Are Darknets All The Same? On Darknet Visibility for Security Monitoring

Francesca Soro, **Idilio Drago**, **Martino Trevisan**, **Marco Mellia**
Politecnico di Torino
first.last@polito.it

João Ceron, **José J. Santanna**
University of Twente
j.m.ceron,j.j.santanna@utwente.nl

***Abstract*—Darknets are sets of IP addresses that are advertised but do not host any client or server. By passively recording the incoming packets, they assist network monitoring activities. Since packets they receive are unsolicited by definition, darknets help to spot misconfigurations as well as important security events, such as the appearance and spread of botnets, DDoS attacks using spoofed IP address, etc. A number of organizations worldwide deploys darknets, ranging from a few dozens of IP addresses to large /8 networks. We here investigate how similar is the visibility of different darknets. By relying on traffic from three darknets deployed in different contintents, we evaluate their exposure in terms of observed events given their allocated IP addresses. The latter is particularly relevant considering the shortage of IPv4 addresses on the Internet. Our results suggest that some well-known facts about darknet visibility seem invariant across deployments, such as the most commonly contacted ports. However, size and location matter. We find significant differences in the observed traffic from darknets deployed in different IP ranges as well as according to the size of the IP range allocated for the monitoring.**

***Index Terms*—Network telescopes, darknets, sinks, darkspaces.**

I. Introduction

Darknets[1] have been used for years as a source of information for cybersecurity [1]. A darknet is a set of IP addresses advertised by routing protocols, however without hosting any device. All traffic reaching the darknet remains unanswered and, by definition, is considered unsolicited. A monitoring probe listens to the darknet traffic, processing it in search for signals of new threats, misconfigurations and possibly sources/victims of attacks.

Years of experience running darknets have shown that three main types of traffic reach such networks [2]: (1) networks scans, both malicious (e.g., by botnets) and legitimate (e.g., by crawlers); (2) backscattering, i.e., deflected traffic received because someone contacted a host spoofing the source IP address belonging to the darknet; (3) traffic due to misconfigured devices or mistyped IP addresses.

Darknets have been used for a number of tasks [1], including (i) the investigation of malware spread [3] and Internet scans [4]–[6]; (ii) the estimation of DDoS frequency and volumes [7]–[9]; (iii) the analysis of Internet censorship [10]; (iv) the estimation of IPv4 address space utilization [11]. Different players deploy darknet infrastructure, from the large-scale projects run by the CAIDA/UCSD[2] and Merit [12] (each relying on a $/8$ IP range) to "sparse" darknets (also called greynets) run by companies[3] and academics [13]. The latter are characterized by a limited number of IP addresses that are distributed across different IP ranges. Several deployment strategies are thus available, and knowing the trade-offs is important for increasing the visibility of events while reducing the allocation of addresses for darknets, particularly relevant given the shortage of IPv4 addresses.

This paper investigates and revisits these questions for understanding how the visibility of darknets varies according to the IP range, size and location of the darknet. We capture traffic simultaneously for 1 month in three darknets, deployed in the Netherlands (a $/15$ network), in Brazil (a $/19$ network) and in Italy (3 $/24$ networks). We contrast the traffic reaching each network, highlighting the mostly seen protocols. We confirm that the size of the darknet matters, and quantify how the visibility is affected by the number of IP addresses allocated for the monitoring. We show that the Autonomous Systems (AS) and countries originating the traffic present significant differences according to the IP range where the darknet is deployed as well as the considered time period. All in all, results show that darknet traffic must be used with care to support security tasks, since the picture obtained in one darknet may not reflect other darknets or the attacks seen on production networks.

We are not the first to study trade-offs in darknet deployment strategies. Since seminal works on Internet Background Radiation [14], [15] (i.e., traffic seen in darknets), authors question the impact of darknet size, an analysis that has been revisited some years later [12] and repeated for IPv6 [16]. Recently, authors of [2] compare traffic observed in CAIDA's and Merit's darknets. Other authors have focused on distributed darknets [13], [17]–[19]. All acknowledge that darknets deployed at different IP blocks and networks observe different events. These works are however aged given the significant changes on the Internet in the last decade. We reappraisal this analysis with current traffic, shedding light on the coverage of interesting events according to parameters of the darknet deployment.

II. Methodology

We rely on data from three darknets composed by IPv4 addresses allocated in two continents. The first darknet resides

[1]They are also called network telescopes, Internet sinks and darkspaces.

[2]https://www.caida.org/projects/network_telescope/

[3]https://greynoise.io/

978-1-7281-1434-7/19/$31.00 © 2019 IEEE

TABLE I: Datasets and percentage packets per protocol.

	Size	Volume	TCP	UDP	ICMP	Other
BR	/19	2.5 GB/day	95.16%	4.39%	0.44%	0.01%
NL	/15	30 GB/day	93.69%	5.72%	0.59%	0.00%
IT	$3 \times /24$	420 MB/day	95.71%	3.89%	0.39%	0.01%

in Brazil, formed by a /19 network allocated by LACNIC (hereafter called *BR*). The second one is formed by a /15 network allocated by RIPE NCC in the Netherlands (hereafter called *NL*). The third one is formed by three /24 networks, with non-continuous addresses, hosted at the Politecnico di Torino in Italy (hereafter called *IT*). This latter is particularly interesting, since the addresses have been allocated for production traffic until recently. IPv4 prefixes are kept private following requests of the research institutions running the networks. Given the different size of the darknets, most of the analyses of Section III are restricted to smaller subnets of *BR* and *NL* darknets (hereafter referred to as *NLs* and *BRs*, respectively) to allow a fair comparison with *IT*.

In each location a network probe captures the traffic arriving to the allocated address, recording the full packet. The probe obfuscates IPv4 prefixes of the darknets (i.e., destination IP addresses) and sends the data to a Hadoop-based cluster for storage and processing. We perform analyses using data collected during 1 month, from the 1st of January to the 1st of February 2019.

Table I summarizes the dataset and provides a per-protocol breakdown of packets reaching the darknets. The majority of the traffic is represented by TCP ($> 93\%$ of the packets), with UDP ranging in $3.89 - 5.72\%$ and less than 1% for other protocols. No significant difference emerges between the darknets. The general picture is similar to the one reported in [2], even if we find higher percentages of TCP packets in these darknets than what is reported in previous work.

When analyzing the composition of traffic reaching darknets, we will focus on some main *traffic categories*:

• **Scan:** TCP packets with only the `SYN` flag set. To filter occasional scan from actual hosts running extensive scans, we mark as scans only those cases where the sender targets at least $k = 10$ different destination addresses or ports in a one hour time bin;

• **Backscattering:** TCP packets with `SYN+ACK`, `RST`, `ECN`, `RST+ACK` or only `ACK` flags set. Since the darknet does send any packets with `SYN` flag set, these packets are mostly from devices contacted with spoofed source IP addresses;

• **UDP:** UDP traffic, regardless payload or ports;

• **ICMP:** ICMP traffic;

• **Other:** All other cases or protocols. These include `SYN` scan messages sent by occasional scanners (that sent less than $k = 10$ messages in one hour).

III. Comparison of darknet traffic

In this section we provide a comparison across darknets, contrasting traffic composition, temporal patterns, sources and targeted ports.

TABLE II: Summary of the traffic per category.

Type	*NL*/15		*BR*/19		*IT* $3 \times /24$	
	Pkts	IP addr.	Pkts	IP addr.	Pkts	IP addr.
Scan	85.1%	12.5%	84.8%	4.6%	86.9%	3.2%
Back.	3.7%	0.8%	2.3%	0.6%	0.2%	0.2%
UDP	5.7%	10.8%	4.3%	2.3%	3.8%	1.8%
ICMP	0.5%	1.6%	0.5%	0.8%	0.3%	0.6%
Other	4.8%	74.1%	7.8%	91.4%	8.6%	93.9%

A. Traffic types

Table II provides a breakdown of the traffic, showing the percentage of packets and source IP addresses distribution across categories. Considering the share of packets on different categories, the highest share of traffic is constituted by Scan, with small differences among the darknets. The *IT* darknet shows a lower share of backscattering because of middleboxes sitting upstream the darknet, which drop incoming packets with inconsistent TCP flags/handshakes. UDP and ICMP shares are consistent among the three networks.

When comparing source addresses per category, interesting considerations hold. First, notice that Scan traffic is responsible of the majority of volume but it is generated by a small fraction (3.2 to 12.5%) of the senders' IP addresses. Recall that these sources are involved in non-occasional *SYN scans*, given the filters described in Sec. II. Second, there is a larger number of IP addresses sending UDP packets to *NL* than to *BR* and *IT*. Here manual inspection confirms a fact about darknet traffic [2]: There exist few sources that send a lot of UDP packets to targeted IP addresses. If the darknet does not include any of such targeted destinations, it would see less UDP events. Finally, note the large percentages of sources whose traffic lies in the "Other" category. Recall these are occasional scans, where the sender sends only few packets to the darknet. This traffic may be due to misconfigurations, low-rate attacks, or stale information in e.g., P2P protocols.

B. Temporal patterns

We next check whether the traffic reaching different darknets follows similar temporal patterns. Since the darknets have different sizes, both the /15 *NL* and the /19 *BR* darknets have been split into smaller subnets having the same dimension as the Italian one - i.e., $3 \times /24$. In the remainder of this Section, we restrict *all* analyses to 3 Dutch and 3 Brazilian samples, thus allowing a fair comparisons with the Italian one. Figure 1 reports time series of packets (top) and IP sources (bottom) per hour for the most relevant traffic category. The remaining categories are omitted given their lesser contribution to the total amount of traffic, and their noisy temporal pattern.

Scan traffic (Figure 1a) presents no clear temporal pattern and no periodicity. Equally, there is no apparent similarity between *BRs*, *NLs* and *IT*, and traffic peaks do not appear to be simultaneous. Figure 1b shows, instead, a more regular pattern in the number of distinct IP sources per hour. Notice that the number of addresses hitting *IT* is generally higher than the ones hitting *BRs* and *NLs*. We conjecture that this possibly happens because such addresses have been previously

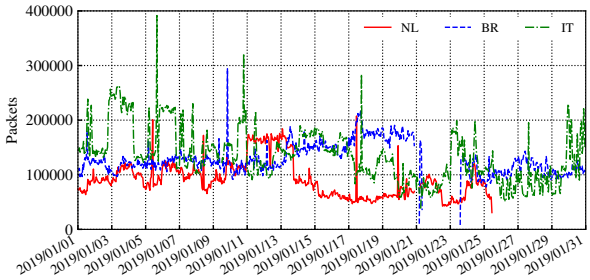

(a) Number of packets

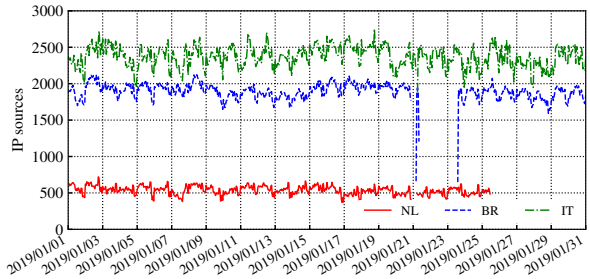

(b) Number of source addresses

Fig. 1: Time series (1h bins) for SCAN packets and sources.

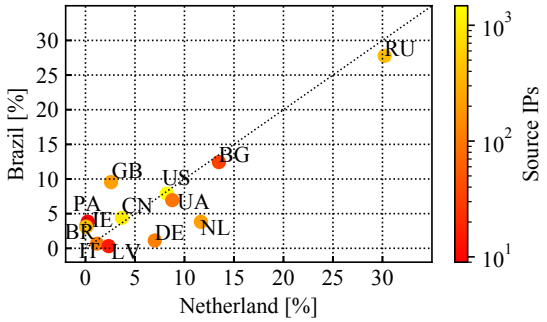

Fig. 2: Top source countries for Scan traffic.

allocated to a production network, and may hence be more known. Similarly, the lowest number of sources is observed on *NLs*, whose addresses have always been allocated as a darknet space. This suggests that the *NLs* darknet may be known to attackers that avoid targeting it. For the sake of brevity, we omit the figures for UDP and backscattering traffic, being significantly more noisy.

We further analyze the time series by calculating the Pearson correlation coefficient between pairs of darknets. Considering the of number of packets, for all traffic categories the pairwise correlation is zero or slightly negative – i.e., time series are uncorrelated. Different considerations hold for the number of sources: The average correlation among all pairs is 0.47 for network scans, 0.31 for backscattering, with *IT* and *BRs* reaching 0.79, 0.87 for UDP and 0.42 for ICMP (again with a peak of 0.79 between *IT* and *BRs*).

TABLE III: Top-10 AS per SCAN traffic – Jan 2019.

BRs			*NLs*			*IT*		
ASN	pkts	IPs	ASN	pkts	IPs	ASN	pkts	IPs
49453	14.8	8	49505	10.57	15	**43350**	22.18	12
57043	10.72	15	202325	9.94	11	204428	7.17	24
202325	6.5	12	204428	7.52	20	58271	7.05	22
58271	5.18	19	58271	6.9	19	**51852**	6.69	5
204428	3.74	18	**201912**	5.8	8	57043	6.28	16
14061	2.75	542	47350	5.07	5	14061	2.80	658
57271	2.51	11	57271	3.38	11	202325	2.75	11
47350	1.86	8	14061	3.03	103	**202425**	2.03	45
50297	1.66	4	**48817**	2.17	8	**206485**	1.74	3
51787	1.64	4	**41390**	1.96	1	49505	1.63	27

C. Origin of Scan traffic

We now focus on Scan packets to check whether sources of traffic are similar across darknets. The same analysis has been conducted for UDP and backscattering as well, but, for the sake of brevity, we consider only on Scan, being historically the most prominent source of attacks. Beside considering source IP addresses, we also map them to the corresponding AS and country with the Maxmind Geo Location database.[4]

Considering IP addresses, we record $27,105$ sources for *BRs*, $29,837$ for *IT* and $4,269$ for *NLs*. As previous works observed, the distribution of packets per IP address is heavy tailed: in our case 95% of the packets are generated by the (i) 22% most active addresses in *BRs*, (ii) by the 18% most active addresses in *NLs* and (iii) by the 23.3% addresses in *IT*. Notice the different order of magnitude in the number of sources observed in *NLs* with respect to the other two darknets. This result leads to the same conjecture reported in Figure 1b: being *NL* addresses allocated in the darknet space from a long time, they may be known as unused and hence less targeted.

Considering source ASes, we find $1,393$ (*BRs*), 142 (*NLs*) and $1,524$ (*IT*) sources, of which 134 are common to all three darknets (i.e., more than the 94% of the addresses targeting *NLs* are seen also in *BRs* and *IT*), while $1,015$ are common between *BRs* and *IT* (i.e., the 72.8% of the ASes seen in *BRs* are also present in *IT*). The top-10 most active ASes are shown in Table III, which reports in bold those that are not common across the three darknets. Table III also shows that the most common ASes generally produce a large percentage of traffic using only a small set of addresses (with AS $14,061$ being the exception). Moreover, rarely a single AS targets all darknet in the same way – e.g., AS $49,450$ is the most active against *NLs*, but it is ranked as last in *IT*, even if the latter is targeted using a wider number of sources.

We finally focus on source countries. In total 133 countries are seen on *IT*, 125 on *BRs*, and only 38 on *NLs* (the latter are all visible also in *IT* and *BRs*). Figure 2 compares the top-10 most seen countries per *BRs* and *NLs*. The scatter plot compares the share of packets from each country, while colors mark the total number of IP addresses observed for the country. The ranks mostly overlap, with 13 countries building the combined lists. Russia is the most popular source for both *BRs* and *NLs*, together with Bulgaria, Great Britain, USA, Ukraine and China.

[4]https://www.maxmind.com/en/home

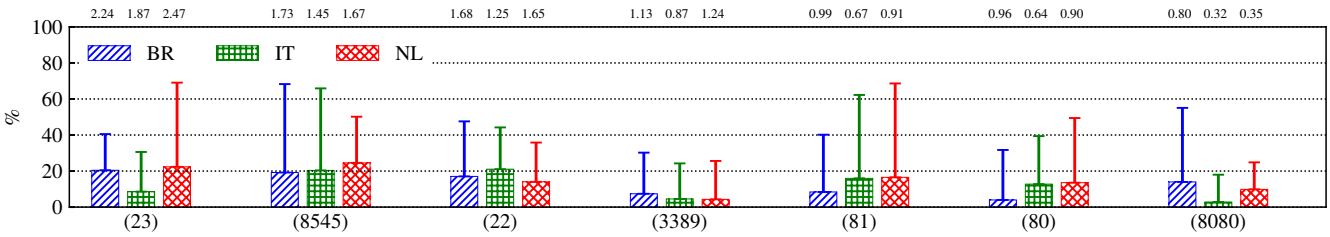

(a) Scan

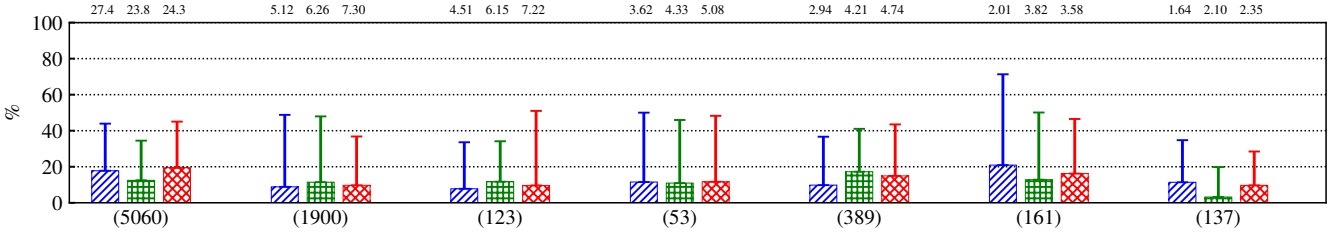

(b) UDP

Fig. 3: Packets due to top-1 (boxes) and top-10 (whiskers) source IP addresses for the 7 most contacted TCP and UDP ports. Numbers in the top x-axis represent the share of the port in the overall number of packets for the given network.

In general, such results confirm a conjecture raised in [2]: The Scan traffic reaching different darknets, while similar, is non-uniform. Finding the root-cause of such differences (e.g., routing configurations, IP ranges, location etc.) is left for future work.

D. Per-port breakdown

We now examine the destination ports of packets reaching the darknets. We restrict our analysis to Scan and UDP traffic since, for backscattering, destination port does not contain useful information.[5] Again, we consider only the three $/24$ *NLs* and *BRs* subnets for a fair comparison with the *IT* darknet.

In Figure 3 we quantify to what extent traffic originates from a small or large set of addresses for the 7 most contacted ports. Boxes represent the share of packets sent by the single most active IP source, while the whiskers represent the share for the top-10 addresses. Numbers in the top of the figure report the overall percentage of traffic to the port in the given network.

Considering Scan (Figure 3a), the most popular ports are associated with services known to be targets of attacks, e.g., telnet (port 23) and ssh (port 22). A significant number of such packets have been linked to attacks targeting IoT devices [20]. The top-10 sources are generally responsible for less than the 40% of the traffic, except in some particular cases (for instance port 81 and 8545, which is targeted by less distributed sources). The single most active source is in most cases generating about the 20% of the traffic alone. Focusing on the upper x-axis, we notice that the volume hitting the ports is similar across darknets.

[5]Backscattering traffic typically comes from victims contacted with spoofed source IP address.

Focusing on UDP (Figure 3b), we see that for almost all ports, about 40% of the packets are related to top-10 IP addresses. Pictures emerging in the three darknets are very similar, with protocols such as SIP (port 5060), NTP (port 123) and UPnP (port 1900) leading the ranks. This is not surprising, as such protocols are well-known targets large-scale attacks and abuses. Again, if we focus on the upper x-axis, we notice that a similar volume of traffic hits the considered ports for the three darknets.

We finally quantify to what extent the traffic sources are shared among darknets. Again, we map each IP address to the corresponding AS, and, separately per port, compute the Jaccard similarity index between the obtained sets of ASes. We consider *NLs* and *BRs* darknets, as they resides in different continents. Figure 4 shows the results. The blue bars represent the average similarity when comparing *BRs* to *NLs* subnets. The red bars serve as baseline, showing the average similarity when comparing the three *NLs* subnets against each other. Considering Scan, Figure 4a shows a Jaccard always around the 80% for *NLs* subnets, as expected. The similarity is generally below 50% when comparing *BRs* to *NLs*. Very interesting is the case of port 8548 (Json-RPC), which shows that scanning attempts on *BRs* and *NLs* descend from exactly the same set of ASes. The same considerations hold for the UDP scenario when comparing the *NLs* subnets among themselves. In Figure 4b, we see that the Jaccard index is generally above 80% when comparing the *NLs* subnets, while, when comparing *BRs* and *NLs*, results vary from port to port (e.g., below 50% for port 137 and port 5060, above 70% for port 53 and 123).

Take Away: *Comparing three different darknets, we observe that: (i) the largest amount of Scan traffic is produced by few source IP addresses, while most of them send only few*

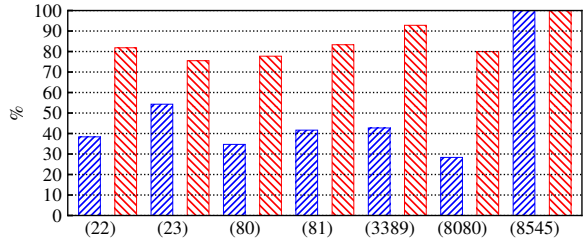

(a) Scan

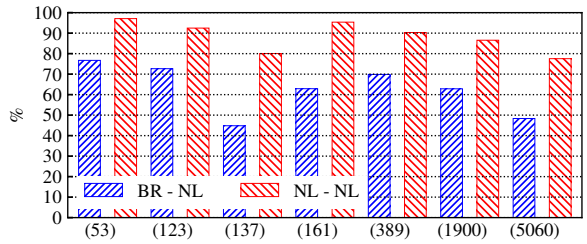

(b) UDP

Fig. 4: Average Jaccard similarity (calculated over sets of ASes) between *BR* and *NL* samples (blue) and among *NL* samples (red) for the top contacted ports.

packets; (ii) The behaviour of sources is particularly similar across darknets for UDP traffic. We register lower similarities for other traffic categories; (iii) Despite similarity, traffic from some sources generates very different volumes in different darknets; (iv) the top contacted ports are similar. For some of them, only few source ASes are behind the traffic. However, the ASes targeting most ports vary considerably.

IV. Effects of darknet size

In this section, we verify how observation period and the darknet size affect the list of observed sources. As in the previous section, we rely on the Jaccard similarity to compare the setups, focusing on source ASes.

A. Observation period

We first analyze the impact of the observation period. For a given darknet setup, we extract the set of ASes observed in a 1-week long period. Then, we again extract the sets of ASes after reducing the observation period to given shorter intervals. The Jaccard similarity is calculated by comparing the sets obtained with short intervals against the one obtained with the 1-week long interval. To increase reliability on results, these steps have been performed multiple times, by sampling 8 /19 subnets from the original /15 *NL* darknet, and 8 /22 subnets from the /19 *BR* darknet. In Figure 5, each data point reports the average Jaccard similarity for the multiple subnets. Figure 5a reports results for *NL*. Separate lines depict results for network scans, backscattering and UDP traffic. The visibility of sources in a darknet is reduced considerably when the observation period is reduced. However, the impact is different for the different traffic types. Considering network scans, the

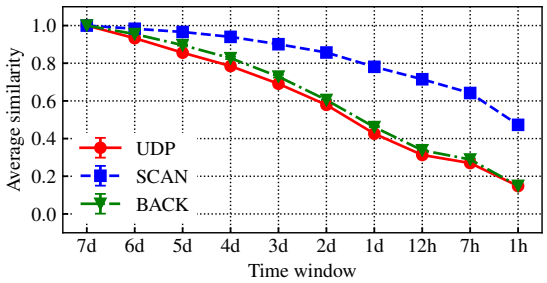

(a) *NL* (/19 samples)

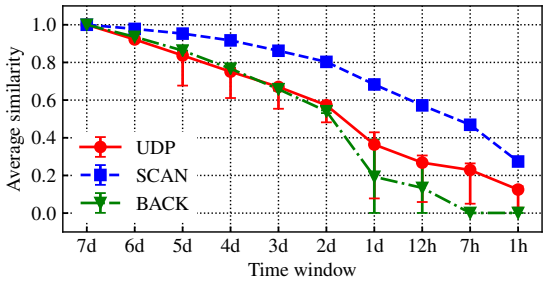

(b) *BR* (/22 samples)

Fig. 5: Average Jaccard similarity when fixing the darknet size and varying observation time.

Jaccard similarity is still at around 0.8 when the observation period is reduced to 1 day (i.e., 80% of the sources seen in one week are visible in one day despite the reduced interval), and around half of the ASes are found if one observes only few hours of traffic. Instead, for UDP and backscattering traffic, the reduction on visibility is much sharper. Already after shrinking the observation period to 2 days, almost half of the ASes are lost. This is also a consequence of the overall volume per traffic type (see Table II): whereas network scans are widespread, the other categories are rarer, which thus need more time to be observed.

Similar considerations hold for *BR* (Figure 5b). Given the smaller dimension of the darknet, the decrease is faster. For network scans, the picture is similar to the *NL* case, with just 30% of the sources found with a 1 hour observation interval. For backscattering, we notice a sudden drop when the observation period is shorter than 2 days. This can be explained by the intrinsic variability of backscattering traffic. Remind that backscattering sources are likely to be victims of attacks with spoofed addresses, typically carried out in a short amount of time. Given the lower volume, UDP decreases faster, too.

B. Darknet size

Finally, we analyze the impact of the darknet size, quantifying to what extent a small darknet observes events also found in larger ones. Remind that a small darknet would require low numbers of IPv4 addresses, thus freeing addresses for production traffic. Figure 6 reports the average similarity obtained when reducing the darknet size. Again, for improving robustness of results, we start by taking samples

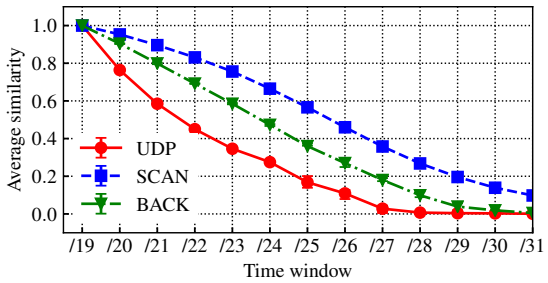

(a) Netherland

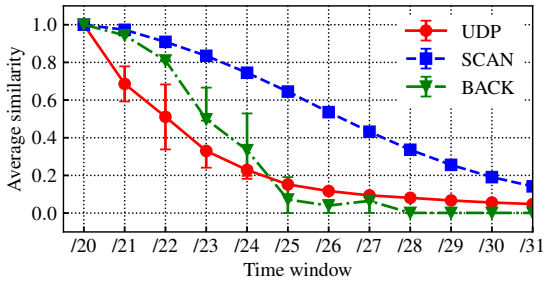

(b) Brazil

Fig. 6: Average Jaccard similarity when fixing observation time and varying the darknet size.

from the original darknets, i.e., 8 /19 subnets for *NL* and 2 /20 for *BR*. Then, for each experiment, we split each of these subnets into smaller subnets.

The figure reports the average similarity when comparing the small subnets to their respective original /19 (*NL*) or /20 (*BR*) subnets. The observation time window is fixed to one week in all cases.

Considering *NL* (Figure 6a), we notice a regular decrease as darknet size shrinks. For network scans, the plot suggests that a /25 network still observes around 60% of the source ASes. To obtain the same result for backscattering and UDP, larger /23 and /21 are needed. In other words, the majority of sources is still visible with a 64-fold reduction in darknet size for scanning, while for backscattering and UDP already a 16-fold and 4-fold size reduction hides 40% of the sources, respectively. Similar considerations hold for *BR* (Figure 6b). For scanning, with a 64-fold size reduction (/26 subnet) more than half source are still present. UDP decreases slightly faster than the *NL* case, with a 8-fold size reduction (/23 subnet) already hiding almost 70% of sources. For backscattering, again we notice a much faster decrease with respect to *NL*, similarly to what emerged from Figure 5b. The /25 subnets lose almost all visibility, confirming that the variability and unpredictability of backscattering traffic requires large darknets to be observed.

Take Away: *Network scans are constant and more prominent, thus easier to monitor. A one order of magnitude reduction in observation time and IP range size removes little of the darknet visibility. For backscattering and UDP, large observation times and IP ranges are needed for a good coverage.*

V. Conclusion

In this paper we compared three darknets deployed at different IP ranges and continents. We confirmed well-known facts about darknet visibility, such as the prevalence of traffic to the ports usually targeted by scans and attacks. Our results also provided new evidences that sources of traffic significantly varies according to the IP range, and the size of the darknet impacts its visibility. As future work, we plan to include greynets in the picture, thus updating our understanding about the differences between darknets and networks of monitoring hosts heavily distributed in the IP range.

References

[1] C. Fachkha and M. Debbabi, "Darknet as a Source of Cyber Intelligence: Survey, Taxonomy, and Characterization," *Commun. Surveys Tuts.*, vol. 18, no. 2, pp. 1197–1227, 2016.

[2] K. Benson, A. Dainotti, K. Claffy, A. C. Snoeren, and M. Kallitsis, "Leveraging Internet Background Radiation for Opportunistic Network Analysis," in *Proc. of the IMC*, 2015, pp. 423–436.

[3] S. Staniford, D. Moore, V. Paxson, and N. Weaver, "The Top Speed of Flash Worms," in *Proc. of the WORM*, 2004.

[4] Z. Durumeric, M. Bailey, and J. A. Halderman, "An Internet-Wide View of Internet-Wide Scanning," in *Proc. of the SEC*, 2014, pp. 65–78.

[5] A. Dainotti, A. King, K. Claffy, F. Papale, and A. Pescape, "Analysis of a "/0" Stealth Scan From a Botnet," *IEEE/ACM Trans. Netw.*, vol. 23, no. 2, pp. 341–354, 2015.

[6] Elias Raftopoulos, Eduard Glatz, Xenofontas Dimitropoulos, and Alberto Dainotti, "How Dangerous Is Internet Scanning? A Measurement Study of the Aftermath of an Internet-Wide Scan," in *Proc. of the TMA*, 2015, pp. 158–172.

[7] D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage, "Inferring Internet Denial-of-Service Activity," *ACM Trans. Comput. Syst.*, vol. 24, no. 2, pp. 115–139, 2006.

[8] C. Fachkha, E. Bou-Harb, and M. Debbabi, "Inferring Distributed Reflection Denial of Service Attacks from Darknet," *Comput. Commun.*, vol. 62, no. C, pp. 59–71, 2015.

[9] M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, and A. Dainotti, "Millions of Targets Under Attack: A Macroscopic Characterization of the DoS Ecosystem," in *Proc. of the IMC*, 2017, pp. 100–113.

[10] A. Dainotti *et al.*, "Analysis of Country-Wide Internet Outages Caused by Censorship," *IEEE/ACM Trans. Netw.*, vol. 22, no. 6, pp. 1964–1977, 2014.

[11] A. Dainotti, K. Benson, A. King, B. Huffaker, E. Glatz, X. Dimitropoulos, P. Richter, A. Finamore, and A. C. Snoeren, "Lost in Space: Improving Inference of IPv4 Address Space Utilization," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1862–1876, 2016.

[12] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston, "Internet Background Radiation Revisited," in *Proc. of the IMC*, 2010, pp. 62–74.

[13] W. Harrop and G. Armitage, "Defining and Evaluating Greynets (Sparse Darknets)," in *Proc. of the LCN*, 2005, pp. 344–350.

[14] D. Moore, C. Shannon, G. M. Voelker, and S. Savage, "Network Telescopes: Technical Report," Tech. Rep., 2004.

[15] J. Fan, J. Xu, M. H. Ammar, and S. B. Moon, "Prefix-Preserving IP Address Anonymization: Measurement-Based Security Evaluation and a New Cryptography-Based Scheme," *Comput. Netw.*, vol. 46, no. 2, pp. 253–272, 2004.

[16] J. Czyz, K. Lady, S. G. Miller, M. Bailey, M. Kallitsis, and M. Karir, "Understanding IPv6 Internet Background Radiation," in *Proc. of the IMC*, 2013, pp. 105–118.

[17] M. Bailey, E. Cooke, F. Jahanian, J. Nazario, and D. Watson, "The Internet Motion Sensor: A Distributed Blackhole Monitoring System," in *Proc. of the NDSS*, 2005, pp. 167–179.

[18] E. Cooke, M. Bailey, Z. M. Mao, D. Watson, F. Jahanian, and D. McPherson, "Toward Understanding Distributed Blackhole Placement," in *Proc. of the WORM*, 2004, pp. 54–64.

[19] M. Bailey, E. Cooke, F. Jahanian, A. Myrick, and S. Sinha, "Practical Darknet Measurement," in *Proc. of the CISS*, 2006, pp. 1496–1501.

[20] L. Metongnon and R. Sadre, "Beyond Telnet: Prevalence of IoT Protocols in Telescope and Honeypot Measurements," in *Proc. of the WTMC*, 2018, pp. 21–26.