Inferring, Characterizing, and Investigating Internet-Scale Malicious IoT Device Activities: A Network Telescope Perspective

Sadegh Torabi[1], Elias Bou-Harb[2], Chadi Assi[1], Mario Galluscio[2], Amine Boukhtouta[1,3], and Mourad Debbabi[1]
[1]Computer Security Lab, Faculty of Engineering and Computer Science, Concordia University, Montreal, QC, Canada
{sa_tora, assi, debbabi, a_boukh}@encs.concordia.ca
[2]Cyber Threat Intelligence Lab, College of Engineering and Computer Science, Florida Atlantic University, Boca Raton, FL, USA
{ebouharb, mgalluscio2016}@fau.edu
[3]Ericsson Research Canada, Saint-Laurent, QC, Canada

***Abstract*—Recent attacks have highlighted the insecurity of the Internet of Things (IoT) paradigm by demonstrating the impacts of leveraging Internet-scale compromised IoT devices. In this paper, we address the lack of IoT-specific empirical data by drawing upon more than 5TB of passive measurements. We devise data-driven methodologies to infer compromised IoT devices and those targeted by denial of service attacks. We perform large-scale characterization analysis of their traffic, as well as explore a public threat repository and an in-house malware database, to underlie their malicious activities. The results expose a significant 26 thousand compromised IoT devices "in the wild," with 40% being active in critical infrastructure. More importantly, we uncover new, previously unreported malware variants that specifically target IoT devices. Our empirical results render a first attempt to highlight the large-scale insecurity of the IoT paradigm, while alarming about the rise of new generations of IoT-centric malware-orchestrated botnets.**

I. Introduction

In recent years, Internet connected devices, or what is currently known as Internet of Things (IoT) devices, have been widely adopted in various parts of our lives. IoT devices and corresponding technologies facilitate efficient data collection, monitoring, and information sharing for consumers (e.g., Internet routers, smart TVs, health monitoring wearables), and Cyber-Physical Systems (CPS) (e.g., power utilities, manufacturing plants, factory automation) [1]. Despite their benefits, the always-connected nature of IoT devices and the inadequate security measures implemented by some manufacturers [2], have turned these devices into attractive targets for cyber adversaries. Furthermore, IoT devices could be used as enablers for orchestrating large-scale attacks towards a variety of targets. The `Mirai` botnet for instance, utilized millions of compromised IoT devices (e.g., CCTV cameras) to launch Distributed Denial of Service (DDoS) attacks on several DNS servers, resulting in service disruption for millions of Internet users across the globe [3]. Very recently, the `Reaper` botnet extended `Mirai` by exploiting IoT-specific vulnerabilities rather than simply guessing credentials [4].

In order to adopt proper mitigation measures and prevent large-scale, IoT-related cyber attacks, security researchers and operators need to assess the magnitude of Internet-scale IoT exploitations, in addition to characterizing and analyzing their malicious activities. Nevertheless, given the lack of empirical data related to IoT devices [5], in addition to their excessive Internet-wide deployments in consumer and CPS, there is an utmost need to explore data-driven methodologies to shed the light and comprehend the characteristics of such compromised IoT devices and their malicious behaviors. To address the lack of knowledge about compromised IoT devices, there is a need to posses an Internet-scale perspective of IoT devices and their unsolicited activities over a period of time. This indeed is quite challenging as it requires authorization from different entities who own and operate these IoT devices in their local realms. Furthermore, monitoring IoT traffic would come with underlying privacy implications. Moreover, there are tremendous variants of IoT devices operating from all around the world and monitoring them would require scalable systems and significant resources.

An effective approach to gain Internet-wide cyber threat intelligence is to study passive measurements gathered using designated sensors or traps that collect traffic from the Internet [6], [7]. These sensors collect traffic targeted towards routable, yet unused Internet Protocol (IP) addresses, which are known as darknets or network telescopes [8]. Characteristically, traffic destined to these inactive hosts is likely to represent suspicious and unsolicited activities. Furthermore, traffic captured at the darknet mainly consists of scanning [9], [10], backscatter traffic resulting from DDoS attacks [11]–[13], and misconfiguration [8], [14]. Therefore, by carefully studying darknet traffic, one can generate useful insights on a portion of unsolicited traffic related to different sources including compromised machines (e.g., malware-infected) and victims of DDoS attacks, to name a few. To this end, in this work, we aim at addressing the problems of inferring Internet-scale compromised IoT devices and analyzing their unsolicited/malicious activities by exploring auxiliary, macroscopic, empirical passive darknet data obtained from a large network telescope. Specifically, we frame the contributions of this paper as follows:

- We draw-upon close to 5TB of recent darknet data and execute correlations with a near real-time IoT database to empirically characterize the magnitude of Internet-scale IoT exploitations in both, consumer and critical CPS realms. The generated insights not only render a first attempt ever to empirically shed the light on the large-scale insecurity of

the IoT paradigm, but are also intended to contribute to operational/actionable cyber security by providing Internet-wide, IoT-tailored notifications of such exploitations, thus permitting rapid remediation.
- We execute a first-of-a-kind, large-scale empirical characterization and analysis of IoT-centric unsolicited activities as perceived by a large network telescope. To this end, we uncover the nature of such traffic, its sources, employed protocols, targeted ports, upon various others. Given the lack of IoT-specific attack signatures, we postulate that the analyzed traffic from this work could be leveraged to design such signatures, in addition to promoting and facilitating further IoT-tailorted forensic investigations by making the captured unsolicited empirical traffic available to the research and operations communities at large.
- Motivated by the rise of new malware families/variants that specifically target and exploit IoT devices such as `Persirai`, `Hajime` and `BrickerBot`, to name a few, we execute non-intrusive correlations between passive measurements and malware threat intelligence to uncover new, previously unreported malware families targeting the IoT paradigm. In this context, we explore a publicly available threat repository and an in-house built malware database facilitated by instrumenting a large corpus of malware samples in a controlled sandbox. The results not only alarm about the severity of this malware issue in the context of the IoT, but also paves the way for future work for addressing the rise of IoT-centric, orchestrated botnets.

The remainder of the paper is organized as follows. Section II reviews the recent literature on various concerned topics to highlight the uniqueness of the proposed work. Section III details the methodology to infer Internet-scale compromised IoT devices by leveraging network telescopes. Section IV performs a large-scale empirical characterization of the generated unsolicited traffic from such IoT devices, putting special emphasis on understanding the nature of the traffic. Section V explores the maliciousness of the identified IoT devices, highlighting their involvement in various misdemeanors as well as pinpointing several newly discovered IoT-specific malware families. Lastly, Section VI provides a discussion on several insightful observations and current work limitations, while Section VII summarizes the outcomes of this work and highlights several topics that pave the way for future work.

II. Related Work

In this section, we review the literature on various concerned topics and highlight the added-value of the proposed work.

IoT security and protocol vulnerabilities. The majority of IoT security research work has been dedicated to synthesizing IoT context-aware permission models. For instance, Yu et al. [5] proposed a policy abstraction language that is capable of capturing relevant environmental IoT contexts, security-relevant details, and cross-device interactions, to vet IoT-specific network activities. Along the same research direction, Jia et al. [15], proposed ContextIoT, a system that is capable of supporting complex IoT-relevant permission models through efficient and usable program-flow and runtime taint analysis. Fernandes et al. [16] proposed a similar program-flow tracking approach that used taint arithmetic to detect policy violations and restrict traffic generated from exploited IoT application. In the context of protocol vulnerabilities, Ur et al. [17] studied numerous types of home automation IoT devices and unveiled various insights with regards to the security and usability of the implemented access control models. Ronen and Shamir [18] demonstrated information leakage attacks by instrumenting a set of IoT smart lights.

IoT data capturing initiatives. Given the rareness of IoT-relevant empirical data, several recent efforts were proposed to collect, curate, and analyze such data. The first IoT tailored honeypot, namely, IoTPOT, was designed and deployed by Pa et al. [19]. IoTPOT emulates Telnet services of various IoT devices running on different CPU architectures. In alternative work, Guarnizo et al. [20] presented the Scalable high-Interaction Honeypot platform for IoT devices (SIPHON). The authors demonstrated how by leveraging worldwide wormholes and few physical devices, they were able to mimic various IoT devices on the Internet and to attract significant malicious traffic.

Network telescope measurements and analysis. The idea of leveraging network telescopes to monitor unused IP addresses for security purposes was first brought to light in the early 1990's by Bellovin for AT&T's Bell Labs Internet-connected computers [21], [22]. Since then, the focus of network telescope studies has shifted several times, closely following the volatile nature of new threat actors. For instance, some of the important contributions include the discovery of the relationship between backscattered traffic and DDoS attacks in 2001 [23], worm propagation analysis between 2003 and 2005 [24], [25], the use of time series and data mining techniques on telescope traffic in 2008 [26], the monitoring of large-scale cyber events through telescopes in 2014 [27], and more recently, the study of amplification DDoS attacks using telescope sensors [28], [29].

This paper compliments the previous contributions by extending network telescope research to particularity address the problem of IoT security, which has yet to be attempted. To this end, the paper develops unique data-driven methodologies to infer and characterize compromised IoT devices, their unsolicited traffic, and their involvement in illicit activities. The paper also sheds light on new, previously undocumented malware families that specifically target IoT devices.

III. Identifying Unsolicited Internet-Scale IoT Devices

We initiate our work by addressing the problems of identifying and characterizing Internet-scale unsolicited IoT devices. We refer to IoT devices as being unsolicited (or compromised) if they were found to be generating any network packets towards the network telescope. Please note that Section IV will detail the nature of such unsolicited traffic and provide an in-depth characterization of its modus-operandi. We herein initially elaborate on the employed datasets and subsequently

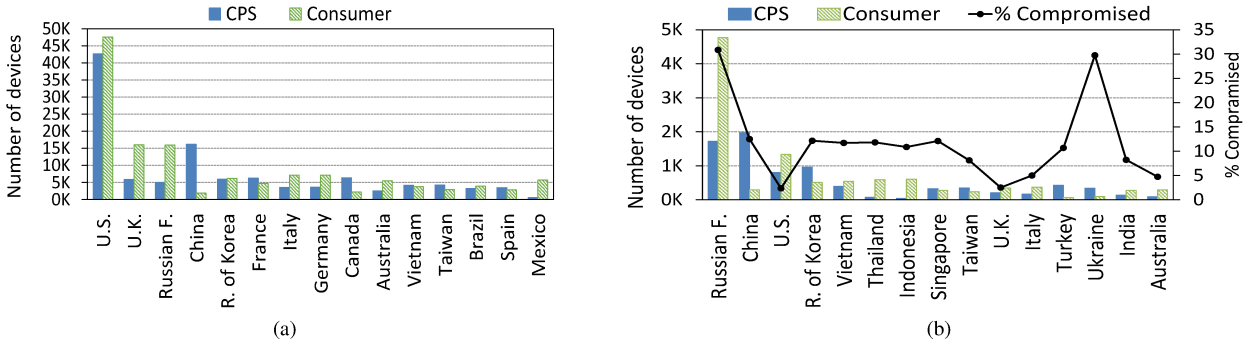

Fig. 1. Top 15 countries hosting the most (a) deployed IoT devices in the obtained data (cumulative percentage 69.3%), and (b) compromised IoT devices (CPS and consumer IoT devices).

provide the methodology and results towards the goal of inferring compromised IoT devices.

A. Obtained Data

1) IoT Device Information: It is quite difficult, if not impossible, to obtain technical information related to Internet-wide IoT devices that have been deployed in consumer and CPS environments due to privacy and logistic reasons. In addition, there is a lack of knowledge about effective fingerprinting approaches for identifying IoT devices by solely observing network traffic. Considering these challenges however, in this work, we leverage a near real-time IoT database provided by `Shodan` [30]. This service executes large-scale active measurements to identify and index Internet-facing IoT devices.

To this end, we obtained information related to 331,000 IoT devices from `Shodan`. These IoT devices, which were deployed in more than 200 countries all around the world, belong to consumer and CPS realms. On one hand, consumer IoT devices represent wireless access points and routers, IP cameras (e.g., webcams and CCTV cameras), printers, network storage media, satellite TV box and digital video recorders (DVRs), and electric hubs/outlets. On the other hand, IoT devices in CPS realms are involved in monitoring, controlling, and managing industrial/automation operations. They represent programmable logic controllers (PLC), remote terminal units (RTU), or other smart equipment that are used in industrial control systems (ICS), supervisory control and data acquisition systems (SCADA), and/or distributed control system (DCS). We obtained information related to approximately 181,000 consumer IoT devices, including, routers (46.9%), printers (29.1%), IP cameras (18.3%), and network storage media (4.6%). The remaining consumer IoT devices accumulate to only 1.1% of the total devices. We also obtained data related to 150,000 IoT devices in CPS that supported 31 industrial/control automation protocols/services. These CPS devices belong to a number of industries including but not limited to: building automation, power generation and distribution, control systems, plant/factory automation, oil and gas transportation, and embedded IoT communications.

As depicted in Figure 1a, the U.S. hosted the largest number of IoT devices (25%), followed by a significantly less number of devices hosted in the U.K. (6%), Russia (5.9%), and China (5%), respectively. Furthermore, by looking at the top 15 countries with the most number of IoT devices (Figure 1a), which account for about 69% of all IoT devices, we noticed that the number of consumer IoT devices were relatively higher than those deployed in CPS for the listed countries except for China, France, Canada, Vietnam, Taiwan, and Spain.

2) Network Telescope Data: Darknet data consists of one-way traffic targeted towards routable, allocated yet unused IP addresses (dark IP addresses). Since these IP addresses are not bound to any services, any traffic targeting them is characteristically unsolicited [8], [14]. Typically, darknet data consists of scanning, backscatter, and misconfiguration traffic [8], [9], [11], [12], [14], [31]. We explored over 5TB of darknet traffic between April 12-18, 2017 (about 80 GB of daily traffic). The darknet traffic is obtained from the UCSD real-time network telescope data maintained by the Center for Applied Internet Data Analysis (CAIDA) [32]. It is one of largest available sources of passive darknet traffic with about 16.7 million globally routed destination IPv4 addresses (i.e., /8 network) capturing over a billion packets every hour. The processed darknet traffic is stored in "flowtuple" files. Each file represents incoming flows towards the darknet that consist of the following flowtuple information: source/destination IP addresses and used ports, protocol, time to live (TTL), TCP flags, IP length, and total number of packets. The daily darknet traffic consists of unique compressed files representing hourly traffic (maximum of 24 files per day). We found that the available data for April 18 was incomplete, with only 15 hours of collected traffic (data might be missing due to technical issues at the telescope). To maintain consistency, we decided to remove the incomplete data from further analysis throughout the paper, resulting in 143 hours of analyzed darknet data that was obtained over 6 days between April 12-17, 2017.

B. Inferring and Characterizing Unsolicited IoT Devices

To infer compromised IoT devices, we executed a correlation algorithm that leverages IP header information to associate the obtained IoT device information with darknet flows. A significant 26,881 IoT devices were found interacting with the darknet, representing relatively more compromised

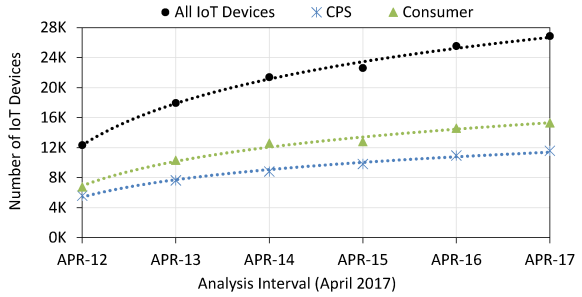

Fig. 2. The cumulative number of daily discovered compromised CPS and consumer IoT devices at the darknet over the 6 days analysis interval.

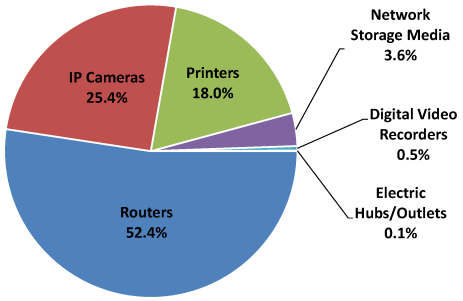

Fig. 3. Percentage of compromised consumer IoT devices by type/category.

consumer IoT devices (57%) than CPS (43%). As shown in Figure 2, slightly over 12,000 (46%) unsolicited IoT devices were correlated with the darknet data at the first day of the analysis (April 12, 2017). For the remaining time period, we discovered an average of about 2,900 newly compromised IoT devices per day. Considering the overall and cumulative numbers of uncovered unsolicited IoT devices, we definitely anticipate that an extended analysis period would result in discovering even more compromised IoT devices.

The compromised IoT devices were located across 161 countries, with the largest number of devices to be hosted in Russia (24.5%), followed by China (8.6%), and the U.S. (8.1%), respectively (Figure 1b). It is worth noting that while the U.S. and the U.K. hosted more number of IoT devices as compared to Russia and China (Figure 1a), the latter countries were found to host a relatively higher number of unsolicited IoT devices, as illustrated in Figure 1b. Furthermore, while Thailand, Indonesia, Singapore, Turkey, Ukraine, and India were not listed among the top 15 hosts with the most deployed IoT devices (Figure 1a), it is interesting to find them among the top 15 countries with the most number of uncovered compromised IoT devices. In fact, Figure 1b illustrates a significant difference in the percentage of unsolicited IoT devices found in Russia (31%) and Ukraine (30%), as compared to countries such as the U.S. (2.4%) and the U.K. (2.5%). While the actual reason behind this significant difference is quite obscured, this might indicate the enforcement of a stronger and more effective IoT security measures and policies in the U.S. and the U.K. in comparison to other countries.

1) Compromised IoT Devices in Consumer Realms: We identified 15,299 unsolicited consumer IoT devices that were correlated with the darknet over the analysis period. These IoT devices were located across 145 countries, with Russia hosting the highest percentage of compromised consumer IoT devices (32%), followed by the U.S. (9%), Indonesia (4%), and Thailand (4%), respectively. These IoT devices were connected to the Internet via 1,762 different Internet Service Providers (ISP), with the Russian "JSC ER-Telecom" hosting the highest percentage of compromised consumer IoT devices (27.6%), as summarized in Table I. In addition, about 52.4% of compromised consumer IoT devices were Internet routers, followed by IP cameras (25.2%), printers (18%), and network storage media (3.6%), respectively. As illustrated in Figure 3, these aforementioned devices accounted for about 99.4% of all consumer IoT devices, while TV boxes/DVRs and electric hubs/outlets represented less than 0.6% of all compromised consumer devices.

TABLE I
TOP 5 ISPS HOSTING THE HIGHEST NUMBER OF COMPROMISED CONSUMER IOT DEVICES.

ISP	Country	Devices	%
JSC ER-Telecom	Russian F.	4,205	**27.6**
PT Telkom	Indonesia	542	3.6
Korea Telecom	R. of Korea	339	2.2
PLDT	Philippine	311	2.0
TOT	Thailand	277	1.8

2) Compromised IoT Devices in CPS Realms: We identified 11,582 compromised IoT devices in CPS environments that were located in 136 countries, with China, Russia, Korea, and the U.S. hosting about 17%, 14.8%, 8.3%, and 6.9% of all the compromised devices respectively. The IP addresses of these devices were associated with 2,279 different ISP across the identified countries. As pinpointed in Table II, "Rostelecom" hosts the highest percentage of compromised IoT devices (about 4%), followed by "Korea Telecom" (3.8%) and "Turk Telekom" (3.2%), respectively.

TABLE II
TOP 5 ISPS HOSTING THE HIGHEST NUMBER OF COMPROMISED IOT DEVICES IN CPS REALMS.

ISP	Country	Devices	%
Rostelecom	Russian F.	461	**4.5**
Korea Telecom	R. of Korea	429	3.8
Turk Telekom	Turkey	347	3.2
HiNet	Taiwan	261	2.5
JSC ER-Telecom	Russian F.	277	1.8

Furthermore, a range of 31 services/protocols were operated by such compromised IoT devices. These services are not mutually exclusive, and therefore, an IoT device in a specific CPS might support one or more of these services. The top 10 operated services/protocols by the most number of unsolicited IoT devices are summarized in Table III. Among all the supported services/protocols, *Telvent OASyS DNA* (20%),

TABLE III
TOP 10 CPS REALMS HOSTING COMPROMISED IOT DEVICES.

Service/Protocol	Common applications	Devices	%
Telvent OASyS DNA	Oil and Gas transportation pipelines and distribution networks	**2,328**	**20.0**
SNC GENe	Control systems	2,126	18.3
Niagara Fox	Building automation systems	1,554	13.4
MQ Telemetry Transport	IoT communications, sensory networks, safety-critical communications	1,497	12.9
Ethernet/IP	Manufacturing automation	1,490	12.8
ABB Ranger	Power generating plants, transmission lines, mining operations, and transportation systems	1,061	9.1
Siemens Spectrum PowerTG	Utility networks	685	5.9
Modbus TCP	Power utilities	639	5.5
Foxboro/Invensys Foxboro	Plant automation systems, flowmeters, single-loop controllers, and product support services	590	5.1
Foundation Fieldbus HSE	Plant and factory automation	354	3.0

which operates in critical oil and gas CPS, and *Niagara Fox* (13.4%), which is common in building automation systems, appear among the most prevalent. CPS hosting compromised IoT devices also include those related to power utilities and manufacturing plants. Having noted this, it is indeed alarming (to say the least) to infer over 11,000 compromised IoT devices operating in such critical and error-sensitive environments.

IV. CHARACTERIZING UNSOLICITED TRAFFIC FROM INTERNET-SCALE IOT DEVICES

The aim of this section is to dissect, thoroughly comprehend, and characterize the unsolicited traffic generated by the inferred compromised IoT devices as perceived by the network telescope. We observed about 141.3M packets that were sent to the darknet from the 26,881 compromised IoT devices (daily $mean = 23.5$M and $\sigma = 0.92$M packets). On average, we captured 10,889 unsolicited IoT devices generating traffic towards the darknet on a daily basis, with slightly larger number of active consumer IoT devices (53%) on a daily basis. In general, consumer IoT devices, which represent 57% of all compromised IoT devices, generated more packets towards the darknet as compared to compromised devices in CPS realms, with approximately 62M packets (daily $mean = 10$M and $\sigma = 1.01$M), and 50M packets (daily $mean = 8.3$M and $\sigma = 1.05$M) for each device type respectively.

Considering the critical CPS contexts in which the compromised IoT devices operate in, it is worrisome to observe their aggressive role in generating significant amount of unsolicited activities. Interestingly, the statistical analysis using a Mann-Whitney U test indicated that the number of packets generated towards the darknet was significantly greater for devices in CPS than for consumer IoT devices ($p < 0.0001$). The higher activities however, might be attributed to the nature of the compromised IoT devices in CPS, which might have access to more powerful processing capabilities, as compared to other IoT devices, which typically have limited processing and memory resources. By contrast, the lower activity rate of compromised consumer IoT devices might be due to the stealthy nature of their generated activities, which aim at maximizing reachability while attempting to avoid detection. In what follows, we further explore the natures and characteristics of unsolicited traffic that have been generated by Internet-scale compromised IoT devices.

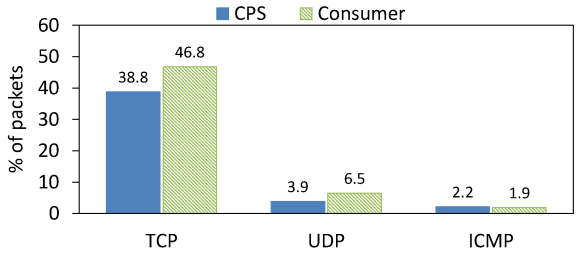

Fig. 4. Percentage of TCP, UDP, and ICMP traffic generated by compromised IoT devices in CPS and consumer realms.

A. Unsolicited UDP Traffic

The analyzed UDP packets represent about 10.4% of all traffic generated by the unsolicited IoT devices, with slightly more UDP packets generated by compromised consumer IoT devices as compared to those in CPS, as illustrated in Figure 4. Indeed, it is well known that UDP packets have been used to scan the Internet for open ports/services [9], [33], in addition to being employed to perform DoS attacks by flooding destination IP addresses or by exploiting open resolvers, causing amplification DoS attacks [29]. Thus, due to the stateless nature of UDP packets, it is quite challenging to classify them into a specific traffic category without further packet inspection. To maintain the focus of this paper, we do not address this challenging objective herein, though we will explore methodologies similar to [13] in future work to achieve this task. Nonetheless, to gain insights related to IoT-generated UDP traffic, we provide an overall characteristic analysis of the observed UDP packets in the following sub-sections.

1) UDP Packets: Overall, we observed about 13M UDP packets generated by a total of 25,242 compromised IoT devices, among which, about 60% were compromised consumer IoT devices, generating 63% of all UDP packets. Such IoT devices also targeted a significantly higher number of ports and destination IP addresses on an hourly average, as compared to those compromised IoT devices deployed in CPS (Figure 5). More specifically, compromised consumer IoT devices targeted an average of about 29,000 ports on more than

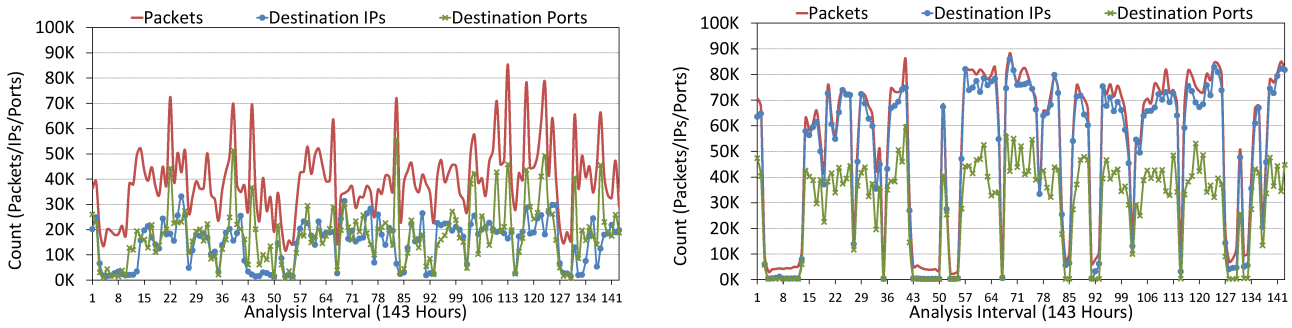

(a) CPS (b) Consumer

Fig. 5. Overall UDP packets sent by compromised (a) CPS and (b) consumer IoT devices to destination IP addresses and ports.

48,000 destination addresses, while compromised IoT devices in CPS targeted less ports (about 18,000) on significantly less number of destination addresses (14,700).

The comparison of the overall behavior of compromised IoT devices in CPS and consumer realms in terms of the generated UDP packets, and targeted destination addresses and ports, illustrates a number of differences. First, the compromised consumer IoT devices were actively sending UDP packets during repeated intervals that lasted for longer hours than those compromised in CPS. Second, the total number of generated UDP packets per hour by compromised consumer IoT devices was very close to the total number of targeted destination IP addresses (Figure 5b), and therefore, very few packets were sent towards each destination IP. We also found a strong positive correlation between the number of targeted ports and destination IP addresses by compromised consumer IoT devices (Pearson's correlation $r = 0.95$ and $p < 0.0001$), which may indicate the effort of such devices to reach a wider range of new destination IP addresses on various ports at each interval. Finally, the compromised CPS devices generated a significantly larger number of UDP packets per hour towards the targeted destinations, with packets possibly sent to a relatively larger number of ports on the same destinations, as illustrated by the recurring spikes in the number of contacted destination ports per hour (Figure 5a).

2) UDP Ports: The compromised IoT devices generated UDP packets towards all available UDP ports (65,535). About 10.7% of all UDP packets were targeting the top 10 ports (Table IV), while the remaining packets (89.3%) were distributed among over 60,000 ports. As shown in Table IV, port 37547 received about 329,000 UDP packets (2.5% of all), followed by port 137 (NetBIOS) and port 53413 with 2.06% and 2.05% of all UDP traffic respectively. In addition, while destination ports 37547, 32124, and 28183 were targeted by more than 9,000 compromised IoT devices, the remaining ports received UDP packets from significantly less number of compromised IoT devices (Table IV). We identified 5 assigned (well-known) services/protocols that correspond to the top 10 targeted ports. Nevertheless, although the remaining ports were not officially assigned to any services/protocols, some of them are known to be associated with known vulnerabilities. For instance, port 37547 has been associated with a backdoor to exploit and control "Netcore/Netis" routers [34].

TABLE IV
TOP 10 TARGETED UDP PROTOCOLS/PORTS.

Protocol/Port	Packets (K)	%	Devices
Not Assigned/37547	329.6	2.52	10,115
NetBIOS/137	269.9	2.06	144
Not Assigned/53413	268.1	2.05	91
Not Assigned/32124	141.2	1.08	9,488
Not Assigned/28183	122.5	0.94	9,710
mDNS/5353	99.4	0.76	165
Not Assigned/4605	50.3	0.38	150
DNS/53	42.6	0.33	158
Teredo/3544	34.4	0.26	226
OpenVPN/1194	34.0	0.26	96

B. Unsolicited Backscatter Traffic

Backscatter traffic, in the context of this work, is a byproduct of (D)DoS attacks that target IoT devices. When a victim IoT device is attacked by a flood of packets generated from spoofed source IP addresses (that happened to be belonging to the network telescope IP space), the device will generate reply packets destined to the darknet, which can then be collected and extracted. These packets are mainly TCP (SYN-ACK and RST) or ICMP reply packets (Echo Reply, Destination Unreachable, Source Quench, Redirect, Time Exceeded, Parameter Problem, Timestamp Reply, Information Reply, or Address Mask Reply) [13].

Our analysis revealed a total of 839 IoT devices that have fallen victims of DoS attacks, with about 10.3M backscatter packets generated towards the darknet (8.2% of total traffic). Approximately half of the victim IoT devices generated less than 170 backscatter packets towards the darknet, while about 17% of the IoT devices generated 10,000 or more backscatter packets (Figure 6). Moreover, only 7 devices generated 100,000 or more backscatter packets, among which 5 of them were likely to be operating in critical CPS realms. In general, about 73% of all backscatter packets were generated by IoT devices in CPS, which represent slightly more than half of the DoS victims (53%). In fact, 5 IoT devices in CPS contributed to about 43% of all backscatter traffic, with two devices

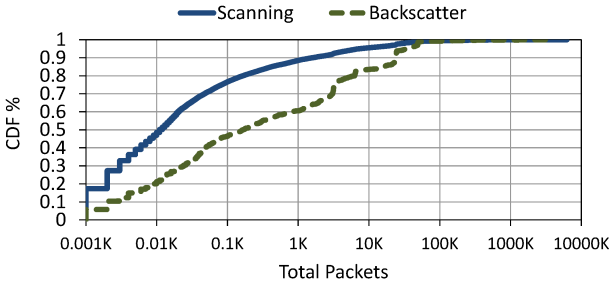

Fig. 6. Distribution of the scanning and backscatter packets generated by compromised IoT devices and DDoS victims, respectively.

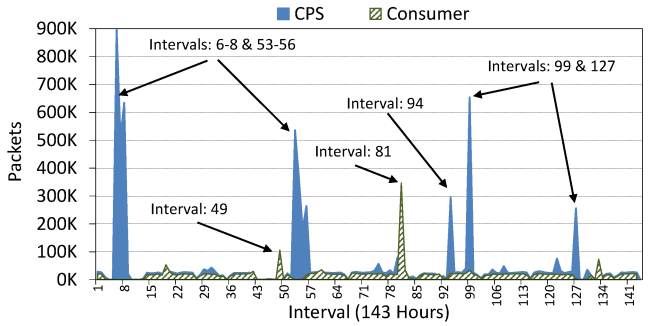

Fig. 7. Distribution of the generated backscatter packets by CPS and consumer IoT devices (143 hours).

that generated about 1.1 and 3.4 million packets respectively. These observations may reflect the nature of the inferred DoS attacks that were focused on target devices in CPS with higher intensity as compared to consumer IoT devices.

1) IoT DoS Victims: By investigating the distribution of backscatter packets as illustrated in Figure 7, we observed few instances with noticeable increase in the number of generated backscatter packets by the IoT devices (e.g., between intervals 6 and 8). These sudden spikes indicate a large magnitude of DoS attacks against CPS and consumer IoT devices during the specified time intervals. It is apparent that IoT devices in CPS realms were attacked more often and with higher intensity as compared to consumer IoT devices. In fact, a conducted Mann-Whitney U test showed a statistically significant difference between the number of generated backscatter packets per hour when comparing IoT devices in CPS and consumer realms ($p < 0.0001$, $U = 6061$, and $Z = -5.95$).

To further investigate targeted IoT devices as DoS victims, we focused at intervals with sudden spikes in the number of backscatter packets. Interestingly, a single victim IoT device generated almost all the packets during every DoS attack interval. For instance, an IoT device in a CPS realm located in China was responsible for more than 99% of all backscatter traffic during intervals 6-8 and 53-55, and about 89% of traffic at interval 56. Similarly, a different CPS device from China was found to be under DoS attacks during intervals 99 and 127, generating about 91% and 97% of all backscatter traffic at those intervals respectively. Both of the aforementioned IoT device operated Ethernet/IP on TCP/UDP port 44818, which is used in manufacturing automation. After some investigations, we inferred that this service is associated with "Rockwell Automation Control Logix PLC" vulnerabilities, which can cause DoS on the targeted IoT devices.[1] Finally, an IoT device in a CPS from Switzerland, which supports Telvent OASyS DNA (used in oil and gas transportation pipelines and distribution networks), contributed towards about 85% of the backscatter traffic at interval 94, indicating another instance of targeted DoS attacks.

Analyzing the DoS events for the consumer IoT devices resulted in similar behavior as the CPS devices. For instance, a printer located in the Netherlands, generated over 104,000 backscatter packets at interval 49, contributing towards 98% of all packets at this interval. In addition, another printer from the U.K. was found to be under targeted DoS attacks as it generated about 85% of all backscatter packets at interval 81.

2) Targeted Countries: The targted IoT devices were found to be located in 80 countries, with China, Singapore, and the U.S. hosting the highest number of DoS IoT victims, respectively (Figure 8a). Moreover, China and the U.S. hosted the most number of targeted IoT devices in CPS realms (103 and 49 devices, respectively), while Singapore and Indonesia hosted the highest number of consumer IoT device victims (64 and 52 devices). From a different perspective, about 52% of all backscatter traffic was generated by IoT devices hosted in China, followed by devices in the U.S. (5.9%) and the U.K. (4.1%), respectively. In addition, we noticed that the U.K, Brazil, Switzerland, and Argentina, were among the top 15 countries with the highest number of generated backscatter packets (Figure 8b), while hosting relatively few victim IoT devices (10, 16, 4, and 5 devices, respectively). This corroborates our previous findings regarding the nature of the observed DoS attacks during the analysis intervals, which represent intensive targeted attacks on a small number of victim IoT devices.

C. Unsolicited Scanning Traffic

Probing traffic generated from unsolicited IoT devices that target the network telescope is an indicator of exploitations of such IoT devices. Such compromised devices would typically be scanning the Internet looking to exploit vulnerable hosts or other IoT devices. In order to identify IoT-generated scanning traffic, we first looked at the remaining non-backscatter ICMP packets, which represented a very small percentage of the total generated traffic by compromised IoT devices (0.23%). More than 99.9% of these packets were ICMP Echo requests, which are typically used for remote network scans (e.g., ping). Moreover, these packets were originated from 56 exploited IoT devices, among which 32 consumer IoT devices generated the majority of the ICMP scanning packets (93%). We also identified slightly over 100M TCP packets that were not classified as backscatter. These TCP packets were mainly TCP-SYN packets (99.97%), which are commonly used for scanning the

[1] https://ics-cert.us-cert.gov/advisories/ICSA-13-011-03

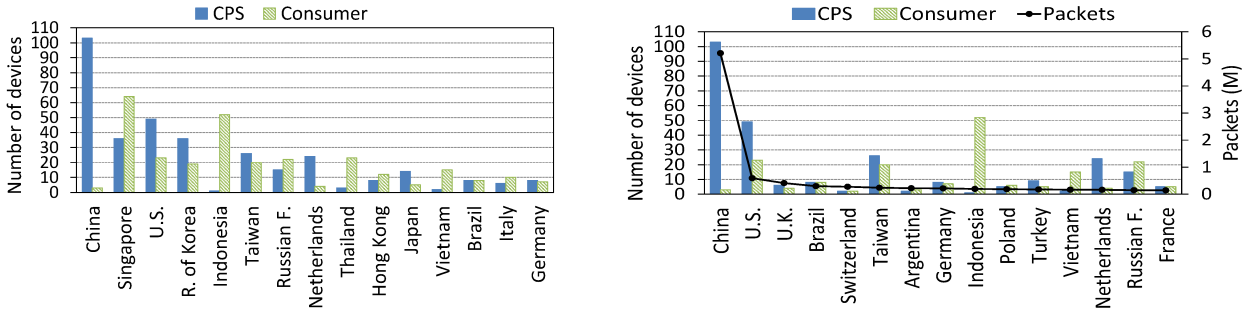

(a) (b)

Fig. 8. Distribution of targeted IoT devices in the top 15 countries with the highest number of (a) DoS victims, and (b) generated backscattered packets.

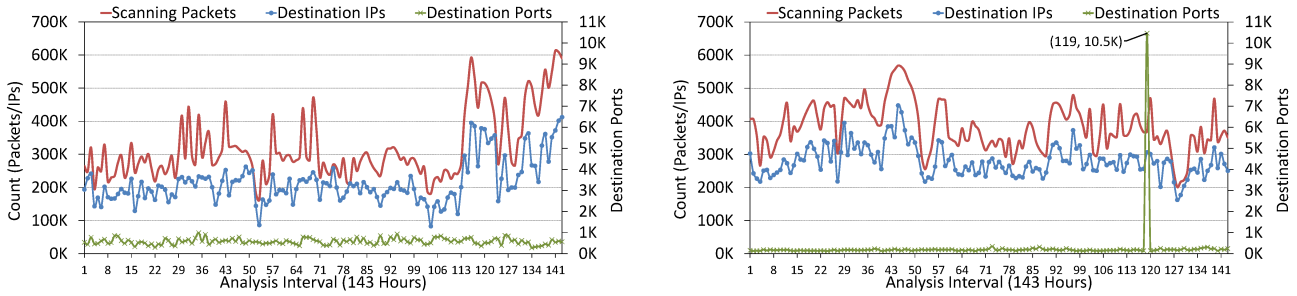

(a) CPS (b) Consumer

Fig. 9. Overall TCP scanning packets generated towards destination IP addresses and ports by exploited (a) CPS and (b) consumer IoT devices.

Internet [9]. The TCP scanning packets were generated by a total of 12,363 compromised IoT devices (55% consumer IoT devices). We illustrate the overall distribution of the TCP scanning packets generated by compromised IoT devices in both, CPS and consumer realms in Figure 9. On average, exploited consumer IoT devices generated more TCP scanning packets per hour, as compared to exploited IoT devices in CPS, with about 382,000 and 318,000 packets for each device type respectively. Nevertheless, while the compromised IoT devices included more consumer IoT devices (55%), the analysis showed no linear correlation between the number of compromised IoT devices and the total generated scanning packets per hour (Pearson's $r \approx 0$ and $p > 0.05$). In addition, exploited consumer IoT devices targeted relatively more destinations per hour, as compared to those deployed in CPS, with an average of 280,000 and 215,000 destinations respectively. Interestingly, while the exploited IoT devices in CPS scanned relatively less number of destinations, they seemed to be scanning a wider range of destination ports as compared to consumer IoT devices, with an average of 576 scanned ports per hour ($min = 271$ and $max = 987$). Exploited consumer IoT devices on the other hand, scanned a smaller range of ports per hour (average of 246 ports), except at interval 119, where a sudden increase in the number of scanned ports is clearly observed (Figure 9b). Investigating the data at interval 119 revealed 734 IoT devices that were generating TCP scanning packets. Among those devices, a single IP camera hosted in the Dominican Republic was responsible for scanning 10,249 ports on 55 destination addresses.

1) Scanned Protocols/Services: We summarize the top 14 protocols/services that received the most scanning activities from the exploited IoT devices in Table V. Telnet received the highest portion of all TCP scanning packets (about 50%), followed by HTTP and SSH, which received significantly less number of packets, with about 9.4% and 7.7% of all TCP scanning packets respectively.

TABLE V
TOP 14 PROTOCOLS/PORTS WITH THE MOST TCP SCANNING PACKETS GENERATED BY EXPLOITED IOT DEVICES (CP=93.3%).

Protocol/Port	Packets (M)	Packets (%)	Consumer (%)	Consumer IP	CPS (%)	CPS IP
Telnet /23/2323/23231	50.08	50.2	63.4	643	36.6	553
HTTP /80/8080/81	9.41	9.4	94.5	1418	5.5	345
SSH /22	7.68	7.7	33.7	64	66.3	80
BackroomNet /3387	6.2	6.2	–	–	100	1
CWMP /7547	4.49	4.5	44.8	169	55.2	244
WSDAPI-S /5358	4.05	4.1	59	94	41	48
MSSQLServer /1433	3.33	3.3	36.2	8	63.8	13
Kerberos /88	2.67	2.7	99	1061	1	23
MS DS /445	2.49	2.5	45.3	43	54.7	330
EtherneIP IO /2222	0.68	0.7	41.6	50	58.4	65
iRDMI /8000	0.67	0.7	98.5	1055	1.5	18
Unassigned /21677	0.57	0.6	0	1	100	87
RDP /3389	0.51	0.5	46.8	42	53.2	61
FTP /21	0.29	0.3	46	20	54	33

Overall, we observed that HTTP, Telnet, Kerberos, and iRDMI, were scanned by a noticeably larger number of

compromised IoT devices as compared to other protocols (Table V). In addition, a significantly larger number of compromised consumer IoT devices were scanning HTTP, Kerberos, and iRDMI protocols, contributing towards the majority of generated TCP scanning packets at these ports (TableV). On the other hand, while only one exploited IoT device in a CPS realm was actively scanning port 3387 (BackroomNet), almost all scanning packets generated towards port 21677 were also found to be generated by compromised CPS IoT devices (negligible TCP traffic was generated by a single compromised consumer IoT device).

The distribution of the TCP scanning packets targeting the top 5 protocols/services is illustrated in Figure 10. It is important to note that most of these protocols were also associated with the recent IoT-initiated cyber attacks (e.g., the `Mirai` botnet and its variations) [3]. In fact, our analysis revealed a number of compromised IoT devices that were actively involved in scanning these protocols. Moreover, these compromised devices were corroborated to perform malicious scanning by comparing them against a publicly available threat repository (`Cymon` [35]), as elaborated in Section V. In what follows, we present further analysis with regards to the top scanned protocol/services.

Telnet. It is clearly observed that Telnet received the highest amount of TCP scanning packets from 1,196 exploited IoT devices. In addition, slightly more compromised consumer IoT devices (54%) were scanning Telnet as compared to those deployed in CPS, generating about 63% of all TCP scans. Moreover, a total of 7 compromised IoT devices contributed towards 55% of all TCP packets targeting Telnet. These exploited IoT devices, which were hosted in different countries, represent three IP cameras, one router, DVR, and printer, and two devices in CPS associated to power utilities and utility networks. Interestingly, these compromised IoT devices were also associated with malicious scanning as indexed by `Cymon`.

SSH. We noticed sudden increases in the overall scanning activities towards SSH at intervals 32 and 69 (Figure 10), with about 242,000 and 253,000 TCP packets generated by compromised IoT devices respectively. Surprisingly, only a hand full of compromised IoT devices, mainly those in CPS, were generating the majority of the TCP scans at these intervals. In particular, two exploited routers hosted in Russia and Australia, and three compromised IoT in CPS (two hosted in China and one in Brazil), generated about 93% of the scans at interval 32. Interestingly, the three exploited IoT in CPS, which generated about 80% of the scanning packets at interval 32, were also found to generate the majority of all scanning traffic at interval 69 (about 90%). In fact, all of the five aforementioned compromised IoT devices were also associated with malicious scanning and/or SSH brute force attacks by `Cymon`.

BackroomNet. We noticed that BackroomNet was scanned by a single compromised IoT in CPS located in Canada, which operated *BACnet/IP* (used in building automation). As shown in Figure 10, the intensive scanning activity started at interval 113 (April 16), generating over 6.2 million packets during

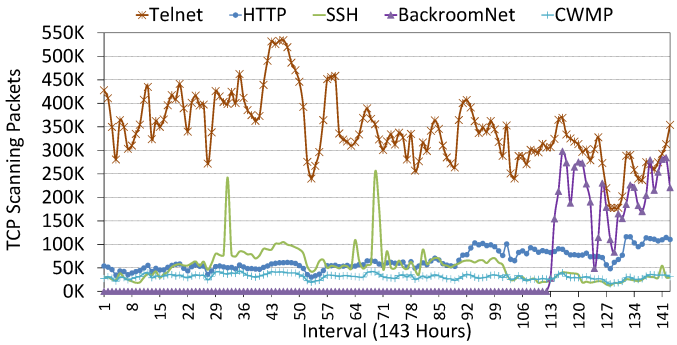

Fig. 10. The distribution of TCP scanning packets generated by exploited IoT devices towards the top 5 targeted protocols/services.

the next 30 hours (average of approximately 200,000 TCP scanning packets per hour). We also compared this suspicious activity against `Cymon`, and confirmed that it was being involved in malicious scanning activities.

HTTP. A total of 1,763 compromised IoT devices scanned HTTP ports, among which about 80% are consumer IoT devices. Moreover, these compromised consumer devices contributed towards the the majority (94.5%) of the scanning packets targeting HTTP ports, with an hourly average of about 62 thousand generated scanning packets from 415 exploited devices. It is interesting to see that despite the gradual increase in the number of generated scanning packets towards HTTP ports after interval 92 (Figure 10), the overall distribution of the scanning packets illustrates a more organized and uniform scanning behavior that does not involve noticeable behavioral changes from the compromised IoT devices. This behavior however, might be resulting from orchestrated stealthy scans generated by compromised IoT devices towards the Internet. Proving this would require further investigations and will be considered for future work.

CWMP. The CPE WAN Management Protocol (CWMP) is a web-based protocol that enables remote configuration and management of routers, gateways, and other IoT devices [36]. More importantly, CWMP was utilized by some variants of the `Mirai` botnet to exploit routers [3]. A total of 413 compromised IoT devices, among which 59% were CPS-related, generated more than 4M TCP scanning packets towards CWMP (Table V). On average, CWMP was scanned by 36 compromised IoT devices, generating more than 31,000 scanning packets per hour. As illustrated in Figure 10, these scans had the least variations in terms of magnitude of the generated packets during the analysis interval. Despite that, we noticed an exploited router, which was located in Australia, to generate relatively more scanning packets (10.6%), as compared to other compromised IoT devices. Moreover, a total of 5 exploited CPS-related IoT devices were also found to generate relatively more packets than those others deployed in other CPS, representing a total of about 25% of all scanning traffic on CWMP. Three of these devices, which supported *Ethernet/IP* (used in manufacturing automation), were hosted

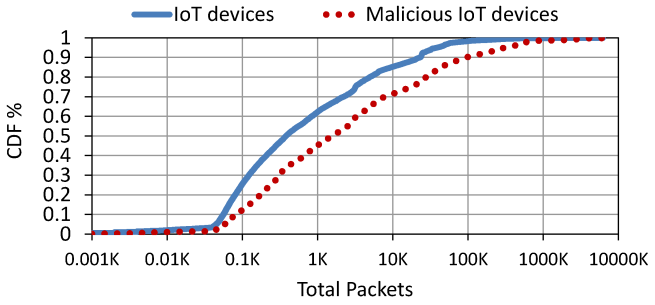

Fig. 11. Distribution of received packets from the top 8,839 IoT devices and the malicious devices flagged by Cymon ($N = 816$).

in Korea. The remaining two devices, which were used in control systems (*SNC GENe*) and oil and gas transportation pipelines and distribution networks (*Telvent OASyS DNA*), were located in China and South Africa, respectively. Finally, all but two of the aforementioned compromised IoT devices were confirmed to be performing malicious scanning of the Internet using Cymon.

V. Analyzing the Maliciousness of Unsolicited Internet-Scale IoT Devices

In this paper, we identified a large number of compromised/exploited IoT devices while characterizing their unsolicited traffic, which pinpointed to some malicious scanning activities. Motivated by these findings, and the plethora of IoT-centric malware that are currently "in the wild" (e.g., Mirai and Hajime), in this section, we aim at exploring the maliciousness of the inferred IoT devices by investigating: (1) whether such IoT devices are involved in other illicit activities, and (2) whether there exists other malware families and variants that could possibly be exploiting such IoT devices.

A. IoT Illicit Activities

To investigate the involvement of such IoT devices in malicious activities, we relied on a publicly available cyber-threat intelligence service provided by Cymon [35]. The latter renders a service to track and aggregate Internet-scale events related to IP addresses and domains, which are involved in malware, phishing, botnets, spamming, DNS blacklisting, scanning, and web attacks. We investigated the malicious activities associated with 8,839 exploited IoT devices, which represent all devices that generated backscatter traffic (839 DoS victims), and the top 4,000 compromised IoT devices with the most generated scanning and UDP packets from each IoT device category (consumer and CPS). As presented in Figure 11, about 10% of the explored IoT devices sent 50 or less packets to the darknet, while only 15% of them sent 10,000 packets or more during the analysis interval. In fact, while less than 2% generated 100,000 packets or more, only 15 devices sent more than 1M packets ($max = 6.25M$ packets). By correlating the explored IoT devices against those IP addresses indexed by Cymon, we uncovered 816 IoT devices (9.2%) that were linked to one or more malicious activities.

TABLE VI
Identified threats summary. Note that the identified threats are not mutually exclusive.

Threat Category	IoT Devices	%
Scanning	786	**96.3**
Miscellaneous (Web/FTP attacks, DNSBL, Malicious domains, VoIP)	574	70.3
Brute force (SSH)	252	30.9
Spam (Mail, IMAP)	227	27.8
Malware (Virus, Worm, Bot/Botnet, Trojan)	117	14.3
Phishing	5	0.6

We amalgamated the identified activities into 6 illicit categories, as summarized in Table VI. It is worthy to note that the threat categories are not mutually exclusive, since different sources of cyber threat intelligence (within Cymon) might flag a given host/IP with multiple malicious activities. The majority of the identified malicious IoT devices were associated with illicit scanning (96.3%). Furthermore, about 70% of the IoT devices were flagged as miscellaneous (e.g., Web attacks), while about 31% and 28% were associated with SSH brute force attacks and spamming respectively. Interestingly, a total of 117 IoT devices were linked to malware-related activities (14.3%), while only 5 devices were associated with phishing activities. Specifically, our findings identified a significant 91 IoT devices operating in various CPS realms that were indeed associated with malware, with the majority (85 devices) to be involved in TCP scanning activities. Furthermore, 26 consumer IoT devices were linked to malware, with 23 devices performing scanning activities. It is interesting to observe that a total of 9 devices (generating the DoS peaks of Figure 7), were found to be related to malware as well.

B. IoT-Centric Malware Families

Given the fact that we have identified 117 IoT devices that were related to malware, we attempted herein to further explore this matter. In this context, we relied on an in-house built database of malware information. The database is an artifact of conducting large-scale dynamic malware instrumentation. Indeed, we have been receiving a malware feed on a daily basis with an average of 30,000 malware samples from ThreatTrack Security.[2] XML reports are produced by analyzing the malware binaries in a controlled environment. It is worthy to mention that these reports contain the executed activities by the malware samples at the network and system levels. On one hand, the network level activities refer to the connections and the exchanged packets, including IP addresses, port numbers, URLs, visited domains and the actual payload data that has been sent. On the other hand, the system level activities constitute the list of Dynamic-link Library (DLL) files that are utilized by the malware, the key registry changes, and the memory usage. The malware database is built by parsing and indexing such XML malware reports. We executed correlations (using IP address information) between all the inferred unsolicited IoT devices from

[2]www.threattrack.com

TABLE VII
IDENTIFIED, PREVIOUSLY UNREPORTED MALWARE FAMILIES EXPLOITING IOT DEVICES

Malware Family
Ramnit
Starman
Kryptik
Nivdort
Razy
Zusy
Bayrod
Artemis
MSIL
Vupa
Allaple

Section III (i.e., 26,881 devices) and the malware database. The outcome is intended to demonstrate if any malware variant has communicated with (possibly exploited) IoT devices.

Our findings revealed 33 domain names and 24 unique malware hashes/variants associated with the identified IoT devices. Given the extracted malware hashes from the malware database, we leveraged `VirusTotal` to unveil 11 malware families that were found to be associated with the IoT devices. The uncovered malware families are summarized in Table VII. While some of the families are already quite popular, such as the `Ramnit` as a backdoor, and `Zusy` for generating email spam, our results demonstrate that new variants of such families are already being empowered to target the IoT paradigm. To the best of our knowledge, such results (*i*) render a first attempt ever to shed the light on IoT-centric malware families by correlating passive measurements and malware samples facilitated through dynamic analysis, (*ii*) highlight on new, previously unreported families (and variants) that have empirically been demonstrated to target the IoT paradigm, and (*iii*) alarm about the rise of new malware variants, which undoubtedly would facilitate the establishment of ever-evolving, IoT-tailored, malware-orchestrated botnets.

VI. DISCUSSION

In this section, we elaborate on a few topics that are worthy of being discussed in the context of the proposed work.

Comprehensively inferring Internet-scale unsolicited IoT devices. While this work leveraged the `Shodan` service to gather a large dataset of IP information related to deployed IoT devices in order to facilitate their correlation with passive measurements, identifying technical information for Internet-wide IoT devices remains a challenging objective. Indeed, without addressing this issue at large, approaches similar to the one presented in this paper would remain partially lacking (at least operationally). In this context, we foresee two approaches moving forward. The first is of technical nature and is rendered by exploring fuzzy matching algorithms and fuzzy hashes/signatures to identify a broader range of IoT devices (previously not indexed by `Shodan`) as perceived by the network telescope by leveraging IoT-relevant darknet traffic (from previously inferred IoT devices). The second is a non-technical approach, requiring ISPs, local IoT operators and industry to collaborate to make such IoT information available. Indeed, the authors are already in touch with `Cisco Systems` to have access to their IoT platform dubbed as `Jasper`, for a larger corpus of IoT device information. The vision is that threat intelligence from the proposed work will be fed back to `Cisco Jasper`.

Malware attribution for tailored remediation. With the continuous rise of new malware variants that specifically target IoT devices in consumer and critical CPS sectors, the objective to attribute such exploitations to certain malware variants is undeniably of high significance. While this work provided initial results from such attempt, further exploration is needed. To this end, we are currently exploring formal correlation approaches between passive measurements and malware network traffic samples to fortify the attribution evidence.

Real-time IoT operational cyber security capabilities and artifacts. IoT-tailored actionable cyber security capabilities are quite lacking and we hope that this work would motivate the research and operational communities to devise such capabilities. In fact, we are currently working to automate the devised methodologies in this work to index, in near real-time, unsolicited Internet-scale IoT devices. We are also working on creating an authenticated API to share IoT-relevant malicious empirical data, IoT-centric attack signatures, and threat intelligence derived from passive measurements with the research community at large.

VII. CONCLUSION

The Internet of Things (IoT) is an emerging paradigm of technical, social, and economic significance. Nevertheless, the initial priorities of IoT vendors have been focused on providing novel functionality, getting products to market sooner, and making IoT devices more accessible and easier to use. Unfortunately, security concerns have not received as much attention. To this end, this paper presented a first empirical look at the magnitude of compromised IoT devices that have been deployed in both consumer and CPS realms. Initially, large-scale correlations between passive measurements and IoT-relevant information is conducted to shed the light on Internet-wide unsolicited IoT devices. Subsequently, empirical measurements, characterization, and analysis is presented to thoroughly investigate IoT-generated unsolicited traffic, including backscattered traffic from IoT devices that have been targeted by DoS attacks, and scanning activities from exploited IoT devices. Finally, an attempt is made to uncover the maliciousness of such unsolicited IoT devices by utilizing a publicly available threat repository and an in-house built malware database. Some of the outcomes include more than 15,000 compromised consumer IoT devices and more than 11,000 compromised IoT devices operating in critical CPS (including oil and gas, manufacturing plants and power utilities). The results also demonstrate the aggressiveness of more than 5,000 compromised IoT devices in CPS in exploiting other services. The outcome also pinpoints the involvement of a large number of IoT devices in malevolent activities as well

as the rise of new malware variants that specifically target the IoT paradigm. Overall, the presented measurements from this work highlight, at large, the insecurity of the IoT paradigm. As for future work, apart from addressing a number of tasks and issues that have been pinpointed throughout this paper, we are working on addressing the challenging problem of identifying and clustering IoT botnets and their illicit activities by solely scrutinizing passive measurements.

ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to the anonymous reviewers for their constructive feedback. The work has been partially supported by Natural Sciences and Engineering Research Council of Canada (NSERC) and Concordia University. This work was also partially supported by a grant from the U.S. National Science Foundation (NSF) (Office of Advanced Cyberinfrastructure (OAC) #1755179).

REFERENCES

[1] Kyoung-Dae Kim and Panganamala R Kumar. Cyber–Physical Systems: A Perspective at the Centennial. *Proceedings of the IEEE 100*, (Special Centennial Issue):1287–1308, 2012.

[2] Jacob Wurm, Khoa Hoang, Orlando Arias, Ahmad-Reza Sadeghi, and Yier Jin. Security Analysis on Consumer and Industrial IoT Devices. In *21st Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 519–524. IEEE, 2016.

[3] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. Understanding the Mirai Botnet. In *26th USENIX Security Symposium*, pages 1093–1110, Vancouver, BC, 2017.

[4] Wired Andy Greenberg. The reaper IoT botnet has already infected a million networks. Online: https://www.wired.com/story/reaper-iot-botnet-infected-million-networks/, 2017.

[5] Tianlong Yu, Vyas Sekar, Srinivasan Seshan, Yuvraj Agarwal, and Chenren Xu. Handling a trillion (unfixable) flaws on a billion devices: Rethinking network security for the Internet-of-Things. In *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, page 5. ACM, 2015.

[6] Craig Labovitz, Abha Ahuja, and Michael Bailey. *Shining Light on Dark Address Space*. Arbor Networks Inc., 2001.

[7] Claude Fachkha, Elias Bou-Harb, Anastasis Keliris, Nasir Memon, and Mustaque Ahamad. Internet-scale Probing of CPS: Inference, Characterization and Orchestration Analysis. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, San Diego, California, 2017.

[8] Claude Fachkha and Mourad Debbabi. Darknet as a Source of Cyber Intelligence: Survey, Taxonomy, and Characterization. *IEEE Communications Surveys & Tutorials*, 18(2):1197–1227, 2016.

[9] Elias Bou-Harb, Mourad Debbabi, and Chadi Assi. On Fingerprinting Probing Activities. *Computers & Security*, 43:35–48, 2014.

[10] Zakir Durumeric, Michael Bailey, and J Alex Halderman. An Internet-Wide View of Internet-Wide Scanning. In *Proceedings of the 23rd USENIX Security Symposium*, pages 65–78, San Diego, CA, 2014.

[11] Nobuaki Furutani, Tao Ban, Junji Nakazato, Jumpei Shimamura, Jun Kitazono, and Seiichi Ozawa. Detection of DDoS Backscatter Based on Traffic Features of Darknet TCP Packets. In *Ninth Asia Joint Conference on Information Security (ASIA JCIS)*, pages 39–43. IEEE, 2014.

[12] Eray Balkanli and A Nur Zincir-Heywood. On the Analysis of Backscatter Traffic. In *Local Computer Networks Workshops (LCN Workshops), 2014 IEEE 39th Conference on*, pages 671–678. IEEE, 2014.

[13] Norbert Blenn, Vincent Ghiëtte, and Christian Doerr. Quantifying the Spectrum of Denial-of-Service Attacks Through Internet Backscatter. In *Proceedings of the 12th International Conference on Availability, Reliability and Security*, ARES '17, pages 21:1–21:10, Reggio Calabria, Italy, 2017.

[14] Eduard Glatz and Xenofontas Dimitropoulos. Classifying Internet One-way Traffic. In *Proceedings of the 2012 Internet Measurement Conference*, IMC '12, pages 37–50, Boston, MA, USA, 2012.

[15] Yunhan Jack Jia, Qi Alfred Chen, Shiqi Wang, Amir Rahmati, Earlence Fernandes, Z Morley Mao, Atul Prakash, and Shanghai JiaoTong Unviersity. ContexIoT: Towards Providing Contextual Integrity to Appified IoT Platforms. In *Proceedings of the 21st Network and Distributed System Security Symposium (NDSS'17)*, 2017.

[16] Earlence Fernandes, Justin Paupore, Amir Rahmati, Daniel Simionato, Mauro Conti, and Atul Prakash. FlowFence: Practical Data Protection for Emerging IoT Application Frameworks. In *USENIX Security Symposium*, 2016.

[17] Blase Ur, Jaeyeon Jung, and Stuart Schechter. The Current State of Access Control for Smart Devices in Homes. In *Workshop on Home Usable Privacy and Security (HUPS)*, 2013.

[18] Eyal Ronen and Adi Shamir. Extended Functionality Attacks on IoT Devices: The Case of Smart Lights. In *IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 3–12. IEEE, 2016.

[19] Yin Minn Pa Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, and Christian Rossow. IoTPOT: A Novel Honeypot for Revealing Current IoT Threats. *Journal of Information Processing*, 24(3):522–533, 2016.

[20] Juan David Guarnizo, Amit Tambe, Suman Sankar Bhunia, Martín Ochoa, Nils Ole Tippenhauer, Asaf Shabtai, and Yuval Elovici. Siphon: Towards Scalable High-Interaction Physical Honeypots. In *Proceedings of the 3rd ACM Workshop on Cyber-Physical System Security*, pages 57–68. ACM, 2017.

[21] Steve Bellovin. There Be Dragons. In *USENIX Summer*, 1992.

[22] Steven M Bellovin. Packets Found on an Internet. *ACM SIGCOMM Computer Communication Review*, 23(3):26–31, 1993.

[23] David Moore, Colleen Shannon, Douglas J Brown, Geoffrey M Voelker, and Stefan Savage. Inferring Internet Denial-of-Service Activity. *ACM Transactions on Computer Systems (TOCS)*, 24(2):115–139, 2006.

[24] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. The Spread of the Sapphire/Slammer Worm. https://www.caida.org/publications/papers/2003/sapphire/sapphire.html, 2003.

[25] Michael Bailey, Evan Cooke, Farnam Jahanian, David Watson, and Jose Nazario. The Blaster Worm: Then and Now. *IEEE Security and Privacy*, 3(4):26–31, July 2005.

[26] Kriangkrai Limthong, Fukuda Kensuke, and Pirawat Watanapongse. Wavelet-based unwanted traffic time series analysis. In *International Conference on Computer and Electrical Engineering (ICCEE)*, pages 445–449. IEEE, 2008.

[27] Alberto Dainotti, Alistair King, kc Claffy, Ferdinando Papale, and Antonio Pescapè. Analysis of a "/0" Stealth Scan from a Botnet. In *Proceedings of the 2012 Internet Measurement Conference*, IMC '12, pages 1–14, 2012.

[28] Claude Fachkha, Elias Bou-Harb, and Mourad Debbabi. Fingerprinting Internet DNS Amplification DDoS Activities. In *the 6th International Conference on New Technologies, Mobility and Security (NTMS)*, pages 1–5. IEEE, 2014.

[29] Christian Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Symposium on Network and Distributed System Security (NDSS)*, 2014.

[30] Shodan. https://www.shodan.io/.

[31] Elias Bou-Harb, Chadi Assi, and Mourad Debbabi. CSC-Detector: A System to Infer Large-Scale Probing Campaigns. *IEEE Transactions on Dependable and Secure Computing*, 2016.

[32] UCSD Real-time Network Telescope Data. provided by the UCSD - Center for Applied Internet Data Analysis. Available online at https://www.impactcybertrust.org/dataset_view?idDataset=206.

[33] Jun Liu and Kensuke Fukuda. Towards a Taxonomy of Darknet Traffic. In *International Wireless Communications and Mobile Computing Conference (IWCMC)*, pages 37–43. IEEE, 2014.

[34] Tim Yeh. Netis Routers Leave Wide Open Backdoor. http://blog.trendmicro.com/trendlabs-security-intelligence/netis-routers-leave-wide-open-backdoor/, August 2014.

[35] Cymon Open Threat Intelligence. https://cymon.io/.

[36] John Blackford and Mike Digdon. TR-069: CPE WAN Management Protocol. https://www.broadband-forum.org/technical/download/TR-069_Amendment-5.pdf, November 2013.