Towards a Taxonomy of Darknet Traffic

Jun Liu
Department of Informatics,
The Graduate University for
Advanced Studies
junliu@nii.ac.jp

Kensuke Fukuda
National Institute of Informatics/
The Graduate University for
Advanced Studies
kensuke@nii.ac.jp

***Abstract*—Darknets can be used to monitor unexpected network traffic destined for allocated but unused IP address blocks, thus providing an effective traffic measurement technique for viewing certain remote network security events. Past works in this field discussed the possible causes (events) of darknet traffic and applied their classification schemes on short-range traces. Our interest lies, however, in how darknets have evolved since those works and the effectiveness of a darknet taxonomy for real long-range traffic. We thus propose a simple but effective taxonomy of darknet traffic, on the basis of observations, and evaluate it on real darknet traces covering six years. The evaluation results show that we can detect and label anomalous events defined by the taxonomy for over 96% of all sources, making the unlabeled source rate extremely low. We also obtain some interesting findings on the evolution of different anomalous events since 2006 (especially in recent years), determine the most appropriate time bin for traffic analysis of our traces, and highlight the general applicability of our taxonomy on different darknet datasets. Finally, we conclude that most sources in our traces are characterized by just one or two events with simple attack mechanisms.**

***Index Terms*—Traffic Analysis, Darknet, Taxonomy.**

I. Introduction

Along with the rapid growth in Internet usage, network security issues have become more difficult to deal with in recent years. By providing an opportunity to view and detect remote network security events, darknets [1] [2] (a.k.a., network telescope [7]) have drawn much attention in the security research community. A darknet consists of globally routable but still unused IP blocks in which little or no legitimate traffic exists. Continual monitoring of such addresses, however, shows that unexpected packets keep arriving at darknets with not low rates from a wide range of sources. These unwanted packets are completely non-productive, since they originate from worm propagation, (D)DoS attacks, Internet outages, network misconfiguration, or other unsolicited events. Darknet traffic can be used to track such security-related activities on a global scale. An analysis of country-wide Internet outages in Egypt and Libya in 2011 based on darknet traffic [3] serves as a best example of this approach.

Past studies [1] [2] have showed the not minor volume of darknet traffic and its great diversity both in terms of the addresses being monitored as well as over time. They also claimed that darknet traffic broadly comes from three types of network events: scanning, backscatter, and misconfiguration. An evaluation has been performed with their dataset, however, this classification did not give clear definitions of events with concrete traffic rules, and a further refinement was required since it was simply based on TCP flags. As described in those works, scanning is largely the result of infected hosts in the Internet attempting to find other vulnerable targets; backscatter most often results from (D)DoS attacks; and misconfiguration generally results from software or hardware errors in network devices.

Some works on one-way traffic analysis helped us better understand darknet traffic considering its unidirectional nature. The one-way traffic analysis tool *iatmon* [16] [18] classifies traffic with two schemes: activity patterns of sessions, created using finite state machine models of host-pair packet-level behavior [17], and packet inter-arrival time (IAT) percentage distributions of sources. The author has evaluated the tool with a half-year trace in 2011. Finding recognizable IAT patterns, however, takes much effort, and the efficiency of classifying long-range darknet traces has not been examined yet.

To examine the evolution of darknet traffic within long-range traffic and avoid the hard work of obtaining IAT patterns, we need a simple but effective taxonomy of darknet traffic. To the best of our knowledge, no such a simple taxonomy has been developed within the research community. In this paper, therefore, we propose a simple taxonomy of darknet traffic, on the basis of observations, and then evaluate the taxonomy on real darknet traces covering six years. Our taxonomy applies concrete traffic rules on source flows generated from our dataset to define five main types of anomalous events we observed: scanning, one flow, backscatter, IP fragment, and small events (see section III for detailed explanations). The evaluation results demonstrate that we can detect and label anomalous events defined by the taxonomy for over 96% of all sources, suggesting an extremely low unlabeled source rate. We obtain some interesting findings on the evolution of different anomalous events since 2006 (especially in recent years), helping to shed light on overall darknet trends. We also determine the most appropriate time bin for our dataset and highlight the general applicability of our taxonomy on different darknet datasets. Finally, we conclude that most sources in our traces are characterized by just one or two events with simple attack mechanisms.

The rest of this paper is organized as follows. Section II summarizes some related work on darknet traffic analysis.

978-1-4799-0959-9/14/$31.00 ©2014 IEEE

Section III introduces a simple but effective taxonomy of darknet traffic with concrete traffic rules. Section IV describes our experiment dataset and shows the evaluation results of our proposal. Section V gives a further discussion of our findings. At last section VI concludes the paper and lists future work.

II. Related work

Since monitoring darknet traffic provides a continual opportunity to view and detect remote network security events, many efforts have been made to build darknet traffic monitoring systems. Among them, three popular systems were proposed in Refs. [7]–[9].

Another work [1] presented the first comprehensive analysis of darknet traffic observed in 2004 at four unused IPv4 network blocks. It showed the great diversity of darknet traffic, both temporally and spatially, and examined the dominant events (i.e., root causes) on popular ports. Six years later, another work attempted to discover how darknet traffic had evolved from 2004 to 2010 [2]. It also introduced a simple categorization of darknet traffic and evaluated it with real traffic. Based on darknet traffic, an analysis of country-wide Internet outages in Egypt and Libya in 2011 was presented in Ref. [3]. In addition, a classification of one-way Internet traffic collected in live networks expanded our understanding of darknet traffic [10].

Many studies have been devoted to characterizing common anomalous events in the Internet. The authors discussed a taxonomy of DDoS attacks and different types of scanning events in Ref. [6]. Ref. [4] provided a longitudinal examination of scanning activities observed at Lawrence Berkeley National Laboratory (LBNL) over 12.5 years. A past work [5] proposed a backscatter analysis technique to infer DoS activity in the Internet. A network activity classification scheme with specification-based finite state machine models of TCP, UDP, and ICMP traffic was introduced in Ref. [17]. By integrating this classification scheme and packet IAT distributions of sources, *iatmon* demonstrated its effectiveness in classifying one-way traffic [16] [18]. More technical details of scanning and backscatter events were covered in Refs. [11] [12].

III. A taxonomy of darknet traffic

In this section we propose a simple taxonomy of darknet traffic. Given that darknet traffic consists of non-productive packets, it is natural to use concrete *traffic rules* to characterize it in terms of a number of *anomalous events*. Table I summarizes the anomalous events in our taxonomy. We base our classification on source flows generated from darknet traffic and explain each of these events in the following subsections.

A. 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 N ($\#portDst \geq N$). 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 $\cup$ FIN-ACK $\cup$ NULL; see Ref. [11] for more details) to be larger than a threshold $R\%$ ($ScanFlagPktRatio \geq R\%$), 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 ($Avg\ \#Pkt\ per\ portDst$).

B. 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 ($\#ipDst \geq N$). 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 in this case. For all three protocols we specify two subcategories (depending on $Avg\ \#Pkt\ per\ ipDst$) for heavy and light attacks.

C. One flow

The notion of one flow characterizes large, repeated traffic ($\#Pkt > N_3$) destined for one destination port ($\#portDst == 1$) in a $(ipSrc, ipDst)$ flow. This happens with both the TCP and UDP protocols. Network misconfiguration is a plausible explanation for this kind of traffic.

D. Backscatter

Backscatter traffic [5] consists of response packets to (D)DoS attacks carried out elsewhere in the Internet. Specifically, attackers somewhere in the Internet 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. For TCP, we use the TCP flags field (SYN-ACK $\cup$ ACK $\cup$ RST $\cup$ RST-ACK; see more details in Ref. [12]) to detect backscatter. For ICMP, we instead consider echo reply ($(Type == 0) \cap (Code == 0)$) and destination unreachable ($Type == 3$) packets (see more details in Ref. [13]). We then count the number of distinct sources ($\#ipSrc == 1$) that send at least one packet belonging to the categories mentioned above as the number of backscatters.

E. IP fragment

IP fragmentation exploits in darknet traffic represent DoS attacks or attempts to defeat packet filter policies. The Rose Attack [15] 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.

TABLE I
A TAXONOMY OF DARKNET TRAFFIC

Event	Category		Traffic rules
Port scan	TCP	Heavy	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg\ \#Pkt\ per\ portDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg\ \#Pkt\ per\ portDst \leq M)$
	UDP	Heavy	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N) \cap (Avg\ \#Pkt\ per\ portDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#ipDst == 1) \cap (\#portDst \geq N) \cap (Avg\ \#Pkt\ per\ portDst \leq M)$
Network scan	TCP	Heavy	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg\ \#Pkt\ per\ ipDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N) \cap (ScanFlagPktRatio \geq R\%) \cap (Avg\ \#Pkt\ per\ ipDst \leq M)$
	UDP	Heavy	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N) \cap (Avg\ \#Pkt\ per\ ipDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#portDst == 1) \cap (\#ipDst \geq N) \cap (Avg\ \#Pkt\ per\ ipDst \leq M)$
	ICMP	Heavy	$(\#ipSrc == 1) \cap (\#ipDst \geq N) \cap ((Type, Code) == (8, 0)) \cap (Avg\ \#Pkt\ per\ ipDst > M)$
		Light	$(\#ipSrc == 1) \cap (\#ipDst \geq N) \cap ((Type, Code) == (8, 0)) \cap (Avg\ \#Pkt\ per\ ipDst \leq M)$
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		$(\#ipSrc == 1) \cap (\#Pkt \geq 1) \cap (TCP\_Flags \in \{SA \cup A \cup R \cup RA\})$
	ICMP		$(\#ipSrc == 1) \cap (\#Pkt \geq 1) \cap (((Type, Code) == (0, 0)) \cup (Type == 3))$
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			$(\#ipSrc == 1) \cap (\#ipDst < N_1) \cap (\#Pkt \leq N_3) \cap ((Type, Code) == (8, 0))$
Other			$Other$
Remark: Our parameter setting $\{N = N_1 = N_2 = 5, R = 50, M = 3, N_3 = 15\}$ was empirically decided according to real traces.			

F. Small SYN

In our real darknet traces, we notice that many sources send a limited number of SYN packets to limited destinations on limited destination ports within a time period, a situation that does not correspond to any of the above events. 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.

G. Small UDP

For “small UDP”, the traffic rules are almost the same as “small SYN” except that we consider UDP packets instead.

H. Small Ping

The “small ping’’ event is also similar to “small SYN” except that 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.

I. Other

Source IP addresses that are not labeled as any of the network events mentioned above fall into this category. Note that the calculation of the unlabeled source rate in our traces is based on this category.

Note here that anomalous events – 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 according to certain specific packets. We emphasize that in our taxonomy one event may overlap others (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 traffic rules for backscatter and IP fragment events: just one backscatter or IP fragmented packet will trigger them. Moreover, our taxonomy allows one source to be characterized by multiple anomalous events; for example, one source send both TCP and ICMP scan packets.

IV. Evaluation

In this section we first introduce our dataset and then present analysis results for real darknet traces covering six years.

A. Dataset

Since Oct. 2006, we have been collecting traffic destined to one /18 allocated but unused IPv4 darknet address block in Japan, with three major data loss time periods: May 27, 2007 to Jun. 28, 2007; Nov. 26, 2010 to Feb. 4, 2011; and Jan. 9, 2012 to Sep. 25, 2012. Since we capture complete packet headers (layers -2, -3, and -4) and only a few payload bytes, our traffic analysis mainly relies on the header information.

To balance the scale of experimental data and corresponding processing time, we select the first seven days' data from every month to experiment with. This data covers 74 weeks (518 days) from October 2006 to November 2013 (not counting the excluded time periods). This long-range data is fairly representative of darknet's overall trends.

We also emphasize that our taxonomy is generally applicable for other darknet datasets, except for the parameter settings, which depend on specific traces.

B. Evaluation result overview

First, we compare our dataset with those used in Ref. [2] to check the similarity of traffic behavior between them. Figure 1 plots time series of (a) the protocol breakdown based on the number of packets and (b) the traffic breakdown defined by the simple categorization introduced in Ref. [2]. A significant change occurred around Oct. 2008 in both plots. Before the change the TCP and UDP traffic were quite similar to each other, but then the TCP traffic (especially TCP-SYN) kept increasing until it accounted for over 70% of the packet volume, thus dominating the complete traffic. As reported in Ref. [2], this change was due to the Conficker worm outbreak in Oct. 2008 [14]. In this regard, our dataset is consistent with the prior ones. Furthermore, our results demonstrate the Conficker worm's great influence on darknet till Nov. 2013.

To detect anomalous events hidden by short time bins, we adopted a longer time bin (24 hours, as explained in subsection IV-C) in our experiment. We first extract daily source IP flows from raw darknet traffic; then, for detection and labeling, we apply the traffic rules for each anomalous event defined in our taxonomy. We emphasize again that the taxonomy allows multiple events for one source IP address. Table II summarizes the results in terms of the percentages of source IP addresses labeled as each type of event. The symbol "-" in this table indicates a percentage of "<0.01%".

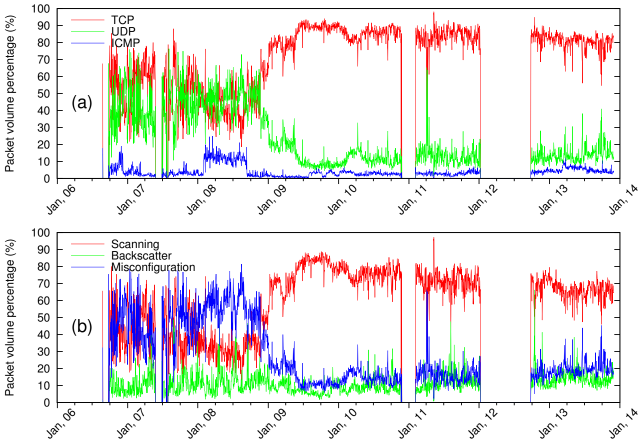

Fig. 1. Time series plot of our dataset from 2006 to 2013, showing traffic breakdown (a) by protocols and (b) by a simple categorization.

Table II clearly demonstrates that throughout the traces small SYN and small UDP events are most popular, with a percentage of at least 10%. This indicates that more sources are likely to send only a few packets destined for a small number of hosts and destination ports within 24 hours. We also notice, however, that the percentage of small SYN events increased a lot in 2008, and then decreased from 69.99% in 2009 to 41.46% in 2013, while the proportion of small UDP events experienced a significant decrease from 2008 to 2009, and maintained a rate of less than 30%. Looking at the proportion of light TCP network scans, we observe a significant increase from 2008 to 2009. Considering that the decrease in small SYN events and increase in light TCP network scan events almost complement each other at the same pace, we conclude 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 those for the other two small events. We confirm that the main reason for these results is the Conficker worm's outbreak in 2008.

As for scanning events, so far they are not popular choices among attackers, except for light TCP and ICMP network scans. 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 less than 2% of all sources in our traces.

Regarding backscatters, in the first three years ICMP backscatter events covered more than 5% of source IP addresses and increased slowly from 0.01% in 2009 to 1.32% in 2013. Compared to ICMP, TCP backscatter events in our traces were relatively stable, ranging from 0.85% to 2.55%. From the backscatter results we conclude that the observed

TABLE II
EVALUATION RESULT OVERVIEW (SOURCE IP ADDRESS %)

Event & Category			2006	2007	2008	2009	2010	2011	2012	2013
Port scan	TCP	Heavy	-	-	0.03	-	0.01	-	-	0.02
		Light	0.01	-	0.03	-	-	-	-	-
	UDP	Heavy	-	-	0.03	-	0.04	0.02	0.01	0.01
		Light	-	0.01	0.01	-	0.02	0.01	0.01	0.01
Network scan	TCP	Heavy	0.05	0.03	0.06	0.03	0.03	0.04	0.06	0.19
		Light	0.61	0.39	4.00	17.39	20.70	21.98	21.37	22.14
	UDP	Heavy	-	0.29	0.01	-	-	-	0.01	0.01
		Light	1.35	0.87	1.25	0.06	0.04	0.07	0.06	0.07
	ICMP	Heavy	-	-	0.01	-	-	-	-	-
		Light	0.14	0.06	0.14	0.01	0.07	0.56	0.63	0.94
One flow	TCP		2.04	1.42	2.58	0.22	0.30	0.20	0.66	0.72
	UDP		1.54	3.23	5.83	0.17	0.57	0.28	0.28	0.32
Backscatter	TCP		1.86	2.39	2.55	0.85	1.01	0.86	1.04	1.16
	ICMP		10.93	15.47	5.03	0.01	0.03	0.04	0.76	1.32
IP fragment			0.05	0.02	0.01	-	-	-	0.01	0.01
Small SYN			42.30	18.53	34.47	69.99	67.38	61.20	54.63	41.46
Small UDP			36.32	54.28	39.50	12.85	10.75	14.23	18.40	27.87
Small Ping			5.32	1.82	3.85	0.18	0.25	1.42	1.64	2.55
Other			0.60	3.09	2.49	0.15	0.34	0.26	1.34	2.34

spoofed-source (D)DoS attacks from our darknet keep relatively inactive in recent years.

As discussed in subsection III-C, one flow events mainly result from network misconfiguration. The highest proportions for both TCP and UDP one flow events appeared in 2008, and both exhibit an overall trend of decreasing since then. By examining the raw packets belonging to one flow events, we find that both single and multiple source ports are possible.

Regarding port usage in port scans, we find that over six years the most popular ports exploited by TCP scanners were 80 (HTTP) and 8080 (HTTP Alternate) while the most popular port for UDP was 28237 (application unknown). Turning to network scan events, we observe that 445 (Microsoft-DS) dominates TCP while 1434 (MS-SQL Monitor) dominates UDP. We also notice that port 53 (DNS) has been popular for UDP network scans in recent years. As expected, port 80 dominates among TCP backscatter source ports, suggesting (D)DoS attacks to web servers. Port 2186 (Guy-Tek Automated Update Application) is the most popular for TCP one flow events, while ports 137 (NetBIOS) and 161 (SNMP) dominate for UDP.

Throughout the six years of data, the proportion of IP fragmentation exploits is almost negligible. Last but not least, we point out that other events maintained a low proportion (the highest is 3.09% in 2007), validating that the proposed taxonomy detects and labels most sources in darknet traffic.

C. Dependency on time bin size

Determining the typical time period for anomalous events in a darknet is crucial to their accurate detection. On the one hand, if we choose a detection bin shorter than the typical time period, for example, some scanning events would likely be miscategorized as small SYN events or just be neglected. On the other hand, a longer bin could lead to redundant packets mixing into specific anomalous events, as well as requiring longer processing time. Thus, to understand this key parameter, we set the time bin to different sizes and experimented on six-day traces in 2007 and 2013.

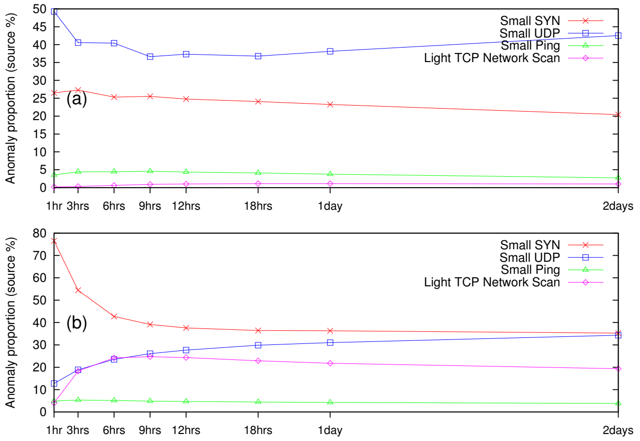

Fig. 2. Dependency on time bin size in terms of anomalous event proportions in (a) 2007 and (b) 2013.

Figures 2(a) and 2(b) plot the results for dependency on time bin size in terms of labeled source proportions before and after the Conficker worm's outbreak in 2008. In Fig. 2(a), small UDP events first decrease then increase a bit 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 much meanwhile. We also notice that with a time bin between six hours and one day both plots show just small fluctuations.

The resuts show that the lower percentage of small SYN events with a time bin between six hours and one day also means a higher detection rate for other types of events, like light TCP network scans. Thus, we conclude that the best time bin for detecting more significant anomalous events in our dataset is between six hours and one day. In fact, for this work we selected one day as the time bin for evaluation.

D. Dependency on darknet space size

To understand how taxonomy parameters influence detection accuracy for darknets with different space sizes, we divide short-range traces into several subnets with different sizes, and then apply our taxonomy on those subnets with the same parameters we used for /18 block before. The results are plotted in Fig. 3.

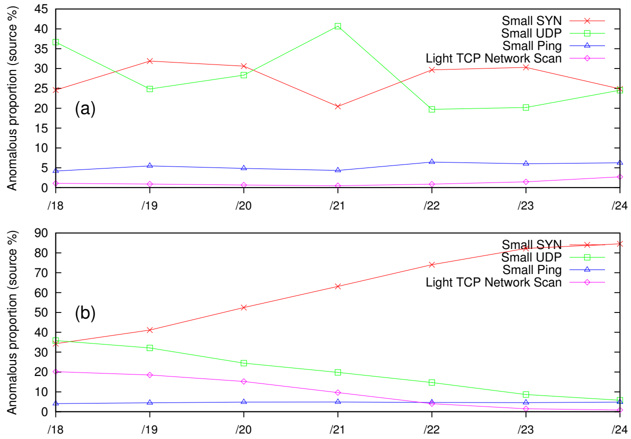

Fig. 3. Dependency on darknet space size in terms of anomalous event proportions in (a) 2007 and (b) 2013.

From Fig. 3(a) we see clearly that our empirical parameters mainly influence small SYN and small UDP events to show fluctuations for darkents with different space sizes before the Conficker worm outbreak. However, Fig. 3(b) shows that small UDP and light TCP network scan events keep decreasing whereas the vast majority of traffic can be characterized as small SYN events with the old parameters for /18 block used as darknet get smaller after the Conficker worm outbreak. This result highlights that our taxonomy is generally applicable to different darknet datasets, though the parameter tuning is required for accurate detection.

E. Diversity of anomalous events per source

Of particular interest to us is whether sources in darknet typically exhibit simple or complicated attack mechanisms. To obtain clues to this issue, we summarize the overall proportions of sources with different numbers of labels since 2006 in Tab. III. From the table we clearly see that sources with one or two labels are the vast majority, together accounting for over 96% of all sources. The extremely high percentage of sources 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, small UDP" and "small SYN, TCP backscatter" dominate among the sources with two labels. These dominant labels and label combinations are quite simple, and do not require deploying complicated mechanisms like those described in Ref. [11].

The results also indicate very few sources with more than two labels. One example of a four-label combination is "small SYN, small UDP, small Ping, TCP backscatter", which is allowed by our taxonomy.

V. Discussion

From Tabs. II and III, we conclude that most sources in our traces send a small number of TCP-SYN or UDP packets destined for small numbers of hosts and ports. We also notice the relative low proportion of one flow events over six years. These two findings suggest that most source IP addresses prefer neither large campaigns nor one-flow attacks. Moreover, the relative stability since 2008 of most of the anomalous events listed in Tab. II proves the persistent influence of the Conficker worm.

We also notice that in our traces over 97% of the sources are characterized by just one simple event and about 1.40% of the sources are labeled by a simple combination of two events, indicating that simple attack mechanisms are most often deployed by the sources in our traces.

Our result for dependency on time bin size shows that a time bin between six hours and one day is the most appropriate for accurate detection of anomalous events in our traces. Furthermore, the dependency on a longer bin suggests that the scanning activities in our traces are more likely to carry out slower probing, in attempting to more easily evade detection. Also, the result for dependency on darknet space size suggests us an importance of appropriate parameter settings corresponding to observed address blocks, especially for light TCP network scan events. The detailed analysis and further improvement will be one of our future work.

VI. Conclusion and Future work

In this paper, we have proposed a simple but effective taxonomy of darknet traffic and analyzed 74 weeks of real traces to evaluate our proposal. The results showed an extremely high detecting and labeling rate, of over 96% of all sources.

Through examining the evolution of anomalous events since 2006 (especially in recent years), we obtained some interesting findings. We highlighted that small SYN and small UDP events have dominated throughout the six years, while light TCP network scans have become more active in recent years. We also confirmed that the Conficker worm maintained its great influence on darknet traffic as of this work. In addition, we concluded that the observed spoofed-source (D)DoS attacks from our dataset and network misconfiguration events have kept relatively inactive in recent years.

We determined that the most appropriate time bin for the analysis of our dataset is between six hours and one day. Also,

TABLE III
OVERALL PROPORTIONS OF SOURCES WITH DIFFERENT NUMBERS OF LABELS SINCE 2006 (%)

#Labels	0	1	2	3	4	≥5
Percentage	0.70	97.83	1.40	0.05	0.01	<0.01

we investigated the dependency on darknet space size and highlighted the need for parameter tuning for different darknet datasets. Furthermore, we emphasized that most sources are characterized by one or two events, and they most often deploy simple attack mechanisms.

In future work, we plan to find a reasonable parameter tuning approach for darknets of different space sizes, and improve our current taxonomy to make it more fine-grained. In addition, we also plan to conduct a quantitative comparison with existing approaches (e.g., *iatmon*) to examine the advantages and disadvantages of our taxonomy.

ACKNOWLEDGMENTS

We thank Johan Mazel, Romain Fontugne, and the anonymous reviewers for their helpful comments and suggestions. This research has been partially supported by the Strategic International Collaborative R&D Promotion Project of the Ministry of Internal Affairs and Communication, Japan, and by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 608533 (NECOMA).

REFERENCES

[1] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of Internet Background Radiation. *IMC'04*, pp. 27–40, 2004.

[2] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston. Internet Background Radiation Revisited. *IMC'10*, pp. 62–74, 2010.

[3] A. Dainotti, C. Squarcella, E. Aben, K. Claffy, and M. Chiesa. Analysis of Country-wide Internet outages Caused by Censorship. *IMC'11*, pp. 1–18, 2011.

[4] M. Allman, V. Paxson, and J. Terrell. A Brief History of Scanning. *IMC'07*, pp. 77–82, 2007.

[5] D. Moore, G. Voelker and S. Savage. Inferring Internet Denial-of-Service Activity. *USENIX Security*, pages 12, 2001.

[6] J. Mirkovic and P. Reiher. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. *ACM Comp. Comm. Rev.*, 34(2):39–54, 2004.

[7] D. Moore, C. Shannon, G. Voelker, and S. Savage. Network Telescopes. *CAIDA Technical Report*, 2004.

[8] V. Yegneswaran, P. Barford, and D. Plonka. On the Design and Use of Internet Sinks for Network Abuse Monitoring. *RAID'04*, pp. 146–165, 2004.

[9] E. Cooke, M. Bailey, D. Watson, F. Jahanian, and J. Nazario. The Internet Motion Sensor–A Distributed Blackhole Monitoring System. *NDSS'05*, pp. 167–179, 2005.

[10] E. Glatz, X. Dimitropoulos. Classifying Internet one-way traffic. *IMC'12*, pp. 37–50, 2012.

[11] G. Lyon. Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning. http://nmap.org/book/man-port-scanning-techniques.html.

[12] Kevin R. Fall, W. Richard Stevens. TCP/IP Illustrated (Volume 1): The Protocols. *Addison-Wesley Longman Publishing Co., Inc.*, 2011.

[13] J. Postel. Internet Control Message Protocol. RFC 792, Sep. 1981.

[14] Protect yourself from the Conficker computer worm. *Microsoft Report*, Apr. 2009. http://blogs.msdn.com/b/mthree/archive/2009/03/31/conficker-033109.aspx.

[15] The Rose Attack Explained. http://www.digital.net/~gandalf/Rose_Frag_Attack_Explained.htm.

[16] N. Brownlee. One-way Traffic Monitoring with Iatmon. *PAM'12*, pp. 179–188, 2012.

[17] J. Treurniet. A Network Activity Classification Schema and Its Application to Scan Detection. *IEEE/ACM Transactions on Networking*, 19(5):1396–1404, 2011.

[18] The *iatmon* tool. http://www.caida.org/tools/measurement/iatmon/.