

An Evaluation of Darknet Traffic Taxonomy

JUN LIU^{1,a)} KENSUKE FUKUDA^{1,2,b)}

Received: May 9, 2017, Accepted: November 7, 2017

Abstract: To enhance Internet security, researchers have largely emphasized diverse cyberspace monitoring approaches to observe cyber attacks and anomalies. Among them darknet provides an effective passive monitoring one. Darknets refer to the globally routable but still unused IP address spaces. They are often used to monitor unexpected incoming network traffic, and serve as an effective network traffic measurement approach for viewing certain remote network security activities. Previous works in this field discussed possible causes (i.e., anomalies) of darknet traffic and applied their classification schemes on short-term traces. Our interest lies, however, in how darknet traffic has evolved and the effectiveness of a darknet traffic taxonomy for longitudinal data. To reach these goals, we propose a simple darknet traffic taxonomy based on network traffic rules, and evaluate it with two darknet traces: one covering 12 years since 2006, while the other covering 11 years since 2007. The evaluation results reveal the effectiveness of this taxonomy: we are able to label over 94% of all source IPs with anomalies defined by the taxonomy, leaving the unlabeled source ratio low. We also examine the evolution of different anomalies since 2006 (especially in recent years), analyze the temporal and spatial dependency and parameter dependency of darknet traffic, and conclude that most sources in the datasets are characterized by just one or two anomalies with simple attack mechanisms. Moreover, we compare the taxonomy with a one-way traffic analysis tool (i.e., *iatmon*) to better understand their differences.

Keywords: evaluation, taxonomy, darknet traffic, anomaly

1. Introduction

Along with the growth in Internet users, network security challenges have become more difficult to deal with. By providing an opportunity to passively monitor them remotely, darknets [25], [32] (a.k.a., network telescopes [23]) have drawn much attention in the research community. A darknet consists of globally routable but still unused IP address blocks in which little or no legitimate traffic exists. Continual monitoring of such address spaces, however, shows that unwanted packets keep arriving with high rates from a wide range of sources all across the Internet. These packets are completely non-productive, since they originate from worm propagation, (D)DoS attacks, network misconfiguration, or other unsolicited activities.

Previous works [25], [32] have showed the not minor volume of darknet traffic and its great diversity both in terms of the address space being monitored as well as over time. Based simply on TCP flags they also classified darknet traffic into three types of network activities: scanning, backscatter, and misconfiguration. Some studies on one-way traffic analysis helped us better understand darknet traffic considering its unidirectional nature. An one-way network traffic analysis tool *iatmon* [6] classifies traffic with two schemes: activity patterns of sessions, created using finite state machine models of host-pair packet-level behavior [30], and packet inter-arrival time (IAT) percentage distributions of sources. The author evaluated the tool with a half-year

darknet trace in 2011. Finding recognizable IAT patterns, however, takes much effort, and the efficiency of classifying real long-range darknet traces has not been examined yet.

To examine the evolution of darknet traffic with longitudinal data and avoid the hard work of obtaining IAT patterns, we need a simple but effective taxonomy of anomalies in darknet traffic. To our best knowledge, no such a simple reference taxonomy has been developed in the research community. In this paper, therefore, we propose a simple taxonomy of anomalies in darknet traffic on the basis of network traffic rules, and then evaluate it with two darknet traffic datasets. the taxonomy applies concrete network traffic rules on source host flows extracted from the datasets to define five main types of anomalies we observed: scanning, one flow, backscatter, IP fragment, and small activities (see Section 3 for more). The evaluation results demonstrate the effectiveness of our proposal: we label anomalous events defined by the taxonomy for over 94% of all sources, suggesting a low unlabeled source rate. We obtain a few interesting findings on the evolution of different anomalies since 2006 (especially in recent years), helping to shed light on overall darknet traffic trends. Specifically, the results showed that most anomalies are characterized by just one or two anomalies. Later we discussed temporal and spatial dependency and parameter dependency of darknet traffic. We also deepen the analysis with other known techniques such as OS fingerprinting, Honeybot, and scanner fingerprints.

Moreover, we conduct a comparison between the taxonomy and *iatmon*. The comparison results highlight the consistency of labeling major anomalies defined by both classification schemes although some *iatmon* output is not accurate due to a bug. In summary, the proposed taxonomy demonstrates its effectiveness

¹ The Graduate University for Advanced Studies, Chiyoda, Tokyo 101–8430, Japan

² National Institute of Informatics, Chiyoda, Tokyo 101–8430, Japan

^{a)} junliu@nii.ac.jp

^{b)} kensuke@nii.ac.jp

with a much simpler classification scheme.

The main contributions of this paper are threefold: (1) Compared with *iatmon*, we propose a simpler state-of-the-art taxonomy of anomalies in darknet traffic based on network traffic rules, and demonstrate its effectiveness through evaluation with two longitudinal datasets. We also release the analysis tool in <http://www.fukuda-lab.org/darknet/index.html> to facilitate its usage by other researchers to perform their related analyses. (2) We leverage the active monitoring technique (Honeypot) to complement darknet. We conclude that traces collected by even a single honeypot server could help to gain further insights of anomalies which cannot be seen only using darknet. (3) To our best knowledge, this is the first time that a taxonomy of anomalies in darknet traffic is evaluated with such a longitudinal traffic analysis. Besides researchers in darknet traffic analysis, researchers who are interested in classifying network traffic, network security scientists and network operators may also use the output of the taxonomy as a reference for their network administration and attack prevention. The output is also useful for network traffic classification after anomaly detection or outlier detection that does not rely on signatures.

2. Related Work

In the past many efforts have been made to build darknet traffic monitoring systems [7], [23], [33].

The work [25] presented a first comprehensive analysis of darknet traffic observed in 2004 at four unused IPv4 address blocks. It showed the great diversity of darknet traffic, both temporally and spatially, and examined the dominant anomalies on popular ports. Another work attempted to discover how darknet traffic had evolved from 2004 to 2010 [32]. Reference [19] focused on the destination port usage in darknet with a dataset covering seven years. An analysis based on darknet of country-wide Internet outages in Egypt and Libya in 2011 was presented in Ref. [8]. In addition, a classification of one-way Internet traffic collected in live networks acts as a reference for darknet traffic analysis [15].

Some studies have been devoted to characterizing common anomalies in the Internet. Reference [12] provided an in-depth survey of darknet traffic analysis. The authors discussed a taxonomy of DDoS attacks and different types of scanning events in Ref. [22]. A more general taxonomy of network and computer attacks was presented in Ref. [17]. Reference [14] discussed adaptive detection for DDoS event in darknet traffic. Reference [2] provided a longitudinal examination of scanning activities observed over 12.5 years. A past work [24] proposed a backscatter analysis technique to infer DoS activity in the Internet. Recently entropy-based metrics in classifying darknet traffic patterns have been studied in Ref. [35]. A network activity classification scheme with specification-based finite state machine models of TCP, UDP, and ICMP traffic was introduced in Ref. [30]. By integrating this classification scheme and packet IAT distributions of sources, *iatmon* demonstrated its effectiveness in classifying one-way traffic [6]. Scan and backscatter analysis and detection were covered in Refs. [5], [21], [28]. Continuous efforts have been made to tackle Conficker through-

out the years [13], [26], [27]. Recently fast Internet-wide scanners [10], [11], [16] are widely used for research purpose and they provide a source for darknet traffic.

Finally, we highlight the differences of this work to the preliminary results [20] in three aspects. First, we have extended the network traffic rules of some types of anomalies, and added three new anomalies to reduce the ambiguity of “Others” anomaly. Second, we also evaluate the taxonomy with another /17 dataset which aggregates 11 years’ traffic. The last work is that we provide the comparison with the output of *iatmon* and demonstrate the effectiveness of the proposed simple classification scheme.

3. A Darknet Traffic Taxonomy

We propose a simple darknet traffic taxonomy that is the extension of Ref. [20]. **Table 1** summarizes all anomalies. To be clear, we claim that the term “flow” used in this paper always means “source host flow,” i.e., the whole incoming darknet traffic originating from a *ipSrc* within a time period (one day in this paper, see Section 5.4.1). Similarly, a (*ipSrc*, *ipDst*) flow denotes the whole darknet traffic from a *ipSrc* to a *ipDst* within one day.

3.1 Port Scan

In a port scan, the attacker sends client request packets to a number of server ports with the goal of finding an active port and then exploiting known vulnerabilities of the service corresponding to that port. Thus, we base our considerations on (*ipSrc*, *ipDst*) pairs and raise a port scan event when the number of distinct destination ports in a (*ipSrc*, *ipDst*) flow exceeds a threshold ($\#portDst \geq N_2$). Note that attackers can perform both TCP and UDP port scans. For TCP we also require the proportion of packets with scan flags (SYN \cup FIN \cup FINACK \cup NULL) to be larger than a threshold ($ScanFlagPktRatio \geq R\%$), in order to ensure that attackers are most likely to attempt to find active destination ports to exploit known vulnerabilities. Moreover, we specify two subcategories characterizing whether the scan traffic is heavy or light, depending on the average number of packets per destination port ($Avg \#Pkt \text{ per } portDst$).

3.2 Network Scan

Unlike a port scan, a network scan attempts to find victims with the same active port and either exploit known vulnerabilities of the service corresponding to that port or just recruit peers for launching larger distributed attacks on as many hosts as possible. We characterize a network scan event as a scan aimed at the same target port ($\#portDst == 1$) from a single source ($\#ipSrc == 1$) and involving several hosts ($\#ipDst \geq N_1$). Network scans can be performed with the TCP, UDP, and ICMP protocols. As with a port scan, we also require $ScanFlagPktRatio \geq R\%$ for TCP. For ICMP, only echo request (Ping) packets ($Type == 8$) \cap ($Code == 0$) are considered. For all three protocols we specify two subcategories (depending on $Avg \#Pkt \text{ per } ipDst$) for heavy and light scans.

3.3 One Flow

The notion of one flow characterizes large traffic ($\#Pkt > N_3$) destined to a destination port ($\#portDst == 1$) in a (*ipSrc*, *ipDst*)

Table 1 A Darknet traffic taxonomy.

Anomaly	Category		Darknet Traffic Rule
Port Scan	TCP	Heavy	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N_2) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg \#Pkt \text{ per } portDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N_2) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg \#Pkt \text{ per } portDst \leq M)$
	UDP	Heavy	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N_2) \cap (Avg \#Pkt \text{ per } portDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N_2) \cap (Avg \#Pkt \text{ per } portDst \leq M)$
Network Scan	TCP	Heavy	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N_1) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg \#Pkt \text{ per } ipDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N_1) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg \#Pkt \text{ per } ipDst \leq M)$
	UDP	Heavy	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N_1) \cap (Avg \#Pkt \text{ per } ipDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N_1) \cap (Avg \#Pkt \text{ per } ipDst \leq M)$
	ICMP	Heavy	$(\#ipSrc == 1) \cap (\#ipDst \geq N_1) \cap ((Type, Code) == (8, 0)) \cap (Avg \#Pkt \text{ per } ipDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#ipDst \geq N_1) \cap ((Type, Code) == (8, 0)) \cap (Avg \#Pkt \text{ per } ipDst \leq M)$
One Flow	TCP		$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst == 1) \cap (\#Pkt > N_3) \cap (Protocol == TCP)$
	UDP		$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst == 1) \cap (\#Pkt > N_3) \cap (Protocol == UDP)$
Backscatter	TCP		$(\#ipSrc == 1) \cap (\#Pkt \geq 1) \cap (TCP_Flags \in \{SA \cup A \cup R \cup RA\})$
	UDP		$(\#ipSrc == 1) \cap (\#Pkt \geq 1) \cap (\#portSrc \in \{53 \cup 123 \cup 137 \cup 161\}) \cap (Protocol == UDP)$
	ICMP		$(\#ipSrc == 1) \cap (\#Pkt \geq 1) \cap (((Type, Code) == (0, 0)) \cup (Type == 3) \cup ((Type, Code) == (11, 0)))$
IP Fragment			$(\#ipSrc == 1) \cap (\#FragmentPkt \geq 1)$
Small SYN			$(\#ipSrc == 1) \cap (\#ipDst < N_1) \cap (\#portDst < N_2) \cap (\#Pkt \leq N_3) \cap (TCP_Flags == S)$
Small UDP			$(\#ipSrc == 1) \cap (\#ipDst < N_1) \cap (\#portDst < N_2) \cap (\#Pkt \leq N_3) \cap (Protocol == UDP)$
Small Ping			$(\#ipSrc == 1) \cap (\#ipDst < N_1) \cap (\#Pkt \leq N_3) \cap ((Type, Code) == (8, 0))$
Others			Including “Other TCP”, “Other UDP”, “Other ICMP” and “Other”

Remark: The parameters $\{N_1 = N_2 = 5, R = 50, M = 3, N_3 = 15\}$ are empirically determined with real data, see Section A.1 for more

flow. This happens with both TCP and UDP protocols. Network misconfiguration is a plausible explanation for this anomaly.

3.4 Backscatter

Backscatter [24] consists of response packets to (D)DoS attacks carried out elsewhere in the Internet. Specifically, attackers forge packets (most often TCP SYN packets) and send those packets to victims to launch (D)DoS attacks while hiding themselves with spoofed source IP addresses. If the spoofed source IP addresses are located inside the darknets, then we get a chance to trap such backscatter traffic. For TCP, we use the TCP flags (SYNACK \cup ACK \cup RST \cup RSTACK) to detect backscatter. For UDP, since we have not found significant source ports in real traces, we decide to label packets with reply in application layer originating from source port 53 (DNS), 123 (NTP), 137 (NetBIOS) and 161 (SNMP) as UDP backscatter traffic. For ICMP, we instead consider echo reply $((Type == 0) \cap (Code == 0))$, destination unreachable $(Type == 3)$ and TTL exceeded $((Type == 11) \cap (Code == 0))$ packets as backscatter.

3.5 IP Fragment

This anomaly represents DoS attacks or attempts to defeat packet filter policies. The Rose Attack [18] is an example of exploiting the IP fragments “Too Many Datagrams,” “Incomplete Datagram,” and “Fragment Too Small.” We count the number of distinct sources $(\#ipSrc == 1)$ that send at least one fragmented packet $(\#FragmentPkt \geq 1)$ as the number of IP fragment events.

3.6 Small SYN

In the real traces, we notice that many sources send a limited number of SYN packets to limited destinations on limited destination ports within a time period. We use the term “small SYN” to characterize this type of event. Specifically, we count the number of distinct sources $(\#ipSrc == 1)$ that send a small number of SYN packets $((\#Pkt \leq N_3) \cap (TCP_Flags == S))$ to a few destinations $(\#ipDst < N_1)$ aimed at a small number of destination ports $(\#portDst < N_2)$ as the number of “small SYN” events.

3.7 Small UDP

The network traffic rules of “small UDP” are almost the same as “small SYN” except UDP packets are considered instead.

3.8 Small Ping

This is also similar to “small SYN” except all packets are ICMP echo requests (Pings). Thus, we count the number of distinct sources ($\#ipSrc == 1$) that send a small number of ICMP echo requests ($(\#Pkt \leq N_3) \cap ((Type, Code) == (8, 0))$) to a few destinations ($\#ipDst < N_1$) as the number of “Small Ping” events.

3.9 Others

Sources not labeled as any of the above anomalies fall into this category. In detail, it includes “other TCP,” “other UDP,” “other ICMP” which indicate the traffic in a source host flow includes unlabeled TCP, UDP, ICMP packets, respectively. For sources still not labeled, we label them as the “Other” anomaly.

Note that anomalies (port scan, network scan, and one flow) cover traffic of more than one protocol originating from one source, while small events (small SYN, small UDP, and small Ping) label one source based on specific packets. We emphasize that one event may overlap others (i.e., some packets can be part of multiple events), but it will never include or be included in other events, except for backscatters and IP fragments. These exceptions are due to the simple network traffic rules for backscatter and IP fragment: just one such packet will trigger them. Moreover, different from *iatmon*, the taxonomy allows one source to be characterized by multiple anomalies; for example, one source which sends both TCP and ICMP scan packets will be labeled as both TCP scan and ICMP scan anomalies.

4. Dataset

We analyze two real darknet traces in this paper. Dataset I continuously collects unexpected packets destined to one /18 allocated but unused IPv4 address block in Japan since Oct. 2006, unfortunately with four major data loss time periods. While dataset II aggregates unsolicited packets propagated to one different /17 IPv4 address block in Japan since Aug. 2007, with five major data loss time periods. The two datasets capture complete packet headers (layer -2, -3, and -4) and payloads. However, because of the low ratio of packets with payloads, the traffic analysis in this paper mainly relies on the header information.

To balance the scale and processing time of real traces, in this paper we just analyze the data of the first seven days per month from dataset I, and the first day per month from dataset II. This data covers 115 weeks from Oct. 2006 to Apr. 2017 (excluding the data loss time periods) in dataset I, and 89 weeks from Aug. 2007 to Apr. 2017 (excluding the data loss time periods) for dataset II.

5. Evaluation

5.1 Longitudinal Analysis

First, we compare datasets I and II with those used in Ref. [32] to check the similarity of network traffic behavior between them. **Figure 1** plots time series of protocol breakdown based on the number of packets for both datasets. A significant change occurred around Nov. 2008 in both plots, and since then TCP traffic (especially TCP SYN) has kept increasing until it accounted for over 70% of the packet volume in both datasets, thus dominating the complete traffic. As reported in Ref. [32], this change was




Fig. 1 The time series plots of packet volume proportion breakdown of (a) dataset I and (b) dataset II by TCP, UDP and ICMP.

confirmed to be Conficker’s outbreak in Nov. 2008 [1].

To label anomalies hidden by short time bins, we adopt a longer time bin (one day, see Section 5.4.1) and use the same empirically determined parameters for the two datasets in the analysis. We first extract daily source host flows from raw darknet traffic; then, we apply the taxonomy to these flows for labeling. We emphasize again that the taxonomy allows multiple labels for one source. **Table 2** summarizes the result for dataset I while **Table 3** is the summary for dataset II in terms of labeled source ratios for each anomaly. The dash “-” in the tables indicates a ratio of “<0.01%.”

The results of Tables 2 and 3 are consistent in general. They demonstrate that in both datasets small SYN and small UDP events are most popular, with a ratio of at least 10%. This indicates that more sources are likely to send only a few packets destined to a small number of hosts and destination ports within one day. An example of this indication is that we find such sources in real traces which have been proved to belong to a /0 stealth scan from a botnet in Feb. 2011 [9]. We also notice, however, that in dataset I the ratio of small SYN events increased rapidly in 2008, and then decreased from 69.77% in 2009 to 37.03% in 2014, while the proportion of small UDP events experienced a significant decrease from 2008 to 2009. We observe a significant increase in light TCP network scans from 2008 to 2009, and it shows that more and more attackers have preferred to apply TCP network scans with light traffic since 2008. The proportion of small Ping events, on the other hand, is much lower than the other two small events. We confirm again that these findings result from Conficker’s outbreak in 2008.

As for scanning events in dataset I, so far they are not popular choices among attackers, except for light TCP and ICMP network scans. We confirm the existence of Internet-wide scanners for research purpose like ZMap [11] and malicious ones like Mirai Botnet [29], [31]. We also find that attackers generally prefer light scanning to heavy scanning because light traffic is more likely to evade detections by intrusion detection systems (IDS). Although light UDP network scans show an overall trend of decreasing, while light ICMP network scans have kept increasing in recent years, together they cover < 2% of sources in the traces.

In dataset I in the first three years ICMP backscatter events covered more than 5% of sources and increased slowly from 0.06%

Table 2 Evaluation Result Overview in Darknet Dataset I (ipSrc %, H: Heavy, L: Light, -: <0.01).

Anomaly			Dataset I											
			06	07	08	09	10	11	12	13	14	15	16	17
Port Scan	TCP	H	-	-	0.03	-	0.01	-	-	0.02	-	-	-	-
		L	0.01	-	0.03	-	-	-	-	-	0.01	0.01	-	0.03
	UDP	H	-	-	0.03	-	0.04	0.02	0.01	0.01	0.01	-	-	-
		L	-	0.01	0.01	-	0.02	0.01	0.01	0.01	-	-	-	-
Network Scan	TCP	H	0.05	0.03	0.06	0.02	0.03	0.04	0.06	0.19	0.63	0.35	1.60	0.84
		L	0.61	0.39	4.00	17.06	20.70	21.98	21.37	22.03	21.88	15.27	31.22	49.24
	UDP	H	-	0.29	0.01	-	-	-	0.01	0.01	0.01	0.01	-	-
		L	1.35	0.87	1.25	0.06	0.04	0.07	0.06	0.07	0.14	0.21	0.56	0.47
	ICMP	H	-	-	0.01	-	-	-	-	-	-	0.01	-	-
		L	0.14	0.06	0.14	0.01	0.07	0.56	0.63	0.97	1.14	0.73	0.27	0.25
One Flow	TCP		1.92	1.32	2.33	0.18	0.19	0.12	0.33	0.25	0.19	5.13	4.70	2.22
	UDP		1.24	1.55	3.53	0.12	0.26	0.15	0.18	0.22	0.22	0.57	0.44	0.16
Backscatter	TCP		1.86	2.39	2.55	0.87	1.01	0.86	1.04	1.13	0.85	1.46	0.57	0.62
	UDP		0.15	0.23	0.38	0.01	0.02	0.03	0.04	0.05	0.09	0.11	0.25	0.01
	ICMP		11.35	15.78	5.55	0.06	0.13	0.17	1.11	1.83	2.02	2.66	0.87	0.46
IP Fragment			0.05	0.02	0.01	-	-	-	0.01	0.01	0.01	0.06	-	-
Small SYN			42.30	18.53	34.47	69.77	67.38	61.20	54.63	40.83	37.03	45.47	51.27	38.98
Small UDP			36.32	54.28	39.50	13.42	10.75	14.23	18.40	28.38	31.31	46.21	28.40	14.95
Small Ping			5.32	1.82	3.85	0.19	0.25	1.42	1.64	2.73	3.84	2.02	0.53	0.43
Other TCP			0.16	0.16	0.32	0.07	0.12	0.09	0.39	0.44	0.41	0.62	1.02	2.72
Other UDP			0.32	4.08	3.72	0.12	0.46	0.19	0.91	1.79	2.41	1.89	0.27	0.11
Other ICMP			0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Other			0.02	0.02	0.06	0.01	0.01	0.02	0.04	0.04	0.03	0.02	0.01	0.06

Table 3 Evaluation Result Overview in Darknet Dataset II (ipSrc %, H: Heavy, L: Light, -: <0.01).

Anomaly			Dataset II										
			07	08	09	10	11	12	13	14	15	16	17
Port Scan	TCP	H	-	0.02	0.01	0.11	0.03	0.02	0.03	-	-	-	-
		L	0.01	0.02	-	0.01	-	0.01	0.01	0.01	-	-	-
	UDP	H	1.34	1.01	0.07	0.32	0.08	0.10	0.05	0.03	0.02	0.01	0.01
		L	0.01	0.01	0.01	0.13	0.02	0.03	0.02	0.02	0.01	-	-
Network Scan	TCP	H	0.48	0.44	0.45	0.44	0.74	0.39	0.42	0.27	0.20	1.80	0.72
		L	4.14	6.45	7.66	14.44	16.85	9.82	4.41	4.34	5.17	33.41	29.67
	UDP	H	0.09	0.08	0.02	0.02	0.02	0.02	0.01	0.01	0.01	0.01	-
		L	0.85	1.09	0.25	0.13	0.22	0.17	0.14	0.19	0.20	0.82	0.23
	ICMP	H	0.02	0.03	0.02	0.01	0.02	0.02	0.02	0.01	0.01	0.01	-
		L	1.88	2.33	0.48	0.30	2.51	1.79	2.38	1.17	0.55	0.35	0.19
One Flow	TCP		1.21	2.19	0.25	0.86	0.27	0.20	0.15	0.14	0.19	0.13	0.06
	UDP		3.66	4.43	0.57	1.79	0.95	0.46	0.70	0.99	0.80	0.22	0.08
Backscatter	TCP		1.06	2.68	2.30	2.09	2.82	1.82	3.10	1.94	0.89	0.79	0.52
	UDP		0.34	0.15	0.06	0.11	0.07	0.05	0.06	0.02	0.03	1.16	0.01
	ICMP		21.19	8.07	0.29	0.69	0.41	2.00	3.05	1.85	1.81	2.60	0.37
IP Fragment			0.01	0.01	0.02	0.01	-	0.01	0.01	0.01	0.01	-	-
Small SYN			14.78	18.44	31.02	21.81	20.63	18.96	12.47	9.20	7.15	24.61	14.14
Small UDP			43.62	47.09	63.84	57.00	52.05	60.48	64.31	75.11	81.76	35.25	53.22
Small Ping			3.52	3.92	0.82	0.72	5.00	3.18	4.92	2.78	1.13	0.50	0.23
Other TCP			0.20	0.43	0.18	0.97	0.27	0.24	0.27	0.61	0.16	0.78	1.36
Other UDP			4.86	4.17	0.58	3.88	0.78	2.67	5.27	3.39	1.89	0.31	0.12
Other ICMP			0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Other			0.03	0.03	0.92	0.24	0.08	0.11	0.21	0.05	0.02	0.02	0.03

in 2009 to 2.66% in 2015. Compared to ICMP, TCP backscatter events in our traces were relatively stable, ranging from 0.62% to 2.55%. We also find that a quite small number of sources are labeled as UDP backscatters throughout both datasets. From the backscatter results we conclude that the observed spoofed-source (D)DoS attacks from the datasets keep relatively inactive in recent years.

One flow events mainly result from network misconfiguration. In dataset I, both TCP and UDP one flow events exhibit an overall trend of decreasing. By examining one flow raw packets, we find that both single and multiple source ports are possible.

We find that the most popular exploited ports by TCP port scanners were 3389 (RDP) and 22 (SSH) while the most popular port for UDP was 58904 (unknown) in dataset I. Turning to network

Table 4 Known Scanner Ratios in 2015 (ipSrc %).

Anomaly	Zmap	masscan
Light TCP Network Scan	0.01	<0.01
Small SYN	0.02	0.01

scan events, we observe that 23 (Telnet) and 445 (Microsoft-DS) dominate TCP while 53413 (unknown) dominates UDP. As expected, port 80 (HTTP) dominates among TCP backscatter source ports, suggesting (D)DoS attacks to web servers. Port 54668 (unknown) is the most popular for TCP one flow events, while ports 137 (NetBIOS) and 53 (DNS) dominate for UDP ones. The detailed destination port usage in the traces has been reported in Ref. [19].

Throughout the 12 years' data in dataset I, the proportion of IP fragmentation exploits is almost negligible. Last but not least, we point out that Others (defined in Section 3.9) maintained a low proportion (the highest is 4.41% in 2007), demonstrating that the taxonomy labels most sources in darknet traffic.


As to dataset II, we find quite similar results in some aspects like small ratios of scanning except light TCP and ICMP network scans. We also find a larger proportion of sources are labeled by backscatters both in TCP and ICMP. However, a notable difference between dataset I and dataset II is that small UDP dominates small SYN in dataset II, which is the opposite case in dataset I. We later confirmed that this is exactly due to the difference in traffic composition: there are larger proportions of UDP sources than TCP ones in dataset II.

5.2 Anomaly Cause Investigation

We are interested in how much some certain known scanners/worms/botnets contribute to the main darknet anomalies. We mainly rely on two more microscopic techniques: (1) known scanner footprints and (2) attacker packet signatures to identify them and investigate their contributions to darknet anomalies.

According to Ref. [11], Internet-wide scanners for research purposes (Zmap and masscan) leave their specific footprints in TCP/IP headers, thus making their identification in darknet easy. **Table 4** shows known scanner ratios using a weekly trace from dataset I in 2015. As listed in the table, in terms of the number of ipSrcs they have little impact on the main darknet anomalies.

Attackers also have their own packet signatures. However, for many worms (e.g., Conficker) their signatures in packet payloads are more accurate than in TCP/IP headers, thus making their identification in darknet difficult. To solve this, we employ a single host belonging to a subnet next to IP address blocks in dataset I as a honeypot to gather further traffic from infected ipSrcs. In total, we obtain over 0.35 million distinct ipSrcs. Then we match these ipSrcs with those appeared in dataset I. Later a careful inspection of packet payloads of suspicious ipSrcs helps identify Conficker according to Ref. [13]. Although the ratio of the matched ipSrcs out of darknet ipSrcs is quite low (less than 0.5% per daily trace), we find that less than 0.1% of the matched ipSrcs are probable to be infected by Conficker. We see that honeypot monitors actively, even a small trace collected on a single honeypot server could help complement darknet and gain new insights which cannot be seen only with darknet.

**Fig. 2** The ratios of Mirai Botnet traffic (dataset I).

As we know, Confickers are highly related to TCP destination port 445. We have examined the ratios of ipSrcs related to 445/TCP in the main darknet anomalies (“Light TCP Network Scan” and “Small SYN”) in 2010 and 2015. Comparing the higher ratios (over 30%) in 2010, the relatively low ratios (around 10%) in 2015 indicate that Conficker is probable to be in a decay phase in 2015. This finding is consistent with the results in Ref. [4].

Recently in the two datasets we have also noticed the outbreak of a significant malicious scanner, i.e., Mirai Botnet [3], [29], [31] since Aug. 2016. Mirai takes the advantage of Telnet vulnerabilities to perform a wide-range scanning. Mirai also sets the destination IP address of its SYN packet to Initial Sequence Number (ISN) and performs a scan against destination ports TCP/23 and TCP/2323. **Figure 2** shows the ratios of Mirai packets and Mirai ipSrcs since its outbreak, and the statistics are consistent in both datasets. In Nov. 2016, Mirai traffic reached its peak ratio with over 60% for both packets and ipSrcs. Further analysis shows that Mirai ipSrcs are mainly labeled as “Light TCP Network Scan” (over 60%) and “Small SYN” anomalies (over 20%), which is expected with the taxonomy.

For identifying the OS distribution of source hosts, we apply a passive OS fingerprinting using p0f [34] for “Small SYN” and “Small UDP” anomalies in dataset I. The result turns out that OS types for > 50% of the two anomalies are unknown, but we still find that Linux (kernel version from 2.2 to 3.2) and Windows (from XP to 8) are the most popular OS types, whereas we hardly see anomalies originating from MacOS or FreeBSD.

5.3 Diversity of Anomalies per ipSrc

Of particular interest to us is whether source IPs in darknet typically exhibit simple or complicated attack mechanisms. To obtain clues to this issue, we summarize the overall proportions of source IPs with different numbers of labels for dataset I. We clearly find that source IPs with one or two labels are the vast majority, together accounting for over 99% of all sources. The extremely high percentage of source IPs with only one label (97.83%) also highlights that most sources are characterized by one simple event. Digging deeper, we find that labels “small SYN,” “light TCP network scan,” and “small UDP” together account for over 95% of the sources with one label, while label combinations “small SYN and small UDP” and “small SYN and TCP backscatter” dominate among the sources with two labels. These straightforward dominant labels and label combinations indicate

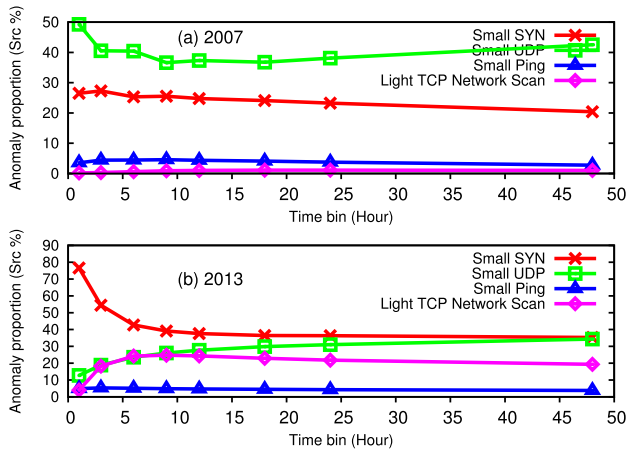


Fig. 3 Plots of time bin size dependency in terms of anomaly proportion using dataset I real traces in (a) 2007 and (b) 2013.

a low possibility of complicated attack mechanisms [21].

5.4 Time and Space Dependency

5.4.1 Time Bin Size Dependency

Determining the typical lifetime for anomalies in a darknet is crucial to their accurate labeling. On one hand, if we choose a time bin shorter than the typical lifetime, for example, some scanning events would likely be miscategorized as small SYN events or just be neglected. On the other hand, a longer time bin could lead to redundant packets mixing into specific anomalies, as well as requiring longer processing time. Thus, to understand this key parameter, we set the time bin to different sizes and analyzed sample traces from dataset I in 2007 and 2013.

Figure 3 (a) and (b) plot the results for time bin size dependency in terms of labeled source ratios before and after Conficker's outbreak in 2008. In Fig. 3(a), small UDP events first decrease then increase slightly while small SYN keeps decreasing as time bin gets longer. However, we observe that small SYN events decrease rapidly from one-hour to six-hour time bin whereas light TCP network scan and small UDP events increase meanwhile. We also notice that with a time bin between six hours and one day both plots show just small fluctuations.

The results show that the lower percentage of small SYN events with a time bin between six hours and one day also means a higher label rate for other types of events, like light TCP network scans. Thus, we conclude that the best time bin for detecting more significant anomalies in the dataset is between six hours and one day (one day used in this work).

5.4.2 Darknet Space Size Dependency

To understand how taxonomy influence labeling accuracy for darknets with different space sizes, we divide sample traces from dataset I into a few subtraces with different subnetwork sizes, and then apply the taxonomy on them with the same parameters used for /18 block before. The results are plotted in **Fig. 4**.

From Fig. 4(a) we see clearly that the parameter of darknet space size is independent of the detectability before Conficker's outbreak. However, Fig. 4(b) shows that small UDP and light TCP network scan events keep decreasing whereas the vast majority of traffic are characterized as small SYN events with the same parameters used for /18 block as darknet gets smaller after




Fig. 4 Plots of darknet space size dependency in terms of anomaly proportion using dataset I real traces in (a) 2007 and (b) 2013.

Conficker's outbreak. As the subnet gets smaller, the number of packets arrived at darknet is not enough for labeling of anomalies like network scans, thus leading to small SYN ratio becoming higher and higher. From this result obviously we understand that a small darknet space size (e.g., /24) is not large enough for labeling anomalies like network scans. Therefore we conclude that darknet space size strongly affects its detectability even for a long time window (i.e., one day). However, we highlight that our taxonomy is generally applicable to different darknets, though parameter tuning is required for accurate labeling.

5.5 Comparison with *iatmon*

To understand the differences between the taxonomy and *iatmon*, we compare them using sample traces from both datasets.

A confusion matrix based on #*ipSrc* is shown in **Table 5**. The output types of *iatmon* (*iatmon* Src Type #) indicate: (0) TCP port scan; (1) UDP port scan; (2) TCP network scan; (3) UDP network scan; (4) ICMP only; (5) TCP one flow; (6) UDP one flow; (7) Backscatter; (8) 1 or 2 packets; (9) both TCP and UDP; (10) TCP unknown; (11) UDP unknown; (12) μ Torrent; (13) Conficker P2P; (14) Unclassified. The numbers separated by the slash “/” in each cell represent the number of sources labeled by both the taxonomy and *iatmon* in dataset I and II, respectively.

Note that the confusion matrix is generated with the output of *iatmon* in which we fixed a bug: the bug leads *iatmon* to incorrectly classify “(1) UDP port scan” as “(0) TCP port scan.”

From this table we see that anomalies in *iatmon* like “TCP and UDP port scan,” “TCP and UDP network scan,” “TCP and UDP one flow,” and “Backscatter” are consistent with the corresponding ones in the taxonomy. Moreover the taxonomy classifies all scans into “Heavy” and “Light” categories.

We notice that the “(4) ICMP only” anomaly of *iatmon* can be further divided into “ICMP network scan” and “ICMP backscatter” anomalies defined in the taxonomy. We also find that the majority of “(7) Backscatter” in *iatmon* is confirmed to be TCP backscatter. Moreover, the “TCP unknown” and “UDP unknown” are labeled as many types of anomalies in the taxonomy. Also, we notice that there is not even single source labeled by “(9) both TCP and UDP” in *iatmon*. “(12) μ Torrent” and “(13) Conficker P2P” detection is based on matching documented TCP/IP header

Table 5 A confusion matrix generated with sample traces in the two datasets (X/Y: #ipSrc in dataset I/II).

iatmon ipSrc Type #			0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Port Scan	TCP	H	11/1	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	36/8	66/35	3/1	0/0	0/1
		L	7/3	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	35/22	2/1	0/0	0/0	0/0
	UDP	H	0/0	15/45	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	13/0	85/46	149/72	0/0	3/0
		L	0/0	41/41	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	4/3	24/11	36/9	0/0	1/5
Network Scan	TCP	H	0/0	0/0	1K/2K	0/0	0/0	1/0	0/0	0/0	0/0	0/0	68/55	7/1	0/0	1/0	1/1
		L	0/0	0/0	389K/16K	0/0	0/0	0/0	0/0	1/0	1/0	0/0	2K/1K	16/10	0/0	11/4	31/16
	UDP	H	0/0	0/0	0/0	49/28	0/0	0/0	0/0	5/0	0/0	0/0	3/1	21/17	0/0	1/0	1/0
		L	0/0	0/0	0/0	3K/325	0/0	0/0	0/0	2/0	0/0	0/0	89/69	128/39	0/0	0/0	72/20
	ICMP	H	0/0	0/0	0/0	0/0	46/40	0/0	0/0	0/0	0/0	0/0	3/1	0/2	0/0	0/0	3/3
		L	0/0	0/0	0/0	0/0	8K/6K	0/0	0/0	0/0	0/0	0/0	118/67	52/12	0/0	0/0	202/112
One Flow	TCP		0/0	0/0	0/0	0/0	0/0	4K/193	0/0	15/1	0/0	0/0	514/89	109/141	4/2	3/1	7/2
	UDP		0/0	0/0	3/0	0/0	0/0	0/0	5K/6K	3/2	2/0	0/0	128/11	282/228	350/159	0/0	8/4
Backscatter	TCP		89/2K	0/0	4K/434	0/0	0/0	5K/2K	0/0	3K/2K	1K/400	0/0	715/2K	120/25	7/0	1K/114	77/31
	UDP		0/0	0/0	0/0	0/0	0/0	0/0	9/0	15/7	0/0	0/0	0/0	4/3	0/0	0/0	0/0
	ICMP		0/0	0/0	0/0	0/0	4K/4K	0/0	6/0	0/0	560/444	0/0	79/60	595/1K	0/0	0/0	290/432
IP Fragment			1/0	0/0	0/1	0/0	0/0	13/0	5/5	0/0	3/0	0/0	15/0	11/8	0/0	2/0	0/0
Small SYN			288/72	0/0	445K/16K	1/0	0/0	74K/13K	4/0	2/2	574K/4K	0/0	5K/1K	1K/1K	206/66	7K/1K	179/84
Small UDP			0/0	769/1K	17/0	1K/188	0/0	2/0	42K/58K	4/1	19K/17K	0/0	4K/1K	9K/13K	4K/2K	5K/1K	388/522
Small Ping			0/0	0/0	6/0	0/0	5K/3K	0/0	0/0	0/0	6K/4K	0/0	593/202	99/62	0/0	0/0	154/54
Other TCP			158/86	0/0	371/244	0/0	0/0	31/85	0/0	8/0	12/13	0/0	1K/1K	0/0	0/0	13/18	34/0
Other UDP			0/0	270/852	0/0	159/99	0/0	0/0	91/0	0/0	3/0	0/0	0/0	9K/21K	2K/587	27/1	1/1
Other ICMP			0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0	0/0
Other			0/0	0/0	0/0	1/0	211/399	0/0	0/0	0/0	3/5	0/0	52/17	484/256	24/8	17/4	70/147

patterns, but we mainly focus on macroscopic behaviors of longitudinal darknets, thus omitting rules to detect such specific anomalies. We emphasize that the large number of sources labeled as “Others,” “small SYN,” “small UDP” and “small Ping” is due to different parameters used in the taxonomy and *iatmon*, and the fact that the taxonomy allows one source to have multiple labels whereas *iatmon* labels each source only once. *iatmon* outputs one label for a single ipSrc flow within a given time interval. However, a proper time interval for *iatmon* requires careful tuning because we do not know active period of each flow in advance. Differently, the taxonomy could output multiple labels without caring much about the time interval.

We briefly compare the speed of the taxonomy and *iatmon*. For a daily trace of 500 MB, it takes less than 30 seconds for *iatmon* to complete analysis, whereas the taxonomy takes 1 minute. We claim that this is mainly due to that *iatmon* core functions are written in C while the taxonomy codes are written in Python.

In summary, the taxonomy can obtain a similar labeling effectiveness as *iatmon* with a much simpler classification scheme.

6. Conclusion

In this paper, we proposed a simple but effective darknet traffic taxonomy based on network traffic rules and analyzed longitudinal traces from two datasets to evaluate this proposal. The results showed an extremely high ipSrc labeling rate of over 94%.

Through examining the evolution of anomalies, we obtained a few interesting findings. We highlighted that small SYN and small UDP anomalies dominated throughout both traces, while light TCP network scans have become more active in recent years. We also confirmed that Conficker worm has lost its great influence on darknet traffic as of this work. In addition, we concluded that the observed spoofed-source (D)DoS attacks from the

datasets and network misconfiguration events have kept relatively inactive recently. We gained further insights with honeypot to complement darknet in anomaly cause investigation. We demonstrated that darknet is an effective approach for passively monitoring Internet-wide scanners for research or malicious purposes.

We determined that the most appropriate time bin for the analysis of the datasets is between six hours and one day. Also, we investigated the dependency on darknet space size and highlighted the general applicability of the taxonomy for different darknets with appropriate parameters. We examined the impacts of different parameters on the taxonomy as well in Section A.1. Furthermore, we emphasized that most sources are characterized by one or two anomalies, and they most often deploy simple attack mechanisms.

Finally, by comparing with *iatmon*, we concluded that the taxonomy can achieve a similar effectiveness for labeling darknet traffic anomalies with a much simpler classification scheme.

Acknowledgments We thank Johan Mazel and Romain Fontugne for their comments and suggestions. This research has been partially supported by JSPS KAKENHI Grant Number 15H02699 and 15KK0019, and MIC and FP7 under grant agreement No.608533 (NECOMA).

References

- [1] Aben, E.: Conficker/Conflicker/Downadup as seen from the UCSD network telescope, CAIDA Technical Report (2009).
- [2] Allman, M., Paxson, V. and Terrell, J.: A brief history of scanning, *IMC'07*, pp.77–82 (2007).
- [3] Antonakakis, M. et al.: Understanding the Mirai Botnet, *USENIX Security'17*, pp.1093–1110 (2017).
- [4] Asghari, H., Ciere, M. and van Eeten, M.J.: Post-Mortem of a Zombie: Conficker Cleanup After Six Years, *USENIX Security'15*, pp.1–16 (2015).
- [5] Balkanli, E. and Zincir-Heywood, A.N.: On the analysis of backscatter traffic, *LCN 2014 Workshops*, pp.671–678 (2014).

- [6] Brownlee, N.: One-way traffic monitoring with *iatmon*, *PAM'12*, pp.179–188 (2012).
- [7] Cooke, E., Bailey, M., Watson, D., Jahanian, F. and Nazario, J.: The Internet motion sensor - A distributed blackhole monitoring system, *NDSS'05*, pp.167–179 (2005).
- [8] Dainotti, A., Squarcella, C., Aben, E., Claffy, K. and Chiesa, M.: Analysis of country-wide Internet outages caused by censorship, *IMC'11*, pp.1–18 (2011).
- [9] Dainotti, A., King, A., Claffy, K., Papale, F. and Pescapé, A.: Analysis of a /0 stealth scan from a botnet, *IMC'12*, pp.1–14 (2012).
- [10] Durumeric, Z., Bailey, M. and Halderman, J.A.: An Internet-wide View of Internet-wide Scanning, *SEC'14*, pp.65–78 (2014).
- [11] Durumeric, Z., Wustrow, E. and Halderman, J.A.: ZMap: Fast Internet-wide Scanning and Its Security Applications, *USENIX Security'13*, pp.605–620 (2013).
- [12] Fachkha, C. and Debbabi, M.: Darknet as a Source of Cyber Intelligence: Survey, Taxonomy, and Characterization, *IEEE Communications Surveys Tutorials*, Vol.18, No.2, pp.1197–1227 (2016).
- [13] Felix Leder, T.W.: Know Your Enemy: Containing Conficker, available from (<https://www.honeynet.org/papers/conficker>)
- [14] Furutani, N., Kitazono, J., Ozawa, S., Ban, T., Nakazato, J. and Shimamura, J.: Adaptive DDoS-Event Detection from Big Darknet Traffic Data, *ICONIP 2015*, pp.376–383 (2015).
- [15] Glatz, E. and Dimitropoulos, X.: Classifying Internet one-way traffic, *IMC'12*, pp.37–50 (2012).
- [16] Graham, R.: Masscan Github Repository, available from (<https://github.com/robertdavidgraham/masscan>)
- [17] Hansman, S. and Hunt, R.: A taxonomy of network and computer attacks, *Computers & Security*, Vol.24, No.1, pp.31–43 (2005).
- [18] Hollis, K.: The Rose Attack Explained, available from (<http://www.digital.net/~gandalf/Rose-Frag-Attack-Explained.htm>)
- [19] Liu, J. and Fukuda, K.: On destination port usage in darknet, *IEICE General Conference 2014*, p.553 (2014).
- [20] Liu, J. and Fukuda, K.: Towards a Taxonomy of Darknet Traffic, *TRAC'14*, pp.37–43 (2014).
- [21] Lyon, G.: Nmap network scanning: The official Nmap project guide to network discovery and security scanning, available from (<http://nmap.org/book/man-port-scanning-techniques.html>)
- [22] Mirkovic, J. and Reiher, P.: A taxonomy of DDoS attack and DDoS defense mechanisms, *ACM SIGCOMM Comput. Commun. Rev.*, Vol.34, No.2, pp.39–53 (2004).
- [23] Moore, D., Shannon, C., Voelker, G. and Savage, S.: Network telescopes, *CAIDA Tech. Rep.* (2004).
- [24] Moore, D., Voelker, G. and Savage, S.: Inferring Internet denial-of-service activity, *ACM Trans. Computer Systems*, Vol.24, No.2, pp.115–139 (2006).
- [25] Pang, R., Yegneswaran, V., Barford, P., Paxson, V. and Peterson, L.: Characteristics of Internet background radiation, *IMC'04*, pp.27–40 (2004).
- [26] Piscitello, D.: Conficker Summary and Review, available from (<https://www.icann.org/en/system/files/files/conficker-summary-review-07may10-en.pdf>)
- [27] Porras, P., Saidi, H. and Yegneswaran, V.: A Foray into Conficker's Logic and Rendezvous Points, *LEET'09*, p.7 (2009).
- [28] Rossow, C.: Amplification Hell: Revisiting Network Protocols for DDoS Abuse, *NDSS'14* (2014).
- [29] Savage, K.: A Post-Mortem on the Mirai Botnet: Part 2: Analyzing the Attack, available from (<https://www.pwnieexpress.com/blog/mirai-botnet-part-2>) (2016).
- [30] Treurniet, J.: A network activity classification schema and its application to scan detection, *IEEE/ACM Trans. Networking*, Vol.19, No.5, pp.1396–1404 (2011).
- [31] Williams, C.: Today the web was broken by countless hacked devices - your 60-second summary, available from (http://www.theregister.co.uk/2016/10/21/dyn_dns_ddos_explained/) (2016).
- [32] Wustrow, E., Karir, M., Bailey, M., Jahanian, F. and Huston, G.: Internet background radiation revisited, *IMC'10*, pp.62–74 (2010).
- [33] Yegneswaran, V., Barford, P. and Plonka, D.: On the design and use of Internet sinks for network abuse monitoring, *RAID'04*, pp.146–165 (2004).
- [34] Zalewski, M.: p0f v3 (version 3.08b), available from (<http://lcamtuf.coredump.cx/p0f3/>)
- [35] Zseby, T., Brownlee, N., King, A. and Claffy, K.: Nightlights: Entropy-based metrics for classifying darkspace traffic patterns, *PAM'14*, pp.275–277 (2014).

Appendix

A.1 Parameter Dependency of the Taxonomy

To examine the impacts of different parameters on the results of the darknet traffic taxonomy, we first set the parameters to empirical values shown in Table 1, then conduct a series of experiments by changing only a single parameter each time to check its parameter dependency. **Figure A.1** shows the changed labeled ratios of major anomalies in the taxonomy as a single parameter changes. Since N_1 and N_2 play as a boundary between “Small SYN” and “Light TCP Network Scan,” their ratios exhibit a significant contrary trend as N_1 and N_2 change in Fig. A.1 (a). The ratio of “Small UDP” stays quite stable in the four plots, and it shows only a little increment as N_1 and N_2 increase in Fig. A.1 (a) and as N_3 gets bigger in Fig. A.1 (b). We also find that as parameter R changes, the ratios of major anomalies are almost the same. We claim that this result is consistent with the simple attack mechanism of darknet ipSrcs mentioned in Section 5.3, i.e., once a ipSrc sends SYN packets, it is highly probable that it hardly sends other kinds of packets and most of its packets are likely to be SYN packets. As shown in Fig. A.1 (d), parameter M is the boundary between heavy and light scans, thus it is expected that the ratio of “Light TCP Network Scan” also grows with M 's growth. From the figure, we conclude that in general R and M have a limited impact on the results of the taxonomy, while N_1 and N_2 and N_3 are much more powerful in labeling the major anomalies.

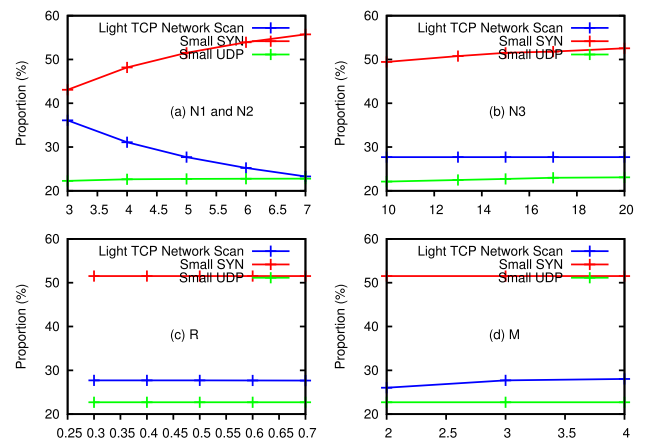


Fig. A.1 Plots of parameter dependency of the taxonomy in terms of anomaly proportion using sample traces from dataset I, (a) N_1 and N_2 , (b) N_3 , (c) R and (d) M .



Jun Liu received his Bachelor's degree in Computer Science from the University of Science and Technology of China (USTC) in 2012. He is now a Ph.D. candidate in the Department of Informatics at the Graduate University for Advanced Studies (SOKENDAI). His current research interests are Internet traffic mea-

surement and analysis, data mining and machine learning applications in big data.



Kensuke Fukuda is an associate professor at the National Institute of Informatics. He received his Ph.D. degree in Computer Science from Keio University in 1999. He was a visiting scholar at Boston University in 2002, and a visiting scholar at the University of Southern California/Information Sciences Institute

in 2014–2015. He was also a researcher of PRESTO JST in 2008–2012. His current research interests are Internet traffic measurement and analysis, intelligent network control architectures, and the scientific aspects of networks.