot score of 0.5, which in principle makes them more suspicious than our set-up, but indicates potentially that either such a setting is not uncommon (if not honeypot), or that other people have been less stealthy in their Honeypot setup than us.

6. RELATED WORK

Honeypots are a common measure to understand attacker activities in computer networks. The authors of [13] provide a taxonomy of honeypots, and differentiate between low and high interaction honeypots. Both low and high interaction honeypots are compared in [7], and the authors conclude that high interaction honeypots provide more insights on attacker behavior than low interaction ones.

The Honeynet Project [33] is a well-established project that focuses on monitoring and analyzing attacks to complement intrusion detection tools. The Honeynet project does not use emulation, and instead leverages real systems and applications. For that reason, the Honeynet project is an example of high interaction honeypots.

In [15], the authors describe honeynets that can be used to increase security in a large computer network. In short, honeynets are clusters of honeypots.

A low interaction honeypot system is implemented in [22]. The honeypot monitors behaviors and learns the advanced attacks that may not be detected by IDS tools. The session of the attacker is redirected to the honeypot system, which then serves requests from the attacker. For that purpose, the authors provided service daemons and a fake shell so that the attacker is not able to discover the system as honeypot.

In [4], the *DecoyPort* system was proposed which redirects attackers towards honeypots. The system creates *DecoyPorts* on active computers, those ports are not used by real services. Whenever any query or request comes to the ports, the system diverts the same to the honeypot. The system does not only act as a port forwarder but also is capable of controlling network loads caused by attackers. With this high interaction honeypot deployment, authors noted that possibility of attack may increase.

In [11], the authors described the implementation and deployment of a honeypot based on a number of real, vulnerable web applications. They hosted all the web applications in seven isolated virtual machines running on a VMWare Server. In order to limit the attack surface, the authors let the exposed services run as a non privileged user. They then analyzed the collected data to study attackers' behavior on the web applications during pre and post exploitation. In contrast, we work with real physical IoT devices to set up a high-interaction honeypot.

Among the related work on general honeypots, we found only one that focuses on IoT [27]. In that work, the authors present a low interaction honeypot for IP cameras and Digital Video Recorders (DVR). The authors emulated *Telnet* services for those devices. No real devices were used to deploy the honeypot. The authors' goal was to capture telnet based attacks and analyze the same with respect to the concerned IoT devices.

Another interesting aspect has been described in [29]. The authors discuss the necessity of deploying distributed honeypots for monitoring local trends of attack and location specific attack behavior. The authors deployed low interaction honeypots in different countries with active participation from various academic, industrial and governmental organizations. These honeypots emulate various operating systems. The authors have shown the similarities and differences in attacks on honeypots in different locations.

A game theoretic model has been proposed in [28]. That paper models a scenario when an attacker decides which server to attack in a network. The authors describe a honeypot selection game which will lure attackers to attack a honeypot instead of an actual server. They studied and analyzed attacker behavior empirically. They have shown that in addition to honeypot design, the strategies should also be emphasized.

Another recent paper [24] proposes a Bayesian game to defend against attacks in a honeypot enabled IoT network. In that paper, various game scenarios are described depending on the changing strategies of both the attacker and defender. The authors perform systematic mathematical analysis of the games and evaluate the Bayesian model through simulation.

Compared to existing literature, our work focuses on a high interaction distributed honeypot consisting of real IoT devices. That setup allows to capture and analyze real traffic and attacks across different geolocations and measure attacker behavior with respect to different locations. To the

best of our knowledge, we are the first to propose and implement such an architecture.

7. CONCLUSIONS

In this paper, we have proposed a design for a distributed and scalable high-interaction IoT honeypot. This design allows portraying physical IoT devices in a single lab as being geographically distributed by establishing tunnels between the public IP addresses and the physical devices. It also allows to collect traffic for further processing and analysis. We have implemented a prototype of this design using five IP cameras, one NVR, making available 85 IoT devices to the attackers worldwide. Our implementation has allowed us to gather several GBs of traffic data that we have preliminary analyzed particularly from the point of view of locations with higher attractiveness to the attacker.

In the data gathered we have particularly noticed more than 400 brute force attempts to bypass the authentication of our devices that have led to 42 successful logins into the devices. We have observed an overwhelming amount of traffic to our devices but comparatively not so many interactions with the device's web interface. We conjecture that this is due to little automation in attack tools to perform such attacks, and comparatively less amount of manual attacks compared to automated attacks. Among all the deployed devices, one particular model of D-Link camera got around 60% of the total traffic, which we conjecture is due to recently published vulnerabilities for this model. Devices in Frankfurt and Singapore, attracted the highest amount of traffic, confirming our hypothesis that *location matters*.

We have inducted other IoT devices like IP printer into the high interaction honeypot and tried to understand attackers' interaction. Moreover, we also deployed a low interaction honeypot using IP camera emulators to compare attackers' behavior with that of a high interaction honeypot. We observed that attackers are more interested in real IP cameras than emulators or other devices.

In the future, we plan to add more devices to our Honeypot, such as home appliances and perform analysis in more depth with the collected data as well as to share our data with other researchers interested in performing such analyses. Apart from this, we also plan to incorporate intrusion prevention systems to prevent substantial damage to our devices. Besides, we plan to integrate the SIPHON framework with the IoT testbed presented in [32]. By incorporating SIPHON with the IoT testbed, we aim at achieving the following two goals. First, we can use the SIPHON as one of the security testing mechanisms of the IoT testbed. This can be achieved by exposing any IoT device that is tested in the testbed to the Internet via the SIPHON framework as a honeypot, and by using the measurement (monitoring) modules in the testbed we can detect attacks originating from the Internet and potential weaknesses of the IoT device. Secondly, we can use the IoT testbed as part of the SIPHON mechanism. By placing an IoT device in the testbed, we can use the stimulation modules of the testbed to increase the realism of the SIPHON honeypot.

8. REFERENCES

- [1] DLink dcs-930l camera vulnerability.
<http://securityaffairs.co/wordpress/49143/breaking-news/d-link.html>. Accessed: 2016-08-10.

- [2] Masscan the internet port scanner. <http://tools.kali.org/information-gathering/masscan>. Accessed: 2016-08-10.
- [3] 2005. ITU report : The Internet of Things.
- [4] *The DecoyPort: Redirecting Hackers to Honeypots*. Springer Berlin Heidelberg, September 2007.
- [5] 2016. Gartner report Forecast: IoT Security, Worldwide.
- [6] 2016. IDC report Internet of Things: Security Practices.
- [7] Eric Alata, Vincent Nicomette, Marc Dacier, Matthieu Herrb, et al. Lessons learned from the deployment of a high-interaction honeypot. *arXiv preprint arXiv:0704.0858*, 2007.
- [8] Eugene Albin. *A comparative analysis of the snort and suricata intrusion-detection systems*. PhD thesis, Naval Postgraduate School, CA, USA, 2011.
- [9] E. Androulaki, C. Soriente, L. Malisa, and S. Capkun. Enforcing location and time-based access control on cloud-stored data. In *Proceedings of Conference on Distributed Computing Systems (ICDCS)*, June 2014.
- [10] Roland Bodenhein, Jonathan Butts, Stephen Dunlap, and Barry Mullins. Evaluation of the ability of the shodan search engine to identify internet-facing industrial control devices. *International Journal of Critical Infrastructure Protection*, 7(2):114–123, 2014.
- [11] Davide Canali and Davide Balzarotti. Behind the scenes of online attacks: an analysis of exploitation behaviors on the web. In *20th Annual Network & Distributed System Security Symposium (NDSS)*, 2013.
- [12] Dyn attack 2016. <http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/>. Accessed: 2016-12-06.
- [13] Wenjun Fan, Zhihui Du, and David Fernández. Taxonomy of honeynet solutions. In *SAI Intelligent Systems Conference (IntelliSys), 2015*, pages 1002–1009. IEEE, 2015.
- [14] Aurélien Francillon, Boris Danev, and Srdjan Capkun. Relay attacks on passive keyless entry and start systems in modern cars. In *Proc. Network and Distributed System Security Symp. (NDSS)*, 2011.
- [15] Julian B. Grizzard, Sven Krasser, and Henry L. Owen. The Use of Honeynets to Increase Computer Network Security and User Awareness. *Journal of Security Education*, 1(2-3):23–37, 2005.
- [16] M. Guri, Y. Poliak, B. Shapira, and Y. Elovici. JoKER: Trusted detection of kernel rootkits in android devices via JTAG interface. In *Proceedings of Trustcom*, volume 1, pages 65–73, Aug 2015.
- [17] Philip Hane. IPWHOIS: A library for RDAP (HTTP) lookups. <https://pypi.python.org/pypi/ipwhois>, 2015.
- [18] Thorsten Holz, Markus Engelberth, and Felix Freiling. Learning more about the underground economy: A case-study of keyloggers and dropzones. In *European Symposium on Research in Computer Security*, pages 1–18. Springer, 2009.
- [19] Y-C Hu, Adrian Perrig, and David B Johnson. Packet leashes: a defense against wormhole attacks in wireless networks. In *Proc. of the IEEE Conference on Computer Communication (InfoCom)*, volume 3, pages 1976–1986. IEEE, 2003.
- [20] Yih-Chun Hu, Adrian Perrig, and David B Johnson. Wormhole detection in wireless ad hoc networks. Technical Report Tech. Rep. TR01-384, Department of Computer Science, Rice University, 2002.
- [21] Yih-Chun Hu, Adrian Perrig, and David B Johnson. Wormhole attacks in wireless networks. *IEEE journal on selected areas in communications*, 24(2):370–380, 2006.
- [22] Miyoung Kim, Misun Kim, and Youngsong Mun. Design and implementation of the honeypot system with focusing on the session redirection. In *International Conference on Computational Science and Its Applications*, pages 262–269. Springer, 2004.
- [23] I Kotuliak, P Rybár, and P Trúchly. Performance comparison of ipsec and tls based vpn technologies. In *Proceedings of Conference on Emerging eLearning Technologies and Applications (ICETA)*, pages 217–221. IEEE, 2011.
- [24] Q. D. La, T. Quek, J. Lee, S. Jin, and H. Zhu. Deceptive attack and defense game in honeypot-enabled networks for the internet of things. *IEEE Internet of Things Journal*, PP(99):1–1, 2016.
- [25] John C Matherly. SHODAN the computer search engine. <https://www.shodan.io>. Accessed: 2016-08-01.
- [26] Mirai malware 2016. <http://blog.malwaremustdie.org/2016/08/mmd-0056-2016-linuxmirai-just.html>. Accessed: 2016-12-04.
- [27] Yin Minn Pa Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, and Christian Rossow. IoTPOT: Analysing the Rise of IoT Compromises. In *9th USENIX Workshop on Offensive Technologies (WOOT)*. USENIX Association, 2015.
- [28] Radek Píbil, Viliam Lisý, Christopher Kiekintveld, Branislav Bošanský, and Michal Pěchouček. *Game Theoretic Model of Strategic Honeypot Selection in Computer Networks*, pages 201–220. Springer Berlin Heidelberg, November 2012.
- [29] F Pouget, M Dacier, and VH Pham. on the advantages of deploying a large scale distributed honeypot platform. In *Proceedings of the E-Crime and Computer Evidence Conference*, 2005.
- [30] Niels Provos. A virtual honeypot framework. In *Proc. of the USENIX Security Symposium*, 2004.
- [31] Niels Provos and Thorsten Holz. *Virtual honeypots: from botnet tracking to intrusion detection*. Pearson Education, 2007.
- [32] Shachar Siboni, Asaf Shabtai, Nils Ole Tippenhauer, Jemin Lee, and Yuval Elovici. Advanced security testbed framework for wearable iot devices. *ACM Transactions on Internet Technology (TOIT)*, 16(4):26, 2016.
- [33] Lance Spitzner. The honeynet project: Trapping the hackers. *IEEE Security & Privacy*, 1(2):15–23, 2003.
- [34] Weizhe Zhang and Baosheng Qu. Security architecture of the internet of things oriented to perceptual layer. *International Journal on Computer, Consumer and Control (IJ3C)*, 2(2):37–45, 2013.