A Survey of Botnets and the Threats they Pose to the Cyber Landscape

Alex J. Pawelczyk
Cheriton School of Computer Science
University of Waterloo
Waterloo, ON
alex.pawelczyk@uwaterloo.ca

***Abstract*— The cyber landscape is a dynamically evolving environment that provides users with a powerful platform for rapid communication and information sharing. However, various forms of network attacks, such as botnets, can disrupt this service and pose serious threats to cybersecurity. A botnet is a network of compromised devices that can perform a variety of large-scale illegal activities, including Distributed Denial of Service attacks, click fraud, data theft, spam, and social media manipulation. The network architectures and technologies that compose different botnets vary from case to case, and there are many motivations for using a botnet in a cyberattack. Similar to legitimate corporations, a successful botnet operation follows a highly-structured business model to help maximize revenue. As technology evolves, botnets will continue to show increasing levels of anonymity and sophistication, making them difficult to detect and disable. This paper provides an in-depth analysis of different botnet attacks, architectures, detection mechanisms, trends, and the business model behind a botnet. The aim of this paper is to provide readers with a view on botnets from both economic and technical perspectives, present relevant examples, and raise awareness about the threats posed by botnets.**

Keywords— botnet, cybersecurity, cyberattack

I. Introduction

The expected growth of internet-capable devices through future technology and the internet of things (IoT) gives rise to an increase in threats from botnets. The term *botnet* refers to a network of end-hosts, called *bots*, which receive and execute a set of commands from a human operator, known as a *botmaster* [1], [2]. Each individual bot is a program that runs on some host device, and by using a *command and control* (C&C) server, a botmaster can remotely direct a collection of bots to automatically execute a set of predefined functions [3]. Botnets are powerful tools for conducting cyberattacks because bots can be in various locations across the planet, they are not physically owned by a botmaster, and the distributed nature of botnets allows them to operate with a high degree of anonymity. Many botnet C&Cs are based on the Internet Relay Chat (IRC) protocol, which provides a centralized command and control mechanism in which a botmaster can directly communicate with bots in real-time. Other C&Cs use the HyperText Transfer Protocol (HTTP), resulting in a centralized channel where bots periodically contact the C&C server(s) to indirectly obtain their commands from a botmaster [4].

While the individual devices that comprise a botnet may be of low computational power, the combination of millions of distributed devices is very powerful when all the devices are connected and performing tasks in unison. A bot can take the form of any device that is infected by malware, and most of the time, the owner of an infected device does not know that their device has been compromised [5]. Thus, certain botnets are composed of millions of devices, and the computational power from this kind of network can have devastating impacts on the target of an attack, such as crashing a server or financial losses. This paper analyzes the threats, architectures, detection mechanisms, and economic impacts of botnets. Moreover, a detailed case study of the Mirai botnet is presented, and future trends in botnet technologies are identified.

The rest of this paper is structured as follows. Section II identifies some of the previous research that has been conducted on botnets. Section III backgrounds the various forms of network attacks that can be executed by a botnet, and Section IV describes three categories of botnet archtiecture that exist today. Following that, Section V explains some of the detection mechanisms that can be used to identify and shut down a botnet. Next, Section VI provides an analysis of the Mirai botnet, a real world example of a successful IoT botnet. After that, Section VII delves into the business model behind a botnet, and lastly, Section VIII discusses some of the recent trends with the use of botnets.

II. Related Work

To date, there have been many papers written to help readers understand the botnet phenomenon. This section reviews some of the related literature on the topic of botnets.

Lange and Kettani [1] provide a review of botnet evolution, trends, and mitigations. They discuss what a botnet is, the types of attacks that can be performed, the types of botnet architectures that exist today, recent trends and developments in botnets, and different mitigations that can be taken to prevent, detect, or shut down a botnet.

Feily et al. [3] provides an overview of botnets, specifically aimed at their characteristics, life-cycle, and detection. The authors identify four techniques for detecting botnets: signature-based detection, anomaly-based detection, DNS-based detection, and mining-based detection. Their analysis then compares different detection techniques based on their abilities to detect unknown bots, detect botnets with encrypted C&C channels, provide real-time botnet detection, not have any dependencies on botnet protocol or structure, and have a low false-positive rate.

Other papers try to understand how botnets work by studying successful implementations of botnets in the past. Kolias et al. [9] studies the Mirai botnet, one that was used in many successful Distributed Denial of Service attacks. The authors detail the major components that form the Mirai botnet, how it operates and communicates, different botnet variants that are similar to Mirai, and the lessons learned from the Mirai attacks. Etaher et al. [12] focuses on ZeuS, a

malicious botnet that is used to target the financial sector by stealing online banking credentials. They describe the history of ZeuS, how ZeuS functions, the components that make up ZeuS, and the different variants of ZeuS.

A strong motivation behind botnet attacks stems from the economic gain for the botmaster. Putman et al. [37] provides an analysis of the economic structure that is required to support a botnet. Their research focuses on analyzing the business model of a botnet and determining the revenue stream of a botnet owner. Sood et al. [5] provides insight into how much money a botnet attack can cost a business. The authors focus on HTTP-based financial botnets and discuss some of the security solutions that can help mitigate the techniques used by these botnets.

While other papers focus on single characteristics of botnets, this paper studies the broader topic, combining the different but related aspects of botnets into one document. Lange and Kettani [1] provide examples of botnet attacks, architectures, detection mechanisms, and trends; however, this paper goes beyond these topics by also providing a detailed examination of the Mirai botnet and the business model of a botnet, which is omitted from previous work. Conversely, Putman et al. [37] analyzes the economics and business model of a botnet, but they provide little insight into botnet architectures, detection mechanisms, or trends. This paper studies different aspects of botnets to provide readers with a better understanding of botnets as a whole, ranging from structural components to economics.

III. Types of Attacks

A 2019 report from Cybersecurity Ventures [45] predicts that cybercrime will cost the world over $6 trillion annually by 2021, up from $3 trillion in 2015. Botnets are a major contributor to these costs because their high computational power and distributed nature makes them effective tools for conducting various illegal attacks, including Distributed Denial of Service, click fraud, data theft, spam, and social media manipulation.

A. Distributed Denial of Service (DDoS)

A major threat that is posed by botnets is the ability to perform massive DDoS attacks against any person or company that uses the Internet. DDoS is a coordinated attack in which an attacker installs malware programs on multiple machines to gain their control, and then uses these compromised hosts to send attack packets to a victim without their knowledge. The primary goal of a DDoS attack is to disrupt a network so that it cannot provide any services to legitimate users, and the severity of a DDoS attack depends on the intensity of the attack packets, along with the number of hosts used for attacking a victim, where more hosts result in a stronger attack [6].

To perform a DDoS attack, attackers typically follow four basic steps. First, they scan a network to obtain information on potential vulnerable hosts that can later be used to launch an attack. Next, the attackers compromise and install malware into the target hosts, gaining control of the devices. Following that, an attack command is sent to all the compromised devices in the botnet, instructing them to send attack packets with specified intensities to a victim. Finally, maintenance is performed to remove all records or history files from memory [6].

Due to the flexibility and power of botnet technologies, botnets are becoming an integral component of most sophisticated DDoS attacks. Botnets are a useful tool for launching DDoS attacks because forming a network with a large number of compromised hosts allows for the quick generation of a powerful flooding attack [6]. Moreover, the distributed nature and scalability of a botnet makes it difficult to find the identity of the actual attacker. Botnets can also use common protocols, such as IRC or HTTP, to bypass security mechanisms, and it is difficult to detect a botnet in real time because botnet behavior is similar to regular network traffic [6].

Using botnets to perform DDoS attacks makes it very difficult for a victim to defend themselves. For example, one technique to defend a network from a DDoS attack is to block all incoming packets coming from the source IP address of an attacker. The problem with this defense mechanism is that due to the large number of devices with unique IP addresses in the botnet, it is very difficult to manually block the malicious traffic based on source address alone [1]. DDoS attacks also have many different variations that make them complicated to defend against. One of these variations is a reflector attack, where compromised hosts send request packets with a spoofed IP address to a victim. This overwhelms the machine of the victim, and because an IP traceback cannot be performed on a reflector attack, tracking down the attacking machines becomes more difficult [7].

Botnets play a vital role in the execution of a sophisticated DDoS attack, and the effects of one of these attacks can cause significant financial damage to a company. A 2017 Kaspersky Lab study [8] found that the average cost of a DDoS attack on an enterprise was $2.3 million per attack, while the average cost per attack for a small business was $123K. The costs of DDoS attacks have significantly increased since 2016, where the average costs per attack for an enterprise and small business were $1.6 million and $106K, respectively [8]. As botnet technologies continue to advance, it is likely that these costs will continue to rise, and coupled with unquantifiable impacts such as reputational damage, a DDoS attack can have a devastating impact on an organization.

One example of a successful, large-scale DDoS attack occurred in September 2016, when the Mirai botnet launched an attack against the website of computer security consultant Brian Krebs. This attack hit Kreb's website with 620 Gbps of traffic, which is "many orders of magnitude more traffic than is typically needed to knock most sites offline" [9]. At about the same time, another DDoS attack using Mirai malware was launched against the French webhost and cloud service provider OVH, peaking at 1.1 Tbps of traffic [9]. In October 2016, the Mirai botnet was also responsible for a DDoS attack against Dyn, the primary Domain Name System (DNS) provider in the United States. This attack resulted in hundreds of websites being taken down for several hours, including Twitter, Netflix, Reddit, and GitHub [9]. Furthermore, Dyn lost nearly 8% of its customers following the Mirai attack [19].

Botnets have also been used in DDoS attacks against the Internet's DNS. On October 21, 2002, attackers used a botnet to send massive amounts of ICMP ping messages to each of the 13 DNS root IP addresses [46]. Fortunately, this attack caused minimal damage because many of the DNS root

servers were protected by packet filters that automatically blocked all ICMP ping messages that were directed at them. Furthermore, most local DNS servers cache the IP addresses of top-level-domain servers, allowing the query process to often bypass the DNS root servers [46].

B. Click Fraud

The industry of online advertising is expected to grow as more devices become connected to the Internet; thus, the use of botnets to perform click fraud is a major threat to the advertising ecosystem. A common business model used in online advertising is one in which ad brokers pay money to advertisers to display ads, and the amount payed to the advertisers is determined by the number of users who click on the ad. Botnets threaten this model because they can be used to automate this process and generate massive amounts of falsified clicks.

Similar to a DDoS attack, the cost of botnet-driven click fraud can cause significant financial damage to a company. A recent study by CHEQ [10] predicts that ad fraud will directly cost advertisers $23 billion in 2019. Coupled with social and economic impacts, such as a lack of trust towards advertisers and a lower return on investment from advertising dollars spent, the total costs of ad fraud are expected to reach $30 billion [10]. One of the most successful ad fraud campaigns to date was launched in 2016 by the Russian botnet Methbot. By using an army of automated web browsers run from fraudulently acquired IP addresses, Methbot could "watch" as many as 300 million video ads per day on falsified websites designed to look like premium publisher inventory. During the Methbot operation, 6,111 premium domains were targeted and spoofed, enabling the botnet to generate between $3 and $5 million per day in counterfeit inventory [11].

C. Data Theft

Botnets are also capable of stealing personal and financial information from their victims. By using keyloggers and screenshots, a botnet can gain access to social security numbers, credit card numbers, banking information, and other private data from their victims [12]. This data is then exploited for the financial gain of the botmaster. ZeuS is an example of a botnet that penetrates large numbers of computers to steal data by logging keystrokes and spreading copies of itself to other computers via instant messaging and email. ZeuS ended up growing to include over 3.5 million devices, and its success lead to the rise of other variants of botnets with similar intentions [12].

Yahos is a botnet that was designed to exploit the Facebook platform and steal victims' Facebook credentials, credit card information, and in some cases, spam their Facebook friends. Consisting of 11 million infected devices, Yahos was able to cause $850 million in wordwide losses by conducting fraudulent transactions globally [5].

D. Spam & Ransomware

Another common attack involves sending spam mail to target inboxes. In 2018, spam accounted for 52.48% of global email traffic, and is considered to be a significant threat to cybersecurity [13]. Spam has the potential to be dangerous because it can carry malware, scam victims out of money, and inhibit the performance of computers. Furthermore, it is very difficult to stop a botnet from carrying out a spam campaign. Traditional defense mechanisms use content filters and DNS blacklisting techniques to block traffic coming from IP addresses associated with spam [1]. However, botnets circumvent this technique by sending low volumes of spam from each individual bot, making it difficult to blacklist the individual machines that form a botnet in an accurate and timely manner [14].

The Necurs botnet, one of the most active distributors of malware in 2016, was responsible for a spam campaign on November 24, 2016, where 2.3 million spam emails were sent to various users of the Internet in one day [15]. This number suggests that a botnets can play a vital role in having a spam campaign reach as many users as possible, increasing the probability of compromising a device.

A recent trend in botnet-driven spam campaigns involves a combination of social engineering tactics and ransomware. The U.S. Department of Justice (DOJ) has described ransomware as a new business model for cybercrime and a global phenomenon, becoming the fastest growing form of cybercrime that exists today [45]. Ransomware is an extortion malware that encrypts the files of a victim and holds the files hostage until a ransom is paid in exchange for a decryption key [43]. Different from other forms of malware, the effects of a ransomware attack can only be reversed via the cryptographic keys of a malicious actor. Some progagation methods for ransomware rely on social engineering tactics to convince a user to click on a malicious link or download a malicious file, and the typical targets of an attack are unsophisticated users that are unlikely to follow security best practices, such as routine data backups [43].

A botnet-driven spam campaign that uses social engineering tactics and sends millions of emails containing ransomware poses a serious threat to users of the Internet. Researchers at IBM security tracked email spam trends and discovered that 40% of spam emails in 2016 contained ransomware, an increase of 6,000% since 2015, when less than 1% of spam emails contained ransomware [44]. Botnets are an effective tool for disseminating massive amounts of spam emails, and the costs from a successful ransomware attack are significant. In 2017, global ransomware damage costs were predicted to exceed $5 billion, a 1,500% increase since 2015 [45]. Cybersecurity Ventures [45] estimates that ransomware damages will cost the world $11.5 billion in 2019, and $20 billion in 2021. The increase in costs are linked to an increase in ransomware attacks, where further estimates show that a business will fall victim to a ransomware attack every 14 seconds by 2019, and every 11 seconds by 2021 [45].

E. Social Media Manipulation

A recent development in botnet attacks comes in the form of social media manipulation. Botnets can be used to generate massive amounts of spam traffic on social media platforms, but another threat comes from a botnet being able to secretly manipulate public opinion. For example, a botnet can be used to like and retweet certain content on Twitter that conforms to the agenda of the botmaster. This type of manipulation was evident during the 2016 U.S. presidential election, where foreign actors in Russia used botnets to tweet divisive material that supported either candidate [1]. Although it is uncertain if these botnets succeeded in manipulating voter opinion, there was a clear attempt by a foreign actor to interfere with the election process.

Social media manipulation has the potential to be the most dangerous type of botnet attack because it opens the door for cyberwarfare between political candidates. Social bots are unique because they are designed to look and behave like the normal users of a platform, making them effective weapons for spreading computational propaganda. Similar to using common advertising and campaign techniques, social botnets provide a new platform to political actors for spreading a message to massive amounts of people. A candidate that uses a social botnet to spread an agenda is likely to reach a larger amount of people than one who only uses traditional advertising and campaign tactics. Since the ultimate goal of any candidate is to garner the most votes from a population, the future political landscape could see both foreign and domestic political actors financially investing into botnets, creating an advantage for the candidate with the most powerful botnet.

IV. Botnet Architectures

Regardless of the motivation of a botmaster or the type of attack that is conducted by a botnet, the architecture of a botnet generally falls into three categories: centralized, peer-to-peer, or hybrid [1].

A. Centralized

Fig. 1 depicts a centralized architecture for botnets, where bots receive commands from a dedicated central C&C server, and common protocols, such as HTTP or IRC, are used to conduct communication [1], [17]. A centralized architecture provides efficient communication between bots and the C&C server, is easy to set up, and has high scalability. However, one of the main disadvantages of a centralized botnet is the possibility of a single point of failure. If the location of the C&C server is determined, then it can either be shut down completely, or it can be taken over to strengthen a rival botnet [1]. Shutting down the C&C server will take down the entire botnet; thus, multiple C&C servers can be used to increase the robustness of a botnet [16].

Other disadvantages of using a centralized architecture include the necessity of hard-coding the addresses of C&C servers into the botnet, and a C&C server can be detected by observing the network traffic of a bot [17]. The centralized architecture can be further sub-divided into star topologies, where each bot is directly connected to the C&C server, and a hierarchical topology, which contains multiple proxy servers between the C&C server and the bots to add an extra layer of obfuscation [1].

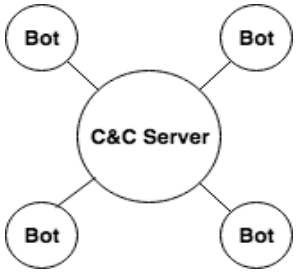

Fig. 1. Centralized botnet with one C&C server and multiple bots [17].

B. Peer-to-Peer

To avoid the vulnerability of a single point of failure, a botmaster may choose to use a peer-to-peer (P2P) architecture instead of a centralized approach. In a P2P botnet, every bot has the potential to be a C&C server, all bots are connected to at least one other bot, bots communicate without passing through a dedicated server, and the commands from a botmaster can only reach the entire botnet if each bot can relay the commands to directly connected bots [17]. However, a P2P botnet presents a tradeoff between latency and overhead.

For example, Fig. 2 shows a fully meshed botnet, where each bot is connected to every other bot in the botnet. This ensures a lower communication latency because less bots are required to broadcast a message to the entire botnet. A fully meshed botnet also has high robustness, since removing a random number of bots from the botnet does not inhibit the communication among other bots. However, the low communication latency and high robustness of a fully meshed botnet comes at the cost of higher overhead and lower scalability, due to an increase in network connections [1], [17]. Conversely, if a bot is only connected to a few of its peers, then there is higher latency due to the amount of time it takes to pass along a command to distant bots, but there is less overhead. Most P2P botnets are not fully meshed because the number of needed connections increases the visibility of the botnet, and the addition or removal of bots due to changes in the network requires a large amount of message coordination [17].

One of the main advantages of using a P2P architecture over a centralized architecture is the lack of a single point of failure. The highly decentralized structure of a P2P architecture also increases the difficulty of detecting and tracking a P2P botnet. P2P architectures are also cost effective, since they do not normally require significant server infrastructure and server bandwidth. A botmaster may also choose to implement a P2P architecture instead of a centralized topology because hijacking a bot in a P2P botnet cannot reveal the identity of a botmaster [16]. However, a P2P botnet is likely to have higher communication latency than a centralized botnet, and implementing a P2P botnet is more difficult to do than one that uses a centralized architecture [16].

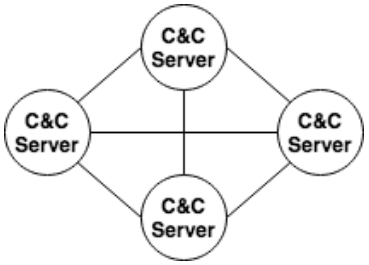

Fig. 2. Fully meshed P2P botnet where there is no dedicated C&C server, and every bot can send commands [17].

C. Hybrid

Both centralized and P2P architectures exhibit weaknesses that can be detrimental to the success of a botnet, and these weaknesses can be remedied by using a hybrid approach. Fig. 3 presents a hybrid architecture that contains a dedicated C&C server, a proxy layer of bots connected in a P2P architecture, and a third layer of bots that execute the commands from a botmaster. The third layer of worker bots helps lower the visibility of the P2P

proxy layer, though at the cost of higher communication latency, additional layers can also be added to increase the protection of the C&C server against detection [17]. A hybrid botnet can also be implemented by using a centralized architecture for one part of botnet communication and a P2P architecture for another part. For example, a P2P section of a botnet can be used to bypass the DNS, while a centralized section is used for all other communication [17].

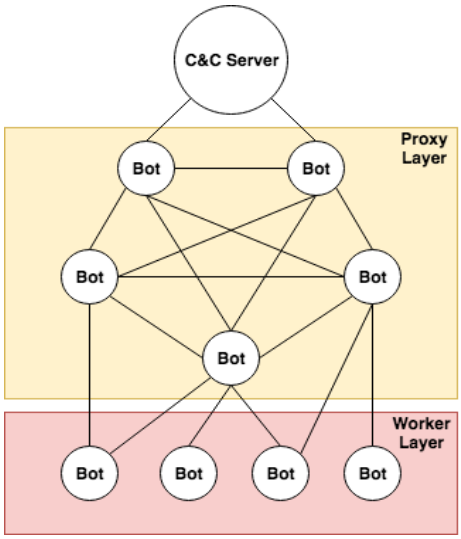

Fig. 3. Hybrid botnet that combines elements from centralized and P2P architectures [17].

V. Botnet Detection Techniques

Over the past two decades, botnet detection and tracking has been a major area of research in network security. Many approaches have been proposed for botnet detection, but the three main approaches involve honeypots, passive network traffic monitoring, and machine learning [3].

A. Honeypot-Based Detection

Spitzner [23] describes a honeypot as a unique security resource that is purposefully set up to be attacked, gaining value the more often a threat, such as a botnet, uses it. A honeypot has little value if no interactions with an attacker occur. However, if a botnet infects and frequently interacts with a honeypot, then researchers can quickly gather data on botnet characteristics. This information can then be used to create signatures that identify botnet traffic, or it can help researchers gain a better understanding of how botnets operate [1]. Honeypots can also be used to slow down or stop an automated attack, and they help with capturing new exploits to gather intelligence on emerging threats. By definition, a honeypot should not see any activity, so one of the advantages of using a honeypot is that any interaction can be assumed to be unauthorized or malicious. Other advantages of using a honeypot include the collection of small data sets, reduced false positives, the capabilities of detecting unknown attacks and capturing encrypted activity, flexibility, and a requirement of minimal resources [23].

B. Passive Network Traffic Monitoring

Another approach for botnet detection is based on passive network traffic monitoring and analysis. Monitoring network traffic has proven to be an effective technique for identifying the existence of botnets, and this detection mechanism can be further sub-divided into signature-based detection, anomaly-based detection, and DNS-based detection [3].

1) Signature-Based Detection: Having prior knowledge about the signatures of existing botnets helps simplify the botnet detection process [3]. For example, by storing signatures of malicious botnets in a database, it is possible to analyze shared malcious packet characteristics, such as addresses or content, and cross-reference these with the known botnet traffic signatures in the database [1]. However, the dependancy on known signatures results in an inability to detect unknown botnets. An example of an intrusion detection system (IDS) that uses signature-based detection is Snort [3]. Snort monitors network traffic to find signs of intrusion, and like most IDS systems, it is configured with a set of rules or signatures to log suspicious traffic.

2) Anomaly-Based Detection: Botnets can also be detected by monitoring network traffic for several anomalies, including high network latency, high volumes of traffic, traffic on unusual ports, and unusual system behavior [3]. Contrary to signature-based detection, anomaly-based detecton techniques are capable of identifying unknown botnets. However, there is a possiblity of generating a false negative. For example, if a botnet exists, but that botnet has yet to be used for an attack, then no anomolies exist, and an IDS that uses anomaly-based techniques will fail to correctly detect the botnet [3].

3) DNS-based Detection: DNS-based detection techniques use the DNS information that is generated by bots to help identify a botnet [3]. To be able to access a C&C server and receive commands from a botmaster, bots make DNS queries to locate the respective C&C server, making it possible to detect botnet DNS traffic and DNS traffic anomalies. DNS-based detection can also reveal the location of a C&C server, helps with identifying the botmaster who controls a botnet, and correctly identifying one bot can lead to the detection of an entire botnet [24].

C. Detection Using Machine Learning

A recent trend in developing botnet detection mechanisms involves a combination of classic botnet detection techniques and ones that use machine learning. Both supervised and unsupervised algorithms can be used to make intelligent decisions from collected packet characteristics, and the use of detection techniques based on machine learning has the potential to outperform traditional detection methods [1]. An advantage of using machine learning is the ability to acquire attribute features from a large amount of data, helping solve classification, clustering, and dimension reduction problems [25].

In terms of botnet detection, machine learning is adventageous because these methods are capable of handling new data and making inferences from inconspicuous patterns that might be missed by a human [1]. However, machine learning requires a large amount of data to make accurate predictions, and it is difficult to obtain accurate and extensive datasets regarding botnets. To detect botnets using machine learning, it is essential to have a large amount of network layer data, including network packet and network flow data [25]. For detection techniques that are based on botnet behaviors, abnormal log data is also needed for the

models to be effective. Without this information, a detection method that is based on machine learning will fail to be effective.

VI. The Mirai Botnet

A prominent example of a recent, succesful botnet is the Mirai botnet. First identified in August 2016, Mirai, along with similar variants, have been the driving force behind some of the most successful DDoS attacks in history [9]. Mirai works by targeting IoT devices, such as webcams, DVRs, and routers, that run some version of BusyBox, a software suite that provides several Unix utilities in a single executable file [9]. Once these devices are infected, Mirai deduces the administrative credentials of other IoT devices through a brute-force attack that uses a small dictionary of potential username-password pairs [9]. The success of the Mirai botnet implies that companies will continue to suffer immense consequences if better security protocols are not implemented in IoT devices.

A. Structual Components

A Mirai botnet consists of bots, a C&C server, a loader, and a report server [9]. Similar to a typical botnet, Mirai bots are devices that are infected by malware, and each bot is responsible for spreading Mirai malware to other devices and executing commands from a botmaster on a target of attack. The C&C server provides a botmaster with a centralized management interface where the conditions of the botnet can be checked, and new DDoS attacks can be orchestrated. Unlike most botnets, communication among Mirai bots and the C&C server is conducted through the Tor network, a free and open-source software that enables anonymous communication [9]. The loader is responsible for distributing executables to various platforms, including ARM, MTPS, and x86, by directly communicating with new victims of an attack. Finally, the report server maintains a database which contains details about the vulnerable devices that are in the botnet, including their stolen login credentials [9], [18]. The report server receives these details from a bot whenever a vulnerable device is identified.

B. Phases of a Mirai Attack

Fig. 4 summarizes the phases of a Miri attack. To find potentially vulnerable devices, Mirai scans random public IP addresses through certain TCP ports (i.e. 22, 23, 32, 2222, and 2323), and then it attempts a brute force attack by using a dictionary of factory default usernames and passwords of IoT vendors [9], [18]. If the correct credentials of a device are identified by a bot, then it will forward various device characteristics to the report server through a different port. Next, via the C&C server, the botmaster will frequently communicate with the report server to obtain information about newly compromised devices, along with the current status of the botnet.

The botmaster then chooses devices to infect, and an infect command will be sent to the loader. After that, the loader will log into a target device and instruct it to download and execute the corresponding binary version of the malware, while also shutting down different points of intrusion, such as Telnet and Secure Shell (SSH) services [9]. At this point, the newly compromised device can communicate with the C&C server to receive attack commands. Once the botmaster decides that there is a sufficient number of bots to conduct an attack, a command will be sent through the C&C server that instructs the entire botnet to begin an attack on a target server. This command contains certain parameters related to the attack, including the type of attack (e.g. Generic Routing Encapsulation (GRE), TCP, HTTP, IP, UDP, and DNS flooding), the duration of an attack, and the IP addresses of the bots and the target server [9], [18]. Upon receiving this command, the bots begin to attack the target server.

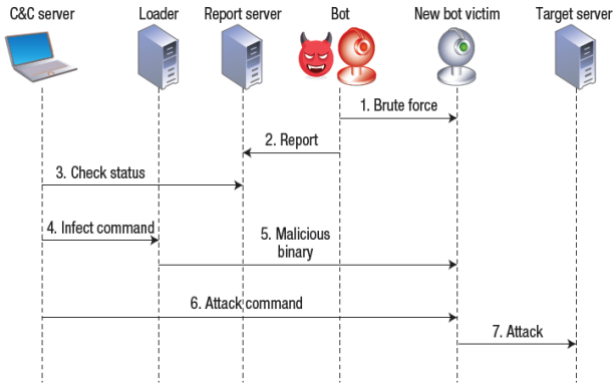

Fig. 4. Phases of a Mirai attack [9].

C. Why Mirai was Successful

The main reason for the success of Mirai comes from the premise of targeting IoT devices. Along with a shortage of legislation, many profit-driven IoT vendors neglect security considerations in favor of user-friendliness and usability, leading to the design of potentially vulnerable IoT devices that can be exploited with little or no effort [19]. In the case of Mirai, the naive approach of conducting a brute force attack by using a dictionary of factory default usernames and passwords was enough to create a massive army of compromised IoT devices. Along with a lack of security considerations from IoT vendors, many IoT devices are constantly running, suffer from poor maintenance, and are powerful enough to produce DDoS attack traffic which is comparable to that of modern desktop systems [9]. Furthermore, infections are more likely to go unnoticed due to minimal user interaction with popular IoT devices. Mirai, along with many of its variants, targeted these vulnerabilities when executing successful attacks.

D. What Could Have Been Done to Stop Mirai

One of the best defense mechanisms against Mirai and its variants is to implement robust security standards for IoT devices and distributors. Mirai took advantage of the fact that many IoT devices use default user credentials, and to this date, this issue remains largely unsolved. Although IoT manufacturers are aware of this flaw, they continue to make consumers take responsibility for changing the default credentials and updating the firmware of their devices. In a world where roughly 48% of consumers are unaware that their connected devices could be used to conduct a cyber attack and close to 40% of consumers never perform firmware updates, placing these responsibilities on the consumers will leave the door open for future botnet-driven attacks [19]. Therefore, IoT vendors should take the responsibility of better securing their products and providing

automatic security updates to all of their devices. Solutions that require manual intervention, such as frequently changing passwords, are unrealistic in the IoT realm because IoT devices must be self-regulating, and trusting the average consumer to practice good security behaviors is a poor solution to the problem [9].

Assuming that IoT vendors fail to improve their security practices, another approach that could be used to prevent a Mirai attack is application whitelisting. An application whitelist is a list of applications and application components (i.e. libraries, configuration files, etc.) that are authorized to be present or active on a host, based on some well-defined baseline [20]. Whitelisting programs then use the whitelist to control which applications are permitted to be installed or executed on a host device. Gopal et al. [18] examined Mirai malware source code and suggest that application whitelisting can be effective in combating IoT malware. Some of the challenges for addressing IoT security requirements include limited storage, power, and computational capabilities of IoT devices [19]. However, due the limited number of applications, low overhead, and low maintenance needed for application whitelisting, it can be used as an effective first line of defense for all IoT devices [18].

Along with application whitelisting, there are a few other techniques that could have been used to stop Mirai. One of the capabilities of Mirai is being able to detect and remove other instances of the malware on an infected device by checking a predefined port and killing the process that is holding it [21]. The same mechanism could be used to automatically detect and eliminate Mirai malware as soon as it starts running on a device. Because the communication between bots and the C&C or reporting server is not encrypted, another approach for shutting down a Mirai botnet involves creating IDS signatures for all parts of the Mirai operation, and using these signatures to detect and shut down the botnet [21].

Meidan et al. [22] proposed N-BaIoT, a network-based approach for IoT that uses deep learning techniques to perform anomoly detection. Within their evaluation, N-BaIot was successful in detecting every attack that was launced by each IoT device in their test set (true positive rate of 100%), had a false positive rate of 0.007 ± 0.01, and the average time to detect an attack was 174 ± 212 ms [22].

VII. Business Model of a Botnet

One of the primary motivations for using a botnet is to generate revenue, and similar to most corporations, a highly structured business model is required to accomplish this task.

A. Life Cycle of a Botnet

To gain a better understanding of the costs that are involved in setting up and maintaining a botnet, it is essential to understand the phases of the botnet life cycle. Rodriguez-Gomez et al. [36] proposes the six phases of the botnet life cycle: conception, recruitment, interaction, marketing, attack execution, and attack success.

The conception phase focuses on the motivation for setting up a botnet. Rodriguez-Gomez et al. [36] argues that the five incentives for a botmaster to setup and maintain a botnet are money, entertainment, ego, cause, and social status, with the primary motivation being financial gain. Regardless of the motivation of a botmaster, the next step in the conception phase is to design the botnet, with specific focus on what architecture will be used (i.e. centralized, P2P, or hybrid). After the botmaster finishes designing the botnet, the final step of the conception phase is to implement the bot code by following a traditional software development process [36].

The second phase of the botnet life cycle involves the recruitment of bots that can be used in a botnet. The botmaster either infects vulnerable devices with botnet malware, or a third party is payed to carry out the infections for the botmaster [37]. There are various ways in which a user can mistakenly compromise their device, including the execution of an attached file in a fake email, downloading a malicious resource from a P2P network, or clicking on a malicious link on a social media platform. Because the power of a botnet is highly dependant on its size, the ultimate goal of a botmaster during the recrutiment phase is to infect as many vulnerable devices as possible.

Following the recrutiment phase, the interaction phase consists of all of the internal and external interactions performed during the botnet operation. Internal interactions are those that are carried out between members of the botnet, including ones from the botmaster to the bots, from the bots to the botmaster, or between bots themselves. External interactions involve the communication that takes place between a member of the botnet and a non-compromised host, typically for the purpose of accessing common services offered in the Internet [36].

The fourth phase of the botnet life cycle revolves around the marketing of a botnet to help the botmaster generate revenue. Revenue can be generated by selling the source code of the botnet malware, renting out a botnet, or renting out the services that a botnet can provide [37].

Once the marketing phase is complete, the botnet is ready to execute an attack. Two powerful features of botnets are their massive size and computational power, which help a botmaster carry out attacks such as DDoS, click fraud, data theft, and spam. Depending on the motivation of a botmaster, the success of a botnet can be determined by the amount of revenue generated, or it can be based on the amount of time that the services of a company are disrupted.

B. Botnet Supply Chain

The development of botnets is a complex process which involves the production and availability of highly structured software architectures, with the goal of having efficient dissemination and monetization [38]. As time goes on, botnets continue to evolve towards a more structured approach. Many key players with different sets of skills, independent revenues, and capabilities to act simultaneously on both the licit and illicit markets are crucial components to the success of a botnet. Bottazzi et al. [38] defines the six categories of the botnet supply chain, ranging from development to utilization.

1) Research and Development: Many organizations are involved in the continuous search for new exploits and vulnerabilities in software that can help create more sophisticated botnets. These organizations also play a role in the development of new malware, and they can sell or rent their expertise on computer systems and software.

2) Money Transfers: Generating revenue on the underground market requires a high level of anonimity, and certain organizations offer secure and anonymous payment services, such as Automated Clearing House (ACH) transactions and wire transfers. For example, money mules are third parties that accept fraudulent ACH transfers, and they play a vital role in the monetization of botnets [39].

3) C&C Bulletproof Hosting: Organizations involved in Bulletproof Hosting offer web-based storage for the stolen information of botnet owners, such as banking or login credentials [37]. These organizations can also host the C&C server that allows a botmaster to successfully carry out an attack.

4) Pay-Per-Install Distribution Model: Even the most sophisticated forms of malware are deemed useless unless they can be spread to numerous devices across the world. Thus, certain organizations offer a pay-per-install (PPI) distribution service, where a botmaster can pay to have their malware spread, providing a commision to the organization per infected device [37]. The PPI distribution model is one of the most used methods for distributing malware, and it is estimated that twelve of the twenty most prevalent malware families are spread using this model [37].

5) Botnet Owners: The owners of a botnet are responsible for the illegal activities of carrying out attacks on their victims. They can also sell the botnet as a whole or rent certain services that a botnet provides to a third party.

6) Pay-Per-Use Attackers: If a person does not own a botnet, then they can rent the services of a botnet from its owner on a pay-per-use basis. Pay-per-use attackers require a minimal amount of technical experience to launch a botnet attack.

C. Botnet Economics

Modern day botmasters and attackers are primarily motivated by profits, rather than creating havoc. One of the essential economic principles states that rational people think at the margin. This suggests that when making economic decisions, people compare costs and benefits, and will only conduct tasks when the benefits of doing them outweigh the costs [39]. To generate the maximum amount of revenue, botmasters must consider the optimal size of botnets, the effective size of bot rental, ways to save money during botnet development and deployment, reducing botnet time-to-market, increasing botnet flexibility, and investing in a highly specialized staff [38], [39]. Moreover, to stay ahead of the continuously evolving defense mechanisms made by IT security stakeholders, a successful business plan of a botnet should not exceed two years.

Section II of this paper presented various attacks that can be executed by a botnet, and the amount of revenue generated by each attack varies. For instance, attacks such as spam and DDoS are considered to be among the least profitable attacks because of their high maintenance costs and "noisiness"; thus, they are typically conducted in the final moments of the life of a botnet [37], [38]. VDoS is an example of an organization that provided a DDoS-for-hire service, and between July 2014 to July 2016, they were generating a median revenue stream of $25,985 per month [40]. Click fraud is a more profitable attack compared to DDoS and spam, and Methbot is an example of a successful botnet that was able to generate between $3 and $5 million dollars per day in counterfeit inventory [11].

Data theft is considered to be the most profitable form of a botnet attack. Eurograbber is an example of a sophisticated botnet that was able to steal an estimated 36 million Euros from more than 30,000 bank customers from multiple banks across Europe [41]. Both corporate and private banking users were targeted, and Eurograbber was capable of illicitly transfering funds out of their accounts in amounts ranging from 500 to 250,000 Euros each [41]. Botnets are clearly a hugely profitable undertaking for those who are successful, and the temptation of these profits shows why botmasters are willing to risk prosecution in order to run their operation.

VIII. Recent Botnet Trends And Developments

As more research is conducted to detect and shut down botnets, botnet developers will focus on using new technologies to stay ahead of IT security stakeholders. The emergence of IoT devices, smartphones, and social media provides new opportunities for botnet attacks.

A. Botnets and the IoT

Recent trends in the technological market show that IoT-based services are experiencing exponential economic growth, expecting to contribute about 1.1-2.5 trillion USD towards the global economy by 2020 [26]. Experts also predict that there will be over 30 billion IoT connected devices in 2020, compared to 9.9 million in 2013 [26]. However, the increasing number of insecure IoT devices with high computational power and diverse locations attracts malicious actors who seek to create large-scale botnets [27], [28]. Most IoT devices are manufactured with an emphasis on quick deployment and convenience, with very little thought going towards the security challenges and requisite threats posed by those devices. This lack of security consideration makes IoT devices excellent prospects for the formation of a botnet, having been used in some of the most prominent botnet attacks in history.

As discussed in Section VI, Mirai was an IoT botnet that was responsible for carrying out some of the most successful DDoS attacks in history. However, researchers at Avast, a Czech cybersecurity software company, recently discovered Torii, a newer and more advanced IoT botnet that has been running since December 2017 [42]. Torii tries to be stealthy and persistent upon infecting an IoT device, and it has yet to be used in any attacks.

Torii also provides a diverse set of features for stealing sensitive information, a modular architecture that allows for the fetching and execution of other commands and executables, and powerful encryption techniques to help prevent detection [42]. Furthermore, Torii can infect a wide range of IoT devices, and it provides support for various kinds of target architectures, such as MIPS, ARM, x86, x64, PowerPC, and SuperH [42]. Torii is a clear example of the evolution of IoT malware, and it shows how criminal actors take advantage of the weak security protocols in the IoT market.

B. Mobile Botnets

Smartphones are a growing commodity for the global market, and their mass and rapid adoption opens the door for malicious actors to create botnets containing them. Because

of the high usage, convenience, and mobility of smartphones, mobile security has become a globally critical issue [28]. Mobile botnets are different from traditional botnets because they do not need to propogate using centralized infrastructures, they can use P2P wireless links to compromise nodes within close proximity, and they can organically evolve via data forwarding [29]. One way to form a mobile is by botnet propagation through infrastructures, where malware sends copies of itself via short or multimedia message services, or malware advertises its applications on mobile markets. Another way for a mobile botnet to form is by proximity infection, where a compromised node sends malware to nearby nodes via peer-to-peer wireless links [29].

Botnet propagation through infrastructures is faster than forming a botnet through proximity infection, but increasingly enhanced security systems, such as Google's Android kill switch, can stop propagation [29]. Hence, an alternative for forming a mobile botnet is to use proximity infection because due to the nature of decentralized infection and a dynamic network topology, infecting nearby nodes can easily persist and remain undetected. After a mobile botnet is formed, it can either have individual or global impacts. For example, access can be blocked to an individual device, or the globally distributed devices can be used to perform a DDoS attack.

C. Social Media Bots

Social media platforms are powerful tools that connect millions of people across the globe, and they play a major role in how information, ideas, news, and opinions are spread across society. Society has a heavy reliance on social media, and social bots can take advantage of this by targetting certain users to promote a specific rhetoric. A social bot refers to a social media account, controled by some version of automated software, that uses algorithms to generate content and establish interactions in a human-like behavior [30]. Contrary to normal users who gain access to a platform via front-end websites, social bots can directly access a website through a mainline, code-to-code connection by using the application programming interface (API) of a particular platform [31]. As time goes on, the presence of bots on social media sites is likely to rise. It is estimated that between 9% and 15% of all Twitter accounts are social bots, which is equivalent to approximately 48 million accounts [35]. Although a social bot can perform many useful functions, such as summarizing news articles or generating useful statistics, a growing number of these bots are being used for malicious purposes.

One motivation for creating a social botnet could be to artificially inflate support for a political candidate, having the potential to undermine the democracy of a nation by influencing the outcome of an election [32]. The 2016 US Presidential Election saw the use of social bots to disseminate politically motivated rumors, share misinformaton, and provide US voters with direct links to political news and information from Russian sources [31]. Furthermore, both presidential candidates used social botnets to spread their political agendas on Twitter, where the largest pro-Trump botnet consisted of 944 bots, compared with 264 bots in the largest pro-Clinton botnet [47]. The pro-Trump botnet was also more centralized and interconnected than the pro-Clinton botnet, suggesting a higher degree of strategic organization.

A strategy used by both botnets was to tweet about their respective candidate's victory during the presidential debates, but on average, the pro-Trump botnet released seven tweets for every one tweet from the pro-Clinton botnet [47]. Howard et al. [33] also points out that in the week leading up to the 2016 US election, the average levels of polarizing political news and misinformation on Twitter were higher in swing states than in uncontested states. This suggests that political actors were specifically trying to influence the opinion of voters in important swing states, where large numbers of votes in the Electoral College were undecided.

Past social science studies have shown that social bots can have a political impact by attacking journalists and discrediting political leaders, but it is still unclear if they succeed in changing voter opinion [31]. Nonetheless, combining the dissemination of misinformation with social bots that promote content in a preprogrammed way gives political actors a powerful set of tools for spreading computational propaganda [33].

Along with their use in the political landscape, social bots can also be used to promote terrorist propaganda and recruitment, manipulate the stock market, and scam users of a particular platform [30], [34]. These malicious activities have a profound impact on society, and social media companies need to take responsibility for securing their platforms and shutting down social botnets. After the 2016 US presidential election, Facebook disabled over 1 billion fake accounts, and its safety and security team doubled in size to more than 20,000 people to handle content in 50 languages [47]. Many social media companies are also investing in machine learning and artificial intelligence that can automatically detect and remove fake news and other illicit activity. As malicious actors come up with new tricks and tactics for using social botnets, it is crucial for social media companies to focus their attention on security and qucikly mitigate these threats.

IX. Conclusions

Botnets pose a serious threat to the individuals, businesses, and governments that rely on the Internet for their day-to-day operations. This paper details different methods of attack performed by botnets, including the theft of private and financial information, committing ad fraud, disseminating spam, and disrupting the services provided by legitimate companies. The primary motivation for carrying out a botnet-driven attack depends on the botmaster, often stemming from economic gain. Similar to legitimate corporations, botnet owners follow a highly-structured business model to help maximize revenue. Although many botnet attacks are illegal, the massive profits that can be generated from a botnet attack are a key factor in why botnet owners risk prosecution to run their operations.

Recent technological advancements in the IoT, mobile phones, and social media present new possibilites for botmasters to exploit vulnerabilities in these platforms and develop new botnets to conduct attacks. In particular, social media manipulation through the use of social botnets poses a new threat in the political landscape, in which political actors attempt to manipulate public opinion through social media and undermine the democracy of a nation. Future elections could see political candidates investing into the latest and most powerful botnet technologies to help garner support for their candidacy. Having the most powerful and far-reaching

botnet will imporve the chances of a candidate winning an election, and cyberwarfare will play a major role in future elections.

Botnet detection and mitigation is the most important area for future research on the topic of botnets. One of the most prominent detection mechanisms involves a combination of classical detection techniques with machine learning, where both supervised and unsupervised learning algorithms can be used to make intelligent decisions based on network flow characteristics. Botnet detection techniques that use machine learning have the potential to outperform traditional methods, but a major disadvantage of this technique is the reliance on large quantities of data, which is difficult to obtain.

Providing cybersecurity education to individuals and corporations that use the Internet is another crucial aspect to the defense against botnets. People are the weakest link in the security chain, and even the latest technologies and layers of security protection can be broken by human error. Coupled with the fact that employee training in recognizing and defending against cyber attacks is the most under spent sector of the cybersecurity industry, the recent success of botnets and other forms of attacks is no surprise [45]. In 2018, organizations from all over the world invested into either training employees on security for the first time, or improving older programs by implementing more robust programs in security awareness and phishing simulations [45]. Estimates show that global spending on security awareness training is predicted to reach $10 billion by 2027, a significant increase from the $1 billion in 2014 [45]. Future research can be conducted to determine if the increase in spending on security awareness results in a decrease of successful cyberattacks.

Along with investing into their employees, social media companies and IoT vendors need to implement protocols that better secure their platforms and devices, rather than relying on consumers to practice good security. Government legislation that mandates a focus on security could also help mitigate the threats that are posed by botnets. For example, having a law that mandates all IoT devices to be protected with a unique username and password combination is a trivial solution that could have stopped the Mirai attacks. Ultimately, the only way to defeat botnets requires cooperation from individuals, corporations, and legislators to make cybersecurity the primary focus of all devices and services that use the Internet.

ACKNOWLEDGMENT

I would like to express my very great appreciation to Dr. Lesley Istead for her valuable and constructive suggestions during the planning and development of this research work. I would also like to thank David Radke and Michael Morgan for their help during the revision process of this paper.

REFERENCES

[1] T. Lange and H. Kettani, "On Security Threats of Botnets to Cyber Systems," *2019 6th International Conference on Signal Processing and Integrated Networks (SPIN)*, Noida, India, 2019, pp. 176-183.

[2] Moheeb Abu Rajab, Jay Zarfoss, Fabian Monrose, and Andreas Terzis, "A multifaceted approach to understanding the botnet phenomenon," *2006 6th ACM SIGCOMM conference on Internet measurement* (IMC '06). ACM, New York, NY, USA, pp. 41-52.

[3] M. Feily, A. Shahrestani and S. Ramadass, "A Survey of Botnet and Botnet Detection," *2009 Third International Conference on Emerging Security Information, Systems and Technologies*, Athens, Glyfada, 2009, pp. 268-273.

[4] G. Gu, J. Zhang, and W. Lee, "BotSniffer: Detecting Botnet Command and Control Channels in Network Traffic," 2008 NDSS.

[5] A. K. Sood, S. Zeadally and R. J. Enbody, "An Empirical Study of HTTP-based Financial Botnets," in *IEEE Transactions on Dependable and Secure Computing*, vol. 13, no. 2, pp. 236-251, 1 March-April 2016.

[6] N. Hoque, D. K. Bhattacharyya and J. K. Kalita, "Botnet in DDoS Attacks: Trends and Challenges," in *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, pp. 2242-2270, Fourthquarter 2015.

[7] B. Nagpal, P. Sharma, N. Chauhan and A. Panesar, "DDoS tools: Classification, analysis and comparison," *2015 2nd International Conference on Computing for Sustainable Global Development (INDIACom)*, New Delhi, 2015, pp. 342-346.

[8] "DDoS Breach Costs Rise to over $2M for Enterprises finds Kaspersky Lab Report," *usa.kaspersky.com*. [Online]. Available: https://usa.kaspersky.com/about/press-releases/2018_ddos-breach-costs-rise-to-over-2m-for-enterprises-finds-kaspersky-lab-report. [Accessed: 26-Nov-2019].

[9] C. Kolias, G. Kambourakis, A. Stavrou and J. Voas, "DDoS in the IoT: Mirai and Other Botnets," in *Computer*, vol. 50, no. 7, pp. 80-84, 2017.

[10] "The economic cost of bad actors on the internet: ad fraud," CHEQ, 2019.

[11] W. O. Inc., "The Methbot Operation," *White Ops - Detect and defeat the web's most evasive bots.* [Online]. Available: https://www.whiteops.com/methbot. [Accessed: 26-Nov-2019].

[12] N. Etaher, G. R. S. Weir and M. Alazab, "From ZeuS to Zitmo: Trends in Banking Malware," *2015 IEEE Trustcom/BigDataSE/ISPA*, Helsinki, 2015, pp. 1386-1391.

[13] M. Vergelis and T. Sidorina, "Spam and phishing in 2018," *Securelist English*. [Online]. Available: https://securelist.com/spam-and-phishing-in-2018/89701/. [Accessed: 26-Nov-2019].

[14] A. Pathak, Q. Feng, Y. C. Hu, Z. M. Mao, and S. Ranjan, "Botnet spam campaigns can be long lasting: evidence, implications, and analysis," *2009 11th International Joint Conference on Measurement and Modeling of Computer Systems*, Seattle, WA, USA, 2009, pp. 13-24.

[15] S. Aimoto *et al.*, "Internet Security Threat Report," Symantec Corp., vol. 22, April 2017, pp. 41-42.

[16] N. S. Raghava, D. Sahgal and S. Chandna, "Classification of Botnet Detection Based on Botnet Architechture," *2012 International Conference on Communication Systems and Network Technologies*, Rajkot, 2012, pp. 569-572.

[17] G. Vormayr, T. Zseby and J. Fabini, "Botnet Communication Patterns," in *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2768-2796, Fourthquarter 2017.

[18] T. S. Gopal, M. Meerolla, G. Jyostna, P. Reddy Lakshmi Eswari and E. Magesh, "Mitigating Mirai Malware Spreading in IoT Environment," *2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, Bangalore, 2018, pp. 2226-2230.

[19] N. Neshenko, E. Bou-Harb, J. Crichigno, G. Kaddoum and N. Ghani, "Demystifying IoT Security: An Exhaustive Survey on IoT Vulnerabilities and a First Empirical Look on Internet-Scale IoT Exploitations," in *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2702-2733, thirdquarter 2019.

[20] A. Sedgewick, M. P. Souppaya, and K. A. Scarfone, "Guide to Application Whitelisting," *National Institute of Standards and Technology*, Oct. 2015.

[21] H. Sinanović and S. Mrdovic, "Analysis of Mirai malicious software," 2017 25th International Conference on Software, Telecommunications and Computer Networks (SoftCOM), Split, 2017, pp. 1-5.

[22] Y. Meidan *et al.*, "N-BaIoT—Network-Based Detection of IoT Botnet Attacks Using Deep Autoencoders," in *IEEE Pervasive Computing*, vol. 17, no. 3, pp. 12-22, Jul.-Sep. 2018.

[23] L. Spitzner, "Honeypots: catching the insider threat," *19th Annual Computer Security Applications Conference, 2003. Proceedings.*, Las Vegas, NV, USA, 2003, pp. 170-179.

[24] N. Kaur and M. Singh, "Botnet and botnet detection techniques in cyber realm," *2016 International Conference on Inventive Computation Technologies (ICICT)*, Coimbatore, 2016, pp. 1-7.

[25] X. Dong, J. Hu and Y. Cui, "Overview of Botnet Detection Based on Machine Learning," *2018 3rd International Conference on Mechanical, Control and Computer Engineering (ICMCCE)*, Huhhot, 2018, pp. 476-479.

[26] I. Makhdoom, M. Abolhasan, J. Lipman, R. P. Liu and W. Ni, "Anatomy of Threats to the Internet of Things," in *IEEE Communications Surveys & Tutorials*, vol. 21, no. 2, pp. 1636-1675, Secondquarter 2019.

[27] E. Bertino and N. Islam, "Botnets and Internet of Things Security," in *Computer*, vol. 50, no. 2, pp. 76-79, Feb. 2017.

[28] M. Eslahi, R. Salleh and N. B. Anuar, "MoBots: A new generation of botnets on mobile devices and networks," *2012 International Symposium on Computer Applications and Industrial Electronics (ISCAIE)*, Kota Kinabalu, 2012, pp. 262-266.

[29] Z. Lu, W. Wang and C. Wang, "On the Evolution and Impact of Mobile Botnets in Wireless Networks," in *IEEE Transactions on Mobile Computing*, vol. 15, no. 9, pp. 2304-2316, 1 Sept. 2016.

[30] O. Varol, E. Ferrara, C. A. Davis, F. Menczer, A. Flammini, "Online human-bot interactions: detection, estimation, and characterization," in *Proc. of the Eleventh International AAAI Conference on Web and Social Media (ICWSM)*, Montreal, Canada, May 15-18, 2017, pp. 280-289.

[31] P. N. Howard, S. Woolley, and R. Calo, "Algorithms, bots, and political communication in the US 2016 election: The challenge of automated political communication for election law and administration," *Journal of Information Technology & Politics*, vol. 15, no. 2, pp. 81–93, Mar. 2018.

[32] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, "The rise of social bots," *Communications of the ACM*, vol. 59, no. 7, pp. 96–104, 2016.

[33] P. N. Howard, B. Kollanyi, S. Bradshaw, L. Neudert, "Social media, news and political information during the US Election: was polarizing content concentrated in swing states?," *Computational Propaganda Research Project*, vol. 2017.8, Sep. 2017.

[34] J. Wright and O. Anise, "Don't @ me: hunting twitter bots at scale," Duo Security Inc., version 1.0., 2018.

[35] E. Alothali, N. Zaki, E. A. Mohamed and H. Alashwal, "Detecting Social Bots on Twitter: A Literature Review," *2018 International Conference on Innovations in Information Technology (IIT)*, Al Ain, 2018, pp. 175-180.

[36] R. A. Rodríguez-Gómez, G. Maciá-Fernández and P. García-Teodoro, "Analysis of botnets through life-cycle," *Proceedings of the International Conference on Security and Cryptography*, Seville, 2011, pp. 257-262.

[37] C. G. J. Putman, Abhishta and L. J. M. Nieuwenhuis, "Business Model of a Botnet," *2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*, Cambridge, 2018, pp. 441-445.

[38] Giovanni Bottazzi and Gianluigi Me. "The Botnet Revenue Model," Proceedings of the 7th International Conference on Security of Information and Networks - SIN '14 (2014), pp. 459–465.

[39] Z. Li, Q. Liao, and A. Striegel, "Botnet Economics: Uncertainty Matters," *Managing Information Risk and the Economics of Security*, pp. 245–267, 2008.

[40] R. Brunt, P. Pandey, and D. McCoy, "Booted: an analysis of a payment intervention on a DDoS-for-hire service," *16th Annual Workshop on the Economics of Information Security*, San Diego, June 26-27, 2017.

[41] E. Kalige and D. Burkey, "A case study of Eurograbber: how 36 million euros was stolen via malware," Check Point Software Technologies Ltd., Dec. 2012.

[42] J. Kroustek, V. Iliushin, A. Shirokova, J. Neduchal, and M. Hron, "New Torii Botnet uncovered, more sophisticated than Mirai," *Avast*, 27-Sep-2018. [Online]. Available: https://blog.avast.com/new-torii-botnet-threat-research. [Accessed: 27-Nov-2019].

[43] N. Scaife, H. Carter, P. Traynor, and K. R. B. Butler, "CryptoLock (and Drop It): Stopping Ransomware Attacks on User Data," *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*, 2016.

[44] "Ransomware Damage Report: 2017 Edition," *Herjavec Group*, 08-Aug-2018. [Online]. Available: https://www.herjavecgroup.com/ransomware-damage-report-2017-edition/. [Accessed: 01-Dec-2019].

[45] S. Morgan, "2019 official annual cybercrime report," Cybersecurity Ventures, 2019.

[46] J. F. Kurose and K. W. Ross, *Computer networking: a top-down approach*, 7th ed. Hoboken: Pearson, 2017.

[47] "How political campaigns weaponize social media bots," *IEEE Spectrum: Technology, Engineering, and Science News*. [Online]. Available: https://spectrum.ieee.org/computing/software/how-political-campaigns-weaponize-social-media-bots. [Accessed: 01-Dec-2019].