Computers & Security 128 (2023) 103123

Contents lists available at ScienceDirect

Computers & Security

journal homepage: www.elsevier.com/locate/cose

A Comprehensive Survey of Recent Internet Measurement Techniques for Cyber Security

Morteza Safaei Pour[a,*], Christelle Nader[b], Kurt Friday[b], Elias Bou-Harb[b]

[a] San Diego State University, California, USA
[b] The Cyber Center for Security and Analytics, University of Texas at San Antonio, Texas, USA

ARTICLE INFO

Article history:
Received 18 August 2022
Revised 15 November 2022
Accepted 24 January 2023
Available online 27 January 2023

Keywords:
Internet measurement
Cyber security
Large-scale analysis
Security threats

ABSTRACT

As the Internet has transformed into a critical infrastructure, society has become more vulnerable to its security flaws. Despite substantial efforts to address many of these vulnerabilities by industry, government, and academia, cyber security attacks continue to increase in intensity, diversity, and impact. Thus, it becomes intuitive to investigate the current cyber security threats, assess the extent to which corresponding defenses have been deployed, and evaluate the effectiveness of risk mitigation efforts. Addressing these issues in a sound manner requires large-scale empirical data to be collected and analyzed via numerous Internet measurement techniques. Although such measurements can generate comprehensive and reliable insights, doing so encompasses complex procedures involving the development of novel methodologies to ensure accuracy and completeness. Therefore, a systematic examination of recently developed Internet measurement approaches for cyber security must be conducted to enable thorough studies that employ several vantage points, correlate multiple data sources, and potentially leverage past successful techniques for more recent issues. Unfortunately, performing such an examination is challenging, as the literature is highly scattered. In large part, this is due to each research effort only focusing on a small portion of the many constituent parts of the Internet measurement domain. Moreover, to the best of our knowledge, no studies have offered an in-depth examination of this critical research domain in order to promote future advancements. To bridge these gaps, we explore all pertinent facets of utilizing Internet measurement techniques for cyber security, ranging from threats within specific application domains to threats themselves. We provide a taxonomy of cyber security-related Internet measurement studies across two dimensions. One dimension relates to the many vertical layers (and components) of the Internet ecosystem, while the other relates to internal normal functions vs. the negative impact of external parties in the Internet and physical world. A comprehensive comparison of the gathered studies is also offered in terms of measurement technique, scope, measurement size, vantage size, and the analysis approach that was leveraged. Finally, a discussion of the roadblocks to performing effective Internet measurements and possible future research directions is elaborated.

1. Introduction: Why measure the Internet?

The Internet is a complex, decentralized and dynamic system with many components and diverse features. It is made up of self-contained networks that communicate with one another using the Internet Protocol (IP) and packet switching (Park, 2005). Due to the Internet's scattered structure, assessing any element of it on a global scale is challenging. At the start of 2021, the number of unique autonomous networks in the Internet's routing system topped 99,378 (Maigron, 2020). By 2023, there will be 5.3 billion Internet users worldwide (66% of the world population), which is higher than 3.9 billion (51% of the world population) in 2018. Furthermore, the number of devices linked to IP networks will exceed three times the world's population. The number of networked devices is also expected to reach 29.3 billion, which is a large increase from 18.4 billion devices accounted for in 2018. In addition, 14.7B Machine-to-Machine (M2M) connections will have been established (Cisco, 2020).

Numerous facets of the Internet's operation and utilization are opaque or constantly changing, and may only be understood by measurement. Moreover, a wide variety of services, ranging from online entertainment to mission-critical, are entirely reliant upon

* Corresponding author. Tel.: +1 619 594 2138.
E-mail addresses: msafaeipour@sdsu.edu, morteza.safaeipour@gmail.com (M. Safaei Pour).

https://doi.org/10.1016/j.cose.2023.103123

the Internet as a foundation and communication medium. Thus, for the purpose of dependability, security, and quality of service, it is essential to continuously monitor performance metrics and network settings, as well as to execute various tests, assessments, configurations, and management tasks.

Regarding the Internet as a complex interaction of numerous simple systems and protocols (Park and Williger, 2005), Internet measurement is a set of methods of large-scale and in-action (remote) collection of measurable data from the Internet to quantitatively describe the structure (individual systems and protocols), their interaction, and use of the Internet (the interrelation between the Internet and the physical world). The three fundamental aspects of Internet measurement research, especially for cyber security, are its empirical basis, large scale, and in-the-wild data collection. More specifically, the scale of data collection procedures must be large enough to produce a representative sample size. Additionally, the empirical nature of Internet measurement is tightly coupled with capturing real data in the wild since it is not feasible to stop or disrupt the Internet's normal operation for this purpose.

Applications of internet measurement can be effectively categorized into three primary groups. The first of these categories studies the protocols and services used on the Internet and how they have adapted to the rapid evolution of the Internet. Since the Internet's widespread adoption in the 1990s, some key protocols have undergone modest revisions to address bugs, enhance the quality of various services, enforce policies, or comply with new criteria such as privacy and security. Among such revisions, HTTP incorporated a few additional headers and methods, Transport Layer Security (TLS) underwent gradual enhancements, Transmission Control Protocol (TCP) adapted congestion management, and Domain Name System (DNS) added capabilities such as Domain Name System Security Extensions (DNSSEC). Additionally, the complexity of protocol implementations has also increased due to the advent of network layer independence and the flexibility to modify Internet protocols at any layer. Moreover, new protocols (e.g., Remote Desktop Protocol (RDP), Virtual Private Networks (VPN)) and standards are regularly being deployed in order to add support for various technologies as their demand increases (e.g., blockchain, IoT, 5G).

In the second category, Internet measurement is leveraged to investigate cyberspace's security. While the applications encompassed by this category are vast, many pertain to examining the security associated with new protocol implementations. Such new protocol implementations are often shipped with vulnerabilities, particularly during the early stages of adoption, which attackers often exploit. Moreover, it may take several years for the extent of these vulnerabilities to be fully realized and appropriately patched. To illustrate, there are protocols (e.g., DNS, Network Time Protocol (NTP)) and services (e.g., Memcached) that have been in use for many years that, as a result of their insecure deployment, enable attackers to amplify Denial of Service (DoS) attacks by several orders of magnitude. Furthermore, ubiquitous TLS vulnerabilities can jeopardize users' privacy and lead to data breaches (Holz et al., 2016). In addition to assessing protocol-related security, Internet measurement is also valuable for investigating a plethora of attacks at scale, including Denial of Service (DoS), Distributed DoS (DDoS), botnets, ransomware, and phishing, among several others.

Finally, in the third category, Internet measurement is mapped to real-world events to assess one's impact on the other. This is possible because of the extent to which the Internet now pervades all facets of human civilization and has inseparably been linked to society and cyberspace. Consequently, any notable occurrence in one will inevitably affect the other. As a result, Internet measurement can be an effective tool for investigating the impact of societal, political, and natural events on the Internet ecosystem. Additionally, Internet measurement can also be leveraged to determine the Internet's resilience to unforeseen real-world developments and the resultant modifications that may need to be integrated. Some instances of unforeseen real-world developments are when both Egypt's and Libya's governments shut down the Internet in 2011 Dainotti et al. (2011), as well as when governments have elected to impose restrictions on the Internet (Marczak et al., 2015).

Ultimately, Internet measurement plays a vital role in analyzing the spread of known vulnerabilities, detecting emerging threats, and tracking the evolution of attackers' activities. It can also be leveraged to garner public attention by shedding light on the magnitude of these issues, such as the inadequacies of their current solutions, despite the cost of these solutions continuing to rise. Indeed, cyber security has emerged as a major issue and will remain as such for many years to come, and Internet measurement offers a viable means of enhancing it. To promote this enhancement given the wide spectrum of Internet measurement, this survey will systematically detail the collection, use cases, and research of Internet measurement data.

Given the scarcity of real-world data in cyber security research, leveraging passive and active measurement techniques in conjunction with robust analytic methods can offer insight on a subject's security posture. Additionally, the Internet is expanding at a rapid pace across a variety of dimensions, including protocols, devices, applications, technologies, platforms, users, and threats. As a result, the behavior of a particular system is dependent on a variety of variables that might alter the measurement. Furthermore, a number of facets of the Internet are constantly evolving. Due to the enormous size of the Internet and its dynamic nature, making definitive statements regarding the Internet's eventual behavior without conducting several experiments and measurements is largely an impossibility. For example, little is known about a data stream that a recipient can attribute directly to the suspected source, as part of the data may have been changed during transit. Internet measurement is usually compared to astronomy in that it includes making remote observations in order to gain a better understanding of how a system works. Furthermore, a measurement taken from one vantage point may bear little resemblance to a measurement taken from another vantage point, leading to inconsistent findings (Wan et al., 2020) due to Internet-enforced policies such as political decisions (e.g., censorship Leyba et al. (2019)) or security precautions (e.g., ISPs block Internet Control Message Protocol (ICMP) messages to avoid exposing their infrastructures to external scans).

Meanwhile, the continued attention to the privacy of users and other entities in cyberspace has resulted in a variety of countermeasures to limit information leaks. While such countermeasures are commendable and necessary, they undoubtedly increase the complexities associated with the collection and assessment of empirical cyber security data. Therefore, conducting Internet measurement is often not straightforward and requires devising innovative techniques to confirm its correctness and completeness. In turn, a systematic review of the devised Internet measurement techniques for cyber security aids researchers in performing comprehensive analysis by highlighting a number of different vantage points that can be leveraged and successful techniques that could be applied to new topics. To the best of our knowledge, this survey is the first systematic review that studies empirical large-scale Internet data collection for cyber security purposes. In this survey, we make the following contributions.

1. A taxonomy of cyber security-focused Internet measurement studies is provided for different components of the Internet ecosystem to pave the way for researchers that are concentrating on a particular subject (Sections 4 through 7). These sections detail the encompassed papers based on the security issues that the works address within a given domain.

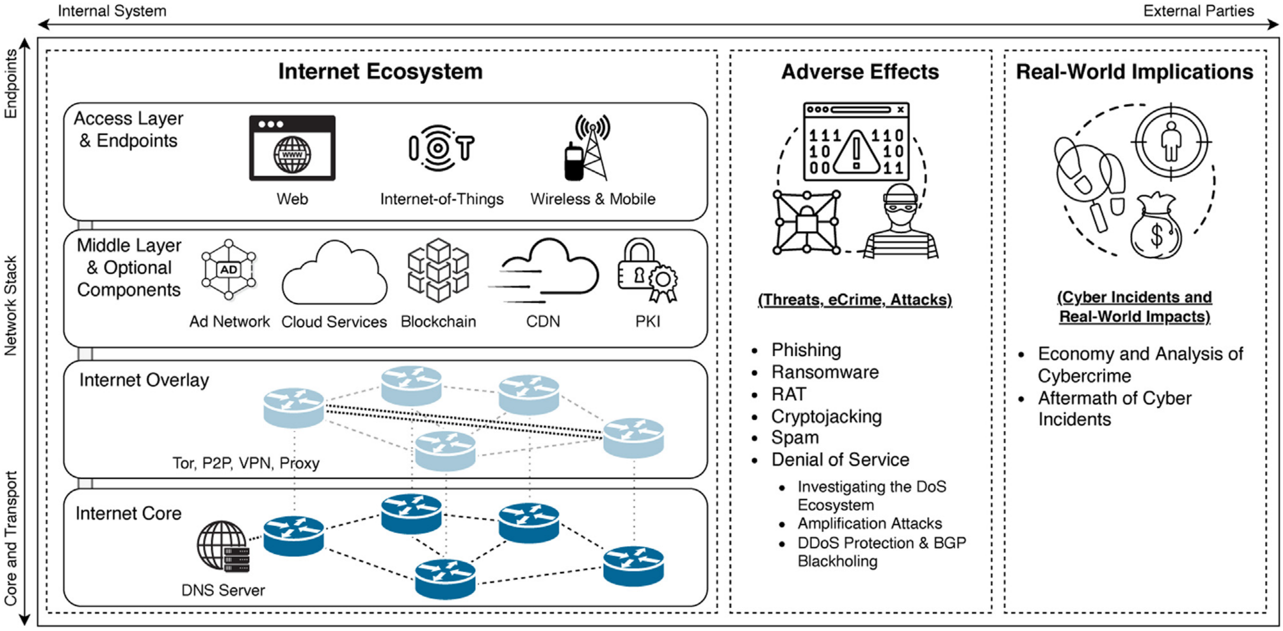

Fig. 1. Holistic view over the Internet measurement for cyber security

2. A review of macroscopic data collection and analysis procedures employed in different cyberspace threat categories is offered (Section 8). This section details publications that investigate a particular cyber attack at scale.
3. A discussion of the impediments to Internet measurement and possible future research directions are provided.

In Section 2, we elaborate on the literature search methodology, covered journals, and conference proceedings. Subsequently, Section 3 reviews all the published surveys on related topics and highlights the differences with this review paper. Figure 1 depicts a comprehensive picture of Internet measurement for cybersecurity, categorizing the studies according to two dimensions. One dimension relates to the many vertical levels of the Internet ecosystem, while the other relates to internal normal functions vs. the negative impact of external parties in the Internet and physical world. We classified the Internet ecosystem and its major components into four main sections (layers) on the basis of their similarity and function within the technology stack. However, components from different layers also depend on one another, and a technology is comprised of multiple components. Consequently, Internet measurement studies frequently explore a particular topic from multiple vantage points. These points of view may be obtained from different components and layers. For instance, encrypted DNS research can utilize techniques from the Internet overlay layer, PKI in the middle layer and optional components layer, and possibly the access layer for a more thorough analysis, despite the fact that DNS component remains the primary focus. Within each layer of the Internet, measuring techniques share similarities. Therefore, the topics related to Internet core architecture and services are elaborated upon in Section 4. Section 5 covers areas of research related to Internet overlay systems and middleboxes. Next, Section 6 reviews topics related to middle layers such as content delivery networks, blockchain, as well as optional components such as advertising networks and web Public Key Infrastructure (PKI). Works pertaining to last mile Internet, access layer and endpoints are explored in Section 7.

In Section 8, we focus on the adverse effects of attackers that are detrimental to the Internet and review the associated papers. Internet measurement papers that study cyber incidents and real-world implications are detailed in Section 9. Finally, Section 10 pinpoints future trends and the challenges associated with conducting Internet measurements.

In this article, we examine the publications with regard to measurement method, analytic method, measurement size, and scope in each of their respective contexts. Internet measuring methods are typically classified as passive, active, or hybrid. In the passive technique, data is gathered through passive observation, which frequently necessitates the deployment of monitoring instrumentation at a network node that can monitor traffic. Typically, this requires cooperation with commercial operators. Active probing, on the other hand, injects traffic or initiates an action in order to collect a response and observe its effect. Active probing is preferred by third-party researchers since it may be performed from the network's edge and requires minimal network access, but it can only infer a limited number of network characteristics. Moreover, passive approaches present privacy concerns for network users and operators. We chose the term 'hybrid' to refer to a multi-step process that combines both approaches to develop more complex measurement methodologies. For instance, passive observation of one data source followed by active probing of another data type based on the outcomes of the passive observation. The analysis method refers to the techniques used to analyze the gathered data. We identified four important categories, including Heuristic, Statistics, Machine Learning (ML), and Graph Theory. Measurement studies must indicate their scope, often known as their focus group. Comparing a study that investigates a phenomenon within an ISP to an Internet-wide measurement study, the latter may require a different methodology and have a higher level of complexity. Major categories of scope include Internet-wide, Country-level, ISP-level, campus-level, and lab settings, among others. The measurement size is an additional criterion for contrasting Internet measurement studies and comparing their comprehensiveness, since it indicates the quantity of data collected on the respective subject. In the summary of the reviewed studies, we specify the measurement period and the overall size of the gathered artifacts.

2. Literature Search Methodology

A literature search for this survey covered many journals and conference proceedings, as demonstrated in Table 1. Besides the conferences and journals dedicated specifically to this field such as

Table 1
List of Literature Search Methodologies (Journals and transactions are appeared in bold.)

Topic	**Conferences & Journals**
Security	IEEE Symposium on Security and Privacy – ACM Conference on Computer and Communications Security – USENIX Security Symposium – The Network and Distributed System Security (NDSS) Symposium – International Symposium on Research in Attacks, Intrusions and Defenses (RAID)– Annual Computer Security Applications Conference (ACSAC) – Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) – ACM ASIA Conference on Computer and Communications Security (ASIACCS) – IEEE European Symposium on Security and Privacy (EuroS&P) – Privacy Enhancing Technologies Symposium (PETS) – International Conference on Availability, Reliability and Security (ARES) – **IEEE Transactions on Information Forensics and Security – IEEE Transactions on Dependable and Secure Computing – ACM Transactions on Privacy and Security – Elsevier Computer & Security**
Internet Measurement	ACM Internet Measurement Conference – Passive and Active Measurement Conference – Network Traffic Measurement and Analysis Conference – The Web Conference (WWW)
Networking	IEEE International Conference on Computer Communications – ACM SIGCOMM – IEEE Global Communications Conference – **Transactions on Networking– IEEE Journal on Selected Areas in Communications**
Surveys	**IEEE Communication Surveys and Tutorials– ACM Computing Surveys**

ACM Internet Measurement Conference (IMC), Passive and Active Measurement Conference (PAM), The Network Traffic Measurement and Analysis Conference (TMA), and World Wide Web Conference (WWW), the topics covered in this survey often appeared in other journals and conferences on computer networks, communications, network security, and information systems.

Although no survey was found to discuss Internet measurement for cyber security, we reviewed survey-oriented journals like IEEE Communication Surveys and Tutorials and ACM Computing Surveys. Further, we used Google Scholar, IEEE Xplore, and ACM Digital Library to search for related papers in venues enumerated in Table 1 using the queries consisting of ”measurement“ + ”empirical“ + ”cyber security“ with the following associated keywords: ”Internet“, ”network“, ”active“, and ”passive“. The 556 identified publications were then manually screened to reject those that did not fall within the scope of this study, resulting in 337 final papers.

Due to the increasing emergence of new threats and vulnerabilities, as well as the high volume of publications in various covered topics, we only included papers that were published from 2015 up to 2022 in this survey, inclusively, in order to focus on the most recent methods and studies. Papers published prior to 2015 maybe be included if they offer fundamental contributions or are still highly relevant.

3. Existing Surveys

To the best of our knowledge, Internet measurement for cyber security as a whole has not been surveyed yet, even though several surveys of particular tasks and use cases were published in recent years.

In 2006, Ziviani (2006) reviewed measurement-based methods related to five topics; bandwidth estimation, traffic matrix estimation, traffic sampling and anomaly diagnosis, network proximity evaluation, and geolocation of Internet hosts. Bou-Harb et al. (2013) discussed the nature of cyber scanning, in addition to the approaches and strategies utilized. The authors provided a classification for 19 cyber scanning techniques. The survey covered techniques adopted by adversaries and botnets for wide-range cyber scanning. Bajpai and Schönwälder (2015) reviewed Internet performance measurement platforms based on their scale, coverage, timeline, deployed metrics and measurement tools, architecture, and overall research impact. Further, the authors provided a classification based on use-case deployment, namely, fixed-line access measurements, mobile access measurements, and operational support. Fachkha and Debbabi (2015) provided a comprehensive survey on darknet (Internet telescope) in 2015, which concerned three main darknet-related categories: deployment, traffic analysis, and visualization. In fact, the darknet has been used extensively to extract insights on Internet-wide probes and scanning activities, DDoS attacks, and other occurrences such as Internet outages and societal/political events.

Subsequently, Aleroud and Zhou (2017) published a survey on phishing threats in 2017. The authors not only investigated phishing attacks and anti-phishing techniques in traditional environments (i.e. emails and websites) but also in new environments such as mobile and social networks. More recently, Zhou et al. (2018) surveyed existing network data collection techniques in 2018. Although this topic is limited in terms of the types of data used (e.g., packets and NetFlow) and size of the network, it shares common techniques and challenges with the Internet measurement topic. Additionally, Torabi et al. (2018b) shed light on the importance of threat detection using passive DNS data, with the aim of providing a systematic review of the previously implemented systems that utilized passive DNS data to detect malicious behavior on the Internet. They discuss the studies with respect to collected data size, DNS data source type, DNS level and experimented machine learning models. Similarly, Zhauniarovich et al. (2018) investigated DNS queries and responses in order to identify malicious domains based on their behavior. They surveyed the existing approaches using several viewpoints: (i) sources of the DNS data, (ii) the data analysis methods, and (iii) the evaluation strategies and metrics.

In 2019, Jing et al. (2018) surveyed data collection and analytics for measuring Internet security. The authors categorized security data into four main classes: (1) packet-level, (2) flow-level, (3) connection-level, and (4) host-level data. Analytic methods applied to DDoS flooding and worm attacks are also elaborated upon in detail. The studies are summarized depending on the type of security data obtained (e.g., number of packets, packet rate, etc.), the analysis method for attack detection, and the various performance metrics (e.g., accuracy, false alarm, etc.).

In the same year, Singh et al. (2019) presented a comprehensive survey on DNS-based botnet detection. The authors offered a new classification of DNS-botnet detection techniques and covered each technique based on the following categories: flows, anomaly, flux, DGA, and bot infection. The studies are reviewed in terms of the techniques used, the whitelist/blacklist criteria, the target botnet, the dataset, and the pros and cons.

Table 2 summarizes the topics and research areas described in the related surveys according to the taxonomy offered in Section 2. This table also highlights how our literature survey differs from the existing ones. After more than two decades since the Internet measurement field has emerged, we have unfortunately found that related literature is rather scattered. In large part, this is due to the field of Internet measurement encompassing many different and equally important aspects in need of investigation. In turn, researchers typically only focus on a subset of them. Consequently, to the best of our knowledge, there are no surveys that offer an extensive view of Internet measurements for cyber security. In ad-

Table 2
A Classification of Related Surveys

	Research Area	Ziviani (2006)	Bou-Harb et al. (2013)	Bajpai and Schönwälder (2015)	Fachkha and Debbabi (2015)	Aleroud and Zhou (2017)	Zhou et al. (2018)	Torabi et al. (2018b)	Zhauniarovich et al. (2018)	Jing et al. (2018)	Singh et al. (2019)	**Ours**
Covered Topics	**Internet Core**	✓	✓	✓			✓	✓	✓	✓	✓	✓
	Internet Overlay											✓
	Optional Services											✓
	Access Layer											✓
	Threats & Attacks				✓	✓		✓	✓	✓	✓	✓
	Real-world Impact											✓
Scope	**Cyber Security**		✓		✓	✓	✓	✓	✓	✓	✓	✓
	Internet-Wide	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓

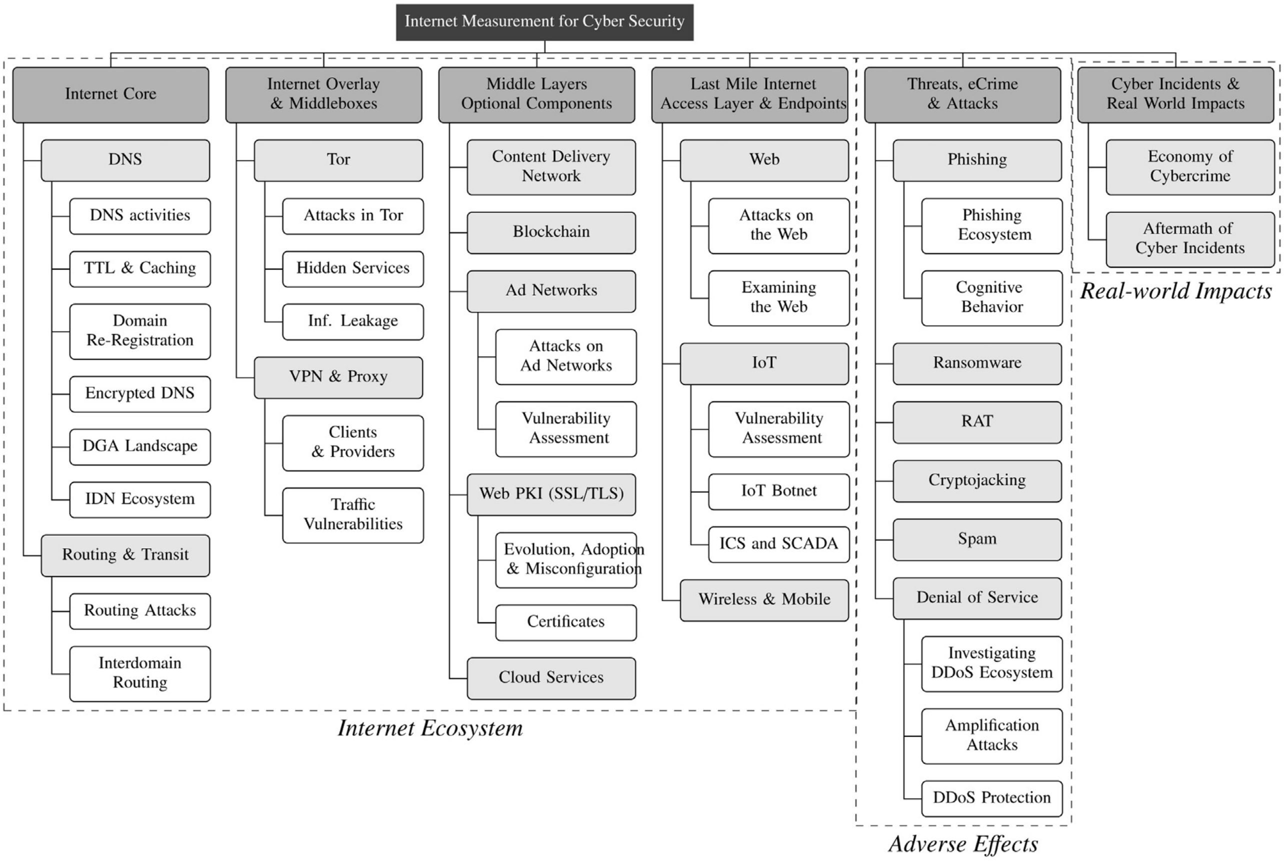

Fig. 2. A taxonomy detailing the application domains of Internet measurement

dition, providing such an overall picture could inform researchers and practitioners that are either new to the topic, wish to expand their knowledge on Internet measurement to associated aspects related to their expertise, or are simply exploring new possible applications for their specialty.

4. Internet Core

In this section, we elaborate on several topics related to Internet core components and services, namely, the DNS ecosystem and routing and transit protocols. Figure 2 presents a categorization of the subjects covered in studies on Internet measuring techniques for analyzing the cyber security of Internet core components.

A number of subjects, such as Internet addressing (IPv4/IPv6) and Internet outages, have some connection to the Internet core section. Internet-wide scanning is often used to comprehend the Internet's topology and security (Durumeric et al., 2013). However, IPv4 Internet scans have been restricted to a subset of services since it is too expensive to scan IPv4 addresses across all ports for all IPv4 services. Therefore, recently, there have been efforts to develop bandwidth-efficient frameworks for discovering services on unusual ports (Izhikevich et al., 2021; 2022b). Further, the trend of IPv6 adoption has continued with many network operators deploying IPv6 to significant parts of their networks (Borgolte et al., 2018b). Considering the vastness of the IPv6 address space, it is not feasible to probe every possible unicast IPv6 address for measurement purposes, such as when conducting classification tasks and other means of comprehending this new address space. Consequently, most of the efforts regarding IPv6 explorations are toward finding efficient ways to probe the IPv6 cyberspace by discovering nominated seeds and other methods in order to generate a target hitlist (Beverly et al., 2018; Borgolte et al., 2018b; Gasser et al., 2018b; 2016; Murdock et al., 2017; Plonka and Berger, 2015). While we will not get into the specifics, it is important to note that this is one method of measuring Internet.

Further, large-scale power outages, underwater cable breaks, bad network administration, security attacks, natural disasters, and government-imposed network outages can all cause network disruptions. These situations may result in the unstable functioning of the Internet, which can jeopardize the dependability and availability of essential services and mission-critical applications. As a result, several research endeavors have focused on outage characterization (Banerjee et al., 2015; Giotsas et al., 2017a; Gunawi et al., 2016; Holterbach et al., 2017; Padmanabhan et al., 2019a; 2019b; Richter et al., 2018; Tao et al., 2019; Zhang et al., 2017), Internet outage detection (Bogutz et al., 2019; Fontugne et al., 2017; Guillot et al., 2019; Shah et al., 2017) and assessing real-world events (Bayat et al., 2021). However, these apply more to the dependability of the Internet than to its security.

4.1. DNS

DNS is a critical Internet protocol that was first specified in 1983 in Request for Comments (RFC) 882 and 883 (Mockapetris, 1983; 1987). Despite that its primary objective of resolving domain names to IP addresses appears to be rather straightforward, the DNS environment is quite complex. It consists of a hierarchy of recursive resolvers and authoritative nameservers that are coupled with a diverse set of query types and DNSSEC. Furthermore, authoritative DNS servers can be configured to behave differently based on latency, geolocation, and content filtering policies. As a result, any measurement used to explore this ecosystem is extremely reliant upon the vantage point(s) from which it is obtained, and it may only provide a partial view of the DNS zone.

Along with the aforementioned complexities, the DNS environment has become a haven for illicit and fraudulent acts due to the multistakeholder governance approach it employs. Such acts pertaining to the name system itself can be categorized by two distinct forms of threats. To begin, because DNS transactions are not verified, DNS may create erroneous mappings of names to IP addresses. Additionally, while the cryptographic DNS zone signature technology has existed for over a decade, it remains unused. Second, the name registration environment promotes extremely opaque name usage, which is not only beneficial for privacy but also can be leveraged by malicious actors.

Typically, passive DNS collects data about the interaction between recursive caching name servers (i.e., resolvers) and authoritative name servers. Data from passive DNS configurations can be used for applications such as tracing names connected with malicious IP addresses. Unfortunately, passive DNS does not offer accurate data over time. Specifically, passive DNS will only record data for domains that are of interest to the clients behind the resolvers that gather passive DNS data. Additionally, passive DNS has no effect on the query's temporal spacing.

The necessity for researchers to obtain dependable DNS data over time prompted the development of a supplementary alternative to passive DNS based on active measurements (van Rijswijk-Deij et al., 2016). With DNS zone files from Top-Level Domains (TLDs) as input, a set selection of queries for each domain in a TLD will be sent once every 24 hours. In fact, the behavior of clients making queries can be controlled via this technique. However, scaling such a method is quite difficult. For instance, the.com domain alone (the Internet's largest TLD) includes over 120M names. An abstract illustration of passive and active DNS measurement are shown in Figure 3. In addition, Table 3 summarizes and compares the papers that will be discussed.

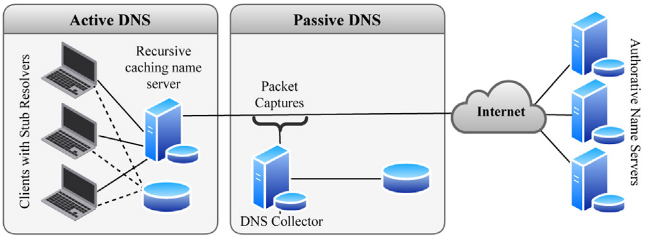

Fig. 3. Abstract view of passive vs active DNS measurement (adapted from van Rijswijk-Deij et al. (2015))

4.1.1. Malicious DNS activities

Passive collection of DNS traffic at recursive DNS, authoritative DNS and TLD servers along with DNS backscatter are typically used to find malicious activities through analysis of queries. Kwon et al. (2016) aspired to utilize DNS for fingerprinting botnets. They leveraged timing information of query generation, regardless of the number of domains and queries and used the Power Spectral Density (PSD) to discover the significant frequencies resulting from periodic DNS queries from botnets. Their system gave a detection accuracy of 95% with a false positive rate of 0.1%, and was able to detect 23 unknown and 26 known botnet groups. DNS is also leveraged as a covert channel by malware authors in order to perform data exfiltration. As a result, Nadler et al. (2019) suggested a machine learning method for identifying classes of covert channels. The authors evaluated their approach on a large-scale recursive DNS server's logs and detected two malware variants.

With malicious domains being key components in a variety of cyber attacks, researchers have been leveraging DNS data as a means of identifying them. Chiba et al. (2016) proposed DomainProfiler, which utilizes DNS logs in order to detect malicious domain names that will likely be abused in the future. DomainProfiler assesses the Temporal Variation Patterns (TVPs) of domain names to extract information regarding how and when a domain was recorded in blacklists and whitelists. The authors showed that their system can predict malicious domain names 220 days in advance with a true positive rate of 0.985. Rahbarinia et al. (2016) presented Segugio, which utilizes passive DNS traffic to detect previously unknown malware-controlled domain names. The approach first constructs a machine-domain bipartite graph of what domains are being queried by machines that are known to be benign or malware-related. Subsequently, Segugio couples this information with a number of other domain-name features (annotated in the graph) to arrive at a probability of a particular domain name being used for C&C or that the machine making the query is infected. The proposed system was implemented in large ISP networks and achieved a true positive rate of up to 94%. Despite success in malicious domain detection, this success has largely been predicated on local features (e.g., domain name patterns, daily similarity, average TTL, etc.), which can potentially be modified by attackers. To address this issue, Khalil et al. (2016) proposed an approach that analyzes global associations among domains. The authors primarily leveraged the expectation that multiple malicious domains will end up being hosted using the same IPs over time and vice versa. Based off this intuition, a graph-based inference technique over associated domains (i.e., with common IPs) is constructed to enable the discovery of a large set of unknown malicious domains by using a small set of domains known to be malicious. The approach was evalu-

Table 3
Summary of the DNS papers

Topic	Papers	Measur. Method	Analysis Method	Scope	Measurement Size
Malicious DNS activities	Kwon et al. (2016)	Passive	Statistics	Internet-wide	~100M queries, ~10M domains
	Nadler et al. (2019)	Passive	ML	Internet-wide	47M queries/hour
	Chiba et al. (2016)	Hybrid	Statistics & ML	Internet-wide	>55M domains
	Rahbarinia et al. (2016)	Passive	Graph Theory & ML	ISP-Level & Country-Level	~101.4M domains
	Khalil et al. (2016)	Passive	Graph Theory	Internet-wide	152K domains
	Kountouras et al. (2016)	Hybrid	Heuristic	Campus-Level, Internet-wide	~1.5B domains
	Fukuda et al. (2017)	Passive	ML	Country-Level, TLD	7.11×10^{11} queries
TTL and Caching	Moura et al. (2019)	Hybrid	Heuristic	Country-Level	58M domains
	Al-Dalky et al. (2019)	Passive	Heuristic	Country-, Campus-Level	>5M resolvers
	Kührer et al. (2015)	Hybrid	Heuristic	Internet-wide	155 domains
Malicious Domain Re-Registration	Liu et al. (2016)	Active	Heuristic	Internet-wide	Top 10k Alexa ~570k subdomains
	Vissers et al. (2017)	Active	Heuristic	Internet-wide	Top 10K NS domains
	Pearce et al. (2017)	Hybrid	Heuristic	Internet-wide	4.2M open resolvers
	Lever et al. (2016)	Passive	Heuristic	ISP-Level	179M expired domains
	Lauinger et al. (2016)	Passive	Heuristic	Internet-wide	7.4M domains
Encrypted DNS	Chung et al. (2017)	Hybrid	Heuristic	Internet-wide	59K DNS resolvers, >147M domains
	Lu et al. (2019)	Hybrid	Heuristic	Internet-wide & ISP-Level	122k IP addresses in 166 countries
	Hoang et al. (2020)	Active	Heuristic, Statistics	Internet-wide	13.6M domains
DGA Landscape	Plohmann et al. (2016)	Hybrid	Heuristic	Internet-wide	>18M DGA domains
	Stevanovic et al. (2015)	Passive	Graph Theory	Regional ISP-Level	13.93M domains
IDN & Typosquatting Domains	Liu et al. (2018)	Passive	Heuristic	Internet-wide	1472k IDNs from 56 TLD zone files
	Quinkert et al. (2019)	Hybrid	Heuristic	Internet-wide	Alexa Top 10K, 3k IDN
	Suzuki et al. (2019)	Hybrid, crowdsource	Heuristic, Statistics	Internet-wide	141M domain
	Le Pochat et al. (2019)	Hybrid	Heuristic	Internet-wide	15k IDNs
	Chiba et al. (2019)	Passive, crowdsource	Heuristic	Internet-wide	4.4M registered IDNs, from 570 TLDs
	Yazdani et al. (2020)	Hybrid	Heuristic	Internet-wide	7.5M IDNs
	Agten et al. (2015)	Active	Heuristic	Internet-wide	Top 500 Alexa domains
	Dam et al. (2019)	Hybrid	Heuristic	Internet-wide	485k domains
	Khan et al. (2015)	Hybrid	Heuristic, Statistics	Internet-wide, campus-, enterprise-level	~54.7M domains

ated on a public passive DNS database and reached a true positive rate of up to 95%.

Kountouras et al. (2016) took a different approach to facilitating the malicious domain detection effort by offering a large-scale and freely-available means of aggregating active DNS data. The proposed system, referred to as Thales, accomplishes this aim by generating numerous DNS queries from a list of publicly accessible sources of collected domain names (e.g., the Alexa list, various TLD zone files, etc.). From the resultant dataset, the authors identified domain names that ultimately ended up consuming over 75% of public blacklists several weeks later. Additionally, it was demonstrated that malicious campaigns can be fingerprinted solely using such active DNS data.

Further, Fukuda et al. (2017) endeavored to identify malicious events from DNS backscatter. By leveraging machine learning, they were able to identify several scanning trends and bursts corresponding to Heartbleed. They also showed that the classification of these queries enabled the fingerprinting of several network events with a precision of 70-80%.

4.1.2. TTL and Caching

DNS Time To Live (TTL) is a setting that determines how long a DNS resolver should cache a query. Moura et al. (2019) executed several controlled experiments on captured passive and active DNS data to demonstrate how setting TTL values affect operational networks and attack mitigation. The authors also enumerated considerations and recommendations for selecting appropriate TTLs. For example, DNS load balancing or DDoS mitigation may require short TTLs of 5 or 15 minutes, whereas other implementations may benefit from TTLs of a few hours.

Al-Dalky et al. (2019) studied the behavior of recursive resolvers that adopted the EDNS0-Client-Subnet (ECS) extension, as well as their caching implementations. This extension includes end-user subnet information in DNS queries to offer good approximations of end-user locations, which enables authoritative DNS servers to better assign users to the nearest edge servers in CDNs. To this end, the authors assessed both the authoritative DNS servers of a major CDN and a busy DNS resolution service. Their analysis uncovered both deviations from the expected caching behavior and detrimental actions (e.g., defying RFC recommendations), which can ultimately lead to the violation of users' privacy, a decreased effectiveness of DNS caching, and a dramatic reduction of the benefits of ECS.

Millions of recursive DNS resolvers are still open to the public, which makes them vulnerable to cache snooping and DDoS amplification attacks. Kührer et al. (2015) shed light on this phenomenon by analyzing it from two different angles: (1) assessing changes over time and classifying such resolvers by their device type and software version, and (2) measuring the authenticity of open resolvers‘ responses from a client viewpoint. The authors found that some open DNS resolvers manipulate DNS resolutions to censor communication channels, inject ads, perform fishing attacks, serve malicious files, or perform malicious redirection. The amplification vulnerabilities of the DNS protocol were also shown, as well as its lack of verification mechanisms at the application-level that allows clients to be redirected to suspicious content.

4.1.3. Malicious Domain Re-Registration

Anyone who re-registers an expired domain implicitly inherits the domain's previous residual trust. This occurrence is leveraged by adversaries to exploit both users and systems. To determine the extent of such DNS-related vulnerabilities and abuse, researchers usually concentrate their study on the top domains (e.g., the Alexa top 10k) and then collect thorough replies and test results. Liu et al. (2016) shed some light on dangling DNS records (records pointing to resources that are not available) and demonstrated how attackers can even have these compromised domains signed with a Certificate Authority (CA). In addition, they found 467 exploitable dangling DNS records in 277 Alexa top 10k domains and 52 edu zones. Adversaries can also hijack domains via exploiting configuration issues and hardware errors in nameservers in order to seize control over their requests. Vissers et al. (2017) studied this phenomenon by performing a large-scale analysis of 10K popular domains. Subsequently, they mapped out existing abuse and vulnerable entities, and uncovered 52.8M domains that were targeted by nameserver bisquatters who respond with rogue IP addresses. They also identified 1.28M domains that are at risk for DoS attacks due to outdated nameservers. In addition to hijacking, DNS manipulation also encompasses DNS poisoning or redirection. To this extent, Pearce et al. (2017) proposed Iris, a system for measuring the global manipulation of DNS resolutions. It relies on a variety of metrics associated with the consistency of resolutions with their control group and others that can be independently verified with external data sources (e.g., the HTTPS certificate infrastructure) in order to determine if manipulation has occurred.

Lever et al. (2016) proposed Alembic, a lightweight algorithm that detects potential domain ownership changes from passive DNS. The authors measured the scope and growth of this threat over six years and identified 27,758 domains from public blacklists and 238,279 domains that expired and were subsequently resolved by malware. Several of the instances of abuse that they uncovered were previously unidentified, including an expired Advanced Persistent Threat (APT) domain. Lauinger et al. (2016) offered insight pertaining to the re-registration phenomenon via analyzing the WHOIS databases. Upon overcoming related shortcomings (e.g., inconsistent data formats, the rate limits imposed on the number of lookups, and ambiguities in the data), the authors collected 7.4M com, net, org, biz, and name domains that were about to be deleted over the course of ten months. Survival analysis revealed that several re-reservations happened soon after the domain's deletion, especially when the domains were older. Also, it is challenging to predict the time between the domain's expiration date and its deletion.

4.1.4. Encrypted DNS Protocols

DNSSEC authenticates the resolution of IP addresses with a cryptographic signature, which ensures the validity of the answers provided by the DNS server. Chung et al. (2017) performed a 21-month study on DNSSEC's PKI management, as well as its deployment and maintenance. This empirical analysis was carried out on data from all DNSSEC-enabled subdomains under the com, org, and net TLDs. In addition, active measurements were conducted on more than 59K DNS resolvers worldwide to evaluate resolver-side validation. Ultimately, widespread DNSSEC infrastructure mismanagement was identified, such as 31% of domains supporting DNSSEC fail to publish all relevant records required for validation. Moreover, of the 82% of resolvers that requested DNSSEC records, only 12% of them attempted to validate the response.

Protocols offering encryption for DNS-related traffic (e.g. DNS-over-TLS (DoT), DNS-over-HTTPS (DoH), and Encrypted Server Name Indication (ESNI)) have also been introduced, which can mitigate threats against plaintext DNS packets. To this extent, Lu et al. (2019) found that encrypted DNS queries are less likely to be disrupted by in-path interception than traditional DNS queries. However, they also discovered that 25% of DNS service providers use invalid SSL certificates, and that far more users rely on traditional DNS than its encrypted counterparts.

4.1.5. DGA Landscape

Domain Generation Algorithms (DGA) are algorithms that are utilized by some malware families to periodically generate a large number of domain names that can function as rendezvous points with Command and Control (C&C) servers. Passive DNS is the primary data source for analyzing and detecting DGA domains. Plohmann et al. (2016) analyzed 43 such malware families and variations to conduct a comprehensive assessment of the DGA landscape. They proposed a DGA taxonomy, which they utilized to characterize and compare the attributes of these malware families. They examined the registration status of over 18 million DGA domains and found that pre-compiling future DGA domain names can accurately identify malware campaigns. Stevanovic et al. (2015) endeavored to facilitate the detection of such malicious domains via machine learning by presenting a semi-manual procedure for labeling DNS queries that are used to reach potentially malicious servers by fast changing domain names and/or IP addresses. The authors used DNS traffic from an Internet Service Provider (ISP) to evaluate their approach and showed the importance of domain/IP address blacklists and whitelists for DNS traffic labeling. However, it was also highlighted that blindly relying on these blacklists and whitelists may result in misleading conclusions about the analyzed DNS traffic.

4.1.6. IDN & Typosquatting Domains

Internationalized Domain Names (IDNs) containing non-ASCII characters and were installed in DNS over 15 years ago. It enables IDN homograph attacks, which is caused by the fact that many different characters look alike but have a different assigned code. Liu et al. (2018) found that approximately 1.4 million IDNs were registered under more than 700 registrars, and east Asian regions have exhibited significant growth in IDN registration. In terms of IDN abuse, this research endeavor uncovered that URL blacklists contain over 6K IDNs. Additionally, 1,516 and 1,497 IDNs were identified as having strong visual (i.e., homograph domains) and semantic similarities, respectively, to well-known brand domains (e.g., apple.com). Quinkert et al. (2019) aspired to detect homograph IDNs and monitor their activity. The authors implemented a two-phase procedure consisting of (1) detecting the homograph IDNs and then (2) collecting metadata about them on a daily basis. The first phase was conducted for a period of eight months and resulted in the discovery of around 3,000 homograph domains targeting both technology and financial organizations. The second phase was performed for over five months and uncovered several cases of scamming and phishing, with some of these cases being active for months.

Suzuki et al. (2019) designed ShamFinder, which is an automated approach for detecting IDN homographs. ShamFinder was used to conduct a large-scale measurement study in order to better understand the IDN homographs that exist in the wild. As a result, the authors were able to offer insights on effective countermeasures against IDN homograph attacks. Chiba et al. (2019) proposed DomainScouter, a new system that can detect various homograph IDNs via a calculated deception score, which is the estimated probability of a user being deceived by a given IDN. Using approximately 4.4M registered IDNs under 570 TLDs, DomainScouter detected 8,284 homograph IDNs. Of these detected homographs, many were previously unknown and targeting non-English brands. Yazdani et al. (2020) applied homoglyph tables to generate ASCII counterparts of IDNs in order to detect suspicious homograph domain names. The approach was performed on data from the OpenINTEL platform, which takes daily active DNS measurements of over 65% of the DNS namespace. The authors also extended existing homoglyph tables, which allowed them to detect roughly 2.97 times more homograph domains on average than would be possible with existing homoglyph tables. While the approach is only designed to fingerprint suspicious domains, it detected domains 21 days before they appeared in the blacklists measured by OpenINTEL, on average.

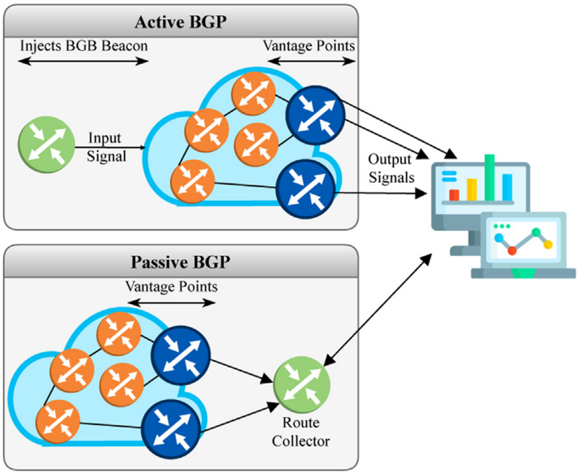

Fig. 4. Abstract view of active and passive BGP measurement

In contrast to homograph domains, typosquatting is when adversaries register domains that are a mistype of a popular domain name. To this end, Agten et al. (2015) undertook a seven-month analysis of typosquatting domains by visiting the Alexa top 500 domains. The authors perform their analysis by clustering the content of potential typosquatting web pages. The study revealed that typosquatters were targeting 95% of the 500 domains studied, and the majority of such domains do not employ their own typosquatting domains as a means of defense. Dam et al. (2019) performed a large-scale investigation of a particular type of scam implemented on typosquatting domains, namely, the pop-up scam. This scam displays a JavaScript alert box, which attracts users' attention by blocking a user interface element. By leveraging the Alexa top 1M list, the authors found a total of 9,857 pop-up notifications on 8,255 different sites, with the majority being displayed to one specific HTTP user agent only. A distribution of this scam across different languages was also given. Khan et al. (2015) endeavored to assess the typosquatting cybercrime related harms to users via proposing a methodology called intent inference. They revealed that typosquatting costs the average user 1.3 seconds per typosquatting occurrence versus obtaining a browser error page. Moreover, it was also discovered that legitimate sites lose about 5% of traffic over an unregistered typo.

4.2. Routing

BGP is the Internet's global routing protocol, which is used by autonomous networks (also known as autonomous systems or ASes) to communicate network topology information. Since BGP was created without security in mind, network operators have been subjected to adversarial attacks on routing (e.g., traffic hijacking). Thus, for more than two decades, the Internet engineering community has continuously attempted to build and deploy interdomain routing security mechanisms. However, finding an efficient method of processing large amounts of distributed or live BGP measurement data is extremely challenging. Such mechanisms include passive and active BGP measurement techniques, as shown in Figure 4. Furthermore, many routing anomalies are caused by inaccurate BGP prefix announcements (e.g., operational interruptions) rather than routing attacks, which complicates the research.

In this subsection we explored the empirical studies about the routing attacks, securing Interdomain routing, and BGP communities. In addition, subsection 8.6.3 provides a comprehensive overview of BPG-based DDoS defense approaches, such as blackholing.

The growing use of encrypted network traffic has become a double-edged sword. On one edge, it ensures safe data transfer, prevents eavesdropping, and increases the reliability of communicating hosts. Yet on the other edge, it makes essential network traffic control tasks (e.g., traffic classification and host recognition) more difficult. Therefore, monitoring, identifying, and classifying encrypted communication has caught the attention of researchers in recent years. In general, these researchers' works can be categorized as fingerprinting encrypted traffic (Aceto et al., 2018; Anderson and McGrew, 2020; Di Martino et al., 2019; van Ede et al., 2020; Husák et al., 2015; Kohls et al., 2019; Mavroudis and Hayes, 2020; Panchenko et al., 2016; Rimmer et al., 2018; Taylor et al., 2017; 2016) or insight generation via network traffic analysis (Al-Naami et al., 2016; Anderson et al., 2018; Frolov and Wustrow, 2019; Holz et al., 2016; Li et al., 2017; Msadek et al., 2019; Siby et al., 2020; Zhang et al., 2018). However, we will not discuss the specifics of these articles because their subject matter is not directly connected to cyber security.

4.2.1. ***Routing Attacks***

Birge-Lee et al. (2019) introduced Surgical Interception using COmmunities attacks (SICO), which launch interception attacks on BGP communities in order to study an adversary's attack and ensure a route to the victim. The authors showed that SICO attacks can target specific IP addresses in order to reduce attack costs. In addition, SICO attacks were executed on the real Internet to evaluate their effectiveness and feasibility. Finally, the authors analyzed the Internet topology and found that at least 83% of multi-homed ASes have the capacity to launch such attacks. Schlamp et al. (2016) devised a novel formalization of Internet routing and applied it to routing anomalies in order to establish a comprehensive attacker model. The authors used this model to classify attacks and evaluate the extensiveness of their impact. Additionally, a new effective approach that investigates hijacking alarms was proposed and implemented, called the Hijacking Event Analysis Program (HEAP). Due to the increase in use of BGP poisoning, Smith et al. (2020) evaluated the feasibility of poisoning in practice. BGP poisoning makes use of the routing protocol between ASes to reroute upstream networks' return paths onto new paths. This route hijacking can be exploited for congestion control, censorship, geopolitical boundaries, etc. By using a multi-country and multi-router, Internet-scale measurement infrastructure, the authors captured and analyzed over 1,400 instances of BGP poisoning across thousands of ASes. They conducted a detailed analysis of the performance of steering paths and the graph-theoretic aspects of available paths. With this data, a re-evaluation of the simulated systems was performed. Cho et al. (2019) leveraged supervised learning classifiers to detect hijacking attacks in order to document the different types of such attacks. The authors introduced four categories of BGP hijacks: (1) typos, (2) prepending mistakes, (3) origin changes, and (4) forged AS paths. By leveraging AS hegemony (i.e. a measure of dependency in AS relationships), they efficiently identified forged AS paths. Additionally, they used heuristic approaches to identify typos and AS prepending mistakes. Overall, the proposed approach correctly classified the authors' collected ground truth dataset into the aforementioned four categories with 95.71% accuracy.

Testart et al. (2019) took on a new perspective on BGP hijacking activity by introducing and tracking long-term routing behavior of serial hijackers. They built a ground truth dataset by extracting information from network operator mailing lists. Then, based on this dataset, they shed some light on serial hijacking and how such networks differ from legitimate networks. They leveraged features that captured these differences and used machine learning models to automatically detect ASes that were subjected to serial hijacking.

The authors found a wide range of malicious activity, misconfiguration, and benign hijacking activity.

Further, researchers found that blocks of IP addresses are being stolen by BGP hijackers in order to launch spam campaigns. Since BGP hijacking attacks in the wild are rarely documented or reported, Vervier et al. (2015) analyzed 18 months of data collected by an infrastructure built for examining the frequency of intentional stealthy BGP hijacks on the Internet. The authors found more than 2,000 malicious hijacks. Due to a lack of ground truth, ISP assistance was required to confirm that some of these hijacks indeed occurred.

4.2.2. Securing Interdomain Routing

There are two primary mechanisms used to secure interdomain routing: origin authentication with Resource Public Key Infrastructure (RPKI) and path validation with BGPsec. However, BGPsec is a long way from being adopted due to several issues such as replacing routing infrastructure and the overhead of online cryptography.

To this end, Iamartino et al. (2015) proposed a measurement-based approach for prefix origin validation in order to mitigate routing hijacks that are often due to BGP misconfiguration. This approach is based on an RPKI in order to reveal mis-origination problems that occur in RPKI repositories of Internet registries. This work used data from public BGP monitors (routeview) over a 2.5-year period to identify the main causes of errors in these registries and to investigate possible solutions. The impact that invalid origins have on routing tables and the traffic at the edge of a large research network was also shown. Cohen et al. (2016) proposed a deployable extension to RPKI called path-end validation to supplement the shortcomings of BGPsec. This approach does not entail in changing the routing infrastructure nor online cryptographic operations. Moreover, the authors showed that path-end validation has significant benefits even when partially deployed. Gilad et al. (2018) proposed DISCO, an automated system for certifying ownership over IP prefixes to address RPKI's deployment challenges. The authors demonstrated that DISCO can be deployed on the Internet as currently constructed. Additionally, it was also shown that DISCO provides a secure path against prefix and subprefix hijacks. Finally, the authors noted that DISCO is a first step towards mitigating more sophisticated attacks on BGP.

Birge-Lee et al. (2018) performed the first real-world demonstration of BGP attacks that can obtain bogus certificates from top CAs. The authors leveraged a dataset of 1.8M certificates in order to assess the vulnerability of PKI and found that an adversary is capable of obtaining bogus certificates for the majority of the domains. Two countermeasures for securing PKI against these attacks were proposed and verified: (1) CAs that verify domains for several vantage points in order to make it harder to launch a successful attack, and (2) a BGP monitoring system for CAs that detects suspicious BGP routes. Sermpezis et al. (2018) proposed Automatic and Real-Time dEtection and MItigation System (ARTEMIS), a defense approach against BGP hijacking. The authors based their method on accurate and fast detection performed by the AS by leveraging public BGP monitoring services that offer real-time streaming. The authors showed that ARTEMIS enables both fast and flexible mitigation of hijacking events. For example, it can neutralize prefix hijacking within a minute.

5. Internet Overlay and Middleboxes

In this section, we cover topics related to Internet overlay and middleboxes (as shown in Figure 2), namely, Tor, VPN and Proxies. A synopsis of the publications pertinent to these subjects is also provided.

5.1. Tor

The Onion Router, commonly referred to as Tor, is a peer-to-peer overlay routing network that ensures the privacy of data during transmission between the source and destination. Tor also provides low-latency communication and, consequently, seeks to strike a balance between both anonymity and performance. There are several research efforts devoted to characterizing the Tor network (Cangialosi et al., 2015; Jansen et al., 2018; Mani et al., 2018). Furthermore, in this section, we group the studies of this network into the following three major classes: (i) attacks in Tor, (ii) Tor hidden services, and (iii) information leakage in the Tor environment.

5.1.1. Attacks in Tor

Iacovazzi et al. (2019) proposed DUSTER, which is an active traffic analysis attack based on flow watermarking. DUSTER exploits a vulnerability in Tor's congestion control mechanism to link a Tor onion service to the target's real IP address. The watermarking system embeds a watermark at the entry relay of a Tor circuit, which is propagated throughout the Tor network and can be identified by the modified Tor relays that are in proximity of the onion service. Several experiments were performed on the real Tor network and showed that their approach gives a true positive rate upwards of 94%. They also discussed a solution to mitigate this and other traffic analysis attacks that exploit Tor's congestion control. Further, Jansen et al. (2018) explored traffic analysis attacks on Tor that are solely conducted with middle relays. Nasr et al. (2018) explored flow correlation attacks on Tor and designed DeepCorr, a system that outperforms the state-of-the-art in correlating Tor connections. The authors leveraged a deep learning architecture in order to learn a flow correlation function specifically tailored to Tor's complex network. By only collecting about 900 packets of each target Tor flow (roughly 900KB of Tor data), DeepCorr provided a flow correlation accuracy of 96%.

Sun et al. (2015) presented a series of new active routing attacks, called Raptor, that can be launched by ASes in order to compromise user anonymity. First, AS-level adversaries can exploit the asymmetric nature of Internet routing to observe at least one direction of user traffic at both ends of the communication. Second, such adversaries can exploit Internet routing churn to lie on the BGP pathways for more users over time. Third, they can tamper with Internet routing by utilizing BGP hijacking (to find users utilizing specific Tor entry nodes) and interceptions (for traffic analysis). The authors evaluated their attacks on the live Tor network. In continuation, Sun et al. (2017) presented a Tor entry relay selection algorithm that mitigates such attacks. They demonstrated that their algorithm improves the security of Tor clients by 36% on average, and ASes with high Tor bandwidth can be less resilient to active routing attacks than ASes with lower Tor bandwidth.

Singh et al. (2017) studied abusive traffic on Tor such as spamming, vulnerability scanning, scraping, and other undesired behavior used by online service providers to discriminate against Tor users. The authors utilized several data sources for their study, such as email complaints sent to exit operators, commercial IP blacklists, web page crawls via Tor, and privacy-sensitive measurements of their own Tor exit nodes. The authors also developed techniques to classify email complaints, as well as an interactive crawler to identify subtle discrimination. Finally, they found that although conservative exit policies are ineffective against the blacklisting of exit relays, it is possible to block them using privacy-sensitive techniques.

5.1.2. Tor hidden services and topology

Since Tor hidden services are part of the dark Web, they can be used for many illicit activities, as they provide a certain anonymity for its users. In contrast to the exploration and analysis of Web graphs that has been widely researched over the years, there is little information on the topology of the Tor Web graph and reliable ways for detecting Tor hidden services. Bernaschi et al. (2017) addressed such lack of information by presenting a study on: (1) automatic Tor Web exploration and data collection approaches, (2) the adoption of metrics to evaluate Tor data, (3) an in-depth analysis of the hidden services graphs, and (4) a correlation analysis of hidden services' semantics and topologies. They provided interesting insights and considerations pertaining to the Tor Web organization. In a continuation, Bernaschi et al. (2019) then described the topology of the Tor graph by measuring global and local properties using well-known metrics. They considered three different snapshots that were obtained by crawling Tor three times over a 5-month time frame. By separately studying each snapshot and their shared stable core, the authors assessed the renowned volatility of Tor hidden services, as well as distinguished time-dependent and structural aspects of the Tor graph. They found that the graph of Tor hidden services shows some of the characteristics of Web graphs and that a very high percentage of nodes do not have outbound links.

5.1.3. Information Leakage in Tor

The Tor anonymity network uses the.onion pseudo-top-level domain as its naming convention and to route requests to the Tor hidden services. However, several.onion requests are still observed in the global DNS infrastructure. Mohaisen and Ren (2017) explored this phenomenon. By leveraging two large DNS traces, the authors uncovered high volumes of diverse leakage in the DNS infrastructure that was geographically distributed and targets several types of services. Additionally, several spikes in the volumes of.onion requests were correlated with various global events, such as geopolitical activities. Information leakage also occurs with the linking of Tor hidden services to the surface Web. In response, Zabihimayvan and Doran (2022) considered how this linking can affect the overall hyperlink structure of Tor hidden services, investigated the extent to which Tor hidden services are vulnerable to such leakage, and presented an overall evaluation of the Tor network. To conduct this study, nearly two million pages from 23,145 onion seed addresses were crawled over a three-month period. Ultimately, the authors found that the dark-to-surface reference network is a single massive connected component with a small number of isolated Tor domains. It was also discovered that only a few Tor hidden services are immune against information leakage, while over 90% have at least one link to the surface Web.

5.2. VPN & Proxy

Proxy and Virtual Private Network (VPN) servers function as gateways between users and other servers on the Internet and have a variety of privacy-related applications (e.g., evading server-side blocking, circumventing geographic limitations, etc.). On the other hand, adversaries may utilize these servers to conduct attacks, harvest user data, and inject both advertisements and files. As a result, several commercial and free VPN services have surfaced to offer privacy and anonymity to consumers. Consequently, these services have been the topic of interest in a number of research endeavors. In the following sections, we break down these works into VPN clients/providers and VPN traffic vulnerabilities.

5.2.1. VPN clients & providers

Ikram et al. (2016) provided an analysis of 283 Android VPN apps. These apps were extracted from more than 1.4 million apps on the Google Play Store. Both active and passive measurements were utilized to investigate the behavior, security, and privacy of these apps. The authors found out that several expose users to security and privacy vulnerabilities, such as insecure VPN tunneling protocols and leakage associated with IPv6 and DNS. Several apps were also discovered to be actively performing TLS interception. Additionally, several apps were shown to inject JavaScript to enable tracking, advertising, and redirecting e-commerce traffic to external partners.

Khan et al. (2018) not only tested the privacy of VPN apps but also their infrastructure. To achieve this aim, they designed an active measurement system that evaluated 62 commercial providers. Their system discovered that many VPNs leak user traffic in a variety of ways and several providers transparently proxy traffic. Additionally, it was found that 5-30% of the vantage points (corresponding to 10% of the studied providers) appear hosted on servers not located in the countries advertised to users. Perta et al. (2015) specifically focused on the privacy and anonymity that commercial VPN service providers claim to offer to users. The authors analyzed 14 of the most popular commercial VPN services by inspecting their internals and infrastructure. Their study revealed that most of these VPN services suffered from IPv6 traffic leakage. A sophisticated DNS hijacking attack was also developed to transparently capture all traffic.

Durumeric et al. (2017) presented a study on the prevalence and impact of HTTPS interception. To do so, the authors built a set of heuristics that detect interception and attribute the responsible entity by investigating middleboxes and client-side security software. They found that the intercepted connections use weaker cryptographic algorithms. Additionally, the study revealed that 62% of middlebox connections were less secure and 58% had severe vulnerabilities that allowed later interception.

5.2.2. Proxy traffic manipulation

Detecting traffic manipulation and violations of application-level, end-to-end connectivity on the Internet at scale remains a difficult undertaking. The majority of successful detection approaches require dedicated hardware, user-installed software, or having privileged access to a popular web site. To this end, Chung et al. (2016a) presented an alternative approach for detecting end-to-end violations by developing Luminati, an HTTP/HTTPS proxy service that routes traffic through millions of end hosts and can detect DNS, HTTP, and HTTPS end-to-end violations. The authors evaluated their approach on 1.2M nodes across 14K ASes in 172 countries. Their findings showed that up to 4.8% of the nodes are subject to end-of-end connectivity violations. As HTTP headers are becoming increasingly prone to middlebox manipulation, Tyson et al. (2017) endeavored to comprehend how ASes intercept HTTP headers in the wild. The authors collected data on thousands of networks, which revealed that 25% of the measured ASes modify HTTP headers. Lastly, the different types of manipulation and how they differ from one region to another were analyzed.

Free web proxies are increasingly being utilized for anonymity and censorship circumvention at zero-cost to their users. Perino et al. (2018) provided an in-depth examination of this complex free-proxy ecosystem. First, the authors performed active measurements to detect free proxies, assess their performance, and look for potential malicious activities. Subsequently, passive measurements were conducted to assess the proxys' performance and their usage in the wild. After monitoring up to 180,000 free proxies since January 2017, the authors' analysis showed that less than 2% of the proxies relay traffic on behalf of the users. Additionally, it was revealed that user traffic is largely derived from web browsing. Open HTTP proxies have also been an attractive option, although routing traffic through untrusted third parties can have severe consequences. To this extent, Tsirantonakis et al. (2018) performed a large-scale analysis of open HTTP proxies in order to

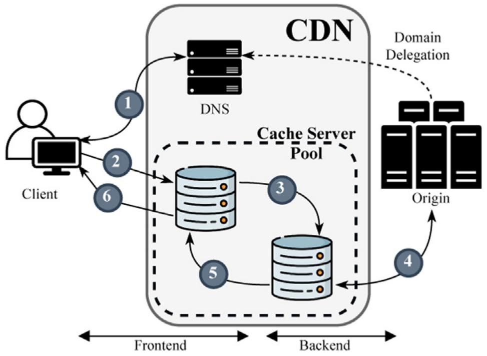

Fig. 5. Overview of CDN architecture (adapted from Jin et al. (2019))

determine the extent to which user traffic is being manipulated while being relayed. The authors' framework not only detects proxies, but is capable of attributing certain traffic modifications to well-defined malicious activities, such as ad injections, user fingerprinting, and redirection to malware landing pages. After applying their approach on a large set of publicly available HTTP proxies, it revealed that 38% of open HTTP proxies modify content. They also found that 5.15% of those proxies modify or inject either malicious or unwanted content.

6. Middle Layers and Optional Components

In this section, we studied several topics related to optional components of the Internet (Figure 2), namely, CDNs, Blockchain, Ad Networks, Web PKI and cloud services.

6.1. Content Delivery Network

A CDN is a distributed infrastructure that efficiently delivers web-related content to end-users. CDNs were originally used to improve the performance, scalability and security of websites. Typically, they are composed of a large number of surrogate servers distributed all around the world. Figure 5 provides an overview of a CDN architecture.

Request-routing techniques are the key component for CDN services, since such techniques are responsible for directing user requests from the original website to the CDN and subsequently to the appropriate surrogates. The following are the three most common request-routing techniques: (1) URL rewriting, which modifies the URL of specific content (e.g. images, css, scripts) in the origin web site; (2) CNAME (canonical name), which is a type of DNS record that links a domain name to another name; (3) Domain hosting, which entails a website using the CDN's DNS server as the authoritative name server for its domain.

The line between CDNs, Cloud-Based Security Providers (CBSPs) and DDoS Protection Services (DPSs) is becoming increasingly blurred. A significant share of CBSPs and DPSs has emerged from CDN providers that started offering security services on top of their existing platform.

Liang et al. (2014) studied how CDN and HTTPS work together and revealed that of the 10,721 HTTPS websites with CDNs, 15% were found to create invalid certificate warnings that violated the HTTPS confidence model. Scott et al. (2016) introduced Satellite, a measurement system for identifying CDN deployments and network interference from only a single measurement node. Satellite detects connected components that represent domains hosted by the same servers by leveraging a bipartite graph, which connects domains queried with IP address responses collected from the open resolvers. Satellite uncovered that 20% of the Alexa top 10K domains are hosted on shared infrastructure, with 10% attributed to CloudFlare accounts.

Chen et al. (2016a) presented and studied forwarding-loop attacks on CDNs. Through building up a large number of requests (or responses) circulating between CDN nodes, forwarding-loop attacks enable attackers to massively consume CDN resources, which results in DoS attacks. The authors measured the vulnerability of several nodes to forwarding-loop attacks in each of the 16 commercial CDNs. The results revealed that all of the studied CDNs are vulnerable to some form of forwarding-loop attack and that their existing defense mechanisms can be bypassed. Vissers et al. (2015) studied origin-exposing attacks amid CBSP solutions. The authors selected five well-known providers for this study: CloudFlare, Incapsula, DOSarrest, Prolexic (PLXedge) and Sucuri (Cloud Proxy). Among the 17,877 tested CBSP-protected websites, 71.5% of them are bypassable and exposed their real IP address through at least one of the eight evaluated attack vectors.

Hao et al. (2018) introduced a new vulnerability of CDNs, called redirection hijacking, that targets the dynamic characteristics of DNS records used for a CDN's request routing. The authors then conducted a large-scale empirical study to investigate security implications of this vulnerability in the DNS-based CDNs. To this end, eight geographically distributed vantage points were configured, and DNS resolution results were retrieved for customer websites hosted in each CDN provider. The results uncovered that in 14 out of 20 popular CDNs, an adversary might be able to manipulate end-user redirection. Jin et al. (2018) conducted a measurement study to investigate the DDoS Protection Service (DPS) usage dynamics and uncover residual resolution attacks. The authors assessed the Alexa top 1M websites for their adoption of DPS and usage behaviors (leave, join, pause, resume, and switch). To this extent, the A records from the nameservers of Cloudflare were retrieved, and the CNAMEs of Incapsula's customers were collected and resolved. As a result, it was revealed that 24.8% of all 3,504 Cloudflare's hidden records and 69% of all 42 Incapsula's point to the real origins, which makes them highly vulnerable to DDoS.

6.2. Blockchain

Blockchain is built on the top of a decentralized Peer-to-Peer (P2P) protocol, which is used to propagate relevant information such as transactions, blocks, and other protocol states. Information about all transactions and blocks is relayed to all peers, and despite inconsistent ordering and partial incompleteness, all honest peers eventually agree on a globally consistent state. Although the blockchain technology and its related services (e.g., cryptocurrency, smart contracts, distributed apps, etc.) have been studied, little attention has been paid to the underlying P2P networks. Indeed, P2P networks are responsible for information propagation and enable maliciousness such as Eclipse attack, DNS hijack, 51% attack, etc. (Saad et al., 2020).

Internet measurement studies focus on analyzing blockchain networks to provide a macroscopic view of active blockchain nodes in order to evaluate their security, reliability, and robustness. Active and passive node scanners can be used to identify the discoverability and reachability of nodes in a blockchain-based network, which allows the mapping of the entire network. Figure 6 depicts an overview of both a passive and active blockchain crawler.

The first approach to scan a network is with the implementation of a so-called *passive crawler* node. Such passive methods refer to injecting a certain number of nodes into different locations within the network that behave like normal nodes. The aim is to accept as many incoming connections as possible to discover nodes by capturing all the traffic sent and received. As opposed to *passive crawlers, active crawlers* provide a more targeted and faster node discovery; however passive approaches have the advantage

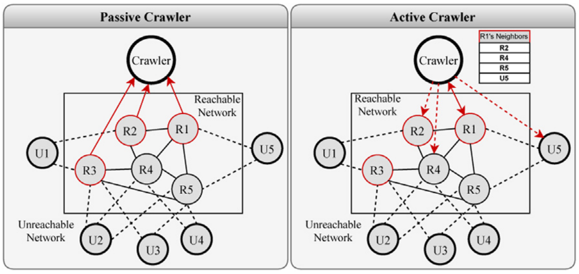

Fig. 6. Blockchain Crawler

of avoiding the additional complexities associated with discovering nodes behind a NAT. Several studies are pertinent to understanding the P2P network, decentralization metrics, and transactions in general (Ben Mariem et al., 2018; Gencer et al., 2018; Kim et al., 2018; Lee et al., 2020; Loe and Quaglia, 2019; Mariem et al., 2020; Wang et al., 2021).

Cheng et al. (2019) designed and implemented an Ethereum honeypot system called Siren that could catch real attacks in the wild. This helped him figure out how the cryptocurrency-stealing attack on Ethereum worked. The developed honeypot listens to the default JSONRPC port (i.e. 8545) and accepts any incoming requests. It is also registered as an Ethereum full node on the Internet and is prepared as a real account that has Ethers. The honeypot captured more than 308M requests from 1,072 distinct IP addresses belonging to 36 attack groups with 59 distinct accounts.

Differently, Hara et al. (2020) concentrated on network communication to understand the behavior of the attacker for transaction manipulation. Hara et al. (2020), on the other hand, focused on network communication to better understand the attacker's behavior that manipulate transactions. The authors conducted detailed analyzes of attacks and attackers by combining three types of data: (1) malicious communication history sent to the simple honeypot installed in 9 countries, (2) information from the Ethereum network and (3) darknet arrival packets. Gao et al. (2019) performed graph analysis of a measurement study related to the Ethereum P2P network. The measurement consisted of two steps: (i) discovering Ethereum nodes, and (2) harvesting the nodes' routing tables. The authors realized that there are a large number of useless and stale nodes, however, Ethereum is still resilient to both random failures and targeted attacks. Roughly 100 abnormal nodes in the network were also detected, which may cause issues in network routing.

6.3. Ad Networks

Businesses use online advertising (e.g., search engine advertising, display advertising, mobile advertising, social media advertising, etc.) to discover new consumers and extend their demographic reach. Given the sheer number of online advertising players, advertising networks function as brokers between website owners and companies to enable efficient and effective ad delivery. Unfortunately, cybercriminals take advantage of this sophisticated structure to engage in illicit activities such as malvertising, click fraud, and auction cheating. To further complicate matters, companies are often aspiring to monitor users while many consumers are employing ad blockers to circumvent such monitoring. In turn, there have been a number of research efforts that have studied the particulars of online advertising, which we subsequently elaborate upon. Table 4 contains a summary and comparison of the papers covered.

In addition, Englehardt and Narayanan (2016); Lerner et al. (2016); Merzdovnik et al. (2017); Pachilakis et al. (2019) provide measurement studies to investigate web tracking in Ad networks which may violate users' privacy. Furthermore, Alrizah et al. (2019); Gugelmann et al. (2015); Iqbal et al. (2017); Malloy et al. (2016); Snyder et al. (2020); Storey et al. (2017) focus on Ad blockers, detecting and measuring anti-ad blocker and privacy intrusive services.

6.3.1. Attacks on Ad Networks

DeBlasio et al. (2017) studied advertiser fraud on Microsoft's Bing search engine platform. Over two years of data pertaining to the advertising campaign's management and the user engagement with the ads (e.g., click-through rates) was used to perform this study. The authors revealed that Bing's policies successfully contained fraud. Despite a significant proportion of new account registrations being fraudulent, Bing successfully prevents the majority from displaying even a single ad.

The majority of malware infections are caused by drive-by downloads or Social Engineering (SE) attacks. To this extent, Nelms et al. (2016) presented the first systematic study of web-based SE attacks designed to entice users into downloading malware. The authors captured live network traffic from more than 2K samples of such attacks in the wild for this study. Analysis of these samples revealed that a large portion of these attacks are delivered from low-tier ad networks, which include bogus updates for Adobe Flash and Java, as well as fake antivirus software. Vadrevu and Perdisci Vadrevu and Perdisci (2019) developed a system based on Chromium and JSgraph for large-scale automated discovery and tracking of Social Engineering Attack Campaigns delivered via Malicious Advertisements (SEACMAs). The authors identified 93,427 different publisher websites that used SEACMAs. From these publishers, detailed logs of the JavaScript code execution related to online ad networks were captured.

Cho et al. (2015) studied the potential security risks of click mobile advertisement. The authors proposed ClickDroid, a tool that automatically simulates click generation attacks on the Android platform. They found that 6 out of the 8 advertising networks (that is Millennial Media, AppLovin, AdFit, MdotM, RevMob, and Cauly Ads) are vulnerable to automated click fraud attacks.

6.3.2. Vulnerability Assessment of Ad Networks

Regarding the rising popularity of Internet video streaming platforms, Rafique et al. (2016) proposed a comprehensive analysis of the Free Live Streaming (FLIS) services ecosystem; they attract millions of users and make use of deceptive advertisements. The authors found that users of FLIS websites are not only exposed to deceptive advertisements but also to malware, malicious browser extensions, and fraudulent scams. Additionally, the authors encountered potential trademark infringements via the abuse of domain names and logos of popular TV channels. Marciel et al. (2016) presented a set of active measurement tools which revealed that, when compared to the other audited portals, YouTube deployed the most effective detection system for fraudulent activity in video ads, although YouTube's system may still be susceptible to simple attacks. It was also discovered that views that were detected as fake and removed from the public view are still monetized.

Starov et al. (2018) analyzed the identifiers that are used by 18 different third-party analytics platforms to discover malicious campaigns amongst them. Subsequently, the authors built a system that automatically detects malicious pages, which enabled the discovery of 11K live domains that use analytics associated with malicious pages. In addition, the proposed approach can be used to improve upon the coverage of existing blacklists in order to identify previously unknown phishing campaigns. Liu et al. (2019) studied how illicit traffic monetization functions on a global scale. Indeed, the authors developed TrafficStop, the first system that passively detects such fraud. By applying more than 231B DNS logs over two weeks, TrafficStop discovered 1,457 fraudulent sites.

Table 4
Summary of the Advertising papers

Topic	Papers	Measur. Method	Analysis Method	Scope	Measurement Size
Attacks on Ad Networks	DeBlasio et al. (2017)	Passive	Heuristic	Internet-wide	~20K advertisers
	Nelms et al. (2016)	Hybrid	Heuristic & ML	Campus-level	2K real-world SE download attacks
	Vadrevu and Perdisci (2019)	Active	Heuristic	Internet-wide	11 low-tier ad networks
	Cho et al. (2015)	Passive	Heuristic	Lab setting	8 ad networks
Vulnerability Assessment of Ad Networks	Rafique et al. (2016)	Hybrid	ML & Heuristic	Internet-wide	5,685 domains
	Marciel et al. (2016)	Hybrid	Statistics, Heuristic	Lab Setting	Top 5 online video portals
	Starov et al. (2018)	Hybrid	Heuristic	Internet-wide	18 third-party analytics platforms
	Liu et al. (2019)	Passive & Active	Statistics	Internet-wide	231B DNS logs
	Subramani et al. (2020)	Hybrid	Statistics	Lab setting	82k second-level domain
	Srinivasan et al. (2018)	Hybrid	Heuristic, Statistics	Internet-wide	23k domains
	Bashir et al. (2019)	Hybrid	Heuristic	Internet-wide	Alexa top-100K

In fact, those sites receive more than 53B DNS requests within a year, which could result in a company losing 53K dollars per day due to this fraudulent traffic. Recently, ad networks have been leveraging new strategies, such as the Web Push technology enabled by modern web browsers, to keep up their revenues. Subramani et al. (2020) proposed PushAdMiner, a system that not only automatically collects Web Push Notifications (WPN) from publisher websites but also finds WPN ads among these notifications to discover malicious WPN-based Ad campaigns. In fact, PushAdMiner identified 572 WPN ad campaigns from a total of 5,143 WPN-based ads. Further, the authors found that 51% of all WPN ads collected were malicious, and that traditional ad-blockers are ineffective against them.

Srinivasan et al. (2018) studied Technical Support Scams (TSS) techniques that combine online abuse and social engineering to exploit services such as sponsored advertisements returned in search engine queries. The authors' study lasted approximately eight months and revealed 9K TSS domains of both passive and aggressive nature. Furthermore, 2,400 domains were discovered that were assisting the detected TSS domains in manipulating organic search results.

In May 2017, the Interactive Advertising Bureau (IAB) Tech Lab introduced the ads.txt standard to help ad buyers verify authorized digital ad sellers. Bashir et al. (2019) presented a 15-month longitudinal study of the ads.txt standard to understand (1) if it provides useful data to researchers and privacy advocates, and (2) if it helps ad buyers in mitigating ad fraud. ads.txt enabled identification over 1K domains that belong to ad exchanges. However, this study also found that while ads.tx data may be useful, it contains errors that should be addressed prior to it being used to mitigate ad fraud.

6.4. Web PKI and SSL/TLS Certificates

Secure communications through HTTPS is becoming more common on the modern web, thanks to Transport Layer Security (TLS) and its predecessor Secure Socket Layer (SSL). SSL is a cryptographic protocol that allows two interacting applications to communicate secretly while still ensuring message integrity. This protocol is based on a secure PKI, which uses digital certificates to verify the ownership of a public key associated with a domain name. SSL is a well-studied protocol, however other protocols that operate on top of it (e.g., HTTPS, websocket secure (WSS)) are more complicated and changing. As a result, a number of noteworthy research efforts have focused on this ecosystem, which we subsequently detail.

6.4.1. Evolution, Adoption, and Misconfiguration

Kotzias et al. (2018) conducted a longitudinal study over a six-year period to study the responses of the TLS ecosystem to high-profile attacks. This study identified a TLS client fingerprint database consisting of 1,684 fingerprints. The authors also observed a significant shift in the ecosystem from 2012, such as changes with cipher suites and TLS extensions offered by clients and accepted by servers. Additionally, some of these changes were able to be correlated with specific attacks on TLS. Lastly, a large number of client software was observed that likely offers unsafe ciphers.

Mayer et al. (2016) assessed the state of the security mechanisms in the email ecosystem. They collected and scanned e-mail related TLS configurations for the entire IPv4 range. To this end, more than 20M IP/port combinations of all related protocols (SMTP, POP3, IMAP) and legacy ports were analyzed. It was shown that securing the server-to-server communication of the scanned protocols is inherently more difficult than securing the client-to-server communication. Lastly, the volatility of TLS certificates and trust anchors in the email ecosystem were analyzed, which uncovered that many steps still need to be taken in order to effectively secure email. Durumeric et al. (2015) reported the global adoption rates of SMTP security extentions, which include STARTTLS, SPF, DKIM, and DMARC. The authors presented data from two perspectives: (i) SMTP server configurations for Alexa top 1M domains and (ii) over a year of SMTP connections to and from Gmail. Analyzing these data revealed that the top email providers (e.g., Gmail, Yahoo, and Outlook) proactively encrypt and authenticate messages. However, only 35% of these providers had successfully configured encryption and 1.1% specified a DMARC authentication policy.

Additionally, Anderson and McGrew (2019) conducted an investigation encompassing what applications are using TLS, how they are using it, and trends in enterprise TLS applications beyond the browser. In order to perform this study, the authors built a system for capturing process information from end hosts and mixed it with network data in order to produce a TLS fingerprint knowledge base. The knowledge base consisted of 471 million labeled and 8 billion unlabeled endpoint TLS sessions obtained from enterprise edge networks in five countries. The authors also incorporated millions of sessions from a malware analysis sandbox, which allowed them to reveal trends in the use of TLS by malware (e.g., adoption of cipher suite randomization). Alternatively, Fadai et al. (2015) examined the trust model used by SSL. In particular, the authors analyzed the number and origin of companies and governmental institutions that are trusted by several operating systems and browser vendors. It was determined that these trusted entities were growing over time. They also correlated Root Certificate Authorities against a variety of trust indexes to assess the trustworthiness of the countries represented by these authorities.

Samarasinghe and Mannan (2019) endeavored to analyze device vulnerabilities based on their certificates and TLS parameters. The authors noticed an increase in TLS adoption from 29.4% in 2016 to 73.7% in 2018. However, some devices were found to still have significant vulnerabilities, such as the use of known private keys, DDLv3, MD5-RSA, and RC4. For this reason, the authors contacted the leading manufacturers of vulnerable devices to convey their findings.

Beurdouche et al. (2015) designed a robust composite state machine that correctly multiplexes between a variety of protocols, extensions, authentication modes, and key exchange methods. Next, popular open source TLS implementations for state machine bugs were tested, which led to the identification and subsequent patching of hidden critical security vulnerabilities, such as the FREAK flaw. With TLS utilization still in its early stages, Mayer and Schmiedecker Mayer and Schmiedecker (2016) aspired to advance the securing of communication content by creating and validating rules used by the HTTPS Everywhere browser extension. The authors first used multiple seeding approaches over a five-month period to obtain results from more than 7,500 websites. TLScompare was then leveraged, which showed that its users had a tendency to disagree with one another. This phenomenon was even found to be true with binary decisions, such as deciding if two websites are similar over ports 80 and 443.

6.4.2. Certificates

Gustafsson et al. (2017) presented the first large-scale characterization of the Certificate Transparency (CT) landscape. Both passive and active measurements were used to highlight similarities and differences in public CT logs, as well as the usage of such logs and the certificates they include. Their analysis showed that popular domains appear to be more willing to pay the extra cost of Extended Validation (EV) certificates. The authors also offered insight pertaining to how these certificates relate to the certificates and keys that are observed in regular web traffic. VanderSloot et al. (2016) delved into different data sources to compare PKI from several perspectives, namely passive monitoring of networks, scanning the most popular domains or the IPv4 space, search engines such as Censys, and CT logs. The authors found that the aggregated CT logs and Censys snapshots encompassed over 99% of all certificates found by any of these techniques. However, such an approach still misses 1.5% of certificates found in a crawl of all domains in.com,.net and.org. The authors demonstrated how combining CT logs and Censys snapshots affected the results of previous studies. Gasser et al. (2018a) also leveraged the information in the CT logs to check if these certificates adhere to the industry's Baseline Requirements and found that 907K certificates violate them. Additionally, the authors used data from active measurements to compare certificates in logs, evaluate CT-specific HTTP headers, and identify non-HTTPS certificates in logs. A combination of passive and active measurements was also performed to analyze CT's gossiping and pollination approaches, which revealed low rates of log inclusion. Furthermore, with the increasing deployment of CT, there have also been concerns about information leakage due to the visibility of certificates in CT logs. To this end, Scheitle et al. (2018) introduced a CT honeypot to better understand this threat. The authors showed that the data from CT logs was being used in order to identify targets for scanning campaigns just minutes after certificate issuance. A methodology that learns and identifies new subdomains from the domains extracted from CT logged certificates was also presented and evaluated.

Since the circumstances of certificate revocation within the SSL ecosystem are still not well understood, Liu et al. (2015) took a closer look at the Web's PKI. The authors performed 74 HTTPS scans of the entire IPv4 address space and found that 8% of the certificates have been revoked. Additionally, it was discovered that obtaining certificate revocation information is often expensive in terms of latency and bandwidth for clients. Finally, the revocation behavior of 30 different combinations of web browsers and operating systems was studied, which showed that browsers often do not check if a certificate has been revoked. Oakes et al. (2019) conducted a longitudinal study of the characteristics of SSL certificate chains that are presented to clients during HTTPS connection setup. The authors leveraged 23B SSL certificate chains collected from a global panel consisting of more than 2M residential client machines for a period of 6 months. As a result, more than 35M unique certificate chains with different relationships were identified at all levels of the PKI hierarchy. Invalid certificate chains were also examined, which revealed that 93% contained an untrusted root certificate. Finally, the unintended behaviors of root certificate deprecation and secure traffic interception were investigated in the data.

Chung et al. (2016b) performed a study of SSL certificates to better understand why the SSL ecosystem of the web is populated by an overabundance of invalid certificates. In particular, the impact these certificates have on security and where they originated from was assessed. By analyzing over 80M certificates, it was discovered that the majority of invalid certificates came from a few categories of end-user devices. Additionally, these invalid certificates had properties that were radically different from their valid counterparts. Moreover, many of these end-user devices reissued invalid certificates on a regular basis. In response, the authors developed new tools to trace such reissues across scans. These tools enabled the tracking of over 6.7M devices.

Roberts et al. (2019) studied domain impersonation attacks that threaten end-to-end authentication. The authors then introduced a new classification of impersonation attacks, referred to as "target embedding". Next, a user-study was performed to better understand the efficacy of target embedding versus other popular impersonation attacks (e.g. typosquatting, combosquatting, and homographs). The study uncovered that target embedding is the most effective attack against modern browsers. The authors also leveraged all HTTPS certificates collected by Censys to conduct longitudinal analysis of how this attack has evolved and who is responsible for issuing impersonating certificates.

6.5. Cloud Services

Infrastructure-as-a-Service (IaaS) and, more broadly, the cloud, have altered the environment of Internet system operations. Large-scale IaaS systems have been introduced for public use by IT giants such as Amazon, Microsoft, and Google. IaaS combines multiple third-party services into a single physical pool, however, sharing physical resources with customers can result in security breaches. While several aspects of IaaS have been explored by a number of researchers, very few measurement studies assess IaaS security and threats.

Miao et al. (2015) presented a large-scale characterization of nine types of inbound and outbound cloud attacks (e.g., DDoS, SQL injection, spam, etc.) from May to December 2013. The authors investigated over 200 TB of NetFlow records collected from dozens of edge routers spread across multiple geographically distributed data centers of a major cloud providers. The study revealed high variations in attack throughput across time and Virtual IPs (VIPs). Most VIPs experienced one attack per day, although there are a few cases of multiple attacks targeting 20 to 60 VIPs simultaneously.

Xu et al. (2015) conducted a measurement study of the co-residence threat in IaaS environments. In particular, the authors investigated Amazon EC2 from the perspectives of virtual machine placement, network management, and Virtual Private Cloud (VPC). ZMap (2022) was used to scan the specified ranges of IP addresses published by EC2 for several well-known ports to acquire a list of live hosts in EC2. Such scanning followed by several experiments (e.g., traceroute) revealed that VPC can prevent the co-residence threat. Finally, the authors presented a technique that achieves co-recovery in VPC and can detect live hosts from their domain names, IP addresses, and mapping between public and private IP addresses.

Borgolte et al. (2018a) conducted a measurement study spanning 120 days that collected all DNS responses in Farsight's passive DNS that contain A records pointing to the Amazon Web Services (AWS) EC2 cloud, the Microsoft Azure cloud, and the Digital Ocean cloud. Next, the authors checked whether the IP addresses respond to ICMP ping requests or 36 of the most frequently used ports within a two-second timeout, in order to identify which IPs are online (i.e., allocated) or offline (i.e., freed). Ultimately, 702,180 unique domains (0.539%) pointed to available or freed IP addresses. In addition, 80.31% of domain migrations were delayed, 17.24% were abandoned, and only 2.45% were auxiliary.

7. Last Mile Internet, Access Layer and Endpoints

In this section, we elaborate on application topics pertaining to web layer, endpoints and last mile Internet. Refer to Figure 2 for the detailed taxonomy. A synopsis of the publications pertinent to these topics is also provided.

There are a few additional topics that may be relevant to this section, including malicious actions in Online Social Networks (OSNs), browser security, and mobile applications. OSNs have become ingrained into our daily life. Consequently, they have attracted the attention of both attackers and the research community. The primary topics of interest include large-scale detection and measurement of fake accounts and impersonation attacks (Boshmaf et al., 2015; Chen and Freire, 2020; Goga et al., 2015; Gong et al., 2019; Mariconti et al., 2017; Yao et al., 2017; Yuan et al., 2019), social botnets and crowd turfing (Besel et al., 2018; Minnich et al., 2017; 2017; Nilizadeh et al., 2017; Stringhini et al., 2015; Zhang et al., 2016). However, we find that the related articles are more concerned with analysis approaches rather than data collection.

Modern web browsers have evolved considerably since being invented for hypermedia dissemination in 1990. Currently, they are relied upon by billions of individuals globally for interfacing with the Web; thus, browsers have garnered substantial attention from researchers in terms of browser extensions (Chen and Kapravelos, 2018; Pantelaios et al., 2020; Starov et al., 2019; Starov and Nikiforakis, 2017), browser security (Jueckstock and Kapravelos, 2019; Snyder et al., 2017), and browser fingerprinting(Gómez-Boix et al., 2018; Laperdrix et al., 2016; Sarker et al., 2020; Vastel et al., 2020). Further, as a consequence of the rapid growth of mobile app ecosystems, the corpus of mobile applications and app markets has been the subject of several research studies (e.g., android apps Allix et al. (2016); Le et al. (2015); Oltrogge et al. (2021); Possemato and Fratantonio (2020); Razaghpanah et al. (2017, 2015); Song and Hengartner (2015); Wang et al. (2019, 2018b), iOS apps Deng et al. (2015); Markert et al. (2020); Orikogbo et al. (2016); Tang et al. (2020), etc.). This involves looking at how apps interact with the operating system through API calls, ensuring secure implementations (e.g., utilizing TLS), and assessing the third-party libraries used in app development (Backes et al., 2016; Chen et al., 2016b; Derr et al., 2017; Tang et al., 2019). Research efforts within this domain often leverage the Google Play Store or Apple App Store, however some studies alternatively limit their analysis to app stores based on vendor, country, or stores provided by third-parties. We observe that the articles in these areas are primarily concerned with bulk software security evaluation techniques rather than Internet measuring and data collection methodologies.

7.1. Web

Web technologies are the primary interface for common users to interact with digital world, and therefore a single website has the potential to impact millions of users. Furthermore, the constant growth of Web technologies has resulted in an evolving attack surface and ever-changing vulnerabilities. For these reasons, researchers have focused their efforts on uncovering the vulnerabilities of the most popular websites (e.g., the Alexa top 1M) by large-scale experimentation or examining trends over time in the Internet Archive.

7.1.1. Attacks on the Web

Lekies et al. (2017) proposed a novel Web attack that can circumvent all of the existing cross-site scripting (XSS) mitigation techniques. This attack abuses script gadgets, which are legitimate JavaScript fragments within an application's code base. Subsequently, the authors empirically studied script gadgets and demonstrated their presence in almost every modern JavaScript framework. Regarding web attack mitigation techniques, Musch et al. (2019a) presented ScriptProject, a non-intrusive transparent protective measure against client-side XSS introduced by third-party scripts. ScriptProject automatically removes unsafe string-to-code conversion capabilites from third-party code. An evaluation of the proposed approach showed that 30% of Alexa top 5K websites could benefit from such protection.

Roth et al. (2020) investigated the impact of script gadgets on Content Security Policy (CSP) at scale. This research endeavor demonstrated that attackers can sideload libraries with known script gadgets as long as the hosting site is whitelisted in the CSP. Additionally, the authors generated sensible CSPs for the Alexa top 10K websites in order to prove that these websites are susceptible to a bypass through script gadget sideloading.

Staicu and Pradel (2018) conducted a large-scale study of Regular expression Denial of Service (ReDoS) vulnerabilities in real-world websites. The authors proposed a novel methodology to analyze the susceptibility of deployed servers to this attack. This study revealed 25 previously unknown ReDoS-based vulnerabilities in popular libraries and tested them against 2,846 of the most popular websites. Of these websites, 339 suffer from at least one of these vulnerabilities.

Ghasemisharif et al. (2018) investigated Single Sign-On (SSO) and its association with account hijacking on the modern Web. The authors introduced novel hijacking attacks that leverage SSO and subsequently evaluated them against 95 Web and mobile services to characterize their severity and stealthy nature. Ultimately, this work uncovered the limitations of the SSO schemes that do not allow users to remediate account hijacking. In response to this shortcoming, an OpenID Connect extension referred to as Single Sign-Off was proposed, which revokes access to all accounts associated with a hijacked account.

Zhang et al. (2019b) investigated click interception practices on the Web present in the Alexa top 250K websites. The authors found that some websites collude with third-party scripts to hijack user clicks for monetization, such as online advertisements. It was also discovered that users can be exposed to malicious content (e.g., scamware) through click interception.

Mirheidari et al. (2020) presented a large-scale study of Web Cache Deception (WCD) in the wild. The study covered 340 high-profile websites from the Alexa top 5K. The authors revealed that WCD vulnerabilities that leak privacy information can be leveraged to conduct severe Web application attacks (e.g., leakage of personal information or security tokens). Additionally, many websites remain susceptible to this attack up to two years after the public disclosure of a WCD vulnerability.

7.1.2. Examining the Web

Stock et al. (2017) examined client-side Web security by studying the Internet Archive. In particular, the authors analyzed code and header information of the 500 most relevant pages for each year between 1997 and 2016. From this data, key client-side technologies (e.g., JavaScript, Flash, Java, and Silverlight) were identified in order to subsequently assess any associated vulnerabilities. The results of this study revealed that client side injection vulnerabilities have risen. Additionally, sites that use HTTPonly cookies are more likely to be vulnerable to XSS. Lauinger et al. (2017) studied the use of outdated JavaScript libraries and found that 37% of the websites evaluated use at least one library with known vulnerabilities. Many of these websites included libraries in an ad hoc and transitive way, which can result in having inconsistent versions of a library in different parts of the same document. It was also discovered that libraries included via code pertaining to adds and tracking are more likely to be vulnerable.

Ikram et al. (2019) performed a large-scale study of dependency chains in the Web. This study was conducted on Alexa top 200K websites and resulted in the following findings: (1) around 50% of first-party websites render content that they did not directly load, (2) 84.91% of websites have short dependency chains (below 3 levels), and (3) 1.2% of the third-party domains that first-party websites import resources from were identified as suspicious by VirusTotal. Chapuis et al. (2020) conducted a large-scale longitudinal study of the use of the Subresource Integrity (SRI) recommendation on the Web. The authors analyzed around 3B URLS crawled on the Web over 3.5 years. The results indicated that SRI is adopted in only around 3.4% of these websites. Moreover, such adoption is highly influenced by the practices of popular library developers and CDN operators. Additionally, the authors noted that SRI is integrated manually, which is a practice that can lead to errors and lacks scalability.

Invernizzi et al. (2016) studied blackhat cloaking techniques that target browsers, networks, or contextual cues in order to detect organic visitors. They first investigated the capabilities of prominent cloaking services by studying several IP blacklists containing over 50M addresses that are tied to the top five search engines, 10 antivirus and security crawlers. As a result, an anti-cloaking system was proposed that could identify split-view content returned to two or more distinct browser profiles with 95.5% accuracy. Kharraz et al. (2018) proposed SurveyLance, a system that leverages machine learning to detect online survey scams. After SurveyLance identified these scams, their capabilities were investigated, their impact on exploited users was quantified, and the connections between the entities involved in them were mapped out. As a result, the authors uncovered 8,623 unique websites involved in online survey scams and that a large number of such scams are easily reachable though the Alexa top 30K websites.

Das et al. (2018) proposed the first large-scale measurement of smartphone sensor API usage and stateless tracking on the mobile Web. The authors developed OpenWPM-Mobile, which is an extension of the OpenWPM Web privacy measurement tool. OpenWPM-Mobile identified that 3,695 of the Alexa top 100K websites accessed one or more sensor APIs. Fingerprinting attempts on mobile platforms were also detected where 63% of the JavaScript programs that accessed motion sensors also engaged in browser fingerprinting. Finally, a large disparity between the intended and actual use of the device sensors was uncovered.

7.2. IoT

Several technical challenges impede data-driven study of Internet of Things (IoT) insecurities at large, including the excessive diversity of IoT devices coupled with their Internet-wide deployments, the lack of IoT-relevant empirical data and the shortage of IoT-specific actionable attack signatures (Mangino et al., 2020). Therefore, researches often deploy specific honeypots to study the threats toward these devices/services. Another approach is running an extensive number of off-the-shelve devices in a lab environment or using a large number of emulated IoT firmware and malware binaries to get insights (Bou-Harb et al., 2016; Dib et al., 2021; Pour et al., 2022; Khoury et al., 2022). Internet scanning on popular IoT-related ports is employed to collect service banners and certificates, as well as to assess vulnerabilities at scale. Further, passive measurement techniques such as an Internet telescope (darknet) (Pour and Bou-Harb, 2019; Torabi et al., 2020), as shown in Figure 7, and ISP level data are valuable resources that require preliminary IoT discovery and fingerprinting ((Feng et al., 2018; Guo and Heidemann, 2018; Perdisci et al., 2020; Saidi et al., 2020; Yan et al., 2019; Yang et al., 2019; Yu et al., 2020)) followed by the analysis of vulnerabilities/threats. Further, there are studies that focused on the automated large scale analysis of IoT firmware (Costin et al., 2016; Zhang et al., 2019a) and malware analysis (Alrawi et al., 2021; Cozzi et al., 2018) to find vulnerabilities in IoT devices. Table 5 summarizes and contrasts the examined publications from a range of perspectives.

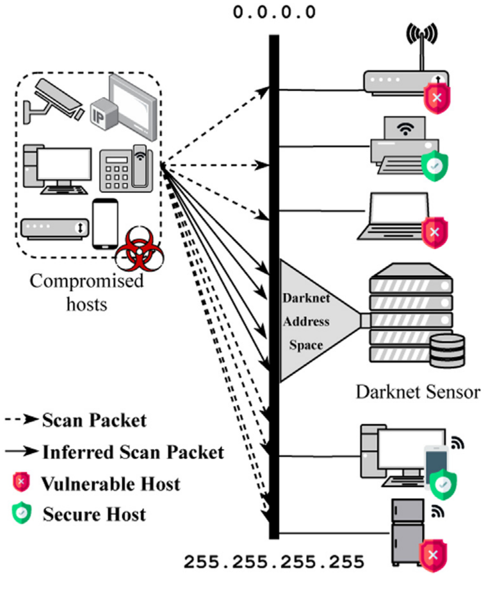

Fig. 7. Darknet (Internet Telescope)

7.2.1. IoT vulnerability Assessment

Scanning the entire Internet can provide significant information for identifying deployed IoT devices with known vulnerabilities. Zhao et al. (2022) presented a large-scale systematic study on the vulnerability of 1,362,906 IoT devices of six different device types. This ten-month-long study provided several insights: (1) 28.25% of the devices suffer from at least one N-days vulnerability, (2) 2,669 of these vulnerable devices may have been compromised by botnets, and (3) 88% of MQTT servers have no password protection. Hastings et al. (2016) measured the actions taken by end users and vendors over time in response to vulnerabilities in network devices associated with weak keys. The authors analyzed public Internet-wide TLS scans performed between 2010 and 2016. From these scans, 81M distinct RSA keys were extracted. Their analysis showed that several vendors never produced a patch and that the number of vulnerable hosts has increased since 2012. Samarasinghe and Mannan (2019) found that some devices still have significant weaknesses, such as the use of known private keys, DDLv3, MD5-RSA, and RC4. The authors compared their findings with the results of their 2016 study that encompassed device vulnerabilities based on certificates and protocol parameters. Kumar et al. (2019) provided

Table 5
Summary of the IoT papers

Topic	Papers	Measur. Method	Analysis Method	Scope	Measurement Size
IoT vulnerability Assessment	Zhao et al. (2022)	Hybrid	Heuristic	Internet-wide	1,362,906 IoT devices
	Hastings et al. (2016)	Passive	Heuristic	Internet-wide	81M distinct RSA keys
	Samarasinghe and Mannan (2019)	Passive	Heuristic	Internet-wide	~8M devices
	Kumar et al. (2019)	Hybrid	Heuristic& ML	Internet-wide	83M devices
	Xu et al. (2018)	Hybrid	Heuristic	Internet-wide	~320K IP cameras
	Ren et al. (2019)	Hybrid& Crowd sourcing	Heuristic	Campus-level	81 devices
	Alrawi et al. (2019)	Passive	Heuristic	Lab setting	45 devices
	Fernandes et al. (2016)	Passive	Heuristic	Lab setting	499 SmartThings apps
	Celik et al. (2018a)	Passive	Heuristic	Lab setting	230 SmartThings apps
	Celik et al. (2018b)	Passive	Heuristic	Lab setting	65 SmartThings apps
	Huang et al. (2020)	Crowd Source & Passive	Heuristic	Internet-wide	54,094 smart home devices
IoT Botnet Analysis	Antonakakis et al. (2017)	Hybrid	Heuristic	Internet-wide	~269K IPs
	Herwig et al. (2019)	Hybrid	Heuristic	Internet-wide	10M unique bot keys
	Tanabe et al. (2020)	Hybrid	Heuristic	Internet-wide	~60k IoT malware samples
	Griffioen and Doerr (2020)	Hybrid	Heuristic	Internet-wide	7,500 IoT honeypots
	Vervier and Shen (2018)	Hybrid	Heuristic	Internet-wide	1.5M unique IP addresses
	Çetin et al. (2019)	Hybrid & Crowd source	Heuristic	ISP-level	~2K IP addresses& 76 subscribers
	Pour et al. (2020)	Hybrid	ML& Statistics	Internet-wide	3.6 TB of darknet traffic
	Torabi et al. (2018a)	Passive	Heuristic	Internet-wide	5TB darknet data
	Pa et al. (2016)	Passive	Heuristic	Internet-wide	87 IP addresses
ICS and SCADA	Fachkha et al. (2017)	Hybrid	Statistics	Internet-wide	50GB darknet data
	Mashima et al. (2019)	Passive	Heuristic	Internet-wide	6GB of ICS network traces
	Husák et al. (2018)	Passive	Heuristic	Internet-wide	16.8 TB of darknet data
	Dahlmanns et al. (2020)	Active	Heuristic	Internet-wide	2k IP addresses
	Vasilomanolakis et al. (2016)	Active	Heuristic	Lab setting	/24 subnet
	Nawrocki et al. (2020)	Passive	Heuristic	ISP and IXP level	329K ICS packets

a large-scale empirical analysis of IoT devices in real-world homes. User-initiated network scans yielded 83M devices. Subsequently, the author investigated security aspects of such devices, including their open services, their weak default credentials, as well as their vulnerabilities to known attacks. It was found that nearly half of the TP-Link home routers have guessable passwords in Eastern Europe and Central Asia. Xu et al. (2018) studied the security of IP cameras without any password protection. They conducted a large-scale empirical investigation of these IP cameras based on insecam.org, the largest online directory for IP cameras. After monitoring the site for 18 days, a comprehensive characteristic analysis of IP cameras was performed. The authors uncovered hidden hosts and services on the internal network where a susceptible IP camera is placed, and then conducted a vulnerability analysis.

Another method involves testing a number of commercially available devices in a laboratory setting. Ren et al. (2019) conducted a set of controlled experiments (34,586) comprising 81 consumer IoT devices in the US and UK in order to investigate the potential privacy risks associated with consumer IoT devices. The authors found that 72 out of 81 devices have at least one destination that is not a first party (i.e., belonging to the device manufacturer). In addition, 56% of the US devices and 83.8% of UK devices where shown to have contacted destinations outside their region. Alrawi et al. (2019) proposed a methodology to analyze the security properties of home-based IoT devices. To identify neglected research areas, 45 devices were evaluated.

Regarding the expansion of app-enabled IoT devices, Fernandes et al. (2016) proposed an in-depth empirical security analysis of an emerging smart home platform. The authors analyzed the SmartThings platform and 499 SmartApps. They discovered two design flaws that lead to significant overprivilege in SmartApps. In fact, it was found that 55% of the apps in the store are overprivileged. Additionally, once the app is downloaded, it is granted full access to the device, even though only limited access is needed. Moreover, four proof-of-concept attacks were constructed that secretly plant and steal door lock codes, disable the vacation mode of a home, and induce a fake fire alarm. In a similar vein, Celik et al. (2018a) presented SAINT, a static taint analysis tool for IoT applications that evaluates the use of sensitive information. Initially, this tool translates IoT source code into Immediate Representations (IR). Then, it performs static analysis to identify sensitive data flows. The authors evaluated SAINT on 230 SmartThings apps and found that 60% of them include sensitive data flows. They also demonstrated their tool on IoTBench, a novel open-source test suite that contains 19 apps with 27 unique data leaks. Celik et al. (2018b) proposed SOTERIA, a static analysis system that validates whether an IoT app or IoT environment adheres to security and functional properties that were identified. This system first translates IoT source code into IR. Next, it extracts a state model from the IR to subsequently verify if this model contains the desired properties. The authors evaluated their system on 65 SmartThings apps and found that 14% of them violated 29% of the properties.

Some efforts rely on crowdsourcing techniques in order to gain knowledge about consumer devices on local networks that are inaccessible via Internet scanning. Huang et al. (2020) developed IoT Inspector, an open-source tool that allows users to observe traffic from smart home devices on their own home networks. In fact, by allowing users to download IoT Inspector, the authors were able to collect (and analyze) labeled network traffic from 54,094 smart home devices. Their analysis showed that many device vendors, such as Amazon and Google, use outdated TLS versions and send unencrypted traffic.

7.2.2. IoT Botnet Analysis

Designing and deploying honeypots is one of the most valuable passive sources for IoT botnet detection and investigation. Pa et al. (2016) proposed IoTPOT, a novel honeypot and sandbox that analyzes Telnet-based attacks against several IoT devices that run on CPU architectures such as ARM, MIPS, and PPC. After evaluating the results of the honeypot and the captured malware samples, the authors showed that at least five distinct DDoS malware families targeted Telnet-enabled IoT devices. It was also revealed that one of the families has evolved to target more devices with different CPU architectures (as many as nine). For six months, Vervier and Shen (2018) used low and high-interaction IoT honeypots to observe the diversity and sophistication of IoT botnets. The authors also discovered that although Mirai is still prevalent, it coexists with other botnets such as Hajime and IoT Reaper. In addition, these botnets were found to be packed with software vulnerabilities that target specific devices in order to have a higher infection rate. Tanabe et al. (2020) presented a comprehensive analysis of IoT botnets based on 23 months of data gathered via honeypots and monitoring botnet infrastructure. The authors focused on three dominant families: Mirai, Bashlite, and Tsunami. Their analysis showed that bots are immediately used after capture and then abandoned. Though IoT botnets appear less advanced than Windows-based botnets, it was discovered that the disposable nature of IoT botnets makes them resistant to blacklisting and C&C take down. Griffioen and Doerr (2020) provided a comprehensive view on Mirai and its many variants.They leveraged 7,500 IoT honeypots as well as the design flaw in Mirai's random number generator to gather insights into Mirai infections worldwide. As a result, botnet networks and particular malware strains were found to be tightly coupled. Additionally, the authors discovered that IoT botnets are not self-sustaining; thus, without continuous pushes from bootstrapping, Mirai and its variants would die out.

The darknet (Internet telescope) is an additional valuable source of data that enables the identification of infected IoT bots through their detected Internet-scanning activity observed on the darknet (Pour et al., 2019, 2021). Antonakakis et al. (2017) proposed a seven-month retrospective analysis of the Mirai botnet's growth. The authors showed its peak of 600K infections and the history of its DDoS victims. Additionally, they analyzed how the botnet emerged, which devices were affected, and how its variants have evolved. Surprisingly, novice malicious techniques were found to compromise enough low-end devices to threaten even the best-defended targets. Finally, several technical and non-technical interventions were recommended to address such risks. Regarding fingerprinting infected IoT devices, Pour et al., 2020; Pour et al. (2019) leveraged macroscopic, passive empirical data provided by network telescopes to shed light on the evolving IoT threat landscape. The authors aimed at classifying compromised IoT devices from one-way network traffic. By analyzing 3.6 TB of darknet traffic along with leveraging active scanning and machine learning techniques, their approach effectively uncovered 440K compromised IoT devices and 350 IoT botnets that were active in the wild. Torabi et al. (2018a) analyzed over 5TB of passive measurements to identify compromised IoT devices as well as those targeted by DoS attacks. Their evaluation exposed 28K compromised IoT devices. They also revealed unreported malware variants that specifically target IoT devices.

Herwig et al. (2019) focused on Hajime IoT botnet by crawling its P2P infrastructure network. The authors found that there are more comprised IoT devices than previously reported. They also discovered that churn is high among IoT devices and that new exploits can quickly increase the size of an IoT botnet.

Çetin et al. (2019) provided an empirical study of IoT malware cleanup in the wild. The authors specifically tried to remove Mirai infections within a medium-sized ISP network. Remediation rates were measured, which revealed that notifying infected customers remediates 92% of infections within 14 days. It was also found that only 5% of the customers who were notified suffered from another infection within five months.

7.2.3. ICS and SCADA

Generally, Industrial Control Systems (ICS) and SCADA devices can be considered subclasses of IoT devices; thus, the same large-scale measurement techniques are often employed by the researchers. However, we chose to separate them into different subsections to put more emphasis on this topic due to their importance and prevalence in the critical infrastructure systems (e.g., Cyber Physical Systems (CPS)).

Mashima et al. (2019) analyzed six months worth of network traces that were collected in low-interaction smart grid honeypot systems deployed in different regions on the Amazon cloud platform. In addition, they discussed the observed attack patterns, as well as other factors such as the correlation among different locations and the dynamics in access sources. Vasilomanolakis et al. (2016) proposed a novel honeypot that detects multi-stage attacks targeting ICS networks. After detecting an attack, the honeypot generates signatures so that Intrusion Detection Systems (IDSs) can prevent attacks with the same signatures. The authors showed that their honeypot detects attacks with good accuracy and that the Bro IDS can successfully use their signatures to prevent future attacks.

Further, Fachkha et al. (2017) exploited passive monitoring from a network telescope with the aim of building broad notions of real CPS maliciousness. The authors inferred and characterized large-scale probing activities that target more than 20 diverse and heavily-employed CPS protocols. Their analysis revealed the presence of 33K probes that target CPS protocols, of which 74% were persistent throughout the analysis. Husák et al. (2018) presented an empirical Internet-wide study on malicious activities generated from and targeted towards critical sectors. The authors leveraged 16.8 TB of darknet data to infer probing activities and DDoS backscatter in critical sectors. In fact, they observed more than 11K probing machines and 300 DDoS attack victims hosted by critical sectors. They also provided insights related to the maliciousness of various business sectors, including the financial sector.

Dahlmanns et al. (2020) leveraged active scanning of the IPv4 address space to study whether Internet-facing OPC UA appliances are configured securely. The authors observed problematic security configurations, such as missing access control, disabled security functionalities, or deprecated cryptographic primitives. Additionally, they found several hundred devices that share the same security certificate, which allows for impersonation attacks. Nawrocki et al. (2020) uncovered unprotected inter-domain ICS traffic at an IXP and ISP level. They leveraged data from honeypots and scans to provide an Internet-wide, in-depth view on ICS communications. They were also able to classify industrial and non-industrial ICS traffic based on cross-correlations with other data sources. In addition, their study showed that ICS systems are controlled remotely without any protective mechanisms, which harms the industrial infrastructures and the Internet.

7.3. Wireless and Mobile

Over the last 20 years, wireless Local Area Networks (LANs) based on 802.11 have grown ubiquitous in workplace and college settings. Additionally, only a small amount of research has been done on SIM-enabled wearable traffic on ISP networks in order to acquire customer insights and better understand the traffic dynamics. The study of mobile devices is becoming more important due to several factors, which includes, the introduction of technologies that enable smart cities (e.g., 5G, 6G, Helium (2021), etc.), Internet connectivity of IoT devices with limited resources (e.g., long-range wide-area network), the continued growth of vehicles connected to the Internet, and the increased demand for micromobility platforms. There is relatively limited research in this topic because of the difficulty in data gathering and the scarcity of data.

Andrade et al. Andrade et al. (2017) conducted network-scale measurements of over 1B radio connections using 1M connected cars on a production cellular network. They found that connected automobiles have a variety of features, including some that are comparable to ordinary smartphones (e.g. overall diurnal pattern) and IoT devices.

Kolamunna et al. (2018) analyzed the network traffic of tens of thousands of SIM-enabled customers of a major European mobile ISP. Over the course of a five-month observation period, they saw

Table 6
Summary of the Phishing papers

Topic	Papers	Measur. Method	Analysis Method	Scope	Measurement Size
Phishing Ecosystem	Tian et al. (2018)	Hybrid	ML	Internet-wide	657K domains, 224M DNS records
	Ho et al. (2019)	Passive	ML	Organization-level	113M emails from 92 organizations
	Han et al. (2016)	Hybrid	Heuristic	Internet-wide	643 unique phishing kits over 5 months
	Oest et al. (2020b)	Passive	Heuristic	Internet-wide	404k phishing URLs, data from paypal
	Zhang et al. (2021)	Hybrid	Heuristic	Internet-wide	112k phishing websites reported to APWG (over 14 months)
	Oest et al. (2019)	Active	Heuristic	Internet-wide	10 anti-phishing platform over 1.5 years. 2,380 phishing sites
	Oest et al. (2020a)	Hybrid	Heuristic	Internet-wide	4k phishing URLs, 2.8k PayPal-branded phishing websites, data from 10 anti-phishing entities
	Maroofi et al. (2020)	Hybrid	Heuristic	Internet-wide	105 phishing websites, data from 7 anti-phishing entities
Victim's Cognitive Behavior	Van Der Heijden and Allodi (2019)	Passive	ML, Statistics, Heuristic	Organization-level	115k emails, 11k alerts (over 10.5 months)
	Simoiu et al. (2020)	Passive	Heuristic	Gmail ecosystem	1.2B phishing and malware attacks, 17M weekly targeted users

a 9% growth in SIM-Enabled wearable users. However, only 34% of these users actually produce any network transaction. They also found that sim-enabled wearable users are much more active in terms of mobility, data consumption, and frequency of app usage.

Zhang et al. (2020) presented a large-scale characterization of the Fake Base Station (FBS) spam ecosystem by leveraging three-months worth of real-world FBS detection results. The authors uncovered the characteristics of FBS spammers, such as business categories, temporal patterns, and spacial patterns. The FBS ecosystem's organization and evasion detection were also investigated. It was found that over 75% of the messages are associated with illegal or fraudulent businesses. Further, 7,884 FBS campaigns were identified via the contact information of spammers.

8. Threats, eCrime and Attacks

In this section, we elaborate on several topics related to threats, namely, Phishing, Ransomware, RAT, Cryptojacking, Spam and Denial of service attacks. We also offer a summary of the papers relevant to these topics.

8.1. Phishing

Phishing attacks are a form of social engineering or online identity theft in which an attacker tricks people into giving away private information. In 2016, the Anti-Phishing Working Group (APWG) recorded over 1.2 million phishing attacks. Such attacks can be carried out via domain squatting techniques such as typosquatting. Typosquatting is the deliberate registration of a domain name that is a misspelling of a well-known domain name. In addition, as characters from different languages can look like each other, IDNs have been used to impersonate popular domains for phishing purposes, i.e., IDN homograph attacks. We direct readers to subsection 4.1.6, which contains a review of IDN-related works, particularly (Hu et al., 2021; Suzuki et al., 2019) and typosquatting-related works such as (Agten et al., 2015; Quinkert et al., 2020; Szurdi et al., 2014) that focus on phishing.

Despite several mitigation efforts such as applying email filters, blocking and taking down phishing websites, or using browser plugins that notify users when they are being redirected to potentially malicious pages, phishing attacks are far from being solved. Table 6 summarizes and contrasts the articles that will be discussed.

8.1.1. Phishing Ecosystem

Tian et al. (2018) introduced SquatPhi, an end-to-end measurement framework to detect squatting phishing domains at both the domain and content levels. The authors tested 224M DNS records for five types of squatting and discovered 657K domains that are potentially impersonating 702 popular brands (e.g., Facebook, Paypal). They also built a machine learning classifier to detect phishing sites. The classifier identified 1,175 squatting pages (857 web pages and 908 mobile pages). It was revealed that over 90% of the phishing websites successfully evaded well known blacklists such as VirusTotal, PhishTank, and eCrimeX.

Ho et al. (2019) investigated lateral phishing attacks (i.e., phishing emails are sent to other users using a compromised enterprise account). The authors leveraged a large dataset consisting of emails sent by more than 100M employees from 92 enterprise organizations. It was found that 14% of randomly sampled organizations experienced at least one incident within a seven-month timespan. It was also found that targeted messages are used in 7% of attacks.

In a more advanced measurement technique, Han et al. (2016) designed a web honeypot system to attract real attackers into installing phishing kits in a compromised web application. According to the findings, 643 unique phishing kits were discovered. Phishing kits were only operational for less than 10 days after being installed, where most of them only gather a small amount of user credentials. The authors also found that attackers use infected websites to load a vast number of phishing kits in a shot-and-forget strategy in order to quickly switch to new phishing pages as the old ones became blacklisted. In a continuation, Oest et al. (2020b) introduced Golden Hour, a framework that passively measures victim traffic to phishing pages. The authors found that the typical phishing attack lasts 21 hours from the first to the last victim's visit, and that anti-phishing entities identify each attack nine hours after the first victim's visit. Because attackers make a lot of money during this period, it was dubbed golden hours.

Several research efforts utilized well-known crowdsourced phishing datasets (e.g., APWG). Zhang et al. (2021) presented CrawlPhish, a framework that collects the source code of reported phishing websites and automatically detects the client-side cloaking evasion techniques that are used. CrawlPhish analyzed 112,005 websites over a fourteen-month period and uncovered 35,067 (31.3%) cloaking usages. It was also revealed that, by the end of the study, the cloaking usages increased from 23.32% to 33.70%.

In addition, a number of studies concentrate on anti-phishing blacklists and evasion tactics. Oest et al. (2019) presented PhishFarm, a scalable framework to methodically test the resilience of anti-phishing entities and browser blacklists. It was revealed that current infrastructures allow some phishing sites to go unnoticed, leaving them accessible to victims. It was also discovered that phishers who employ simple cloaking techniques have reduced the likelihood of being blacklisted by 55% on average. Oest et al. (2020a) proposed PhishTime, a framework that continuously identifies unmitigated phishing websites in the wild and generates longitudinal experiments to measure the given ecosystem's protections. In six experimental deployments over a nine-month period, the authors reported 2,862 new phishing websites and evaluated the performance and consistency of blacklists. The results revealed that the average response time against unsophisticated phishing websites is 55 minutes, whereas the average response time of phishing websites that use common evasion techniques is 2 hours and 58 minutes. Maroofi et al. (2020) looked at how anti-phishing entities fare against three advanced anti-analysis approaches based on human verification: Google reCAPTCHA, alert boxes, and session-based evasion. The authors found that only Google Safe Browsing detected all the reported URLs protected by alert boxes. It was also identified that none of the entities could detect phishing URLs armed with Google reCAPTcha, making it currently the most effective phishing content protection solution accessible to malicious actors.

8.1.2. Victim's Cognitive Behavior

Van Der Heijden and Allodi (2019) conducted a measurement study on the effect of cognitive vulnerability triggered in phishing emails. The authors collected 115,698 reported emails between February 1st, 2018 and 15 December, 2018 from Org (a large financial organization in Europe with more than 8M customers). The authors also gathered 11,936 alerts for malicious links that Org's phishing response team detected. The analysis results revealed that consistency and scarcity exercised a clear positive effect on the number of generated clicks. Simoiu et al. (2020) analyzed over 1.2B email-based phishing and malware assaults against Gmail users to see what factors put people at risk of being attacked. The authors discovered that assault campaigns are often short-lived and, at first look, appear to indiscriminately target consumers on a worldwide scale. However, by modeling the distribution of targeted individuals, it was discovered that a person's demographics, geography, email usage behaviors, and security posture had a major impact on the chance of an attack.

8.2. Ransomware

Some Papers were dedicated to understanding ransomware attacks in scale from ransom payment aspects. Huang et al. (2018) created an end-to-end framework to analyze cybercriminal operations that have adopted Bitcoin as their payment channel. The authors tracked the financial transactions, from the moment victims acquire bitcoins to when operators cash them out. Akcora (2020) proposed a novel, efficient, and tractable framework that can automatically predict new ransomware transactions in a ransomware family. The authors found that their approach has a higher accuracy than existing heuristic and ML-based procedures. Meanwhile, other papers developed novel ransomware detection methods which lead to uncovering new insights. Kolodenker et al. (2017) developed PayBreak, an automated proactive defense mechanism against ransomware. This approach relies on the fact that secure file encryption uses hybrid encryption with symmetric session keys on the victim computer. The authors evaluated PayBreak against 20 successful families of real-world ransomware and were able to restore the files of 12 of these families. Moussaileb et al. (2018) proposed a graph-based ransomware detection. The authors leveraged more than 700 active samples of ransomware that were analyzed in their environment. It was discovered that the per-thread file system traversal is enough to reveal malicious activities.

8.3. RAT

RATs are a type of malware that grant an attacker direct, interactive access to a victim's computer, which enable the adversary to steal personal information, spy on the victim in real-time via the camera and microphone, or harass the victim verbally via the speaker. Farinholt et al. (2017) studied DarkComet, a popular commercial RAT. By utilizing 19,109 malware samples, the authors monitored the behaviors of each of the samples in their honeypot environment. The operator behavior on 2,747 interactive sessions that was captured during the experiment was reported. Operators were found to engage with victims through remote desktop, capture video, audio, and keystrokes. Operators were also found to exfiltrate files and credentials during these sessions. Rezaeirad et al. (2018) reported the attackers and the victims of two popular RATs, namely, njRAT and DarkComet. The authors first leveraged VirusTotal to collect all instances of these RATs and identify the domain names of their controllers. The expired domains were then extracted to determine the victims of such attacks. The results showed that over 99% of the 828,137 IP addresses that were connected to the author's sinkhole were not real victims. Farinholt et al. (2020) presented a longitudinal study of the DarkComet RAT ecosystem. The authors leveraged 6,620 victim log databases from DarkComet controllers during a five-year period and proposed novel techniques to track RAT controllers across hostname changes. The analysis showed that there were at least 57,805 victims of DarkComet over this period. It was also found that 69 new victims were being infected every day, where many of their keystrokes, activities, and webcams have been caught, recorded, and observed.

8.4. Cryptojacking

Illicit crypto-mining enables criminals to mine cryptocurrency using resources stolen from victims. Apart from binary-based crypto-mining malware that exploits the host device's resources, the advent of memory-bound cryptocurrencies such as Monero and Coinhive has made the deployment of mining code in browser-based JavaScript a viable alternative to specialized mining rigs. Certain websites mine cryptocurrency without the customers' awareness in lieu of showing advertisements. This strategy of monetizing websites has enticed both website owners and criminals to look for new revenue streams.

Some research focused on the top Alexa websites to determine the prevalence of cryptojacking on the web. Rüth et al. (2018) inspected 137M.com/.net/.org and Alexa top 1M domains for mining code, and proposing a new fingerprinting method. It was revealed that the prevalence of cryptojacking is less than 0.08% in such domains. Coinhive was identified as the largest cryptojacking provider, which is used by 75% of the mining sites. Further investigation clarified that the authors contributed 1.18% of the blocks in the Monero blockchain, which is worth 150K USD per month. Similarly, Konoth et al. (2018) performed analysis on Alexa top 1M websites. By leveraging 28 Coinhive-like services, 20 active cryptomining campaigns were identified. The authors also discussed how current heuristics based on CPU usage as well as blacklisting approaches are insufficient. Finally, MineSweeper, a detection technique that is based on intrinsic characteristics of cryptomine code, was proposed. Bijmans et al. (2019a) performed an Internet-scale crawling using the Minesweeper tool and investigated about 20% of 1,136 TLDs (48.9M websites). Overall, the authors identified 204 cryptojacking campaigns in the wild. It was also discovered that attackers spread cryptojacker scripts over a large number of domains by using third-party software (e.g., WordPress). Musch et al. (2019b) proposed a three-phase analysis approach in the Alexa 1M websites. Their findings showed that cryptojacking is common, where 1 out of 500 sites host a mining script. The authors then performed a measurement of code characteristics, an estimate of expected mining revenue, and an evaluation of the current blacklist-based countermeasures to gain insight into the cryptojacking characteristics. Hong et al. (2018) proposed CMTracker, a behavior-based tool with two runtime profilers that automatically detects and tracks cryptocurrency mining scripts and their related domains. By leveraging 853,936 popular webpages, CMTracker successfully discovered 2,770 unique cryptojacking samples. The authors found that attackers frequently update their attack domains on the order of days to evade detection. Further, attackers use many other evasion techniques such as code obfuscation and limiting their CPU usage.

Pastrana and Suarez-Tangil (2019) studied crypto-mining malware during a twelve-year period by investigating about 4.5M malware samples (1.2M malicious miners) pulled from malware feeds such as VirusTotal, Palo Alto network, virus share, and others. The authors employed static and dynamic analysis to extract information from the samples, such as wallet identifiers and mining pools. The results revealed that Monero is currently the preferred currency used by criminals who have accumulated a large amount of earnings from illicit mining.

Bijmans et al. (2019b) reported a new attack vector in cryptojacking where cyber criminals use a firmware vulnerability in MikroTik routers to rewrite outgoing user traffic and embed cryptomining code in every outgoing web connection. The authors monitored the activities of Netflows recorded in a Tier 1 network, semiweekly crawls, and network telescope traffic over a ten-month period. The findings showed that 1.4M routers were under the adversaries' control, which consisted of approximately 70% of all Mikrotik devices deployed worldwide.

8.5. Spam

Spam emails have an impact on millions of users, waste important resources, and are a drain on email systems. Spam has long been a popular technique for criminals to do unlawful operations on the Internet, such as stealing sensitive information, selling counterfeit goods, and distributing malware, among other maliciousness. As a result, spam emails include valuable cyber security information.

Liao et al. (2016) studied how long-tail Search Engine Optimization (SEO) spam is implemented on cloud hosting platforms. The authors leveraged 3,186 cloud directories and 318,470 doorway pages to characterize long-tail SEO spam's abusive behavior. The results showed that 6% of the doorway pages made it into the top ten search results for poisoned long-tail keywords. It was also discovered that these doorway pages are able to monetize traffic and manage the cloud platform's countermeasures. Dinh et al. (2015) proposed a framework that detects, analyzes, and investigates spam campaigns from 678,095 messages and 91,370 unique IP addresses. This framework not only identified spam campaigns in real-time but also labeled, scored, and collected a multitude of information about them. Gupta et al. (2018b) leveraged 23M posts over five different OSNs to characterize cross-platform spam campaigns that use phone numbers for monetization purposes. The authors found that although Indonesian campaigns create the highest volume ($3.2M posts), only 1.6% of the accounts that are involved in such campaigns have been suspended. Additionally, if intelligence was shared throughout the OSNs, roughly 35K victims and $8.8 million may have been saved. Gupta et al. (2018a) studied multiple spam campaigns, including tech support scams. The authors leveraged approximately 70M tweets from 2.5M user accounts that contain over 5,786 phone numbers during a fourteen-month period to measure the properties of these campaigns. First, the data collection technique is able to identify tweets containing phone numbers. Second, they were able to cluster tweets that are a part of Outgoing Phone Communication (OPC) attack campaigns. It was also found that only about 3.5% of the accounts involved in the top 15 spam campaigns were suspended by Twitter.

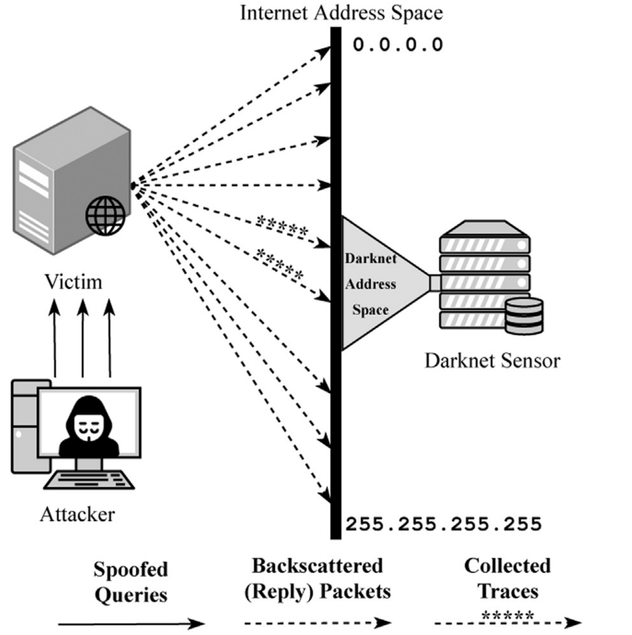

Fig. 8. Leveraging darknet (Internet Telescope) for DDoS-related measurements. (adapted from Fachkha and Debbabi (2015))

8.6. Denial of Service

DoS attacks are becoming one of the biggest threats to Internet stability and reliability. Moreover, DDoS attacks have become commoditized, allowing abusive subscribers to cheaply extort, harass, and threaten businesses. In addition, amplification DDoS attacks have gained increasing popularity by abusing so-called amplifiers (or reflectors) to exhaust the bandwidth of a victim. One common DDoS-related measurement technique is leveraging backscatters at the darknet (Fachkha et al., 2015) as shown in Figure 8. BGP blackholing is one of the main mitigation techniques, which discards traffic bound for the given prefix. Table 7 summarizes and contrasts the papers encompassed herein that study this topic.

8.6.1. Investigating the DoS Ecosystem

Blenn et al. (2017) used passive traffic from the /16 network telescope to research the attack size and duration of DDoS attacks over the span of 26 months, which revealed new insights into emerging challenges and the progression of adversarial tactics over time. Despite the fact that the media reports new milestones of DDoS attack sizes on a daily basis, the vast majority of such attacks are small and can be defended with readily accessible technologies and applications. Jonker et al. (2017) introduced a new framework for the macroscopic characterization of DoS ecosystem (e.g., attacks, attack targets, and mitigation behaviors) by correlating diverse sets of the global Internet infrastructure's measurement data. The authors found that more than one third of the total number of IPv4 network blocks that are operating on the Internet were targeted by over 20M DoS attacks during the study period. They observed that after an attack, 4.3% of attack targets migrate to a DPS.

Wang et al. (2018a) used both active and passive measuring techniques to provide a complete picture of both attackers and vic-

Table 7
Summary of the DDoS papers

Topic	Papers	Measur. Method	Analysis Method	Scope	Measurement Size
Investigating the DoS Ecosystem	Blenn et al. (2017)	Passive	Heuristic	Internet-wide	/16 Internet telescope data over 26 months
	Jonker et al. (2017)	Hybrid	Heuristic	Internet-wide	1M events targeted at > 2M /24 network blocks
	Wang et al. (2018a)	Hybrid	Heuristic	Internet-wide	50k DDoS from 674 botnets (7 months)
	Moura et al. (2018)	Hybrid	Heuristic	Internet-wide	Leveraging 9k RIPE Atlas probes, +15K vantage points,.nl and the root DNS zone
	Abhishta et al. (2019)	Active	Heuristic	Internet-wide	NS1 and Dyn DNS service providers, OpenINTEL active DNS
Amplification Attacks	Krämer et al. (2015)	Hybrid	Heuristic	Internet-wide	>1.5M attacks from 21 globally-distributed AmpPot (4 months)
	Krupp et al. (2016)	Hybrid	Heuristic	Internet-wide	1.3M amplification attacks (over 23 weeks), 48 AmpPot instances
	Rossow (2014)	Hybrid	Heuristic	Internet-wide, ISP-level	130 real-world DRDoS attacks, 14 protocols, large ISP (1M users), a /17 and a /27 network telescope
	MacFarland et al. (2017)	Active	Heuristic	Internet-wide	130M DNS domains and 1M unique DNS servers
	Bushart and Rossow (2018)	Active	Heuristic	Internet-wide	entire Ipv4 and 4,170k DNS resolvers
	Rytilahti and Holz (2016)	Hybrid	Heuristic	Internet-wide	47,900 responsive hosts
	Bock et al. (2021)	Hybrid	Heuristic, ML	Internet-wide	184 IP addresses and 1,052 URLs from the Quack dataset
	Moon et al. (2021)	Active	Heuristic	Internet-wide, Lab settings	10K sampled servers using 31 nodes from CloudLab
DDoS Protection	Giotsas et al. (2017b)	Hybrid	Heuristic	ISP/IXP-level	12,940 IP peers
	Jonker et al. (2018)	Passive, Hybrid	Heuristic	Internet-wide	28.14M DoS, 1.30M blackholing using /8 Internet telescope, Ampbot honeypot, BGP data
	Jonker et al. (2016)	Active	Heuristic	Internet-wide (9 major DPS)	Alexa top 1M over 1.5 years
	Dietzel et al. (2016)	Passive	Heuristic	IXP-wide	22,994 blackholing BGP announcements from a large IXP
	Nawrocki et al. (2019)	Passive	Heuristic	IXP-level	590M sampled flows from control & data plane IXP (104 days)

tims. At first malware families (used to launch a range of DDoS attacks) are reverse-engineered and assigned to a known malware family. Then, hosts participating in the specific botnet are enumerated, monitored over time, and their activities are evaluated by connecting with infrastructure infected by that malware family (e.g. the C&C). They characterized 50,704 different DDoS attacks launched by 674 different botnets from 23 different botnet families directly observed in a seven-month period. During this measurement period, 9,026 victim IPs belonging to 1,074 organizations in 186 countries were identified. The authors discovered that the geospatial distribution of the attacking sources follows certain patterns, which enables the accurate source prediction of future attacks.

Moura et al. (2018) conducted a controlled measurement experiment and assessed the resilience of the DNS resolution system during DDoS attacks on authoritative servers to gain insights about the role of DNS caching, retries, and use of multiple DNS recursive resolvers. The authors uncovered that in about 30% of the time, clients do not benefit from caching. However, it was also discovered that DNS caching and retries provide a high degree of client user resilience during DDoS attacks. Abhishta et al. (2019) analyzed two DDoS attack events on NS1 and Dyn on May 16th, 2016 and October 21st, 2016, respectively. They discovered that a significant number of customers who were solely utilizing Dyn's or NS1's Multicast DNS (mDNS) service switched to non-exclusive use following the attacks. Further, the bulk of newly non-exclusive customers began using an mDNS service provider as a secondary DNS to help minimize the chance of downtime.

8.6.2. Amplification Attacks

Krämer et al. (2015) developed AmpPot, a novel honeypot that imitates services that are considered to be vulnerable to amplification attacks, such as DNS and NTP. They observed more than 1.5M attacks between February and May 2015 by deploying 21 globally-distributed AmpPot instances. The researchers found that the vast majority of attacks are brief, and most victims are only targeted once. Using AmpPot, Krupp et al. (2016) carried out a measurement study that took a snapshot of 1,351,852 amplification attacks detected by honeypots over the course of 23 weeks and attribute such attacks to the infrastructures that triggered them. The authors devised a method for embedding a fingerprint on scanners that conduct reconnaissance for amplification attacks, which allows them to trace subsequent attacks back to the scanner. This method lead to the identification of scanners that were involved in 58% of all attacks with more than 99.9% confidence. Further investigation revealed that only 20 scanners were responsible for almost half of the attacks.

Rossow (2014) revisited popular protocols of network services, online games, and P2P filesharing networks to assess their security against Distributed Reflection Denial of Service (DRDoS) attacks. The authors found 14 protocols to be susceptible to bandwidth amplification that could multiply traffic by up to a factor of 4670. Additionally, millions of public hosts could be abused as amplifiers. MacFarland et al. (2017) studied the threat potential associated with DNS amplification attacks that rely on using authoritative servers as amplifiers. The authors found that in certain data sets, fewer than 3.8% of authoritative servers are responsible for the largest amplification factors. The results also showed that rate-limiting is only utilized by 10.23% of servers in practice. Bushart and Rossow (2018) proposed DNS Unchained, an application-layer DoS attack against core DNS infrastructure that uses amplification. The authors carefully chained CNAME records and forced resolvers to perform deep name resolutions to achieve an attack amplification of 8.51. In addition, 178,508 potential amplifiers were identified, where 74.3% of them could be used in such an attack because of the way they cache records with low TTL values. Finally, several countermeasures were suggested for DNS servers to limit the impact of DNS chaining attacks. Rytilahti and Holz (2016) demonstrated how the same features used for NTP-based attacks can be used to obtain a global picture of ongoing attacks on the Internet. The authors found that only a fraction of all the vulnerable services are susceptible to attacks and attack tracking. Many known vulnerable hosts that remained unused because of their small response sizes were also discovered. Bock et al. (2021) demonstrated that non-trivial TCP-based amplification is conceivable, and that it can be more effective than the commonly used UDP-based amplification. The authors found and maximized the efficacy of new TCP-based reflection amplification attacks using a novel application of a recent genetic algorithm, and presented several packet sequences that lead network middleboxes to respond with more packets than what they sent. In addition, the entire IPv4 Internet was surveyed and hundreds of thousands of IP addresses with amplification factors greater than 100x were discovered. Several open questions regarding DoS attacks were also investigated. Finally, network factors that caused some TCP-based attacks to be so effective were discussed.

Moon et al. (2021) presented AmpMap, a lightweight Internet health monitoring service that systematically and continuously quantifies DDoS attacks in order to facilitate mitigation efforts. They found that relying on previous recommendations to block or rate-limit specific queries is not enough, as such recommendations can result in many other amplification-inducing query patterns being missed. Instead, the authors developed an efficient approach to search across the space for protocol headers and servers.

8.6.3. DDoS Protection & BGP Blackholing

BGP blackholing is an operational countermeasure that utilizes the characteristics of BGP to mitigate various forms of maliciousness, such as DoS attacks and spam. BGP blackholing is implemented using the BGP communities attribute, which is a BGP extension that enables the exchange of additional information between BGP peers. BGP blackholing uses a particular set of BGP community tags to request an upstream provider (ISP) or IXP to filter traffic to a certain destination prefix (Giotsas et al., 2017b). Giotsas et al. Giotsas et al. (2017b) were able to uncover that hundreds of networks (including about 50 IXPs and large transit providers) offer blackholing services to their customers, peers, and members. Upon leveraging both active measurements and passive datasets, it was found that blackholing is highly effective. Jonker et al. (2018) conducted a global-scale study to discover the operational deployment of blackholing as a DoS mitigation strategy by comparing DoS attacks with BGP blackholing events. For this purpose, the authors leveraged DoS attacks identified from a network telescope and DoS honeypot, and compared these attacks with a set of blackholing events inferred from BGP routing data over 1100 days. According to the findings, blackholing is a fast and effective way to mitigate such attacks. Additionally, 44% of attacks that are blackholed are mitigated within one minute, and 85% are mitigated within ten minutes.

Leveraging active DNS measurement, Jonker et al. (2016) studied the adoption of cloud-based DPSs by analyzing nine major DPSs globally. The findings showed a 1.24x increase in DPS adoption over 1.5 years, a notable trend in comparison to the overall expansion of the global domain name system. Large hosting providers and domainers are the key drivers of this pattern, as they turn DDoS protection on and off for lots of domains on a regular basis.

Some studies utilize data from large IXPs. Dietzel et al. (2016) looked at the prevalence of blackholing used by the IXP members and its efficacy. The authors collected five-minute snapshots of routing and traffic measurements over a three-month period (from December 2014) from one of the largest IXPs in Europe, which lead to the discovery of 22,994 blackhole BGP announcements. Blackholing was found to succeed in reducing DDoS attack traffic. Nawrocki et al. (2019) correlated 104 days of data and control plane measurements in order to study the collateral damage introduced by blackholing at a big European IXP. The authors discovered that just a third of the 34k apparent blackholing cases are linked to DDoS attack signals. Besides, they realized that blackholing drops just half of unnecessary traffic on average, making it a much less effective method for preventing DDoS attacks than previously thought. In addition, they found 300 blackholing events with traffic containing collateral damage from 1000 detected servers.

9. Cyber Incidents and Real-World Impacts

In this section, we go through three subjects linked to traces of cyber incidents in real world. We begin with an examination of the economics and a study of cybercrime. After that, we go into the aftermath of cyber attacks. Finally, we discuss an assessment of existed security platforms. Refer to Figure 2 for the detailed taxonomy of the application domains. A synopsis of the publications pertinent to these subjects is also provided in this section.

9.1. Economy and Analysis of Cybercrime

Anonymity is a double-edged sword. On one edge it safeguards people's privacy and, in authoritarian societies, fosters freedom of speech. On the other edge, it is used by criminals and even cyberterrorists to conceal their operations. We are facing two identity shielding technologies: (1) the Dark Web and (2) cryptocurrencies. The Dark Web leverages anonymous routing methods (e.g., Tor) to hide users' identities. Cybercriminals also exploit the dark web for the commoditization of instruments and services, which is anticipated to accelerate cybercrime growth. It is also used for illegal advertisement material and hosting C&C servers. Cryptocurrencies (e.g., Bitcoin and Ethereum) enable individuals to perform P2P transactions without a central authority. The Bitcoin blockchain scheme requires people to use an extensive list of aliases to perform transactions safely. This cryptocurrency poses threats to law enforcement as it is decentralized, completely uncontrolled, and both participants in a malicious transaction are shielded behind pseudo-anonymous identities. However, Bitcoin has a property that is unfavorable for cybercriminals: by design, all transactions are public. This helps researchers to retrieve the economic inner workings of cyber-criminal activities through transaction clustering and tracing.

Iliou et al. (2016) presented a crawler capable of traversing both the Surface Web and the Dark Web (i.e. Tor, I2P, and Freenet). The crawler is designed to discover web resources on any given topic, with a particular focus on topics of interest to law enforcement agencies. The experimental examination showed its considerable efficacy, and incorporating an additional classifier further improved effectiveness without posing a bottleneck to its performance. Pastrana et al. (2018) introduced CrimeBot, a tool to scrape underground forums and cybercriminal communities. They created CrimeBB, a dataset of more than 48M posts made from 1M accounts in four different operational forums over the course of a decade. As a case study, they investigated how currencies have evolved over the last 10 years, noted the increase of exchanges involving Amazon gift cards in the last two years, and identified key actors. Lee et al. (2019) introduced MFScope, a Dark Web data collection and analysis platform to study where the money goes from the trading of illicit services and goods by online merchants, as well as how perpetrators acquire their money while minimizing the risk of being tracked. The authors revealed that Bitcoin contains the majority (99.8%) of cryptocurrency addresses among collected addresses, and that more than 80% of Bitcoin addresses are used for illicit intent. By developing a Taint-based Bitcoin flow analysis, around 180M USD was estimated to be the illicit monetary volume. Vu et al. (2020) investigated a dataset of transactions created and completed on the established underground marketplace HackForums. The dataset includes the associated threads and posts from June 2018 to June 2020. The authors observed the growth of users making only one transaction, as well as power-users who make many transactions. Insights were also provided about the sorts of services being exchanged, preferred payment methods, and how users overcome the cold start problem when joining the market without established trust.

Van Wegberg et al. (2018) conducted a measurement study to answer to the main questions regarding the increasing commoditization of cybercrime. The authors found 44K listings and over 564K transactions between 2011 and 2017. It was also discovered that cash-out services contain the most listings and generate the largest revenue. At least 15M was generated between 2011-2017 in overall revenue for cybercrime commodities on online anonymous markets.

Booters are another area of interest to the academics. Karami et al. (2016) analyzed 15 booter services on the dark web and conducted a large-scale payment intervention in partnership with PayPal to assess the impact of this deterrent against their operations. They also shed light on booters technological and economical structure, namely, advertising, attack, hosting, and payment. Krupp et al. (2017) presented several techniques to attribute DDoS amplification attacks to booter services. They proposed a machine learning classifier that leverages features related to a DDoS service, such as the set of reflectors used by that service. In the most challenging real-time attribution scenario, they were able to attribute DNS and NTP attacks with a precision of over 99% and a recall upwards of 69%.

Regarding the ransom payments, Conti et al. (2018) reported the economic impact of ransomware families from the Bitcoin transactions, presented a lightweight framework to analyze Bitcoin addresses managed by malicious actors, and subsequently provided a longitudinal measurement related to twenty recent Bitcoin ransomware instances. The authors began with a set of collected ransom addresses reported by victims and security analysts in the online community, and then used those as initial seeds to identify other addresses that have a similar payment behavior. Further, they considered multi-input transactions and shadow/change addresses. Simoiu et al. (2019) studied the prevalence and characteristics of ransomware attacks on the general population, how users perceive risks and respond to attacks, as well as what proportion of the population pays. The authors developed a proof-of-concept method for risk-assessment that is based on self-reported security habits. By using a representative survey sample of 1,180 American adults, they estimated that 2-3% of the adults were affected over a one-year period (between 2016 and 2017). In addition, the average demanded payment was $530, where only about 4% of the affected users reported paying. Meland et al. (2020) performed a measurement study on the current impact of Ransomware-as-a-Service (RaaS) and the participating actors. They studied RaaS over a period of two years (fall of 2017 to the fall of 2019) within well-known dark web forums and markets. The findings showed that the occurrence of RaaS actually tends to be more modest than suggested in media by security firms.

9.2. Aftermath of Cyber Incidents

Cybercriminals steal login credentials for webmail accounts in order to publish them or sell them on the dark web. Little else is known about what these stolen accounts are utilized for. For this purpose, a number of articles have contributed to a better understanding of how stolen credentials are used. Some studies examine how law enforcement and organizations counteract specific cyber attacks on the Internet, while others examine how individuals patch their computers over time.

Onaolapo et al. (2016) shed some light on the modus operandi of cybercriminals who access stolen Gmail accounts. They developed a methodology that monitors the activity of Gmail users by leaking 100 controlled account credentials on public paste sites and underground forums. Subsequently, they monitored these leaked accounts for a period of seven months. Such monitoring allowed the authors to devise a taxonomy of the behavior of cybercriminals that use stolen Gmail credentials and their malicious activity in order to identify several mitigation techniques. Thomas et al. (2017) presented a longitudinal measurement study of the underground ecosystem that provides stolen credentials and poses a risk to millions of users. The authors found 788K potential victims of off-the-shelf keyloggers, 12.4M potential victims for phishing kits, and 1.9B credential leakages from data breaches. They also found that 7-25% of exposed passwords match a Google account. Peng et al. (2019a) performed an empirical measurement of the transmission and sharing of stolen login credentials. Over the course of five months, they collected more than 179K phishing URLs from 47K live phishing sites. The authors found that third-party servers are often located in a different country as compared to the phishing server, which may create difficulties in taking them down. Golla et al. (2018) presented two user studies on password-reuse attacks by presenting data from 180 users. The first study showed that less than a third of the respondents had intentions to change their passwords. In the second study, the authors synthesized 15 variations of a model notification based on the results from the first study. Then, 588 respondents saw one of 15 variations of the model. The results showed that although the variations' impacts differ slightly, the respondents would still be vulnerable to future password re-use attacks.

Regarding the booter ecosystem, Collier et al. (2019) measured the impact of police interventions (e.g., arrests and website take downs) by looking at usage reports that booters themselves provided and at measurements of reflected UDP DoS attacks. They found that take downs of individual booters precede significant yet short-lived reductions in attacks. They also discovered that the closure of HackForum's booter reduced attacks for 13 weeks globally and the FBI's coordinated operation in December 2018 reduced attacks by a third for at least 10 weeks. Kopp et al. (2019) looked at booter-based DDoS attacks in the wild and the effect of an FBI takedown targeting 15 booter websites in December 2018. The authors purchased Gbps-level attacks against their own infrastructure from four popular booters and collected five months of data to understand spatial and temporal trends of the booters' DDoS traffic. They discovered that the takedown had an immediate impact on DDoS amplification traffic, especially against reflectors, but it has little impact on traffic reaching victims or the amount of attacks detected.

Nappa et al. (2015) leveraged data collected over five years on 8.4M hosts through Symantec's WINE platform and analyzed the deployment of patches on hosts around the world. The authors found that, at most, 14% of vulnerable hosts are patched when exploits are released. The patching rate is affected by user- and application-specific factors such as automated updating mechanisms, which have lower median times to patch.

10. Internet Measurement Impediments and Future Directions

The more intertwined the Internet becomes and the greater the variety of technologies and dependence on numerous components, the more difficult it is for researchers to conduct independent measurements and investigations. To this end, throughout the past few years, there have been a number of studies and workshops devoted to the current difficulties in Internet measuring and related discussions claffy and Clark;(Claffy and Clark, 2020; Claffy et al., 2021; Claffy and Clark, 2019). In this section, we explore the measurement impediment of the Internet, the reproducibility of research and the assessment of security measurement platforms, as well as a few trendy research topics for the future.

10.1. Internet Measurement Impediments

Independent, third-party researchers can only work with the data that is available to them. They have no ability to compel the release of data. Traditionally, the network measurement community has used two methods to gather data: active probing, and passive observation. In addition, a measurement can be performed from the edge or within the Internet. Independent academic researchers most often analyze the Internet from the edge and use active probing. Active probing can only infer a few aspects of the network being probed and may fail to obtain a complete picture of the target, as the Internet's complexity can restrict access to endpoints. Moreover, the set of networks that are accessible via active probing is a small fraction of the networks on the Internet, and they are not where most traffic originates. Passive observation requires deploying monitoring instrumentation at a point in a network where it can observe activities. In this case, the vantage size influences the completeness of the view, particularly when it is collected on edge (e.g., darknet, honeypot etc.). Besides, passive monitoring raises serious concerns about privacy individuals communicating across the network and network operators.

On the other hand, network operators and employees of commercial firms collect extensive data on their own networks, but with restricted access and no corporate interest in sharing. There is no incentive for a commercial network operator to let any unaffiliated party gather data from its network. Sometimes it is illegal to do so, but even if legal barriers are overcome, there is always a risk that data related to a provider's service offering can shed light on aspects of that service that the provider wished to keep secret or that could potentially reflect poorly on the operator. Occasionally, a research team can negotiate a one-time data sharing contract with a commercial company to acquire access to such data for research purposes. In certain cases, the outcomes of these collaborations appear to be noteworthy. Although, concerns regarding scientific objectivity may arise when an associate of a linked company is named as an author. Additionally, since the data is not public, the analysis cannot be checked or reproduced.

Independent research is hindered by current trends that reduce visibility across all Internet layers, from the physical layer to the application. With advances in ISP and cloud connections come measuring challenges. CDNs, like ISPs and cloud providers, operate dense server networks that leverage anycast or DNS-based network traffic redirection. These servers are typically hosted by third-party networks, which obscures the CDN's existence. Wireless cables, WiFi access points, and repeaters are common components of modern home networks. Physical layer and media access control system variability hampers conclusions regarding home network characteristics. The proliferation of wirelessly linked gadgets, such as e-readers and smart household appliances, aggravates the situation. Cloud and ISP actions are obscured by changes at the application layer. DNS resolution is one example of a previously on-premises application that now communicates with cloud-based service endpoints. In fact, the majority of cloud-based services are complicated systems with muddled underlying architectures and linkages. Concerns about privacy are motivating further encryption at the application layer, such as QUIC at the transport layer and TLS wrapping plain text application-layer protocols. The conflict between privacy and measurement research is not new, but scientific evidence has mainly been overlooked in the argument. Allowing researchers to see user endpoints (with authorisation) would allow them to collaborate with users while respecting their privacy. In this regard, there is a growth in the number of works that rely on crowdsourcing approaches in which end users consent to share their own data. Several cellular infrastructure components are obscured by end-to-end Internet measurements (such as radio access and core networks). Extending these insights to a vast number of phones, however, remains a difficulty. Mobile network design (for example, OpenRAN) is becoming more open and interoperable, providing more insight into mobile carrier networks.

10.2. Research Reproducibility and Assessment of Security Measurement Platforms

Existing initiatives to promote reproducibility, public data, and reevaluation of results have had minimal effectiveness. Due to the dynamic nature of the Internet, it is challenging to replicate a previous study. In addition, it is difficult to reuse data collected for a certain study objective for another objective. Different aims frequently necessitate distinct studies and data sets.

Furthermore, threat intelligence data streams are available from a variety of public and private sources. The knowledge of this data, its characterization, and the extent to which it can support its intended purposes is currently restricted. In addition, researchers make extensive use of threat intelligence systems, phishing blacklists, reputation feeds, and online scan engines to label harmful URLs and files. Unfortunately, these systems often function as blackboxes and offer no understanding of how labels are formed or how dependable the outcomes are. This raises the questions of whether these labels are even reliable and whether researchers are making the best use of these platforms. Recent efforts have concentrated on the reliability of commercial threat intelligence (Bouwman et al., 2020), VirusTotal (Peng et al., 2019b; Zhu et al., 2020a; 2020b), third-party blocklists (Li et al., 2021), domain classification platforms (Vallina et al., 2020) and Internet scanning search engines (Zhao et al., 2022); however, researchers must pay greater attention to the aforementioned objectives.

10.3. Future Directions

Here, we first elaborate on possible future directions with respect to core elements in Internet security and resiliency, namely, the DNS ecosystem and routing security. Subsequently, we discuss newly emerging topics that have attracted considerable attention.

Due to the significance of the DNS environment's role in detecting and preventing cyber threats, it deserves special consideration despite a large body of research devoted to it. A topic of interest is indeed the prevalence of domain encryption (e.g., DNS-over-HTTP). ISPs have traditionally (using DHCP) had control over the resolver, so recursive server operators have the ability to ban domains (though skilled users can configure their operating system to use a different resolver). Using DNS-over-HTTP, the browser provider and any native application on a mobile device can determine which recursive resolver is being used. In fact, any application may exercise control over name resolution, e.g., by utilizing a custom resolver or avoiding the DNS entirely, which makes measurements and large-scale evaluations difficult. Moreover, domain name encryption can be a two-edged sword for network administrators who desire complete visibility and control over domain resolutions in their networks. Domain name information extracted from network traffic has been especially valuable to the operation of firewalls, intrusion detection systems, and anti-spam or anti-phishing filters up until now (Hoang et al., 2020). Therefore, the ground truth upon which major threat detection systems rely will shift and new Internet measurement tools need to be developed to overcome technical and performance related challenges (Izhikevich et al., 2022a).

Surprisingly, routing security is one of the areas where there is a plethora of data available. However, increased measurement coverage is just one factor in determining the frequency, reach, and impact of various hijacking types throughout time. A fundamental hurdle in the routing security discussion is the difficulty in differentiating sophisticated traffic engineering (with malevolent intent) from configuration mistakes. After more than two decades of research, there is currently no consensus on the occurrence and efficacy of route hijacking attacks. Further, there is no public source to report which ASes/prefixes are being hijacked. Another line of research may be determining the real extent of harm that hijackers are capable of doing.

Other interesting research directions in this area discussed in Claffy and Clark (2020) are related to MANRS and MANRS+ initiatives. MANRS stands for the Mutually Agree Norms for Routing Security. The project specifies steps that ISPs should take to mitigate vulnerabilities such as route hijacking and leaks. The Internet Society does not collect statistics on its members' compliance with these requirements, nor does it conduct research on the effectiveness of these activities in mitigating risks. Some questions that require more measurement studies are Claffy and Clark (2020): • Are MANRS members directly connected today? • Are participating ISP comply with the requirements? • Which areas of the Internet remain vulnerable to a certain hijack? • Do various types of hijacks spread differently? • Are they all supplied by the same ASes? • Would a more dense deployment of BGP probes enhance the analysis's credibility? Only a non-compliant area may spread a hijack. • How do these areas appear? • Are they multi-homed? • Can we do an analysis of the topology surrounding the infamous serial hijacker ASes? • Are clients of that AS at a lower risk of being hijacked? • Can additional benefits be quantified? • Is it feasible to demonstrate that compliance with MANRS+ benefits the participating AS?

Apart from fundamental Internet system security, the early stages of emerging technology adoption are interesting time to perform Internet measurement studies in order to give security assessments. In this regard, we found a few major themes:

1. **IoT, ICS, and smart cities:** Given that we are still in an era of growing use of smart devices and the creation of more networked devices for the aim of developing smart cities, this issue demands increased attention to regularly examine and analyze the cyber security posture. The variety of IoT devices in terms of vendor, firmware, kind, and application makes fingerprinting, vulnerability detection, and analysis of attacks against them difficult on a broad scale. Currently, many studies focus on a small number of popular gadgets and analyze them in test environments.
2. **Blockchain transactions, P2P networks and decentralized apps:** Large-scale analysis of transactions can uncover illicit money exchange on their network and help with cyber forensics and assessing cyber attack harm, particularly for ransomware. Further, Blockchains are adopted in different application domains as an infrastructure. As an example, Helium is a decentralized blockchain network that leverages a global network of Hotspots to provide long-range connectivity to IoT devices (Jagtap et al., 2021). Analysing blockchain P2P networks along with other measurement datasets can uncover interesting outcomes.
3. **Teleworking:** COVID-19 has demonstrated that dynamics in the area of cyber security have shifted dramatically as a result of enterprises' operations being shifted from the city to the household in telework mode. At home, the cyber security architecture is not as secure as it is in businesses, and if the house is built on a smart infrastructure plan, the cyber security attack surface grows. Additionally, remote work is the norm in today's world and will continue to be a part of our lives. As a result, large-scale measurement and assessment were required to ascertain the cyber security consequences of teleworking. Measuring this transition is difficult. Analysts can establish digital footprints to identify which businesses utilize remote workers, information flow patterns, and economic effect on businesses and employees. They may compare performance and other results by geographic and network region using this data. Stutz et al. (2021) is an example of efforts to address teleworking challenges.
4. **Large-scale machine learning/deep learning model security assessment:** Machine/deep learning models are widely deployed in cloud, software, IoT devices, websites, and smartphone apps. This is mostly because pre-trained models and simple-to-use cloud-based APIs are available. These applications take on a variety of forms and facets, ranging from visual filters to cyber security protections. There is a need to examine them from a variety of angles, particularly those that threaten the security and privacy of users and organizations, in order to determine their current state and potential harm. Recent initiatives include a holistic comparative analysis of widely deployed machine learning models on devices in the wild (Almeida et al., 2021) and an empirical research of machine learning model protection on mobile devices (Sun et al., 2021).
5. **5G, Edge computing and Software Defined Network (SDN):** The advent of 5G, edge computing and SDN has brought with it support for the increase in traffic volume and quality of service demands resulting from the expanse of the Internet and its applications. Notably, while each of the aforementioned technologies aims to increase throughput and reduce latency, 5G and edge computing are also tightly coupled with amplifying Internet traffic rates. Such an increase in traffic rates is inherently problematic for typical Internet measurement techniques running on general purpose CPUs, as their software-based operations may result in undesirable delays (Kfoury et al., 2021), such as for analyzing, detecting or mitigating attacks. In contrast, the Tbps traffic processing capabilities of SDN and its programmable data plane counterpart allow the same operations to be easily performed in real-time. To this extent, how can SDNs and programmable data planes best be leveraged to obtain real-time Internet measurements in order to promptly mitigate the risks associated with privacy violations, more sophisticated botnets and a plethora of other forms of maliciousness that escalate within 5G environments? Additionally, the centralized control planes of SDN and programmable data plane environments inherently offer flexibility and scalability (AlSabeh et al., 2022). In turn, can this flexibility and scalability be harnessed to offer effective cyber security solutions amid the rapidly varying network conditions of edge computing schemes?

11. Conclusion

The transition of the Internet to critical infrastructure has left society more exposed to security issues. Despite significant efforts by a number of parties (e.g., industry, government, academia) to mitigate these issues, cyber security threats are running rampant. Thus, it is necessary to measure the success of both risk reduction initiatives and defenses deployed. Such efforts require the application of unique analysis techniques to large-scale empirical data obtained through Internet measurement methodologies. Third-party researchers either leverage the data they have access to or perform Internet measurements; however, both options are frequently difficult and require individualized approaches to assure accuracy and completeness. Alternatively, researchers can perform more comprehensive investigations by utilizing diverse vantage points, correlating multiple data sources, and perhaps extending past methodologies to new predicaments. Unfortunately, the literature associated with these measurement strategies is scattered, as researchers generally focus on a few key components. To the best of our knowledge, this is the first study that has performed a comprehensive aggregation of such literature in order to examine this vital research field in depth. In particular, we investigated threats inside specific application domains as well as the threats themselves. A taxonomy of cyber security-related Internet measurement studies across multiple application areas was presented to aid academics working in specific domains. Moreover, a review of the macroscopically-acquired data analysis of cyber attacks was presented. Each corresponding study's scope, measurement breadth, and vantage size are compared, along with the analytic approach leveraged. Finally, a discussion of the challenges to Internet measurement was included as well as prospective paths.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgment

The authors are grateful in advance for the tremendous effort that the reviewers have spent in reviewing this work. The work is funded by U.S. National Science Foundations grants #2104273, #2219772 and #2219773.

References

Abhishta, A., van Rijswijk-Deij, R., Nieuwenhuis, L.J., 2019. Measuring the impact of a successful ddos attack on the customer behaviour of managed dns service providers. ACM SIGCOMM Computer Communication Review 48 (5), 70–76.

Aceto, G., Ciuonzo, D., Montieri, A., Pescapé, A., 2018. Mobile encrypted traffic classification using deep learning. In: 2018 Network traffic measurement and analysis conference (TMA). IEEE, pp. 1–8.

Agten, P., Joosen, W., Piessens, F., Nikiforakis, N., 2015. Seven months' worth of mistakes: A longitudinal study of typosquatting abuse. 22nd Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2015.23058.

Akcora, C., 2020. Bitcoinheist: Topological data analysis for ransomware prediction on the bitcoin blockchain. In: Proceedings of the 2020 International Joint Conference on Artificial Intelligence (IJCAI), pp. 4439–4445.

Al-Dalky, R., Rabinovich, M., Schomp, K., 2019. A look at the ecs behavior of dns resolvers. In: Proceedings of the Internet Measurement Conference, pp. 116–129.

Al-Naami, K., Chandra, S., Mustafa, A., Khan, L., Lin, Z., Hamlen, K., Thuraisingham, B., 2016. Adaptive encrypted traffic fingerprinting with bi-directional dependence. In: Proceedings of the 32nd Annual Conference on Computer Security Applications, pp. 177–188.

Aleroud, A., Zhou, L., 2017. Phishing environments, techniques, and countermeasures: A survey. Computers & Security 68, 160–196.

Allix, K., Bissyandé, T.F., Klein, J., Le Traon, Y., 2016. Androzoo: Collecting millions of android apps for the research community. In: 2016 IEEE/ACM 13th Working Conference on Mining Software Repositories (MSR). IEEE, pp. 468–471.

Almeida, M., Laskaridis, S., Mehrotra, A., Dudziak, L., Leontiadis, I., Lane, N.D., 2021. Smart at what cost? characterising mobile deep neural networks in the wild. In: Proceedings of the 21st ACM Internet Measurement Conference, pp. 658–672.

Alrawi, O., Lever, C., Antonakakis, M., Monrose, F., 2019. Sok: Security evaluation of home-based iot deployments. In: 2019 IEEE symposium on security and privacy (sp). IEEE, pp. 1362–1380.

Alrawi, O., Lever, C., Valakuzhy, K., Snow, K., Monrose, F., Antonakakis, M., et al., 2021. The circle of life: A {Large-Scale} study of the {IoT} malware lifecycle. In: 30th USENIX Security Symposium (USENIX Security 21), pp. 3505–3522.

Alrizah, M., Zhu, S., Xing, X., Wang, G., 2019. Errors, misunderstandings, and attacks: Analyzing the crowdsourcing process of ad-blocking systems. In: Proceedings of the Internet Measurement Conference, pp. 230–244.

AlSabeh, A., Khoury, J., Kfoury, E., Crichigno, J., Bou-Harb, E., 2022. A survey on security applications of p4 programmable switches and a stride-based vulnerability assessment. Computer Networks 207, 108800.

Anderson, B., McGrew, D., 2019. Tls beyond the browser: Combining end host and network data to understand application behavior. In: Proceedings of the Internet Measurement Conference, pp. 379–392.

Anderson, B., McGrew, D.. Accurate tls fingerprinting using destination context and knowledge bases.

Anderson, B., Paul, S., McGrew, D., 2018. Deciphering malware's use of tls (without decryption). Journal of Computer Virology and Hacking Techniques 14 (3), 195–211.

Andrade, C.E., Byers, S.D., Gopalakrishnan, V., Halepovic, E., Poole, D.J., Tran, L.K., Volinsky, C.T., 2017. Connected cars in cellular network: a measurement study. In: Proceedings of the 2017 Internet Measurement Conference, pp. 235–241.

Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., Durumeric, Z., Halderman, J.A., Invernizzi, L., Kallitsis, M., et al., 2017. Understanding the mirai botnet. In: 26th {USENIX} security symposium ({USENIX} Security 17), pp. 1093–1110.

Backes, M., Bugiel, S., Derr, E., 2016. Reliable third-party library detection in android and its security applications. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 356–367.

Bajpai, V., Schönwälder, J., 2015. A survey on internet performance measurement platforms and related standardization efforts. IEEE Communications Surveys & Tutorials 17 (3), 1313–1341.

Banerjee, R., Razaghpanah, A., Chiang, L., Mishra, A., Sekar, V., Choi, Y., Gill, P., 2015. Internet outages, the eyewitness accounts: Analysis of the outages mailing list. In: International Conference on Passive and Active Network Measurement. Springer, pp. 206–219.

Bashir, M.A., Arshad, S., Kirda, E., Robertson, W., Wilson, C., 2019. A longitudinal analysis of the ads. txt standard. In: Proceedings of the Internet Measurement Conference, pp. 294–307.

Bayat, N., Mahajan, K., Denton, S., Misra, V., Rubenstein, D., 2021. Down for failure: Active power status monitoring. Future Generation Computer Systems 125, 629–640.

Ben Mariem, S., Casas, P., Donnet, B., 2018. Vivisecting blockchain p2p networks: Unveiling the bitcoin ip network. ACM CoNEXT student workshop.

Bernaschi, M., Celestini, A., Guarino, S., Lombardi, F., 2017. Exploring and analyzing the tor hidden services graph. ACM Transactions on the Web (TWEB) 11 (4), 1–26.

Bernaschi, M., Celestini, A., Guarino, S., Lombardi, F., Mastrostefano, E., 2019. Spiders like onions: on the network of tor hidden services. In: The World Wide Web Conference, pp. 105–115.

Besel, C., Echeverria, J., Zhou, S., 2018. Full cycle analysis of a large-scale botnet attack on twitter. In: 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM). IEEE, pp. 170–177.

Beurdouche, B., Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Kohlweiss, M., Pironti, A., Strub, P.-Y., Zinzindohoue, J.K., 2015. A messy state of the union: Taming the composite state machines of tls. In: 2015 IEEE Symposium on Security and Privacy. IEEE, pp. 535–552.

Beverly, R., Durairajan, R., Plonka, D., Rohrer, J.P., 2018. In the ip of the beholder: Strategies for active ipv6 topology discovery. In: Proceedings of the Internet Measurement Conference 2018, pp. 308–321.

Bijmans, H.L., Booij, T.M., Doerr, C., 2019. Inadvertently making cyber criminals rich: A comprehensive study of cryptojacking campaigns at internet scale. In: 28th {USENIX} Security Symposium ({USENIX} Security 19), pp. 1627–1644.

Bijmans, H.L., Booij, T.M., Doerr, C., 2019. Just the tip of the iceberg: Internet-scale exploitation of routers for cryptojacking. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pp. 449–464.

Birge-Lee, H., Sun, Y., Edmundson, A., Rexford, J., Mittal, P., 2018. Bamboozling certificate authorities with {BGP}. In: 27th {USENIX} Security Symposium ({USENIX} Security 18), pp. 833–849.

Birge-Lee, H., Wang, L., Rexford, J., Mittal, P., 2019. Sico: Surgical interception attacks by manipulating bgp communities. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pp. 431–448.

Blenn, N., Ghiëtte, V., Doerr, C., 2017. Quantifying the spectrum of denial-of-service attacks through internet backscatter. In: Proceedings of the 12th International Conference on Availability, Reliability and Security, pp. 1–10.

Bock, K., Alaraj, A., Fax, Y., Hurley, K., Wustrow, E., Levin, D., 2021. Weaponizing middleboxes for {TCP} reflected amplification. In: 30th {USENIX} Security Symposium ({USENIX} Security 21), pp. 3345–3361.

Bogutz, R., Pradkin, Y., Heidemann, J., 2019. Identifying important internet outages (extended). Technical Report. TR ISI-TR-735, USC/ISI.

Borgolte, K., Fiebig, T., Hao, S., Kruegel, C., Vigna, G., 2018. Cloud strife: Mitigating the security risks of domain-validated certificates. 25th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2018.23327.

Borgolte, K., Hao, S., Fiebig, T., Vigna, G., 2018. Enumerating active ipv6 hosts for large-scale security scans via dnssec-signed reverse zones. In: 2018 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 770–784.

Boshmaf, Y., Logothetis, D., Siganos, G., Lería, J., Lorenzo, J., Ripeanu, M., Beznosov, K., 2015. Integro: Leveraging victim prediction for robust fake account detection in osns. In: 22nd Annual Network and Distributed System Security Symposium, NDSS. Internet Society, pp. 8–11. 10.14722/ndss.2015.23260

Bou-Harb, E., Debbabi, M., Assi, C., 2013. Cyber scanning: a comprehensive survey. Ieee communications surveys & tutorials 16 (3), 1496–1519.

Bou-Harb, E., Debbabi, M., Assi, C., 2016. A novel cyber security capability: Inferring internet-scale infections by correlating malware and probing activities. Computer Networks 94, 327–343.

Bouwman, X., Griffioen, H., Egbers, J., Doerr, C., Klievink, B., van Eeten, M., 2020. A different cup of {TI}? the added value of commercial threat intelligence. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 433–450.

Bushart, J., Rossow, C., 2018. Dns unchained: amplified application-layer dos attacks against dns authoritatives. In: International Symposium on Research in Attacks, Intrusions, and Defenses. Springer, pp. 139–160.

Cangialosi, F., Levin, D., Spring, N., 2015. Ting: Measuring and exploiting latencies between all tor nodes. In: Proceedings of the 2015 Internet Measurement Conference, pp. 289–302.

Celik, Z.B., Babun, L., Sikder, A.K., Aksu, H., Tan, G., McDaniel, P., Uluagac, A.S., 2018. Sensitive information tracking in commodity iot. In: 27th {USENIX} Security Symposium ({USENIX} Security 18), pp. 1687–1704.

Celik, Z.B., McDaniel, P., Tan, G., 2018. Soteria: Automated iot safety and security analysis. In: 2018 {USENIX} Annual Technical Conference ({*USENIX*}{ATC} 18), pp. 147–158.

Çetin, O., Ganán, C., Altena, L., Kasama, T., Inoue, D., Tamiya, K., Tie, Y., Yoshioka, K., van Eeten, M., 2019. Cleaning up the internet of evil things: Real-world evidence on isp and consumer efforts to remove mirai. 26th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2019.23438.

Chapuis, B., Omolola, O., Cherubini, M., Humbert, M., Huguenin, K., 2020. An empirical study of the use of integrity verification mechanisms for web subresources. In: Proceedings of The Web Conference 2020, pp. 34–45.

Chen, J., Zheng, X., Duan, H.-X., Liang, J., Jiang, J., Li, K., Wan, T., Paxson, V., 2016. Forwarding-loop attacks in content delivery networks. 23rd Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2016.23442.

Chen, K., Wang, X., Chen, Y., Wang, P., Lee, Y., Wang, X., Ma, B., Wang, A., Zhang, Y., Zou, W., 2016. Following devil's footprints: Cross-platform analysis of potentially harmful libraries on android and ios. In: 2016 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 357–376.

Chen, Q., Kapravelos, A., 2018. Mystique: Uncovering information leakage from browser extensions. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1687–1700.

Chen, Z., Freire, J., 2020. Proactive discovery of fake news domains from real-time social media feeds. In: Companion Proceedings of the Web Conference 2020, pp. 584–592.

Cheng, Z., Hou, X., Li, R., Zhou, Y., Luo, X., Li, J., Ren, K., 2019. Towards a first step to understand the cryptocurrency stealing attack on ethereum. In: 22nd international Symposium on research in Attacks, Intrusions and Defenses ({RAID} 2019), pp. 47–60.

Chiba, D., Hasegawa, A.A., Koide, T., Sawabe, Y., Goto, S., Akiyama, M., 2019. Domainscouter: Understanding the risks of deceptive idns. In: 22nd International Symposium on Research in Attacks, Intrusions and Defenses ({RAID} 2019), pp. 413–426.

Chiba, D., Yagi, T., Akiyama, M., Shibahara, T., Yada, T., Mori, T., Goto, S., 2016. Domainprofiler: Discovering domain names abused in future. In: 2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). IEEE, pp. 491–502.

Cho, G., Cho, J., Song, Y., Kim, H., 2015. An empirical study of click fraud in mobile advertising networks. In: 2015 10th International Conference on Availability, Reliability and Security. IEEE, pp. 382–388.

Cho, S., Fontugne, R., Cho, K., Dainotti, A., Gill, P., 2019. Bgp hijacking classification. In: 2019 Network Traffic Measurement and Analysis Conference (TMA). IEEE, pp. 25–32.

Chung, T., Choffnes, D., Mislove, A., 2016. Tunneling for transparency: A large-scale analysis of end-to-end violations in the internet. In: Proceedings of the 2016 Internet Measurement Conference, pp. 199–213.

Chung, T., Liu, Y., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Wilson, C., 2016. Measuring and applying invalid ssl certificates: The silent majority. In: Proceedings of the 2016 Internet Measurement Conference, pp. 527–541.

Chung, T., van Rijswijk-Deij, R., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Wilson, C., 2017. A longitudinal, end-to-end view of the {DNSSEC} ecosystem. In: 26th {USENIX} Security Symposium ({USENIX} Security 17), pp. 1307–1322.

Cisco, U., 2020. Cisco annual internet report (2018–2023) white paper.

claffy, k., Clark, D.,. Challenges in measuring the internet for the public interest. https://catalog.caida.org/paper/2022_challenges_measuring_internet_public_interest, Accessed: 2022-11-12.

Claffy, K., Clark, D., 2020. Workshop on internet economics (wie 2019) report. ACM SIGCOMM Computer Communication Review 50 (2), 53–59.

Claffy, K., Clark, D., Heidemann, J., Bustamante, F., Jonker, M., Schulman, A., Zegura, E., 2021. Workshop on overcoming measurement barriers to internet research (wombir 2021) final report. ACM SIGCOMM Computer Communication Review 51 (3), 33–40.

Claffy, K.C., Clark, D., 2019. The 11th workshop on active internet measurements (aims-11) workshop report. ACM SIGCOMM Computer Communication Review 49 (3), 39–43.

Cohen, A., Gilad, Y., Herzberg, A., Schapira, M., 2016. Jumpstarting bgp security with path-end validation. In: Proceedings of the 2016 ACM SIGCOMM Conference, pp. 342–355.

Collier, B., Thomas, D.R., Clayton, R., Hutchings, A., 2019. Booting the booters: Evaluating the effects of police interventions in the market for denial-of-service attacks. In: Proceedings of the Internet Measurement Conference, pp. 50–64.

Conti, M., Gangwal, A., Ruj, S., 2018. On the economic significance of ransomware campaigns: A bitcoin transactions perspective. Computers & Security 79, 162–189.

Costin, A., Zarras, A., Francillon, A., 2016. Automated dynamic firmware analysis at scale: a case study on embedded web interfaces. In: Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, pp. 437–448.

Cozzi, E., Graziano, M., Fratantonio, Y., Balzarotti, D., 2018. Understanding linux malware. In: 2018 IEEE symposium on security and privacy (SP). IEEE, pp. 161–175.

Dahlmanns, M., Lohmöller, J., Fink, I.B., Pennekamp, J., Wehrle, K., Henze, M., 2020. Easing the conscience with opc ua: An internet-wide study on insecure deployments. In: Proceedings of the ACM Internet Measurement Conference, pp. 101–110.

Dainotti, A., Squarcella, C., Aben, E., Claffy, K.C., Chiesa, M., Russo, M., Pescapé, A., 2011. Analysis of country-wide internet outages caused by censorship. In: Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference, pp. 1–18.

Dam, T., Klausner, L.D., Buhov, D., Schrittwieser, S., 2019. Large-scale analysis of pop-up scam on typosquatting urls. In: Proceedings of the 14th International Conference on Availability, Reliability and Security, pp. 1–9.

Das, A., Acar, G., Borisov, N., Pradeep, A., 2018. The web's sixth sense: A study of scripts accessing smartphone sensors. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1515–1532.

DeBlasio, J., Guha, S., Voelker, G.M., Snoeren, A.C., 2017. Exploring the dynamics of search advertiser fraud. In: Proceedings of the 2017 Internet Measurement Conference, pp. 157–170.

Deng, Z., Saltaformaggio, B., Zhang, X., Xu, D., 2015. iris: Vetting private api abuse in ios applications. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 44–56.

Derr, E., Bugiel, S., Fahl, S., Acar, Y., Backes, M., 2017. Keep me updated: An empirical study of third-party library updatability on android. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 2187–2200.

Di Martino, M., Quax, P., Lamotte, W., 2019. Realistically fingerprinting social media webpages in https traffic. In: Proceedings of the 14th International Conference on Availability, Reliability and Security, pp. 1–10.

Dib, M., Torabi, S., Bou-Harb, E., Assi, C., 2021. A multi-dimensional deep learning framework for iot malware classification and family attribution. IEEE Transactions on Network and Service Management 18 (2), 1165–1177.

Dietzel, C., Feldmann, A., King, T., 2016. Blackholing at ixps: On the effectiveness of ddos mitigation in the wild. In: International Conference on Passive and Active Network Measurement. Springer, pp. 319–332.

Dinh, S., Azeb, T., Fortin, F., Mouheb, D., Debbabi, M., 2015. Spam campaign detection, analysis, and investigation. Digital Investigation 12, S12–S21.

Durumeric, Z., Adrian, D., Mirian, A., Kasten, J., Bursztein, E., Lidzborski, N., Thomas, K., Eranti, V., Bailey, M., Halderman, J.A., 2015. Neither snow nor rain nor mitm… an empirical analysis of email delivery security. In: Proceedings of the 2015 Internet Measurement Conference, pp. 27–39.

Durumeric, Z., Ma, Z., Springall, D., Barnes, R., Sullivan, N., Bursztein, E., Bailey, M., Halderman, J.A., Paxson, V., 2017. The security impact of https interception. 24th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2017.23456.

Durumeric, Z., Wustrow, E., Halderman, J.A., 2013. {ZMap}: Fast internet-wide scanning and its security applications. In: 22nd USENIX Security Symposium (USENIX Security 13), pp. 605–620.

van Ede, T., Bortolameotti, R., Continella, A., Ren, J., Dubois, D.J., Lindorfer, M., Choffnes, D., van Steen, M., Peter, A., 2020. Flowprint: Semi-supervised mobile-app fingerprinting on encrypted network traffic. 27th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2020.24412.

Englehardt, S., Narayanan, A., 2016. Online tracking: A 1-million-site measurement and analysis. In: Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, pp. 1388–1401.

Fachkha, C., Bou-Harb, E., Debbabi, M., 2015. On the inference and prediction of DDoS campaigns. Wireless Communications and Mobile Computing 15 (6), 1066–1078.

Fachkha, C., Bou-Harb, E., Keliris, A., Memon, N.D., Ahamad, M., 2017. Internet-scale probing of cps: Inference, characterization and orchestration analysis. 24th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2017.23149.

Fachkha, C., Debbabi, M., 2015. Darknet as a source of cyber intelligence: Survey, taxonomy, and characterization. IEEE Communications Surveys & Tutorials 18 (2), 1197–1227.

Fadai, T., Schrittwieser, S., Kieseberg, P., Mulazzani, M., 2015. Trust me, i'm a root ca! analyzing ssl root cas in modern browsers and operating systems. In: 2015 10th International Conference on Availability, Reliability and Security. IEEE, pp. 174–179.

Farinholt, B., Rezaeirad, M., McCoy, D., Levchenko, K., 2020. Dark matter: Uncovering the darkcomet rat ecosystem. In: Proceedings of The Web Conference 2020, pp. 2109–2120.

Farinholt, B., Rezaeirad, M., Pearce, P., Dharmdasani, H., Yin, H., Le Blond, S., McCoy, D., Levchenko, K., 2017. To catch a ratter: Monitoring the behavior of amateur darkcomet rat operators in the wild. In: 2017 IEEE symposium on Security and Privacy (SP). Ieee, pp. 770–787.

Feng, X., Li, Q., Wang, H., Sun, L., 2018. Acquisitional rule-based engine for discovering internet-of-things devices. In: 27th {USENIX} Security Symposium ({USENIX} Security 18), pp. 327–341.

Fernandes, E., Jung, J., Prakash, A., 2016. Security analysis of emerging smart home applications. In: 2016 IEEE symposium on security and privacy (SP). IEEE, pp. 636–654.

Fontugne, R., Pelsser, C., Aben, E., Bush, R., 2017. Pinpointing delay and forwarding anomalies using large-scale traceroute measurements. In: Proceedings of the 2017 Internet Measurement Conference, pp. 15–28.

Frolov, S., Wustrow, E., 2019. The use of tls in censorship circumvention. 26th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2019.23511.

Fukuda, K., Heidemann, J., Qadeer, A., 2017. Detecting malicious activity with dns backscatter over time. IEEE/ACM Transactions on Networking 25 (5), 3203–3218.

Gao, Y., Shi, J., Wang, X., Tan, Q., Zhao, C., Yin, Z., 2019. Topology measurement and analysis on ethereum p2p network. In: 2019 IEEE Symposium on Computers and Communications (ISCC). IEEE, pp. 1–7.

Gasser, O., Hof, B., Helm, M., Korczynski, M., Holz, R., Carle, G., 2018. In log we trust: revealing poor security practices with certificate transparency logs and internet measurements. In: International Conference on Passive and Active Network Measurement. Springer, pp. 173–185.

Gasser, O., Scheitle, Q., Foremski, P., Lone, Q., Korczyński, M., Strowes, S.D., Hendriks, L., Carle, G., 2018. Clusters in the expanse: Understanding and unbiasing ipv6 hitlists. In: Proceedings of the Internet Measurement Conference 2018, pp. 364–378.

Gasser, O., Scheitle, Q., Gebhard, S., Carle, G., 2016. Scanning the ipv6 internet: Towards a comprehensive hitlist. In: Proc. 8th Int. Workshop on Traffic Monitoring and Analysis. http://www.net.in.tum.de/fileadmin/bibtex/publications/papers/gasser2016ipv6hitlist.pdf

Gencer, A.E., Basu, S., Eyal, I., Van Renesse, R., Sirer, E.G., 2018. Decentralization in bitcoin and ethereum networks. In: International Conference on Financial Cryptography and Data Security. Springer, pp. 439–457.

Ghasemisharif, M., Ramesh, A., Checkoway, S., Kanich, C., Polakis, J., 2018. O single sign-off, where art thou? an empirical analysis of single sign-on account hijacking and session management on the web. In: 27th {USENIX} Security Symposium ({USENIX} Security 18), pp. 1475–1492.

Gilad, Y., Hlavacek, T., Herzberg, A., Schapira, M., Shulman, H., 2018. Perfect is the enemy of good: Setting realistic goals for bgp security. In: Proceedings of the 17th ACM Workshop on Hot Topics in Networks, pp. 57–63.

Giotsas, V., Dietzel, C., Smaragdakis, G., Feldmann, A., Berger, A., Aben, E., 2017. Detecting peering infrastructure outages in the wild. In: Proceedings of the conference of the ACM special interest group on data communication, pp. 446–459.

Giotsas, V., Smaragdakis, G., Dietzel, C., Richter, P., Feldmann, A., Berger, A., 2017. Inferring bgp blackholing activity in the internet. In: Proceedings of the 2017 Internet Measurement Conference, pp. 1–14.

Goga, O., Venkatadri, G., Gummadi, K.P., 2015. The doppelgänger bot attack: Exploring identity impersonation in online social networks. In: Proceedings of the 2015 internet measurement conference, pp. 141–153.

Golla, M., Wei, M., Hainline, J., Filipe, L., Dürmuth, M., Redmiles, E., Ur, B., 2018. ” what was that site doing with my facebook password?” designing password-reuse notifications. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1549–1566.

Gómez-Boix, A., Laperdrix, P., Baudry, B., 2018. Hiding in the crowd: an analysis of the effectiveness of browser fingerprinting at large scale. In: Proceedings of the 2018 world wide web conference, pp. 309–318.

Gong, Q., Zhang, J., Chen, Y., Li, Q., Xiao, Y., Wang, X., Hui, P., 2019. Detecting malicious accounts in online developer communities using deep learning. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management, pp. 1251–1260.

Griffioen, H., Doerr, C., 2020. Examining mirai's battle over the internet of things. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, pp. 743–756.

Gugelmann, D., Happe, M., Ager, B., Lenders, V., 2015. An automated approach for complementing ad blockers' blacklists. Proceedings on Privacy Enhancing Technologies 2015 (2), 282–298.

Guillot, A., Fontugne, R., Winter, P., Merindol, P., King, A., Dainotti, A., Pelsser, C., 2019. Chocolatine: Outage detection for internet background radiation. In: 2019 Network Traffic Measurement and Analysis Conference (TMA). IEEE, pp. 1–8.

Gunawi, H.S., Hao, M., Suminto, R.O., Laksono, A., Satria, A.D., Adityatama, J., Eliazar, K.J., 2016. Why does the cloud stop computing? lessons from hundreds of service outages. In: Proceedings of the Seventh ACM Symposium on Cloud Computing, pp. 1–16.

Guo, H., Heidemann, J., 2018. Detecting iot devices in the internet (extended). USC/ISI Technical Report ISI-TR-726 July.

Gupta, P., Perdisci, R., Ahamad, M., 2018. Towards measuring the role of phone numbers in twitter-advertised spam. In: Proceedings of the 2018 on Asia Conference on Computer and Communications Security, pp. 285–296.

Gupta, S., Kuchhal, D., Gupta, P., Ahamad, M., Gupta, M., Kumaraguru, P., 2018. Under the shadow of sunshine: Characterizing spam campaigns abusing phone numbers across online social networks. In: Proceedings of the 10th ACM Conference on Web Science, pp. 67–76.

Gustafsson, J., Overier, G., Arlitt, M., Carlsson, N., 2017. A first look at the ct landscape: Certificate transparency logs in practice. In: International Conference on Passive and Active Network Measurement. Springer, pp. 87–99.

Han, X., Kheir, N., Balzarotti, D., 2016. Phisheye: Live monitoring of sandboxed phishing kits. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1402–1413.

Hao, S., Zhang, Y., Wang, H., Stavrou, A., 2018. End-users get maneuvered: Empirical analysis of redirection hijacking in content delivery networks. In: 27th {USENIX} Security Symposium ({USENIX} Security 18), pp. 1129–1145.

Hara, K., Sato, T., Imamura, M., Omote, K., 2020. Profiling of malicious users targeting ethereum's rpc port using simple honeypots. In: 2020 IEEE International Conference on Blockchain (Blockchain). IEEE, pp. 1–8.

Hastings, M., Fried, J., Heninger, N., 2016. Weak keys remain widespread in network devices. In: Proceedings of the 2016 Internet Measurement Conference, pp. 49–63.

Helium, 2021. Helium - introducing the people's network. https://www.helium.com/.

Herwig, S., Harvey, K., Hughey, G., Roberts, R., Levin, D., 2019. Measurement and analysis of hajime, a peer-to-peer iot botnet. 26th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2019.23488.

Ho, G., Cidon, A., Gavish, L., Schweighauser, M., Paxson, V., Savage, S., Voelker, G.M., Wagner, D., 2019. Detecting and characterizing lateral phishing at scale. In: 28th {USENIX} Security Symposium ({USENIX} Security 19), pp. 1273–1290.

Hoang, N.P., Akhavan Niaki, A., Borisov, N., Gill, P., Polychronakis, M., 2020. Assessing the privacy benefits of domain name encryption. In: Proceedings of the 15th ACM Asia Conference on Computer and Communications Security, pp. 290–304.

Holterbach, T., Vissicchio, S., Dainotti, A., Vanbever, L., 2017. Swift: Predictive fast reroute. In: Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 460–473.

Holz, R., Amann, J., Mehani, O., Wachs, M., Kafaar, M.A., 2016. Tls in the wild—an internet-wide analysis of tls-based protocols for electronic communication. 23rd Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2016.23055.

Hong, G., Yang, Z., Yang, S., Zhang, L., Nan, Y., Zhang, Z., Yang, M., Zhang, Y., Qian, Z., Duan, H., 2018. How you get shot in the back: A systematical study about cryptojacking in the real world. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1701–1713.

Hu, H., Jan, S.T.K., Wang, Y., Wang, G., 2021. Assessing browser-level defense against idn-based phishing. In: 30th USENIX Security Symposium (USENIX Security 21). USENIX Association, Vancouver, B.C., pp. 3739–3756. https://www.usenix.org/conference/usenixsecurity21/presentation/hu-hang

Huang, D.Y., Aliapoulios, M.M., Li, V.G., Invernizzi, L., Bursztein, E., McRoberts, K., Levin, J., Levchenko, K., Snoeren, A.C., McCoy, D., 2018. Tracking ransomware end-to-end. In: 2018 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 618–631.

Huang, D.Y., Apthorpe, N., Li, F., Acar, G., Feamster, N., 2020. Iot inspector: Crowdsourcing labeled network traffic from smart home devices at scale. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 4 (2), 1–21.

Husák, M., Cermák, M., Jirsík, T., Celeda, P., 2015. Network-based https client identification using ssl/tls fingerprinting. In: 2015 10th international conference on availability, reliability and security. IEEE, pp. 389–396.

Husák, M., Neshenko, N., Pour, M.S., Bou-Harb, E., Čeleda, P., 2018. Assessing internet-wide cyber situational awareness of critical sectors. In: Proceedings of the 13th International Conference on Availability, Reliability and Security, pp. 1–6.

Iacovazzi, A., Frassinelli, D., Elovici, Y., 2019. The {DUSTER} attack: Tor onion service attribution based on flow watermarking with track hiding. In: 22nd International Symposium on Research in Attacks, Intrusions and Defenses ({RAID} 2019), pp. 213–225.

Iamartino, D., Pelsser, C., Bush, R., 2015. Measuring bgp route origin registration and validation. In: International Conference on Passive and Active Network Measurement. Springer, pp. 28–40.

Ikram, M., Masood, R., Tyson, G., Kaafar, M.A., Loizon, N., Ensafi, R., 2019. The chain of implicit trust: An analysis of the web third-party resources loading. In: The World Wide Web Conference, pp. 2851–2857.

Ikram, M., Vallina-Rodriguez, N., Seneviratne, S., Kaafar, M.A., Paxson, V., 2016. An analysis of the privacy and security risks of android vpn permission-enabled apps. In: Proceedings of the 2016 Internet Measurement Conference, pp. 349–364.

Iliou, C., Kalpakis, G., Tsikrika, T., Vrochidis, S., Kompatsiaris, I., 2016. Hybrid focused crawling for homemade explosives discovery on surface and dark web. In: 2016 11th International Conference on Availability, Reliability and Security (ARES). IEEE, pp. 229–234.

Invernizzi, L., Thomas, K., Kapravelos, A., Comanescu, O., Picod, J.-M., Bursztein, E., 2016. Cloak of visibility: Detecting when machines browse a different web. In: 2016 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 743–758.

Iqbal, U., Shafiq, Z., Qian, Z., 2017. The ad wars: retrospective measurement and analysis of anti-adblock filter lists. In: Proceedings of the 2017 Internet Measurement Conference, pp. 171–183.

Izhikevich, L., Akiwate, G., Berger, B., Drakontaidis, S., Ascheman, A., Pearce, P., Adrian, D., Durumeric, Z., 2022. Zdns: a fast dns toolkit for internet measurement. In: Proceedings of the 22nd ACM Internet Measurement Conference, pp. 33–43.

Izhikevich, L., Teixeira, R., Durumeric, Z., 2021. {LZR}: Identifying unexpected internet services. In: 30th USENIX Security Symposium (USENIX Security 21), pp. 3111–3128.

Izhikevich, L., Teixeira, R., Durumeric, Z., 2022. Predicting ipv4 services across all ports. In: Proceedings of the ACM SIGCOMM 2022 Conference, pp. 503–515.

Jagtap, D., Yen, A., Wu, H., Schulman, A., Pannuto, P., 2021. Federated infrastructure: usage, patterns, and insights from” the people's network”. In: Proceedings of the 21st ACM Internet Measurement Conference, pp. 22–36.

Jansen, R., Juárez, M., Galvez, R., Elahi, T., Díaz, C., 2018. Inside job: Applying traffic analysis to measure tor from within. In: 25th Annual Network and Distributed System Security Symposium, pp. 1–15.

Jin, L., Hao, S., Wang, H., Cotton, C., 2018. Your remnant tells secret: Residual resolution in ddos protection services. In: 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). IEEE, pp. 362–373.

Jin, L., Hao, S., Wang, H., Cotton, C., 2019. Unveil the hidden presence: Characterizing the backend interface of content delivery networks. In: 2019 IEEE 27th International Conference on Network Protocols (ICNP). IEEE, pp. 1–11.

Jing, X., Yan, Z., Pedrycz, W., 2018. Security data collection and data analytics in the internet: A survey. IEEE Communications Surveys & Tutorials 21 (1), 586–618.

Jonker, M., King, A., Krupp, J., Rossow, C., Sperotto, A., Dainotti, A., 2017. Millions of targets under attack: a macroscopic characterization of the dos ecosystem. In: Proceedings of the 2017 Internet Measurement Conference, pp. 100–113.

Jonker, M., Pras, A., Dainotti, A., Sperotto, A., 2018. A first joint look at dos attacks and bgp blackholing in the wild. In: Proceedings of the Internet Measurement Conference 2018, pp. 457–463.

Jonker, M., Sperotto, A., van Rijswijk-Deij, R., Sadre, R., Pras, A., 2016. Measuring the adoption of ddos protection services. In: Proceedings of the 2016 Internet Measurement Conference, pp. 279–285.

Jueckstock, J., Kapravelos, A., 2019. Visiblev8: In-browser monitoring of javascript in the wild. In: Proceedings of the Internet Measurement Conference, pp. 393–405.

Karami, M., Park, Y., McCoy, D., 2016. Stress testing the booters: Understanding and undermining the business of ddos services. In: Proceedings of the 25th International Conference on World Wide Web, pp. 1033–1043.

Kfoury, E.F., Crichigno, J., Bou-Harb, E., 2021. An exhaustive survey on p4 programmable data plane switches: Taxonomy, applications, challenges, and future trends. IEEE Access 9, 87094–87155.

Khalil, I., Yu, T., Guan, B., 2016. Discovering malicious domains through passive dns data graph analysis. In: Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, pp. 663–674.

Khan, M.T., DeBlasio, J., Voelker, G.M., Snoeren, A.C., Kanich, C., Vallina-Rodriguez, N., 2018. An empirical analysis of the commercial vpn ecosystem. In: Proceedings of the Internet Measurement Conference 2018, pp. 443–456.

Khan, M.T., Huo, X., Li, Z., Kanich, C., 2015. Every second counts: Quantifying the negative externalities of cybercrime via typosquatting. In: 2015 IEEE Symposium on Security and Privacy. IEEE, pp. 135–150.

Kharraz, A., Robertson, W., Kirda, E., 2018. Surveylance: Automatically detecting online survey scams. In: 2018 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 70–86.

Khoury, J., Safaei Pour, M., Bou-Harb, E., 2022. A Near Real-Time Scheme for Collecting and Analyzing IoT Malware Artifacts at Scale. In: Proceedings of the 17th International Conference on Availability, Reliability and Security, pp. 1–11.

Kim, S.K., Ma, Z., Murali, S., Mason, J., Miller, A., Bailey, M., 2018. Measuring ethereum network peers. In: Proceedings of the Internet Measurement Conference 2018, pp. 91–104.

Kohls, K., Rupprecht, D., Holz, T., Pöpper, C., 2019. Lost traffic encryption: fingerprinting lte/4g traffic on layer two. In: Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks, pp. 249–260.

Kolamunna, H., Leontiadis, I., Perino, D., Seneviratne, S., Thilakarathna, K., Seneviratne, A., 2018. A first look at sim-enabled wearables in the wild. In: Proceedings of the Internet Measurement Conference 2018, pp. 77–83.

Kolodenker, E., Koch, W., Stringhini, G., Egele, M., 2017. Paybreak: Defense against cryptographic ransomware. In: Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, pp. 599–611.

Konoth, R.K., Vineti, E., Moonsamy, V., Lindorfer, M., Kruegel, C., Bos, H., Vigna, G., 2018. Minesweeper: An in-depth look into drive-by cryptocurrency mining and its defense. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1714–1730.

Kopp, D., Wichtlhuber, M., Poese, I., Santanna, J., Hohlfeld, O., Dietzel, C., 2019. Ddos hide & seek: On the effectiveness of a booter services takedown. In: Proceedings of the Internet Measurement Conference, pp. 65–72.

Kotzias, P., Razaghpanah, A., Amann, J., Paterson, K.G., Vallina-Rodriguez, N., Caballero, J., 2018. Coming of age: A longitudinal study of tls deployment. In: Proceedings of the Internet Measurement Conference 2018, pp. 415–428.

Kountouras, A., Kintis, P., Lever, C., Chen, Y., Nadji, Y., Dagon, D., Antonakakis, M., Joffe, R., 2016. Enabling network security through active dns datasets. In: International Symposium on Research in Attacks, Intrusions, and Defenses. Springer, pp. 188–208.

Krämer, L., Krupp, J., Makita, D., Nishizoe, T., Koide, T., Yoshioka, K., Rossow, C., 2015. Amppot: Monitoring and defending against amplification ddos attacks. In: International Symposium on Recent Advances in Intrusion Detection. Springer, pp. 615–636.

Krupp, J., Backes, M., Rossow, C., 2016. Identifying the scan and attack infrastructures behind amplification ddos attacks. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1426–1437.

Krupp, J., Karami, M., Rossow, C., McCoy, D., Backes, M., 2017. Linking amplification ddos attacks to booter services. In: International Symposium on Research in Attacks, Intrusions, and Defenses. Springer, pp. 427–449.

Kührer, M., Hupperich, T., Bushart, J., Rossow, C., Holz, T., 2015. Going wild: Large-scale classification of open dns resolvers. In: Proceedings of the 2015 Internet Measurement Conference, pp. 355–368.

Kumar, D., Shen, K., Case, B., Garg, D., Alperovich, G., Kuznetsov, D., Gupta, R., Durumeric, Z., 2019. All things considered: an analysis of iot devices on home networks. In: 28th {USENIX} Security Symposium ({USENIX} Security 19), pp. 1169–1185.

Kwon, J., Lee, J., Lee, H., Perrig, A., 2016. Psybog: A scalable botnet detection method for large-scale dns traffic. Computer Networks 97, 48–73.

Laperdrix, P., Rudametkin, W., Baudry, B., 2016. Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints. In: 2016 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 878–894.

Lauinger, T., Chaabane, A., Arshad, S., Robertson, W., Wilson, C., Kirda, E., 2017. Thou shalt not depend on me: Analysing the use of outdated javascript libraries on the web. 24th Annual Network and Distributed System Security Symposium, NDSS10.14722/ndss.2017.23414

Lauinger, T., Onarlioglu, K., Chaabane, A., Robertson, W., Kirda, E., 2016. Whois lost in translation: (mis) understanding domain name expiration and re-registration. In: Proceedings of the 2016 Internet Measurement Conference, pp. 247–253.

Le, A., Varmarken, J., Langhoff, S., Shuba, A., Gjoka, M., Markopoulou, A., 2015. Antmonitor: A system for monitoring from mobile devices. In: Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data, pp. 15–20.

Le Pochat, V., Van Goethem, T., Joosen, W., 2019. Funny accents: Exploring genuine interest in internationalized domain names. In: International Conference on Passive and Active Network Measurement. Springer, pp. 178–194.

Lee, S., Yoon, C., Kang, H., Kim, Y., Kim, Y., Han, D., Son, S., Shin, S., 2019. Cybercriminal minds: An investigative study of cryptocurrency abuses in the dark web. 26th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2019.23055.

Lee, X.T., Khan, A., Sen Gupta, S., Ong, Y.H., Liu, X., 2020. Measurements, analyses, and insights on the entire ethereum blockchain network. In: Proceedings of The Web Conference 2020, pp. 155–166.

Lekies, S., Kotowicz, K., Groß, S., Vela Nava, E.A., Johns, M., 2017. Code-reuse attacks for the web: Breaking cross-site scripting mitigations via script gadgets. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1709–1723.

Lerner, A., Simpson, A.K., Kohno, T., Roesner, F., 2016. Internet jones and the raiders of the lost trackers: An archaeological study of web tracking from 1996 to 2016. In: 25th {USENIX} Security Symposium ({USENIX} Security 16), pp. 997–1013.

Lever, C., Walls, R., Nadji, Y., Dagon, D., McDaniel, P., Antonakakis, M., 2016. Domain-z: 28 registrations later measuring the exploitation of residual trust in domains. In: 2016 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 691–706.

Leyba, K.G., Edwards, B., Freeman, C., Crandall, J.R., Forrest, S., 2019. Borders and gateways: measuring and analyzing national as chokepoints. In: Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies, pp. 184–194.

Li, F., Razaghpanah, A., Kakhki, A.M., Niaki, A.A., Choffnes, D., Gill, P., Mislove, A., 2017. lib• erate,(n) a library for exposing (traffic-classification) rules and avoiding them efficiently. In: Proceedings of the 2017 Internet Measurement Conference, pp. 128–141.

Li, V.G., Akiwate, G., Levchenko, K., Voelker, G.M., Savage, S., 2021. Clairvoyance: Inferring blocklist use on the internet. In: Passive and Active Measurement: 22nd International Conference, PAM 2021, Virtual Event, March 29–April 1, 2021, Proceedings. Springer Nature, p. 57.

Liang, J., Jiang, J., Duan, H., Li, K., Wan, T., Wu, J., 2014. When https meets cdn: A case of authentication in delegated service. In: 2014 IEEE Symposium on Security and Privacy. IEEE, pp. 67–82.

Liao, X., Liu, C., McCoy, D., Shi, E., Hao, S., Beyah, R., 2016. Characterizing long-tail seo spam on cloud web hosting services. In: Proceedings of the 25th International Conference on World Wide Web, pp. 321–332.

Liu, B., Li, Z., Zong, P., Lu, C., Duan, H., Liu, Y., Alrwais, S., Wang, X., Hao, S., Jia, Y., et al., 2019. Traffickstop: Detecting and measuring illicit traffic monetization through large-scale dns analysis. In: 2019 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE, pp. 560–575.

Liu, B., Lu, C., Li, Z., Liu, Y., Duan, H.-X., Hao, S., Zhang, Z., 2018. A reexamination of internationalized domain names: The good, the bad and the ugly. In: DSN, pp. 654–665.

Liu, D., Hao, S., Wang, H., 2016. All your dns records point to us: Understanding the security threats of dangling dns records. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1414–1425.

Liu, Y., Tome, W., Zhang, L., Choffnes, D., Levin, D., Maggs, B., Mislove, A., Schulman, A., Wilson, C., 2015. An end-to-end measurement of certificate revocation in the web's pki. In: Proceedings of the 2015 Internet Measurement Conference, pp. 183–196.

Loe, A.F., Quaglia, E.A., 2019. You shall not join: A measurement study of cryptocurrency peer-to-peer bootstrapping techniques. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pp. 2231–2247.

Lu, C., Liu, B., Li, Z., Hao, S., Duan, H., Zhang, M., Leng, C., Liu, Y., Zhang, Z., Wu, J., 2019. An end-to-end, large-scale measurement of dns-over-encryption: How far have we come? In: Proceedings of the Internet Measurement Conference, pp. 22–35.

MacFarland, D.C., Shue, C.A., Kalafut, A.J., 2017. The best bang for the byte: Characterizing the potential of dns amplification attacks. Computer Networks 116, 12–21.

Maigron, P., 2020. World - autonomous system number statistics - sorted by number. https://www-public.imtbs-tsp.eu/~maigron/RIR_Stats/RIR_Delegations/World/ASN-ByNb.html.

Malloy, M., McNamara, M., Cahn, A., Barford, P., 2016. Ad blockers: Global prevalence and impact. In: Proceedings of the 2016 Internet Measurement Conference, pp. 119–125.

Mangino, A., Pour, M.S., Bou-Harb, E., 2020. Internet-scale insecurity of consumer internet of things: An empirical measurements perspective. ACM Transactions on Management Information Systems (TMIS) 11 (4), 1–24.

Mani, A., Wilson-Brown, T., Jansen, R., Johnson, A., Sherr, M., 2018. Understanding tor usage with privacy-preserving measurement. In: Proceedings of the Internet Measurement Conference 2018, pp. 175–187.

Marciel, M., Cuevas, R., Banchs, A., Gonzalez, R., Traverso, S., Ahmed, M., Azcorra, A., 2016. Understanding the detection of view fraud in video content portals. In: Proceedings of the 25th International Conference on World Wide Web, pp. 357–368.

Marczak, B., Weaver, N., Dalek, J., Ensafi, R., Fifield, D., McKune, S., Rey, A., Scott-Railton, J., Deibert, R., Paxson, V., 2015. An analysis of china's "great cannon". 5th {USENIX} Workshop on Free and Open Communications on the Internet ({FOCI} 15). USENIX Association, Washington, D.C.. https://www.usenix.org/conference/foci15/workshop-program/presentation/marczak

Mariconti, E., Onaolapo, J., Ahmad, S.S., Nikiforou, N., Egele, M., Nikiforakis, N., Stringhini, G., 2017. What's in a name? understanding profile name reuse on twitter. In: Proceedings of the 26th International Conference on World Wide Web, pp. 1161–1170.

Mariem, S.B., Casas, P., Romiti, M., Donnet, B., Stütz, R., Haslhofer, B., 2020. All that glitters is not bitcoin–unveiling the centralized nature of the btc (ip) network. In: NOMS 2020-2020 IEEE/IFIP Network Operations and Management Symposium. IEEE, pp. 1–9.

Markert, P., Bailey, D.V., Golla, M., Dürmuth, M., AviG, A.J., 2020. This pin can be easily guessed: Analyzing the security of smartphone unlock pins. In: 2020 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 286–303.

Maroofi, S., Korczyński, M., Duda, A., 2020. Are you human? resilience of phishing detection to evasion techniques based on human verification. In: Proceedings of the ACM Internet Measurement Conference, pp. 78–86.

Mashima, D., Li, Y., Chen, B., 2019. Who's scanning our smart grid? empirical study on honeypot data. In: 2019 IEEE Global Communications Conference (GLOBECOM). IEEE, pp. 1–6.

Mavroudis, V., Hayes, J.. Adaptive traffic fingerprinting: Large-scale inference under realistic assumptions.

Mayer, W., Schmiedecker, M., 2016. Tlscompare: Crowdsourcing rules for https everywhere. In: Proceedings of the 25th International Conference Companion on World Wide Web, pp. 471–476.

Mayer, W., Zauner, A., Schmiedecker, M., Huber, M., 2016. No need for black chambers: Testing tls in the e-mail ecosystem at large. In: 2016 11th International Conference on Availability, Reliability and Security (ARES). IEEE, pp. 10–20.

Meland, P.H., Bayoumy, Y.F.F., Sindre, G., 2020. The ransomware-as-a-service economy within the darknet. Computers & Security 92, 101762.

Merzdovnik, G., Huber, M., Buhov, D., Nikiforakis, N., Neuner, S., Schmiedecker, M., Weippl, E., 2017. Block me if you can: A large-scale study of tracker-blocking tools. In: 2017 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE, pp. 319–333.

Miao, R., Potharaju, R., Yu, M., Jain, N., 2015. The dark menace: Characterizing network-based attacks in the cloud. In: Proceedings of the 2015 Internet Measurement Conference, pp. 169–182.

Minnich, A., Chavoshi, N., Koutra, D., Mueen, A., 2017. Botwalk: Efficient adaptive exploration of twitter bot networks. In: Proceedings of the 2017 IEEE/ACM international conference on advances in social networks analysis and mining 2017, pp. 467–474.

Mirheidari, S.A., Arshad, S., Onarlioglu, K., Crispo, B., Kirda, E., Robertson, W., 2020. Cached and confused: Web cache deception in the wild. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 665–682.

Mockapetris, P., 1983. Domain names: Implementation specification. RFC 883. RFC Editor.

Mockapetris, P. V., 1987. Rfc1034: Domain names-concepts and facilities.

Mohaisen, A., Ren, K., 2017. Leakage of. onion at the dns root: Measurements, causes, and countermeasures. IEEE/ACM Transactions on Networking 25 (5), 3059–3072.

Moon, S.-J., Yin, Y., Sharma, R.A., Yuan, Y., Spring, J.M., Sekar, V., 2021. Accurately measuring global risk of amplification attacks using ampmap. In: 30th {USENIX} Security Symposium ({USENIX} Security 21), pp. 3881–3898.

Moura, G.C., Heidemann, J., Müller, M., de O. Schmidt, R., Davids, M., 2018. When the dike breaks: Dissecting dns defenses during ddos. In: Proceedings of the Internet Measurement Conference 2018, pp. 8–21.

Moura, G.C., Heidemann, J., Schmidt, R.d.O., Hardaker, W., 2019. Cache me if you can: effects of dns time-to-live. In: Proceedings of the Internet Measurement Conference, pp. 101–115.

Moussaileb, R., Bouget, B., Palisse, A., Le Bouder, H., Cuppens, N., Lanet, J.-L., 2018. Ransomware's early mitigation mechanisms. In: Proceedings of the 13th International Conference on Availability, Reliability and Security, pp. 1–10.

Msadek, N., Soua, R., Engel, T., 2019. Iot device fingerprinting: Machine learning based encrypted traffic analysis. In: 2019 IEEE wireless communications and networking conference (WCNC). IEEE, pp. 1–8.

Murdock, A., Li, F., Bramsen, P., Durumeric, Z., Paxson, V., 2017. Target generation for internet-wide ipv6 scanning. In: Proceedings of the 2017 Internet Measurement Conference, pp. 242–253.

Musch, M., Steffens, M., Roth, S., Stock, B., Johns, M., 2019. Scriptprotect: mitigating unsafe third-party javascript practices. In: Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security, pp. 391–402.

Musch, M., Wressnegger, C., Johns, M., Rieck, K., 2019. Thieves in the browser: Web-based cryptojacking in the wild. In: Proceedings of the 14th International Conference on Availability, Reliability and Security, pp. 1–10.

Nadler, A., Aminov, A., Shabtai, A., 2019. Detection of malicious and low throughput data exfiltration over the dns protocol. Computers & Security 80, 36–53.

Nappa, A., Johnson, R., Bilge, L., Caballero, J., Dumitras, T., 2015. The attack of the clones: A study of the impact of shared code on vulnerability patching. In: 2015 IEEE symposium on security and privacy. IEEE, pp. 692–708.

Nasr, M., Bahramali, A., Houmansadr, A., 2018. Deepcorr: Strong flow correlation attacks on tor using deep learning. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1962–1976.

Nawrocki, M., Blendin, J., Dietzel, C., Schmidt, T.C., Wählisch, M., 2019. Down the black hole: dismantling operational practices of bgp blackholing at ixps. In: Proceedings of the Internet Measurement Conference, pp. 435–448.

Nawrocki, M., Schmidt, T.C., Wählisch, M., 2020. Uncovering vulnerable industrial control systems from the internet core. In: NOMS 2020-2020 IEEE/IFIP Network Operations and Management Symposium. IEEE, pp. 1–9.

Nelms, T., Perdisci, R., Antonakakis, M., Ahamad, M., 2016. Towards measuring and mitigating social engineering software download attacks. In: 25th {USENIX} Security Symposium ({USENIX} Security 16), pp. 773–789.

Nilizadeh, S., Labrèche, F., Sedighian, A., Zand, A., Fernandez, J., Kruegel, C., Stringhini, G., Vigna, G., 2017. Poised: Spotting twitter spam off the beaten paths. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1159–1174.

Oakes, E., Kline, J., Cahn, A., Funkhouser, K., Barford, P., 2019. A residential client-side perspective on ssl certificates. In: 2019 Network Traffic Measurement and Analysis Conference (TMA). IEEE, pp. 185–192.

Oest, A., Safaei, Y., Doupé, A., Ahn, G.-J., Wardman, B., Tyers, K., 2019. Phishfarm: A scalable framework for measuring the effectiveness of evasion techniques against browser phishing blacklists. In: 2019 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 1344–1361.

Oest, A., Safaei, Y., Zhang, P., Wardman, B., Tyers, K., Shoshitaishvili, Y., Doupé, A., 2020. Phishtime: Continuous longitudinal measurement of the effectiveness of anti-phishing blacklists. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 379–396.

Oest, A., Zhang, P., Wardman, B., Nunes, E., Burgis, J., Zand, A., Thomas, K., Doupé, A., Ahn, G.-J., 2020. Sunrise to sunset: Analyzing the end-to-end life cycle and effectiveness of phishing attacks at scale. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 361–377.

Oltrogge, M., Huaman, N., Amft, S., Acar, Y., Backes, M., Fahl, S., 2021. Why eve and mallory still love android: Revisiting tls (in) security in android applications. In: 30th USENIX Security Symposium (USENIX Security 21), pp. 4347–4364.

Onaolapo, J., Mariconti, E., Stringhini, G., 2016. What happens after you are pwnd: Understanding the use of leaked webmail credentials in the wild. In: Proceedings of the 2016 Internet Measurement Conference, pp. 65–79.

Orikogbo, D., Büchler, M., Egele, M., 2016. Crios: Toward large-scale ios application analysis. In: Proceedings of the 6th Workshop on Security and Privacy in Smartphones and Mobile Devices, pp. 33–42.

Pa, Y.M.P., Suzuki, S., Yoshioka, K., Matsumoto, T., Kasama, T., Rossow, C., 2016. Iotpot: A novel honeypot for revealing current iot threats. Journal of Information Processing 24 (3), 522–533.

Pachilakis, M., Papadopoulos, P., Markatos, E.P., Kourtellis, N., 2019. No more chasing waterfalls: a measurement study of the header bidding ad-ecosystem. In: Proceedings of the Internet Measurement Conference, pp. 280–293.

Padmanabhan, R., Schulman, A., Dainotti, A., Levin, D., Spring, N., 2019. How to find correlated internet failures. In: International Conference on Passive and Active Network Measurement. Springer, pp. 210–227.

Padmanabhan, R., Schulman, A., Levin, D., Spring, N., 2019. Residential links under the weather. In: Proceedings of the ACM Special Interest Group on Data Communication, pp. 145–158.

Panchenko, A., Lanze, F., Pennekamp, J., Engel, T., Zinnen, A., Henze, M., Wehrle, K., 2016. Website fingerprinting at internet scale. 23rd Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2016.23477.

Pantelaios, N., Nikiforakis, N., Kapravelos, A., 2020. You've changed: Detecting malicious browser extensions through their update deltas. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, pp. 477–491.

Park, K., 2005. The internet as a complex system. In: The Internet as a Large-Scale Complex System, p. 322.

Park, K., Williger, W., 2005. The Internet As a Large-Scale Complex System. Oxford University Press, Inc., USA.

Pastrana, S., Suarez-Tangil, G., 2019. A first look at the crypto-mining malware ecosystem: A decade of unrestricted wealth. In: Proceedings of the Internet Measurement Conference, pp. 73–86.

Pastrana, S., Thomas, D.R., Hutchings, A., Clayton, R., 2018. Crimebb: Enabling cybercrime research on underground forums at scale. In: Proceedings of the 2018 World Wide Web Conference, pp. 1845–1854.

Pearce, P., Jones, B., Li, F., Ensafi, R., Feamster, N., Weaver, N., Paxson, V., 2017. Global measurement of {DNS} manipulation. In: 26th {USENIX} Security Symposium ({USENIX} Security 17), pp. 307–323.

Peng, P., Xu, C., Quinn, L., Hu, H., Viswanath, B., Wang, G., 2019. What happens after you leak your password: Understanding credential sharing on phishing sites. In: Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security, pp. 181–192.

Peng, P., Yang, L., Song, L., Wang, G., 2019. Opening the blackbox of virustotal: Analyzing online phishing scan engines. In: Proceedings of the Internet Measurement Conference, pp. 478–485.

Perdisci, R., Papastergiou, T., Alrawi, O., Antonakakis, M., 2020. Iotfinder: Efficient large-scale identification of iot devices via passive dns traffic analysis. In: 2020 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE, pp. 474–489.

Perino, D., Varvello, M., Soriente, C., 2018. Proxytorrent: Untangling the free http (s) proxy ecosystem. In: Proceedings of the 2018 World Wide Web Conference, pp. 197–206.

Perta, V.C., Barbera, M.V., Tyson, G., Haddadi, H., Mei, A., 2015. A glance through the vpn looking glass: Ipv6 leakage and dns hijacking in commercial vpn clients. Proceedings on Privacy Enhancing Technologies 2015 (1), 77–91.

Plohmann, D., Yakdan, K., Klatt, M., Bader, J., Gerhards-Padilla, E., 2016. A comprehensive measurement study of domain generating malware. In: 25th {USENIX} Security Symposium ({USENIX} Security 16), pp. 263–278.

Plonka, D., Berger, A., 2015. Temporal and spatial classification of active ipv6 addresses. In: Proceedings of the 2015 Internet Measurement Conference, pp. 509–522.

Possemato, A., Fratantonio, Y., 2020. Towards {HTTPS} everywhere on android: We are not there yet. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 343–360.

Pour, M.S., Mangino, A., Friday, K., Rathbun, M., Bou-Harb, E., Iqbal, F., Shaban, K., Erradi, A., 2019. Data-driven curation, learning and analysis for inferring evolving iot botnets in the wild. In: Proceedings of the 14th International Conference on Availability, Reliability and Security, pp. 1–10.

Pour, M.S., Bou-Harb, E., 2019. Theoretic derivations of scan detection operating on darknet traffic. Computer Communications 147, 111–121.

Pour, M.S., Bou-Harb, E., Varma, K., Neshenko, N., Pados, D.A., Choo, K.K.R., 2019. Comprehending the IoT cyber threat landscape: A data dimensionality reduction technique to infer and characterize Internet-scale IoT probing campaigns. Digital Investigation 28, S40–S49.

Pour, M.S., Khoury, J., Bou-Harb, E., 2022. April. HoneyComb: A Darknet-Centric Proactive Deception Technique For Curating IoT Malware Forensic Artifacts. In: NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium. IEEE, pp. 1–9.

Pour, M.S., Mangino, A., Friday, K., Rathbun, M., Bou-Harb, E., Iqbal, F., Samtani, S., Crichigno, J., Ghani, N., 2020. On data-driven curation, learning, and analysis for inferring evolving internet-of-things (iot) botnets in the wild. Computers & Security 91, 101707.

Pour, M.S., Watson, D., Bou-Harb, E., 2021. Sanitizing the iot cyber security posture: An operational cti feed backed up by internet measurements. IEEE, pp. 497–506.

Quinkert, F., Degeling, M., Blythe, J., Holz, T., 2020. Be the phisher–understanding users' perception of malicious domains. In: Proceedings of the 15th ACM Asia Conference on Computer and Communications Security, pp. 263–276.

Quinkert, F., Lauinger, T., Robertson, W., Kirda, E., Holz, T., 2019. It's not what it looks like: Measuring attacks and defensive registrations of homograph domains. In: 2019 IEEE Conference on Communications and Network Security (CNS). IEEE, pp. 259–267.

Rafique, M.Z., Van Goethem, T., Joosen, W., Huygens, C., Nikiforakis, N., 2016. It's free for a reason: Exploring the ecosystem of free live streaming services. 23rd Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2016.23030.

Rahbarinia, B., Perdisci, R., Antonakakis, M., 2016. Efficient and accurate behavior-based tracking of malware-control domains in large isp networks. ACM Transactions on Privacy and Security (TOPS) 19 (2), 1–31.

Razaghpanah, A., Niaki, A.A., Vallina-Rodriguez, N., Sundaresan, S., Amann, J., Gill, P., 2017. Studying tls usage in android apps. In: Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies, pp. 350–362.

Razaghpanah, A., Vallina-Rodriguez, N., Sundaresan, S., Kreibich, C., Gill, P., Allman, M., Paxson, V.. Haystack: A multi-purpose mobile vantage point in user space.

Ren, J., Dubois, D.J., Choffnes, D., Mandalari, A.M., Kolcun, R., Haddadi, H., 2019. Information exposure from consumer iot devices: A multidimensional, network-informed measurement approach. In: Proceedings of the Internet Measurement Conference, pp. 267–279.

Rezaeirad, M., Farinholt, B., Dharmdasani, H., Pearce, P., Levchenko, K., McCoy, D., 2018. Schrödinger's {RAT}: Profiling the stakeholders in the remote access trojan ecosystem. In: 27th {USENIX} security symposium ({USENIX} Security 18), pp. 1043–1060.

Richter, P., Padmanabhan, R., Spring, N., Berger, A., Clark, D., 2018. Advancing the art of internet edge outage detection. In: Proceedings of the Internet Measurement Conference 2018, pp. 350–363.

van Rijswijk-Deij, R., Jonker, M., Sperotto, A., Pras, A., 2015. The internet of names: A dns big dataset. In: Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, pp. 91–92.

van Rijswijk-Deij, R., Jonker, M., Sperotto, A., Pras, A., 2016. A high-performance, scalable infrastructure for large-scale active dns measurements. IEEE Journal on Selected Areas in Communications 34 (6), 1877–1888.

Rimmer, V., Preuveneers, D., Juarez, M., Van Goethem, T., Joosen, W., 2018. Automated website fingerprinting through deep learning. In: Proceedings 2018 Network and Distributed System Security Symposium. Internet Society doi:10.14722/ndss.2018.23105.

Roberts, R., Goldschlag, Y., Walter, R., Chung, T., Mislove, A., Levin, D., 2019. You are who you appear to be: a longitudinal study of domain impersonation in tls certificates. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pp. 2489–2504.

Rossow, C., 2014. Amplification hell: Revisiting network protocols for ddos abuse. In: 21st Annual Network and Distributed System Security Symposium, NDSS. Internet Society, pp. 1–15. doi:10.14722/ndss.2014.23233.

Roth, S., Backes, M., Stock, B., 2020. Assessing the impact of script gadgets on csp at scale. In: Proceedings of the 15th ACM Asia Conference on Computer and Communications Security, pp. 420–431.

Rüth, J., Zimmermann, T., Wolsing, K., Hohlfeld, O., 2018. Digging into browser-based crypto mining. In: Proceedings of the Internet Measurement Conference 2018, pp. 70–76.

Rytilahti, T., Holz, T., 2016. Poster: The curious case of ntp monlist. European S&P 2016.

Saad, M., Spaulding, J., Njilla, L., Kamhoua, C., Shetty, S., Nyang, D., Mohaisen, D., 2020. Exploring the attack surface of blockchain: A comprehensive survey. IEEE Communications Surveys & Tutorials 22 (3), 1977–2008.

Saidi, S.J., Mandalari, A.M., Kolcun, R., Haddadi, H., Dubois, D.J., Choffnes, D., Smaragdakis, G., Feldmann, A., 2020. A haystack full of needles: Scalable detection of iot devices in the wild. In: Proceedings of the ACM Internet Measurement Conference, pp. 87–100.

Samarasinghe, N., Mannan, M., 2019. Another look at tls ecosystems in networked devices vs. web servers. Computers & Security 80, 1–13.

Sarker, S., Jueckstock, J., Kapravelos, A., 2020. Hiding in plain site: Detecting javascript obfuscation through concealed browser api usage. In: Proceedings of the ACM Internet Measurement Conference, pp. 648–661.

Scheitle, Q., Gasser, O., Nolte, T., Amann, J., Brent, L., Carle, G., Holz, R., Schmidt, T.C., Wählisch, M., 2018. The rise of certificate transparency and its implications on the internet ecosystem. In: Proceedings of the Internet Measurement Conference 2018, pp. 343–349.

Schlamp, J., Holz, R., Jacquemart, Q., Carle, G., Biersack, E.W., 2016. Heap: reliable assessment of bgp hijacking attacks. IEEE Journal on Selected Areas in Communications 34 (6), 1849–1861.

Scott, W., Anderson, T., Kohno, T., Krishnamurthy, A., 2016. Satellite: Joint analysis of cdns and network-level interference. In: 2016 {USENIX} Annual Technical Conference ({USENIX}{ATC} 16), pp. 195–208.

Sermpezis, P., Kotronis, V., Gigis, P., Dimitropoulos, X., Cicalese, D., King, A., Dainotti, A., 2018. Artemis: Neutralizing bgp hijacking within a minute. IEEE/ACM Transactions on Networking 26 (6), 2471–2486.

Shah, A., Fontugne, R., Aben, E., Pelsser, C., Bush, R., 2017. Disco: Fast, good, and cheap outage detection. In: 2017 Network Traffic Measurement and Analysis Conference (TMA). IEEE, pp. 1–9.

Siby, S., Juarez, M., Diaz, C., Vallina-Rodriguez, N., Troncoso, C., 2020. Encrypted dns–> privacy? a traffic analysis perspective. 27th Annual Network and Distributed System Security Symposium, NDSS doi:10.14722/ndss.2020.24301.

Simoiu, C., Bonneau, J., Gates, C., Goel, S., 2019. " i was told to buy a software or lose my computer. i ignored it": A study of ransomware. In: Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019), pp. 155–174.

Simoiu, C., Zand, A., Thomas, K., Bursztein, E., 2020. Who is targeted by email-based phishing and malware? measuring factors that differentiate risk. In: Proceedings of the ACM Internet Measurement Conference, pp. 567–576.

Singh, M., Singh, M., Kaur, S., 2019. Issues and challenges in dns based botnet detection: A survey. Computers & Security 86, 28–52.

Singh, R., Nithyanand, R., Afroz, S., Pearce, P., Tschantz, M.C., Gill, P., Paxson, V., 2017. Characterizing the nature and dynamics of tor exit blocking. In: 26th {USENIX} Security Symposium ({USENIX} Security 17), pp. 325–341.

Smith, J.M., Birkeland, K., McDaniel, T., Schuchard, M., 2020. Withdrawing the bgp re-routing curtain: Understanding the security impact of bgp poisoning via real-world measurements. 27th Annual Network and Distributed System Security Symposium, NDSS doi:10.14722/ndss.2020.24240.

Snyder, P., Taylor, C., Kanich, C., 2017. Most websites don't need to vibrate: A cost-benefit approach to improving browser security. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 179–194.

Snyder, P., Vastel, A., Livshits, B., 2020. Who filters the filters: Understanding the growth, usefulness and efficiency of crowdsourced ad blocking. Proceedings of the ACM on Measurement and Analysis of Computing Systems 4 (2), 1–24.

Song, Y., Hengartner, U., 2015. Privacyguard: A vpn-based platform to detect information leakage on android devices. In: Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices, pp. 15–26.

Srinivasan, B., Kountouras, A., Miramirkhani, N., Alam, M., Nikiforakis, N., Antonakakis, M., Ahamad, M., 2018. Exposing search and advertisement abuse tactics and infrastructure of technical support scammers. In: Proceedings of the 2018 World Wide Web Conference, pp. 319–328.

Staicu, C.-A., Pradel, M., 2018. Freezing the web: A study of redos vulnerabilities in javascript-based web servers. In: 27th {USENIX} Security Symposium ({USENIX} Security 18), pp. 361–376.

Starov, O., Laperdrix, P., Kapravelos, A., Nikiforakis, N., 2019. Unnecessarily identifiable: Quantifying the fingerprintability of browser extensions due to bloat. In: The World Wide Web Conference, pp. 3244–3250.

Starov, O., Nikiforakis, N., 2017. Xhound: Quantifying the fingerprintability of browser extensions. In: 2017 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 941–956.

Starov, O., Zhou, Y., Zhang, X., Miramirkhani, N., Nikiforakis, N., 2018. Betrayed by your dashboard: Discovering malicious campaigns via web analytics. In: Proceedings of the 2018 World Wide Web Conference, pp. 227–236.

Stevanovic, M., Pedersen, J.M., D'Alconzo, A., Ruehrup, S., Berger, A., 2015. On the ground truth problem of malicious dns traffic analysis. computers & security 55, 142–158.

Stock, B., Johns, M., Steffens, M., Backes, M., 2017. How the web tangled itself: Uncovering the history of client-side web (in) security. In: 26th {USENIX} Security Symposium ({USENIX} Security 17), pp. 971–987.

Storey, G., Reisman, D., Mayer, J., Narayanan, A.. The future of ad blocking: An analytical framework and new techniques.

Stringhini, G., Mourlanne, P., Jacob, G., Egele, M., Kruegel, C., Vigna, G., 2015.

{EVILCOHORT}: Detecting communities of malicious accounts on online services. In: 24th {USENIX} Security Symposium ({USENIX} Security 15), pp. 563–578.

Stutz, E., Pradkin, Y., Song, X., Heidemann, J., 2021. Visualizing internet measurements of covid-19 work-from-home. 2021 IEEE International Conference on Big Data (Big Data) 5633–5638.

Subramani, K., Yuan, X., Setayeshfar, O., Vadrevu, P., Lee, K.H., Perdisci, R., 2020. When push comes to ads: Measuring the rise of (malicious) push advertising. In: Proceedings of the ACM Internet Measurement Conference, pp. 724–737.

Sun, Y., Edmundson, A., Feamster, N., Chiang, M., Mittal, P., 2017. Counter-raptor: Safeguarding tor against active routing attacks. In: 2017 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 977–992.

Sun, Y., Edmundson, A., Vanbever, L., Li, O., Rexford, J., Chiang, M., Mittal, P., 2015. {RAPTOR}: Routing attacks on privacy in tor. In: 24th {USENIX} Security Symposium ({USENIX} Security 15), pp. 271–286.

Sun, Z., Sun, R., Lu, L., Mislove, A., 2021. Mind your weight (s): A large-scale study on insufficient machine learning model protection in mobile apps. In: 30th {USENIX} Security Symposium ({USENIX} Security 21), pp. 1955–1972.

Suzuki, H., Chiba, D., Yoneya, Y., Mori, T., Goto, S., 2019. Shamfinder: An automated framework for detecting idn homographs. In: Proceedings of the Internet Measurement Conference, pp. 449–462.

Szurdi, J., Kocso, B., Cseh, G., Spring, J., Felegyhazi, M., Kanich, C., 2014. The long "taile" of typosquatting domain names. In: 23rd {USENIX} Security Symposium ({USENIX} Security 14), pp. 191–206.

Tanabe, R., Tamai, T., Fujita, A., Isawa, R., Yoshioka, K., Matsumoto, T., Gañán, C., van Eeten, M., 2020. Disposable botnets: examining the anatomy of iot botnet infrastructure. In: Proceedings of the 15th International Conference on Availability, Reliability and Security, pp. 1–10.

Tang, Z., Tang, K., Xue, M., Tian, Y., Chen, S., Ikram, M., Wang, T., Zhu, H., 2020. ios, your {OS}, everybody's {OS}: Vetting and analyzing network services of ios applications. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 2415–2432.

Tang, Z., Xue, M., Meng, G., Ying, C., Liu, Y., He, J., Zhu, H., Liu, Y., 2019. Securing android applications via edge assistant third-party library detection. Computers & Security 80, 257–272.

Tao, L., Li, D., Zhang, P., Liang, X., 2019. Probe recommendation algorithm for link delay detection. In: International Conference on Artificial Intelligence and Security. Springer, pp. 369–378.

Taylor, V., Spolaor, R., Conti, M., Martinovic, I., 2017. Robust smartphone app identification via encrypted network traffic analysis. IEEE Transactions on Information Forensics and Security 13 (1), 63–78.

Taylor, V.F., Spolaor, R., Conti, M., Martinovic, I., 2016. Appscanner: Automatic fingerprinting of smartphone apps from encrypted network traffic. In: 2016 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE, pp. 439–454.

Testart, C., Richter, P., King, A., Dainotti, A., Clark, D., 2019. Profiling bgp serial hijackers: capturing persistent misbehavior in the global routing table. In: Proceedings of the Internet Measurement Conference, pp. 420–434.

Thomas, K., Li, F., Zand, A., Barrett, J., Ranieri, J., Invernizzi, L., Markov, Y., Comanescu, O., Eranti, V., Moscicki, A., et al., 2017. Data breaches, phishing, or malware? understanding the risks of stolen credentials. In: Proceedings of the 2017 ACM SIGSAC conference on computer and communications security, pp. 1421–1434.

Tian, K., Jan, S.T., Hu, H., Yao, D., Wang, G., 2018. Needle in a haystack: Tracking down elite phishing domains in the wild. In: Proceedings of the Internet Measurement Conference 2018, pp. 429–442.

Torabi, S., Bou-Harb, E., Assi, C., Galluscio, M., Boukhtouta, A., Debbabi, M., 2018. Inferring, characterizing, and investigating internet-scale malicious iot device activities: A network telescope perspective. In: 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). IEEE, pp. 562–573.

Torabi, S., Bou-Harb, E., Assi, C., Karbab, E.B., Boukhtouta, A., Debbabi, M., 2020. Inferring and investigating IoT-generated scanning campaigns targeting a large network telescope. IEEE Transactions on Dependable and Secure Computing 19 (1), 402–418.

Torabi, S., Boukhtouta, A., Assi, C., Debbabi, M., 2018. Detecting internet abuse by analyzing passive dns traffic: A survey of implemented systems. IEEE Communications Surveys & Tutorials 20 (4), 3389–3415.

Tsirantonakis, G., Ilia, P., Ioannidis, S., Athanasopoulos, E., Polychronakis, M., 2018. A large-scale analysis of content modification by open http proxies. 25th Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2018.23244.

Tyson, G., Huang, S., Cuadrado, F., Castro, I., Perta, V.C., Sathiaseelan, A., Uhlig, S., 2017. Exploring http header manipulation in-the-wild. In: Proceedings of the 26th International Conference on World Wide Web, pp. 451–458.

Vadrevu, P., Perdisci, R., 2019. What you see is not what you get: Discovering and tracking social engineering attack campaigns. In: Proceedings of the Internet Measurement Conference, pp. 308–321.

Vallina, P., Le Pochat, V., Feal, Á., Paraschiv, M., Gamba, J., Burke, T., Hohlfeld, O., Tapiador, J., Vallina-Rodriguez, N., 2020. Mis-shapes, mistakes, misfits: An analysis of domain classification services. In: Proceedings of the ACM Internet Measurement Conference, pp. 598–618.

Van Der Heijden, A., Allodi, L., 2019. Cognitive triaging of phishing attacks. In: 28th {USENIX} Security Symposium ({USENIX} Security 19), pp. 1309–1326.

Van Wegberg, R., Tajalizadehkhoob, S., Soska, K., Akyazi, U., Ganan, C.H., Klievink, B., Christin, N., Van Eeten, M., 2018. Plug and prey? measuring the commoditization of cybercrime via online anonymous markets. In: 27th {USENIX} security symposium ({USENIX} security 18), pp. 1009–1026.

VanderSloot, B., Amann, J., Bernhard, M., Durumeric, Z., Bailey, M., Halderman, J.A., 2016. Towards a complete view of the certificate ecosystem. In: Proceedings of the 2016 Internet Measurement Conference, pp. 543–549.

Vasilomanolakis, E., Srinivasa, S., Cordero, C.G., Mühlhäuser, M., 2016. Multi-stage attack detection and signature generation with ics honeypots. In: NOMS 2016-2016 IEEE/IFIP Network Operations and Management Symposium. IEEE, pp. 1227–1232.

Vastel, A., Rudametkin, W., Rouvoy, R., Blanc, X., 2020. Fp-crawlers: studying the resilience of browser fingerprinting to block crawlers. MADWeb'20-NDSS Workshop on Measurements, Attacks, and Defenses for the Web. Internet Society doi:10.14722/madweb.2020.23010.

Vervier, P.-A., Shen, Y., 2018. Before toasters rise up: A view into the emerging iot threat landscape. In: International Symposium on Research in Attacks, Intrusions, and Defenses. Springer, pp. 556–576.

Vervier, P.-A., Thonnard, O., Dacier, M., 2015. Mind your blocks: On the stealthiness of malicious bgp hijacks. 22nd Annual Network and Distributed System Security Symposium, NDSS. Internet Society doi:10.14722/ndss.2015.23035.

Vissers, T., Barron, T., Van Goethem, T., Joosen, W., Nikiforakis, N., 2017. The wolf of name street: Hijacking domains through their nameservers. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 957–970.

Vissers, T., Van Goethem, T., Joosen, W., Nikiforakis, N., 2015. Maneuvering around clouds: Bypassing cloud-based security providers. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 1530–1541.

Vu, A.V., Hughes, J., Pete, I., Collier, B., Chua, Y.T., Shumailov, I., Hutchings, A., 2020. Turning up the dial: the evolution of a cybercrime market through set-up, stable, and covid-19 eras. In: Proceedings of the ACM Internet Measurement Conference, pp. 551–566.

Wan, G., Izhikevich, L., Adrian, D., Yoshioka, K., Holz, R., Rossow, C., Durumeric, Z., 2020. On the origin of scanning: The impact of location on internet-wide scans. In: Proceedings of the ACM Internet Measurement Conference, pp. 662–679.

Wang, A., Chang, W., Chen, S., Mohaisen, A., 2018. Delving into internet ddos attacks by botnets: characterization and analysis. IEEE/ACM Transactions on Networking 26 (6), 2843–2855.

Wang, H., Li, H., Guo, Y., 2019. Understanding the evolution of mobile app ecosystems: A longitudinal measurement study of google play. In: The World Wide Web Conference, pp. 1988–1999.

Wang, H., Liu, Z., Liang, J., Vallina-Rodriguez, N., Guo, Y., Li, L., Tapiador, J., Cao, J., Xu, G., 2018. Beyond google play: A large-scale comparative study of chinese android app markets. In: Proceedings of the Internet Measurement Conference 2018, pp. 293–307.

Wang, T., Zhao, C., Yang, Q., Zhang, S., Liew, S.C., 2021. Ethna: Analyzing the underlying peer-to-peer network of ethereum blockchain. IEEE Transactions on Network Science and Engineering 8 (3), 2131–2146.

Xu, H., Xu, F., Chen, B., 2018. Internet protocol cameras with no password protection: An empirical investigation. In: International Conference on Passive and Active Network Measurement. Springer, pp. 47–59.

Xu, Z., Wang, H., Wu, Z., 2015. A measurement study on co-residence threat inside the cloud. In: 24th {USENIX} Security Symposium ({USENIX} Security 15), pp. 929–944.

Yan, Z., Lv, S., Zhang, Y., Zhu, H., Sun, L., 2019. Remote fingerprinting on internet-wide printers based on neural network. In: 2019 IEEE Global Communications Conference (GLOBECOM). IEEE, pp. 1–6.

Yang, K., Li, Q., Sun, L., 2019. Towards automatic fingerprinting of iot devices in the cyberspace. Computer Networks 148, 318–327.

Yao, Y., Viswanath, B., Cryan, J., Zheng, H., Zhao, B.Y., 2017. Automated crowdturfing attacks and defenses in online review systems. In: Proceedings of the 2017 ACM SIGSAC conference on computer and communications security, pp. 1143–1158.

Yazdani, R., van der Toorn, O., Sperotto, A., 2020. A case of identity: Detection of suspicious idn homograph domains using active dns measurements. In: 2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW). IEEE, pp. 559–564.

Yu, L., Luo, B., Ma, J., Zhou, Z., Liu, Q., 2020. You are what you broadcast: Identification of mobile and iot devices from (public) wifi. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 55–72.

Yuan, D., Miao, Y., Gong, N.Z., Yang, Z., Li, Q., Song, D., Wang, Q., Liang, X., 2019. Detecting fake accounts in online social networks at the time of registrations. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pp. 1423–1438.

Zabihimayvan, M., Doran, D., 2022. A first look at references from the dark to the surface web world: a case study in tor. International Journal of Information Security 1–17.

Zhang, J., Zhang, R., Zhang, Y., Yan, G., 2016. The rise of social botnets: Attacks and countermeasures. IEEE Transactions on Dependable and Secure Computing 15 (6), 1068–1082.

Zhang, L., Chen, J., Diao, W., Guo, S., Weng, J., Zhang, K., 2019. Cryptorex: Large-scale analysis of cryptographic misuse in iot devices. In: 22nd International Symposium on Research in Attacks, Intrusions and Defenses ({RAID} 2019), pp. 151–164.

Zhang, M., Li, J., Brooks, S., 2017. I-seismograph: Observing, measuring, and analyzing internet earthquakes. IEEE/ACM Transactions on networking 25 (6), 3411–3426.

Zhang, M., Meng, W., Lee, S., Lee, B., Xing, X., 2019. All your clicks belong to me:

investigating click interception on the web. In: 28th {USENIX} Security Symposium ({USENIX} Security 19), pp. 941–957.

Zhang, P., Oest, A., Cho, H., Johnson, R., Wardman, B., Sarker, S., Kpravelos, A., Bao, T., Wang, R., Shoshitaishvili, Y., Doupé, A., Ahn, G.-J., 2021. CrawlPhish: Large-scale Analysis of Client-side Cloaking Techniques in Phishing. In: In Proceedings of the 42nd IEEE Symposium on Security and Privacy (Oakland). IEEE, San Francisco, CA, pp. 1109–1124.

Zhang, W., Meng, Y., Liu, Y., Zhang, X., Zhang, Y., Zhu, H., 2018. Homonit: Monitoring smart home apps from encrypted traffic. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pp. 1074–1088.

Zhang, Y., Liu, B., Lu, C., Li, Z., Duan, H., Hao, S., Liu, M., Liu, Y., Wang, D., Li, Q., 2020. Lies in the air: Characterizing fake-base-station spam ecosystem in china. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, pp. 521–534.

Zhao, B., Ji, S., Lee, W.-H., Lin, C., Weng, H., Wu, J., Zhou, P., Fang, L., Beyah, R., 2022. A large-scale empirical study on thevulnerability of deployed iot devices. IEEE Transactions on Dependable and Secure Computing 19 (03), 1826–1840.

Zhauniarovich, Y., Khalil, I., Yu, T., Dacier, M., 2018. A survey on malicious domains detection through dns data analysis. ACM Computing Surveys (CSUR) 51 (4), 1–36.

Zhou, D., Yan, Z., Fu, Y., Yao, Z., 2018. A survey on network data collection. Journal of Network and Computer Applications 116, 9–23.

Zhu, S., Shi, J., Yang, L., Qin, B., Zhang, Z., Song, L., Wang, G., 2020. Measuring and modeling the label dynamics of online anti-malware engines. In: 29th {USENIX} Security Symposium ({USENIX} Security 20), pp. 2361–2378.

Zhu, S., Zhang, Z., Yang, L., Song, L., Wang, G., 2020. Benchmarking label dynamics of virustotal engines. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, pp. 2081–2083.

Ziviani, A., 2006. An overview of internet measurements: Fundamentals, techniques, and trends. African Journal of Information & Communication Technology 2 (1), 39–49.

ZMap, 2022. GitHub. https://github.com/zmap/zmap

Morteza Safaei Pour received the Ph.D. degree in information systems and cyber security from University of Texas at San Antonio, USA. He is currently an assistant professor at the management information system department at San Diego State University (SDSU). His current research interests are in the areas of Internet measurement, operational cyber security, IoT security, attacks detection and characterization etc.

Christelle Nader is currently pursuing an M.S. in Information Technology at the University of Texas at San Antonio (UTSA) with a concentration in cyber security. She received her B.S. in computer science from Notre Dame University-Louaize (NDU), Lebanon, in 2020. She has executed several projects in the field of IoT security, including fingerprinting NATed IoT devices and detecting infected Internet-scale IoT bots behind a NAT. Her research interests include operational cyber security, IoT security, attack detection and characterization, malware analysis, data science, Internet measurements for cyber security, and big data analytics.

Kurt Friday is pursuing his Ph.D. in Information Technology at the University of Texas at San Antonio (UTSA) with added emphasis on cyber security, artificial intelligence, and information systems. He has a background in computer science, which he obtained at Florida Atlantic University (FAU). He has extensive experience with programmable networking and Software Defined Networking (SDN), and he has developed several security and forensic mechanisms for such networks. His current research interests are in the areas of operational cyber security, attack detection and characterization, Internet measurements for cyber security, malware analysis, digital forensics, and data science.

Elias Bou-Harb is currently the Director of the Cyber Center For Security and Analytics at UTSA, where he leads, co-directs and co-organizes university-wide innovative cyber security research, development and training initiatives. He is also a tenured Associate Professor at the department of Information Systems and Cyber Security specializing in operational cyber security and data science as applicable to national security challenges. Previously, he was a senior research scientist at Carnegie Mellon University (CMU) where he contributed to federally-funded projects related to critical infrastructure security and worked closely with the Software Engineering Institute (SEI). He is also a permanent research scientist at the National Cyber Forensic and Training Alliance (NCFTA) of Canada; an international organization which focuses on the investigation of cyber-crimes impacting citizens and businesses. Dr. Bou-Harb holds a Ph.D. degree in computer science from Concordia University in Montreal, Canada, which was executed in collaboration with Public Safety Canada, Industry Canada and NCFTA Canada. His research and development activities and interests focus on operational cyber security, attacks' detection and characterization, malware investigation, cyber security for critical infrastructure and big data analytics.