Characteristics Comparison between Carpet Bombing-type and Single Target DRDoS Attacks Observed by Honeypot

Mao, Qingxin; Makita, Daisuke; van Eeten, Michel; Yoshioka, Katsunari; Matsumoto, Tsutomu

DOI
10.2197/ipsjjip.32.731

Publication date
2024

Document Version
Final published version

Published in
Journal of Information Processing

Journal of Information Processing Vol.32 731–747 (Sep. 2024)

[DOI: 10.2197/ipsjjip.32.731]

Regular Paper

Characteristics Comparison between Carpet Bombing-type and Single Target DRDoS Attacks Observed by Honeypot

Qingxin Mao[1,a)] Daisuke Makita[2,3] Michel van Eeten[4] Katsunari Yoshioka[1,3]
Tsutomu Matsumoto[1,3,5]

Received: December 7, 2023, Accepted: June 10, 2024

Abstract: Carpet bombing-type DDoS attacks targeting a wide-range network rather than a single IP address have threatened the Internet. Some researchers have investigated the characteristics of single-target DDoS attacks. Still, much less is known about the characteristics of carpet bombing, even the differences between them. In this paper, we profile characteristics of carpet bombing via data from amplification DDoS honeypots and the differences between single-target DRDoS attacks and carpet bombing. We analyze attacks highly concentrated on a specific network on victims, duration, number of packets, ports, and TTLs, and describe the differences between single-target DRDoS attacks and carpet bombing. Our analysis at the level of Autonomous Systems demonstrates that carpet bombing attacks target more hosting networks, including some critical targets, than single-target attacks. We found carpet bombing attacks targeting more "Corporate" networks. We also found that each IP address targeted by carpet bombing receives fewer packets than single-target DRDoS attacks. According to the result of the comparison of attack duration and TTL, carpet bombing lasted longer and referred to having diverse values of TTL in the packets. On the contrary, most single-target DRDoS attacks have a single value of TTL in the packets. This implies carpet bombing has a higher probability of originating from multiple sources. Finally, comparing ports shows that using various ports for Carpet Bombing is highly proportional to single-target DRDoS attacks.

Keywords: DRDoS Attack, Carpet Bombing

1. Introduction

Attacks that disrupt regular communication and impair Internet services, such as Distributed Denial of Service (DDoS) attacks, pose a significant threat to communication service providers, companies offering cloud and network services, educational research institutions, and government agencies [1]. Distributed Reflection Denial of Service (DRDoS) attacks involve malicious actors exploiting improperly exposed devices on the Internet as reflectors (**Fig. 1**). These attackers concentrate a large volume of packets on the target, disrupting the target's ability to provide services effectively [2].

A method for observing DRDoS attacks involves the proposal of DRDoS honeypots (AmpPot) [3], [4], [5], [8], [9]. This observation technique observes DRDoS attacks from the perspective of reflectors by deploying the decoy system, namely DRDoS honeypots, on the Internet. Additionally, other methods have been proposed for detecting and defending against DRDoS attacks [10], [11], [12], [13], [14], [15], [16]. For example, Nawrocki et al. survey common amplification honeypot platforms as well as the underlying methods to infer attack detection thresholds and to extract knowledge from the data [17]. Bekeneva et al. present experiments on DNS attack, NTP attacks and combined DRDoS-attack simulation. They simulated several protection mechanisms as well as a mechanism developed by them, and compared these protection mechanisms for different kinds of attacks [18].

Carpet bombing-type DRDoS attacks [19], [20] (hereinafter referred to as carpet bombing) are a recently emerged method of DRDoS attacks that target a broad range of IP addresses instead of a single IP address (**Fig. 2**). Due to the broad scope of carpet bombing attacks, it is inferred that the impact is more significant than DRDoS attacks targeting a single IP address. While detection methods for DRDoS attacks based on single IP addresses

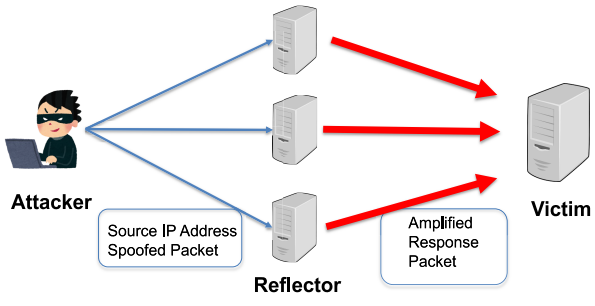

Fig. 1 DRDoS attacks.

[1] Yokohama National University, Yokohama, Kanagawa 240–8501, Japan
[2] National Institute of Information and Communications Technology, Koganei, Tokyo 184–8795, Japan
[3] Institute of Advanced Sciences, Yokohama National University, Yokohama, Kanagawa 240–8501, Japan
[4] Delft University of Technology, Faculty of Technology, Policy and Management, Jaffalaan 5, 2628 BX Delft, Netherlands
[5] Cyber Physical Security Research Center, AIST, Koto, Tokyo 135–0064, Japan
[a)] mao-qingxin-fp@ynu.jp

have been proposed, the distributed nature of the targets in carpet bombing makes it challenging for conventional detection methods to capture the overall picture of carpet bombing [5].

Therefore, in this study, we aim to address this issue: What are the characteristics of carpet bombing? How does it differ from DRDoS attacks targeting a single target? To answer these questions, we aggregate DRDoS attack events observed by honeypots at the network block level, focusing on incidents where the attack targets are extremely concentrated, to conduct a characteristic analysis of carpet bombing. Additionally, we analyze the characteristics of DRDoS attacks targeting a single target and compare them with the characteristics of carpet bombing to reveal the differences.

We conducted this study based on five years of data observed by DRDoS honeypots. Building on the results from our previous paper [reference to the journal paper], we aggregate attack events at the /16 network block level in this research. We define events with an included number of IP addresses ranging from 251 to 256 as carpet bombing, while those with only one IP address are defined as DRDoS attacks targeting a single target. Subsequently, we analyze and compare the characteristics of each, including AS type, AS rank, IP connection type, packet count, attack duration, exploitation status of ports, and types of TTL.

Our contributions are as follows:

(1) As a result of our investigation using ASrank, we found that some carpet bombing attacks targeted more important AS types.
(2) In our examination of IP Connection types, we observed that single-target DRDoS attacks primarily targeted IP addresses belonging to the “Cable/DSL” category, whereas carpet bombing attacks targeted IP addresses in a more diverse range, including the “Corporate” category.
(3) The comparison of packet numbers revealed that, before aggregation, carpet bombing events had fewer packets per event compared to single-target DRDoS attacks. However, after aggregation, it was evident that the packet count for carpet bombing events was significantly higher.
(4) Our comparison of TTL types showed that the sources of carpet bombing attacks were predominantly multi-source.

The structure of this paper is as follows: In Section 2, we introduce the honeypots used for data collection and the aggregation algorithm. Section 3 presents the results of the comparisons, followed by a discussion of the findings in Section 4. Section 5 provides an overview of related research, and finally, Section 6 summarizes the results.

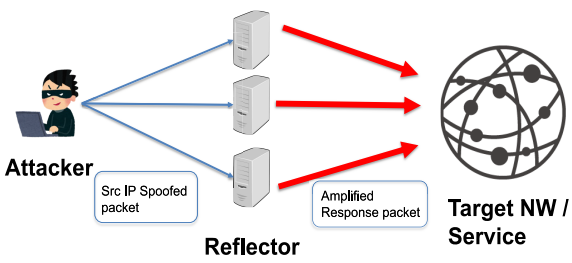

Fig. 2 Carpet bombing.

2. Data Source and Aggregation Algorithm

In this section, we introduce our DRDoS attack observation system called DRDoS honeypots, and the aggregation algorithm used to consolidate attack events.

2.1 DRDoS Honeypot

A DRDoS honeypot (**Fig. 3**) refers to a decoy system deployed on the Internet that conducts observations of DRDoS attacks in IPv4 network space from the perspective of reflectors [4], [5]. The results of observation in IPv6 network space are not the subject of this paper and are a subject for future consideration. The DRDoS honeypots currently operational in our setup consist of two types. The first is the Agnostic Honeypot, which observes across all ports, and the second is the Proxied Honeypot, which observes only on ports commonly exploited. The technology of AmpPot remains unchanged from the RAID2015 paper [5], with only an increase in the number of units. While we cannot disclose the IP addresses of the honeypots, they operate on ISP lines within Japan with fixed IP addresses.

Agnostic Honeypot responds to requests on all ports with a large random response. While this allows for observation across all ports, the non-compliance of Agnostic Honeypot responses with protocols means that attackers can potentially identify it as not a typical reflector by scrutinizing the responses. However, at present, various types of DRDoS attacks have been observed even with this honeypot [4], [5], [21].

Proxied Honeypot focuses on services frequently exploited in DRDoS attacks, running actual services to observe attacks by adversaries exploiting these services [5]. Specifically, it observes attacks abusing the following ports: 17/UDP (Quote of the Day), 19/UDP (Character Generator Protocol), 53/UDP (DNS, Domain Name System), 123/UDP (NTP, Network Time Protocol), 161/UDP (SNMP, Simple Network Management Protocol), 1900/UDP (SSDP, Simple Service Discovery Protocol), 11211/UDP (Memcached).

As the DRDoS honeypot is exposed on the Internet to observe DRDoS attacks, it captures a multitude of other communications beyond just DRDoS attacks, including network scans using tools like Nmap [22] or Zmap [23], as well as attacks targeting device vulnerabilities. To minimize the impact of unrelated communications and extract DRDoS attacks as much as possible from the massive traffic volume, a method has been proposed to define attacks as ‘events’ [5]. This involves grouping a series of pack-

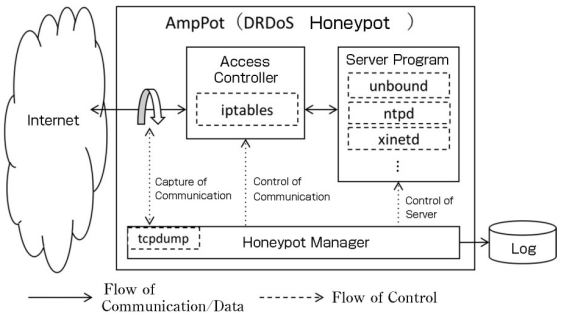

Fig. 3 DRDoS Honeypot.

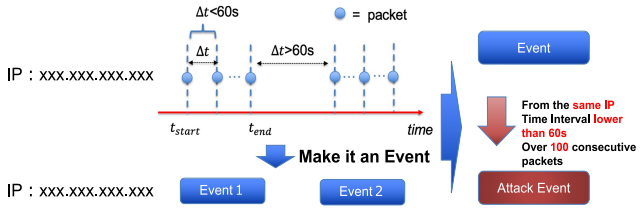

Fig. 4 Definition of attack event.

ets observed by the DRDoS honeypot over time intervals and determining whether they constitute an 'attack event' based on the number of packets in each group. Specifically, the observed packets are organized by destination IP address, and a group of packets with time intervals between them of 60 seconds or less is defined as an 'event.' An 'attack event' is then identified as one event with a packet count of 100 or more (**Fig. 4**).

2.2 Aggregation Algorithm

However, the event definition criteria based on individual IP addresses are challenging to apply to the recently prevalent carpet bombing. First, due to the distributed nature of carpet bombing attacks, the number of attack packets sent to individual destination IP addresses is also dispersed. As a result, dispersed packets may not exceed the recognition threshold for an attack event, leading to a situation where the attack cannot be adequately identified. Additionally, when looking at the entire targeted network, although carpet bombing may be consistently occurring, viewing it on a per-IP address basis may result in larger time intervals between events, causing dispersed packets not to be correctly recognized as attack events. Furthermore, with a multitude of targets in carpet bombing attacks, defining events based on individual IP addresses makes it difficult to capture the overall picture of carpet bombing.

Given the dispersed nature of carpet bombing attacks, the packet count after aggregation is anticipated to be higher than that of single-target DRDoS attacks. Still, there is a significant possibility that the packet count before aggregation may be lower than that of single-target DRDoS attacks. To comprehensively analyze carpet bombing, this study lowers the criteria for identifying attack events from the traditional threshold of over 100 packets to a lower level, thus including potentially carpet bombing-related attacks in the analysis scope. However, events with fewer than 100 packets may not only encompass attack events but also other activities such as network scans. To differentiate between network scan and attack events, we analyze network scan data observed by Yokohama National University's darknet (a network of unused IP addresses that do not respond to incoming packets, and thus the data observed by the darknet is presumed to be network scan activity primarily aimed at discovering devices). Specifically, we analyze network scan data observed by the darknet and elucidate the packet count characteristics of network scans. In particular, packets received by the darknet within intervals of less than 60 seconds are eventized per IP address, and the distribution of packet counts per event is examined. The results of analyzing data from March 2018 are presented in **Figs. 5** and **6**. Figures 5 and 6 are CDFs (Cumulative Distribution Function), with the horizontal axis representing packet count and the vertical axis

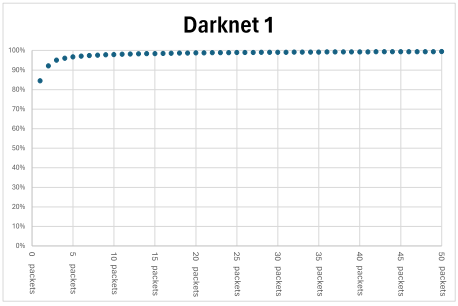

Fig. 5 The distribution of the packet count (Horizontal axis: Packets received by a single darknet IP address within intervals of less than 60 seconds. Vertical axis: Proportion).

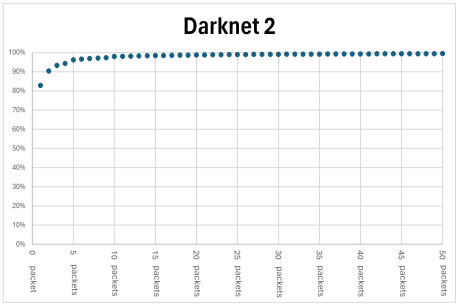

Fig. 6 The distribution of the packet count (Horizontal axis: Packets received by a single darknet IP address within intervals of less than 60 seconds. Vertical axis: Proportion).

representing the packets received by the darknet within intervals of less than 60 seconds. The blue plots represent the proportion corresponding to the packet count. From Figs. 5 and 6, it is evident that over 96% of network scan activities are concentrated within 5 packets or fewer. While there is no ground truth on the boundary between network scans and DRDoS attacks, we set our threshold as 6 to filter out majority of network scans while capturing possibly segmented pieces of carpet bombing attacks. Note that these attack events are eventually aggregated and only aggregation events with targeted IP addresses of 251–256 would be further analyzed as carpet bombing attacks, as those are likely targeting /24 networks as discussed in Section 2.3. Those random scans that happened to exceed the threshold of 6 would be dropped through the aggregation process.

To analyze carpet bombing that concentrates on specific destination networks within a particular period, we propose the following method: aggregating attack events observed during the same time period, where the destination addresses fall within a specific range, and classifying them as a single aggregation event (**Fig. 7**).

(1) Divide IPv4 addresses into fixed network ranges (e.g., /24) and group attack events within the same network.

(2) Among the attack event groups grouped in Step 1, repeatedly process pairs of attack event groups occurring within a fixed time interval until no further aggregation is possible. Sum the packet counts of the aggregated attack events to determine the packet count of the aggregated event. Additionally, set the start time of the aggregation event as the earliest start time among the aggregated attack events and the end time as the latest end time among them.

2.3 The parameters for Aggregation

In this aggregation algorithm, two parameters, namely network range and time interval, are specified for aggregating attack events. Based on the results of the paper published earlier [24], it was observed that many carpet bombing attacks targeted /24 networks, but occasionally there were carpet bombing attacks targeting networks larger than /24. In such cases, aggregating at the /24 level could result in counting a single attack event multiple times, potentially impacting the accuracy of the analysis results. To capture each attack event as comprehensively as possible and maintain the accuracy of the analysis results, we performed aggregation at the /16 network level, which is larger than /24, to eliminate the impact of duplicate counting. However, aggregating at the /16 network level poses the risk of inadvertently aggregating unrelated attacks, leading to a larger number of targeted IP addresses in aggregation events (noise cases). To mitigate the influence of unrelated attacks, it is necessary to understand the distribution of the number of targeted IP addresses in aggregation events when aggregated at the /16 network level and identify the attack ranges where carpet bombing is concentrated.

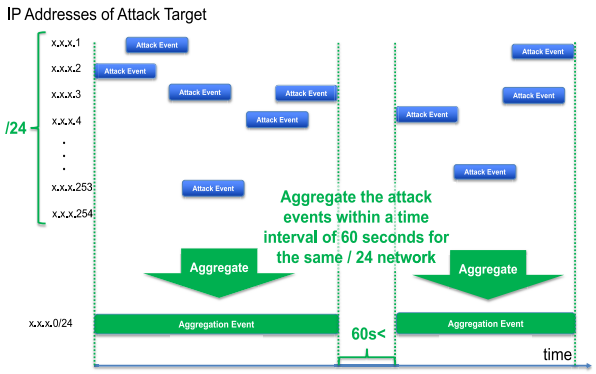

Fig. 7 Specific example of aggregation.

The **Fig. 8** illustrates the distribution of the number of targeted IP addresses in aggregation events aggregated with a /16 network range and a time interval of 1,200 seconds. From Fig. 8, it is evident that aggregation events are primarily concentrated where the number of targeted IP addresses is 255 and 256. This indicates that carpet bombing attacks mainly target /24 network ranges. In this study, we focus on analyzing aggregated events with attack targets ranging from 251 to 256.

However, even with aggregation events having a target count

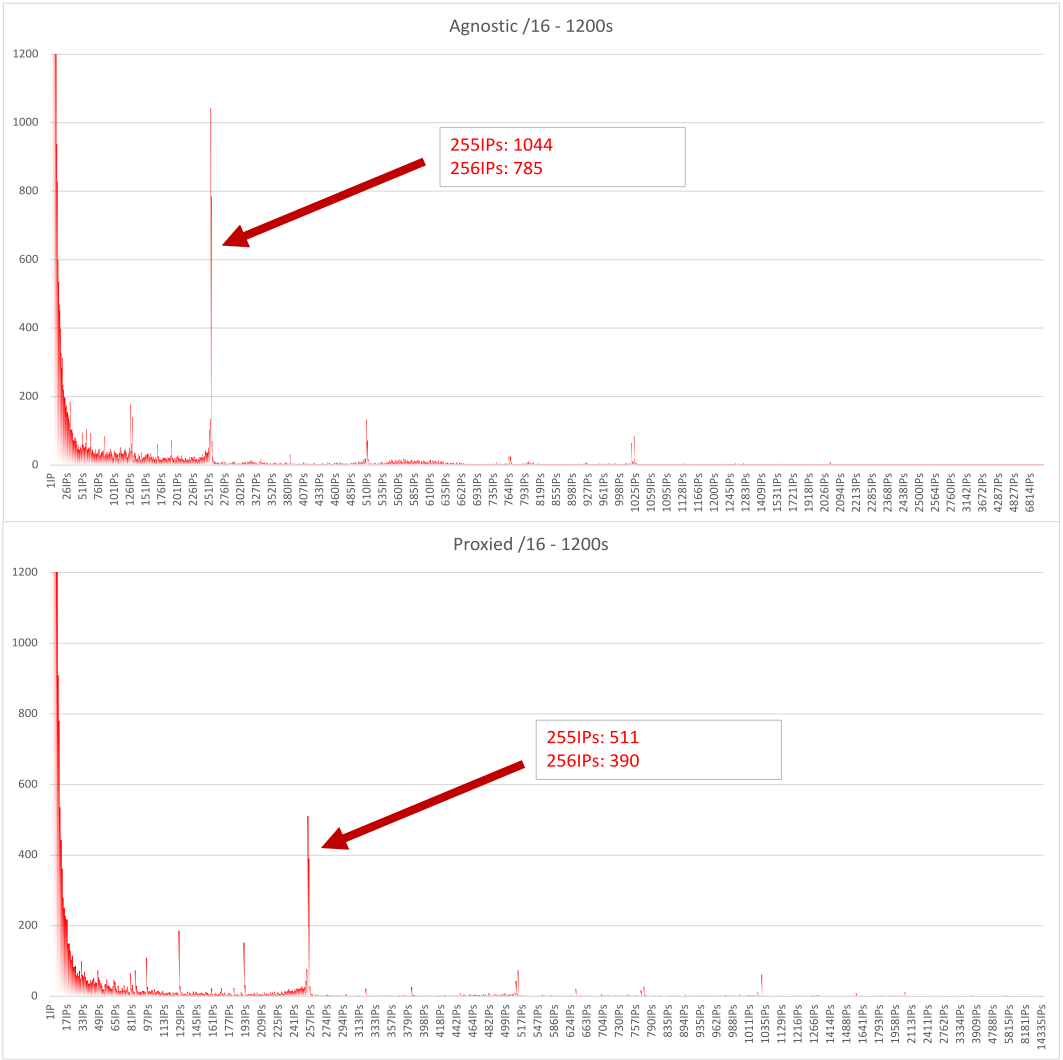

Fig. 8 The distribution of the number of targeted IP address in aggregation events aggregated with a /16 network and a time interval of 1,200 seconds.

Table 2 The results of observation and aggregation.

Period: 2018.3.21 to 2023.3.31		
Honeypot	Agnostic	Proxied
Total Events (Packets ≥ 6)	736,492,155	333,825,377
Agnostic	Before aggregation	After aggregation
Number of single-target DRDoS Attacks	120,823,288	19,549,649
Number of Carpet Bombing	70,764,805	**7,920**
Total	191,588,093	19,557,569
Proxied	Before aggregation	After aggregation
Number of single-target DRDoS Attacks	99,412,930	12,657,577
Number of Carpet Bombing	11,946,936	**6,505**
Total	111,359,866	12,664,082
Agnostic	Number of ASes	Number of IPs
Single-target DRDoS Attacks	37,259	8,583,985
Carpet Bombing	1,368	829,320
Proxied	Number of ASes	Number of IPs
Single-target DRDoS Attacks	34,329	5,521,331
Carpet Bombing	1,467	876,962

Table 1 The distribution of the calculated results.

Agnostic		**Proxied**	
IP range	Count	IP range	Count
256 IPs	4,117	256 IPs	4,139
255 IPs	3,201	255 IPs	1,876
254 IPs	227	254 IPs	226
253 IPs	287	253 IPs	255
252 IPs	17	252 IPs	3
251 IPs	21	251 IPs	6
Total	7,920	Total	6,505

of 251–256, there is a possibility of mistakenly aggregating attacks that do not represent carpet bombing attacks, as they might target multiple networks within the /16 range. To verify this, IP addresses are converted to decimal numbers, and the IP addresses within the same aggregated event are arranged in ascending order. Then, the last IP address in the list subtracts the value of the first IP address, and adding 1 reveals the actual width of the attack range. The distribution of the calculated results is shown in **Table 1**. From Table 1, it is evident that the actual attack range of aggregated events with a target count of 251–256 falls entirely within /24 network ranges. Upon verifying the actual data, there were instances of attacks spanning two /24 networks, but since the IP addresses were contiguous, it can be concluded that there was no misaggregation as described above.

Therefore, in this study, attack events are aggregated at the /16 network level. Additionally, considering the relatively long interval of 1,200 seconds, and based on the results of the previous journal paper, the time interval for aggregating attack events is adjusted to 300 seconds. Finally, considering the above analysis results, this study defines events with an IP address count of 251–256 as carpet bombing attacks, and those with only one IP address as single-target DRDoS attacks. The aggregation algorithm and processing of analysis data were implemented using Python [25] and Bash scripts.

2.4 Results of Observation and Aggregation

The overall situation of observation and aggregation is depicted in **Table 2**. The data used in this study comprises observations from 8 Agnostic Honeypots and 12 Proxied Honeypots, covering the period from March 21, 2018, to March 31, 2023. Examining the event counts before and after aggregation, it is evident that the number of carpet bombing attack events before aggregation is high, while the number of aggregated events after aggregation significantly decreases. From this result, it is inferred that carpet bombings, due to their dispersed targets, are often misidentified under traditional event definitions, leading to the erroneous recognition of a single carpet bombing as multiple attack events.

3. Characteristic comparison

In this section, we describe the procedure for analyzing the data and present the results of the data analysis.

3.1 Methodology

To conduct the analysis, we first extract the following attributes of information from the attack event data: ASN (Autonomous System Number, target of the attack), spoofed source IP address (target of the attack), packet count, attack duration, destination port number, and TTL (Time To Live).

Comparison of AS types and AS rank: To begin with, in order to elucidate which AS are being targeted, we investigate and compare the types of ASs under attack. The honeypots observe a substantial volume of attacks daily. The observed attacks are stored in the database as ‘attack events.’ Before being entered into the database, the source IP addresses are examined using the GeoIP2 ISP Database [26] to retrieve information about the corresponding AS number, which is then included in the attack event. Subsequently, based on this AS number, we query the AS type using ASDB. ASDB is a system that leverages data from established business intelligence databases and machine learning to categorize ASes accurately on a large scale. It achieves a 96% coverage of ASes with 93% accuracy on 17 industry categories and 75% on 95 sub-categories [27].

Secondly, to ascertain which attacks target more significant entities, we utilize CAIDA’s ASrank to investigate the importance of Autonomous Systems (AS). ASrank is CAIDA’s ranking system for Autonomous Systems (AS) and organizations (Orgs) [28]. This ranking is derived from topological data collected by CAIDA’s Archipelago Measurement Infrastructure and Border Gateway Protocol (BGP) routing data collected by the

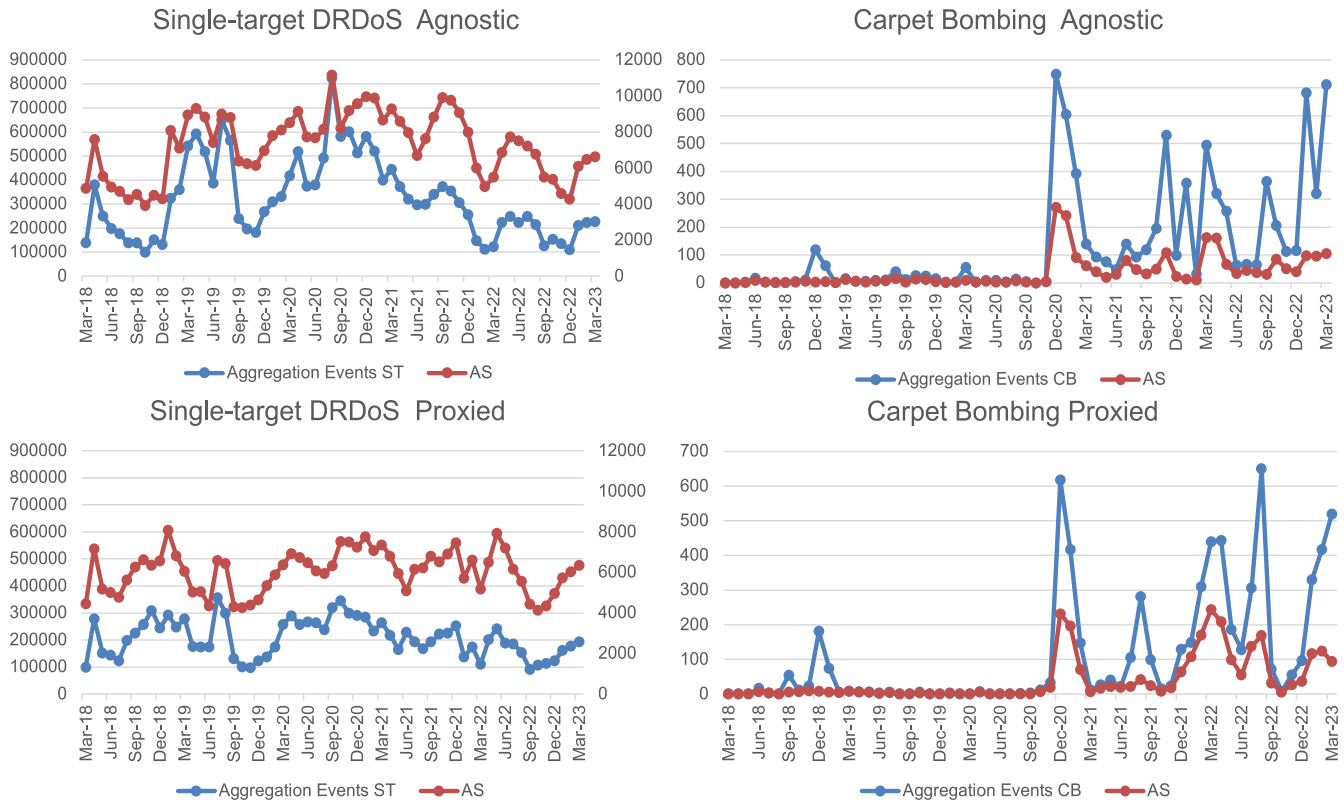

Fig. 9 Monthly aggregation event counts and the number of attacked ASes.

Route Views Project and RIPE NCC. ASes and Orgs are ranked based on their customer cone size, which represents the number of their direct and indirect customers.

Thirdly, to investigate which IP addresses are being targeted, we examine the Connection Type of the attacked IP addresses using the GeoIP2 Connection Type Database [29]. The GeoIP2 Connection Type Database can identify the connection type of visitors based on their IP address, distinguishing between cellular, cable/DSL, and corporate connection speeds.

Fourthly, to examine the volume of attacks and their duration, we compare the distribution of packet counts and attack durations in the attack events before and after aggregation with the distribution of packet counts and attack durations in the aggregated events. This allows us to observe the differences in the results.

Fifthly, to determine whether the sources of the attacks are single or multiple, we extract TTL information from the packets and examine the distribution of TTL types in the aggregated events. Time to Live (TTL) refers to the amount of time or 'hops' that a packet is set to exist inside a network before being discarded by a router. In other words, if all the TTL values of the packets included in a single aggregated event are the same, it can be considered an attack launched by the attacker from a single location. Conversely, if the TTL values are varied, it suggests that the attacker manipulated multiple machines located in different places simultaneously to carry out the attack. Alternatively, it can also be said that the attacker randomized the TTL values of the packets before carrying out the attack to conceal their location.

Finally, to elucidate the state of the services used in the attacks, we examine information about the ports used by carpet bombing and single-target attacks. Furthermore, we look into whether multiple services are used simultaneously in the attacks by examining the distribution of the types of destination port numbers in the aggregated events.

3.2 AS Type, AS Rank

Firstly, we compare the types and importance of targets attacked by carpet bombing and single-target DRDoS attacks. Before delving into the specifics, **Fig. 9** illustrates the monthly aggregation event counts and the number of attacked ASes for these two attack types. Since two types of honeypots are used for attack observation, the data is presented separately for each type. The horizontal axis represents time, while the vertical axis represents the count of aggregation events. Due to the significantly higher number of events for single-target DRDoS attacks, and to avoid overshadowing the AS count on the same axis, the AS count is displayed on a separate axis to the right. Red plots stand for the monthly aggregation event counts and blue plots stand for the number of AS. Figure 9 shows that as the number of single-target DRDoS attacks increases, the corresponding AS count also tends to increase. This suggests that single-target DRDoS attacks do not concentrate on specific ASes. Conversely, during periods of increased carpet bombing attacks, there are instances where the number of targeted ASes does not increase proportionally. This indicates that carpet bombing tends to focus on specific ASes during attacks.

After understanding the temporal trends in the attacks, we compare the types of attacked ASes. **Figures 10** and **11** illustrates the distribution of attacked AS types over time. The horizontal axis represents time, and the vertical axis shows the percentage of each AS type. Regarding the Industry Category, both types

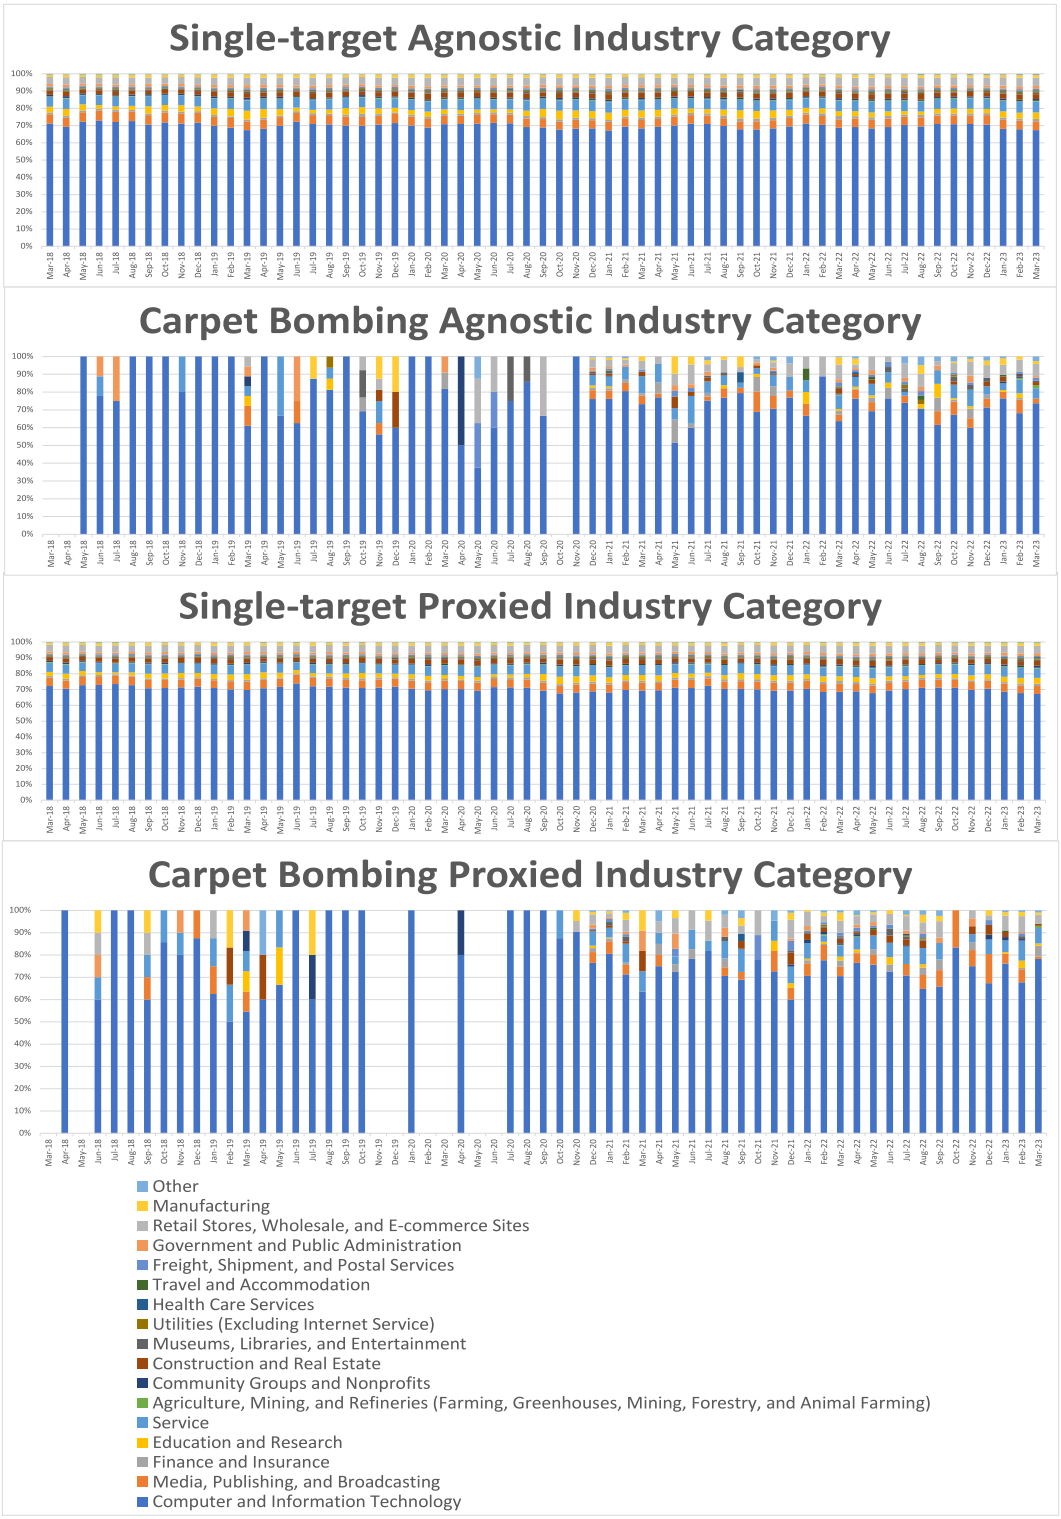

Fig. 10 The distribution of attacked AS types over time - industry categories.

of attacks concentrate heavily on the 'Computer and Information Technology' category, with little significant difference between them. However, when comparing Subcategories, it is observed that single-target DRDoS attacks predominantly target the 'ISP' category, whereas carpet bombing attacks target a greater number of ASes in the 'Hosting' category.

After comparing the types of attacked ASes, we proceed to compare the importance of the targeted ASes. The results from the ASrank analysis are depicted in **Fig. 12**. The horizontal axis represents time, while the vertical axis shows the AS rank, where lower values indicate higher importance. Upon reviewing Fig. 12, it becomes evident that only certain carpet bombing attacks targeted more important ASes. Additionally, it is observed that some attacks targeted ASes of lesser importance.

3.3 IP Connection Type

In this section, we compare the IP Connection Types of the targeted addresses. **Figure 13** illustrates the monthly aggregation

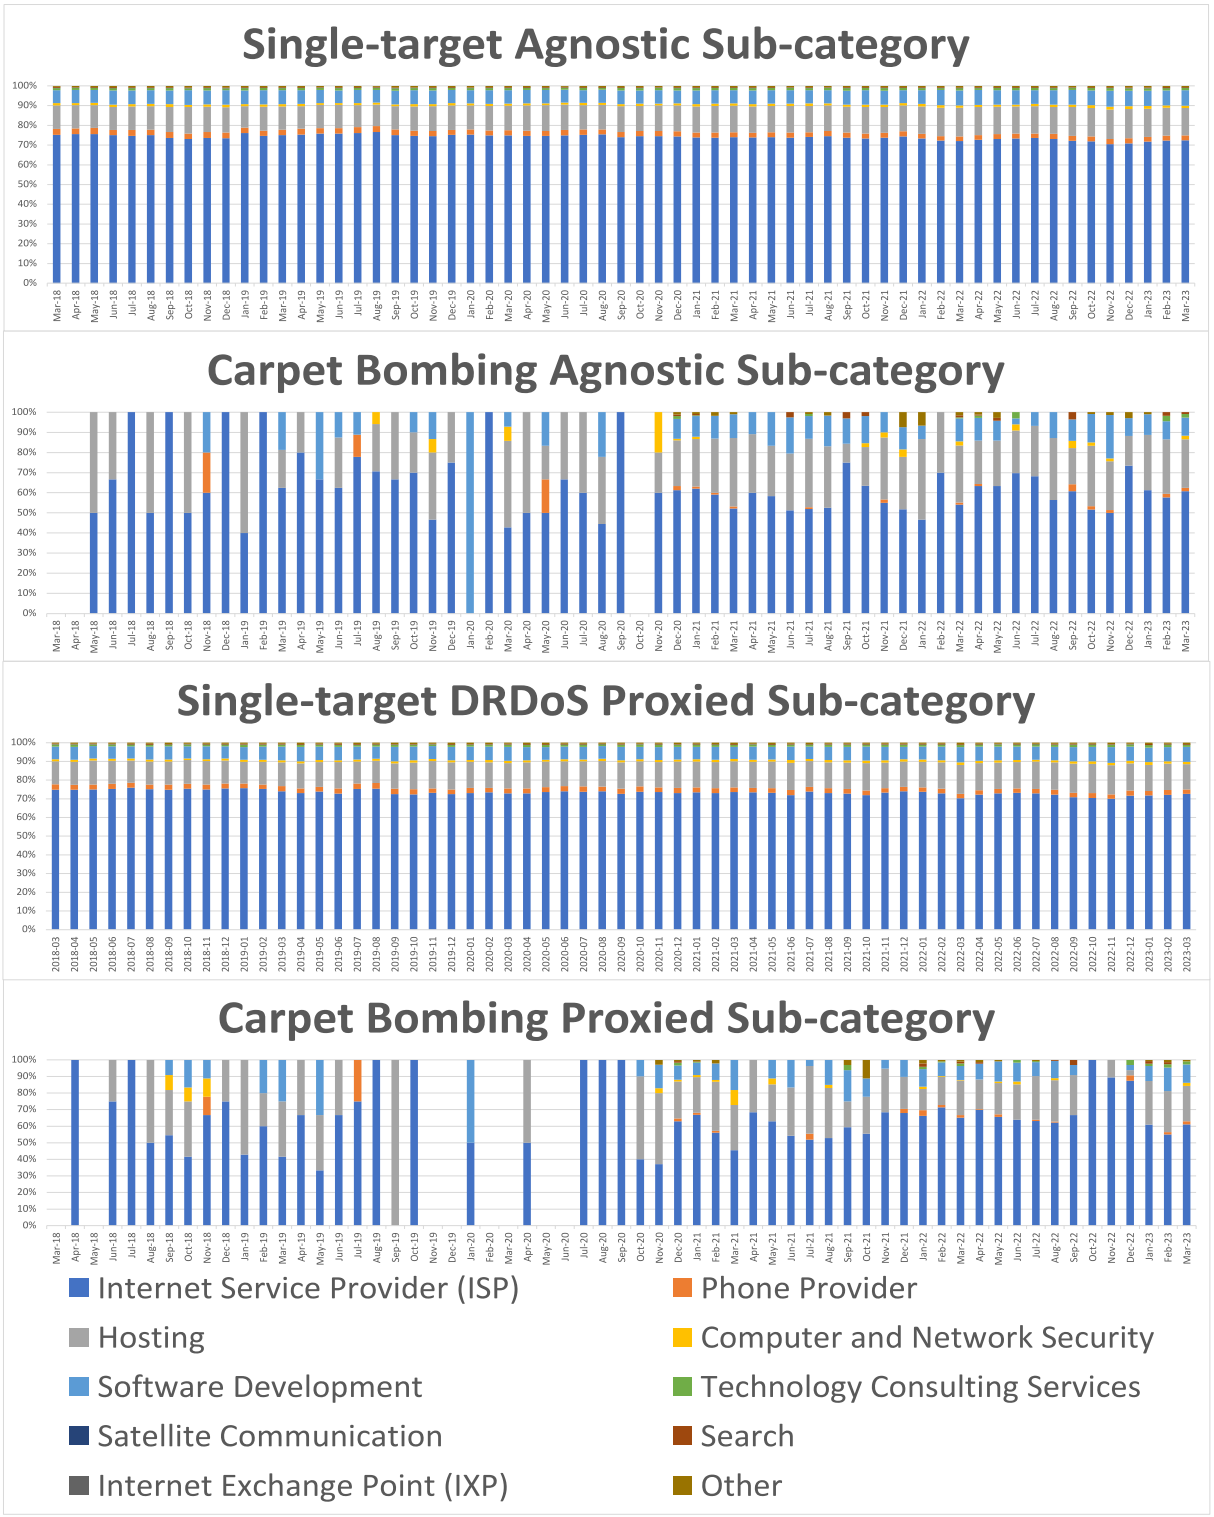

Fig. 11 The distribution of attacked AS types over time - sub-categories.

event counts and the number of attacked IP addresses. The horizontal axis represents time, and the vertical axis shows the aggregation event counts. Due to the significant difference in magnitude between the aggregated event counts and the number of attacked IP addresses in carpet bombing, the number of IP addresses is represented on the right axis. Figure 13 shows that the number of IP addresses affected by a single attack is significantly higher in carpet bombing.

Next, we compare the distribution of Connection Types for the targeted IP addresses. The results of the IP Connection Type comparison are presented in **Fig. 14**. The horizontal axis represents time, and the vertical axis shows the percentage of each category. From Fig. 14, it is evident that the targeted IP addresses in the DRDoS attacks aimed at a single target are concentrated in the 'Cable/DSL' category. In contrast, carpet bombing attacked more IP addresses belonging to the 'Corporate' category.

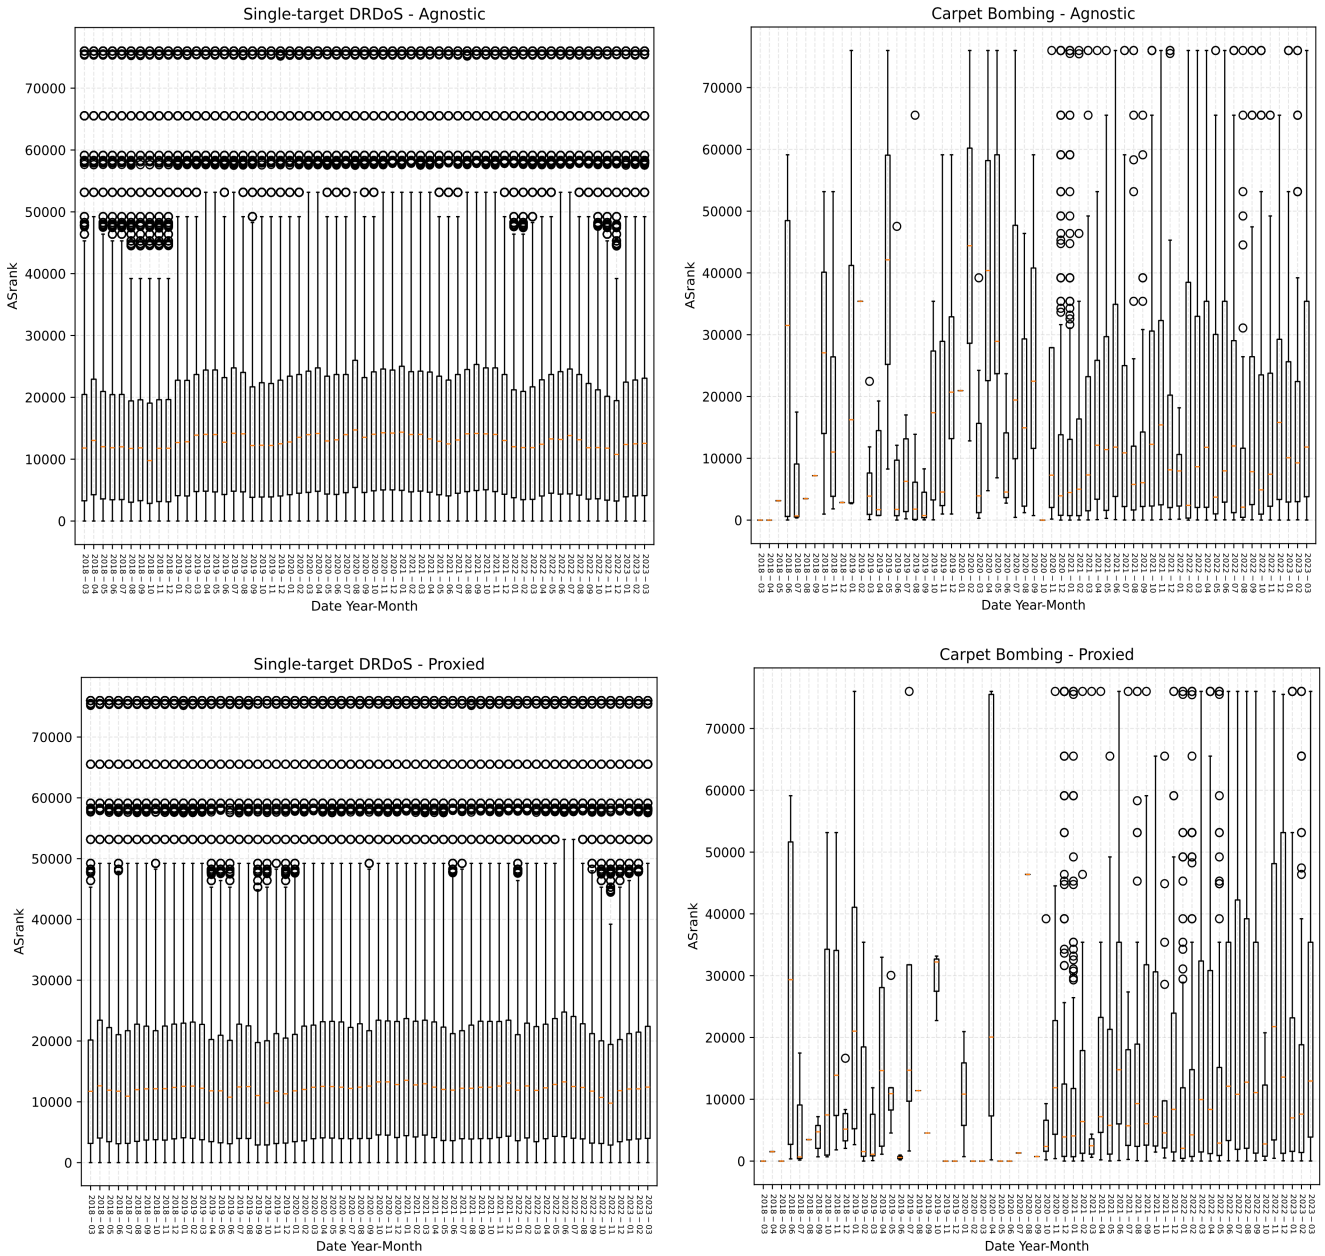

Fig. 12 The result of ASrank analysis.

3.4 Packet Count

In this section, we compare the differences in attack volume. **Figure 15** illustrates the distribution of packet counts before and after aggregation. The horizontal axis represents the sensor numbers, and the vertical axis represents the packet counts. Sensors 204 to 211 correspond to Agnostic Honeypots, while sensors 009 to 020 represent Proxied Honeypots. Examining the packet count distribution before aggregation, it is evident that the packet count for each event in carpet bombing is relatively low. This aligns with our expectations, as the dispersed nature of carpet bombing's attack targets leads to a wide variance in packet counts, making it challenging to recognize using traditional event definitions.

The differing packet count distributions between Agnostic and Proxied Honeypots can be attributed to the distinct ports they monitor. Proxied Honeypots observe only ports frequently exploited for attacks, while Agnostic Honeypots monitor all ports. Consequently, Proxied Honeypots may encounter information from non-monitored ports, affecting the overall packet count trends. Focusing on either type of Honeypot in isolation yields similar packet count distributions, supporting the above inferences.

Upon examining the aggregation results, it becomes evident that the packet count for carpet bombing is significantly higher than that for DRDoS attacks targeting a single target.

3.5 Attack Duration

We present the results for attack duration in **Fig. 16**. From this figure, it becomes apparent that regardless of aggregation, the attack duration for carpet bombing is surprisingly longer than that of DRDoS attacks targeting a single target.

3.6 Types of TTL

In this section, we analyze the distribution of Time to Live (TTL) types. TTL signifies the lifespan of a packet being trans-

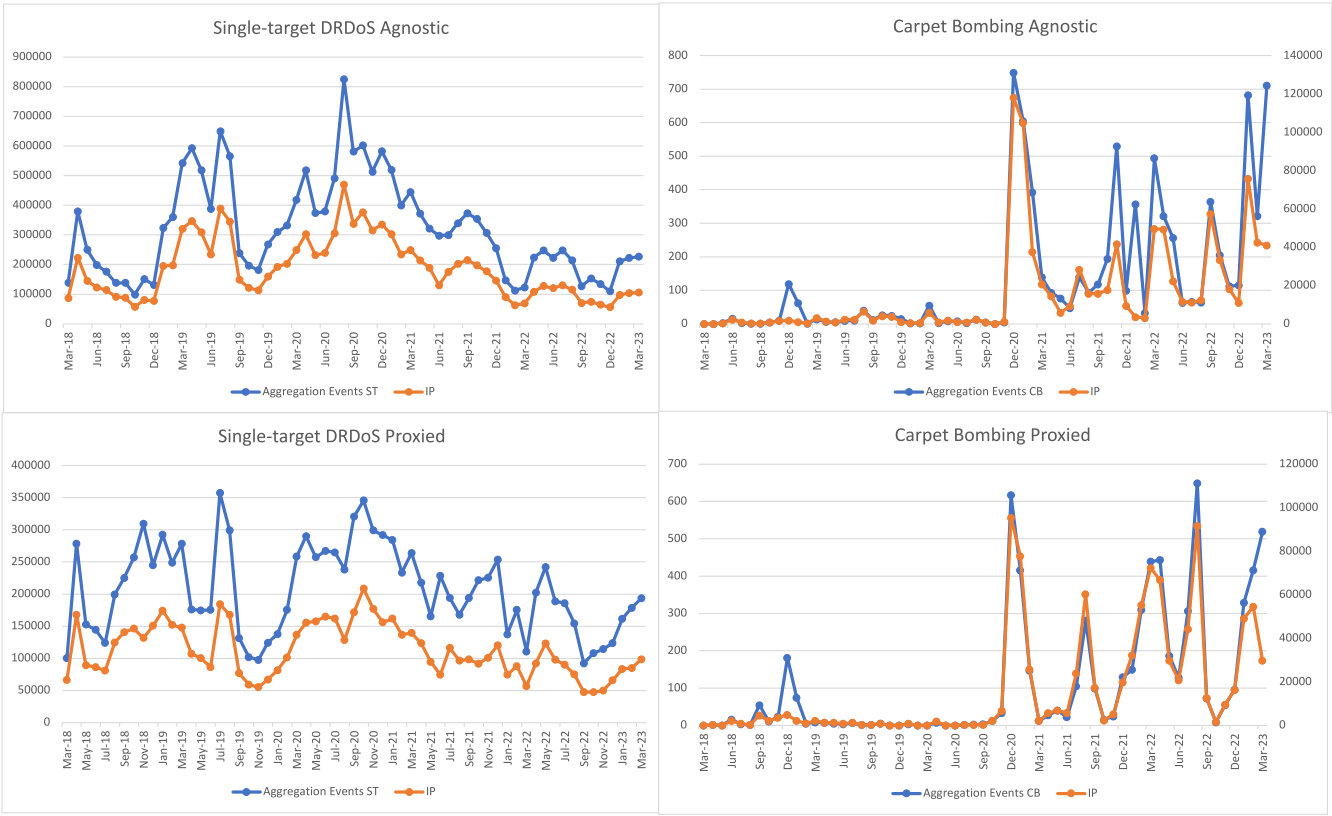

Fig. 13 Monthly aggregation event counts and the number of attacked IP addresses.

ferred over the Internet, and the TTL value decreases by 1 for each transfer. In essence, if the number of TTL types is 1, we can conclude that the attack originated from the same source. If there are multiple TTL types, it indicates a multi-source attack. However, sudden changes in the packet transfer path or alterations in TTL values for some packets due to various reasons may lead to the misinterpretation of attacks coming from the same source as multi-sourced attacks, introducing potential errors. Moreover, since the aggregation combines attacks occurring in the same network block during the same timeframe, any network scans or other communications from the aggregated network block during that period might lead to misinterpretations. To mitigate these effects, when calculating the number of TTL types in an aggregation event, packets with a single TTL value are considered outliers and discarded if they do not constitute at least 1/100 of the overall packet count for the aggregated event. The distribution of TTL types is illustrated in **Table 3**. From Table 3, it is evident that the majority of DRDoS attacks targeting a single target originated from a single source. In contrast, carpet bombing attacks exhibited a higher occurrence of multi-sourced attacks.

3.7 Exploitation status of ports

In this section, we analyze the abuse patterns of ports. The overall port abuse situation is depicted in **Table 4**. As Agnostic Honeypot observes attacks on all ports, it observes more attacks than Proxied Honeypot. Consequently, among the attacks observed by Agnostic Honeypot, it was found that port 389 (LDAP, Lightweight Directory Access Protocol) was the most abused. Among the attacks observed by Proxied Honeypot, port 123 (NTP, Network Time Protocol) was the most abused. The distribution of port types is shown in **Table 5**. Similar to TTL, outliers are removed when calculating the number of port types. A packet using a specific port is considered an outlier and discarded if it does not constitute at least 1/100 of the overall packet count for the aggregation event. Table 5 reveals that most DRDoS attacks targeting a single target utilize a single service, whereas carpet bombing attacks tend to exploit multiple services.

4. Discussion

The primary objective of this study is to uncover the differences in characteristics between carpet bombing and single-target DRDoS attacks. DRDoS attacks have been causing harm to the Internet for many years. While countermeasures against them have become more robust, attack methods are also evolving. Carpet bombing is a newly emerged attack technique in recent years and poses a significant threat to large organizations, yet effective countermeasures have been lacking. The results of this study are expected to make a substantial contribution to addressing the challenges posed by carpet bombing.

4.1 Definition of Carpet Bombing

Carpet Bombing stands for attacks that target many destination IP addresses at once. In this study, we conducted analysis by recognizing attacks that exhibit extreme concentration as carpet bombing. However, attacks that target multiple IP addresses simultaneously without concentrating on a specific target are still poorly understood. For instance, carpet bombing that targeted

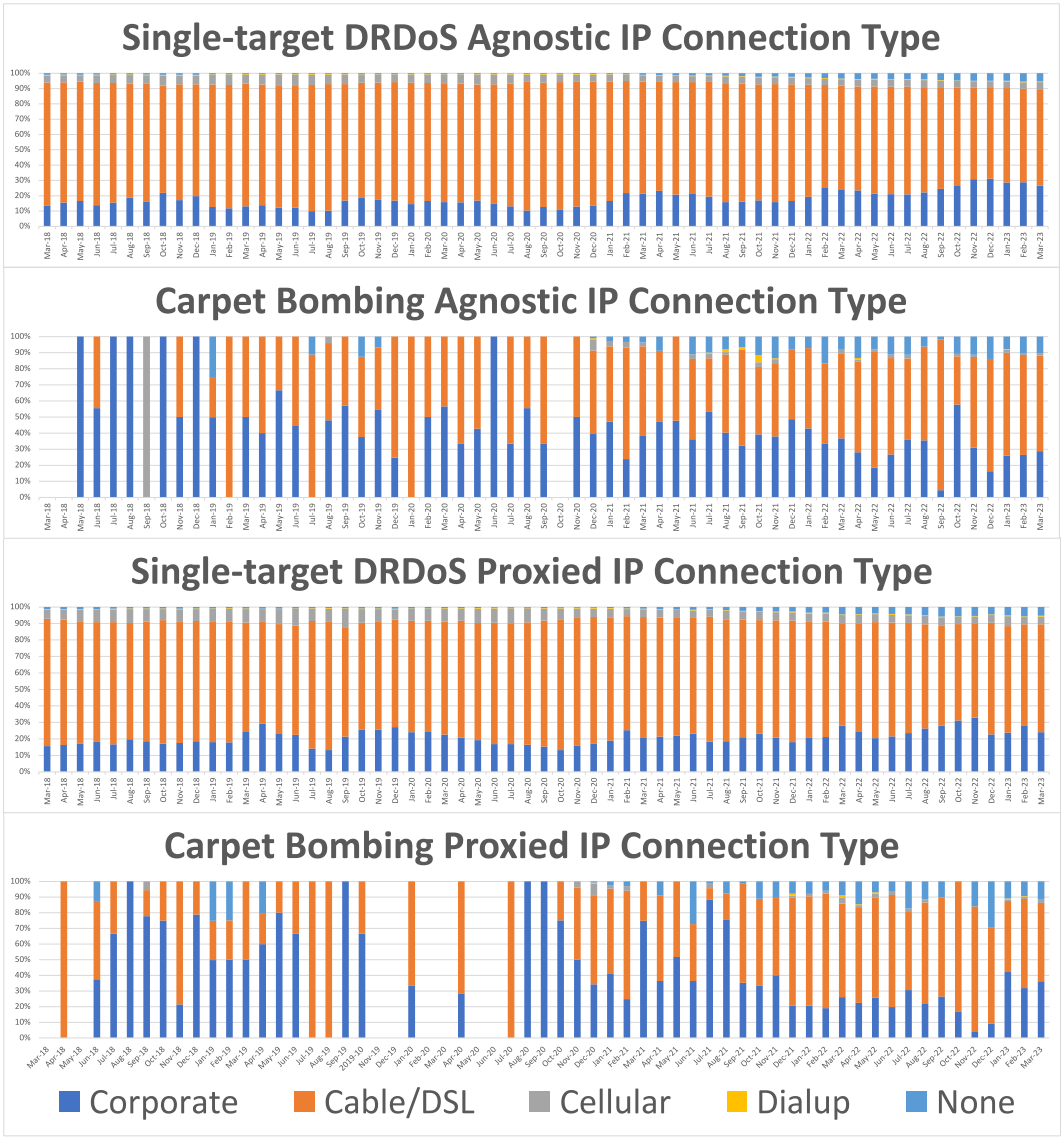

Fig. 14 The results of the IP connection type comparison.

only about half of a /24 network but didn't concentrate much on the target had a significant presence in the aggregated results at the network block level. The data analyzed in this study was aggregated at /16 network block levels and with a time interval of 300 seconds, resulting in aggregation events with 251–256 IPs of attack targets. However, there were also a considerable number of aggregation events with attack target numbers ranging from 2 to 250 and 257 or more. Given the ambiguity in the definition of carpet bombing itself, defining these attacks is challenging. Aggregation events with a small number of attack targets can be referred to as DRDoS attacks targeting multiple entities simultaneously. But how many IP addresses does an attack need to target to be considered carpet bombing? One potential solution to this issue is to perform aggregation based on network blocks that actually exist in the world, without worrying about an organization's network ownership status. Assuming that attackers do not conduct meaningless attacks, they are likely to target network ranges owned by organizations as attack targets. By aggregating based on real-world network blocks, it is possible to minimize the impact from unrelated attacks.

4.2 Results of the Comparison

The analysis results highlight seven differences between carpet bombing and single-target DRDoS attacks. In the comparison of AS types, carpet bombing targeted more 'Hosting' type ASes than single-target DRDoS attacks. In the ASrank comparison, some instances of carpet bombing targeted relatively important entities, while others attacked less significant ASes. Regarding the IP Connection Type, carpet bombing attacked more IP addresses belonging to the 'Corporate' category. In the comparison of packet numbers, the number of pre-aggregation events for carpet bombing was lower, but post-aggregation, the packet numbers for carpet bombing were significantly higher. In the comparison of attack duration, carpet bombing had longer durations both before and after aggregation compared to single-target DRDoS attacks. The TTL comparison revealed that carpet bombing predominantly involved attacks from multiple sources. Finally, in the comparison of service abuse patterns, carpet bombing exhibited a tendency to exploit multiple services simultaneously. These findings are believed to be highly beneficial for the defense against carpet bombing. For instance, the insights can definitely be uti-

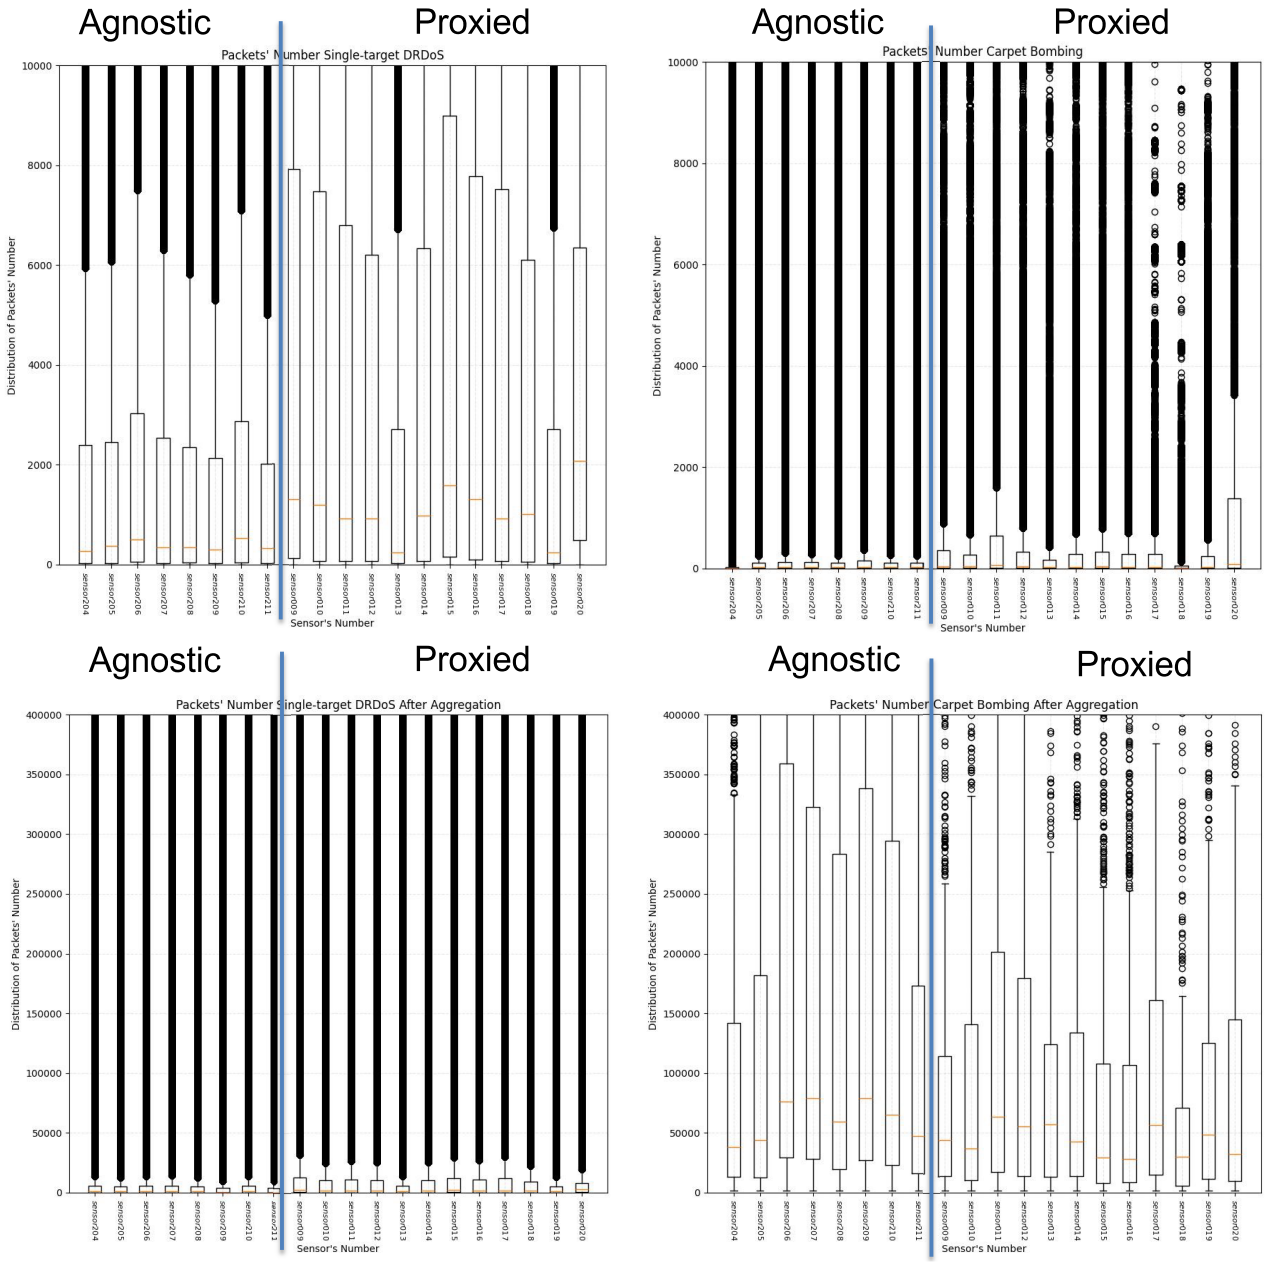

Fig. 15 The distribution of packet counts before and after aggregation.

lized by network operators to recognize the attacks better. Similarly, we can improve the alerting system [6], [7] by Amppot that has been running since 2014, which detects DRDoS attacks from traffic that DRDoS honeypots collect in real-time and sends their alert information to collaborating organizations in order to support early response against DRDoS attacks. You can refer to our amppot website here. We have sent alerts to major Japanese ISPs and over 6,000 network operators worldwide via Shadowserver. Carpet Bombing has been a problem, as a single Carpet Bombing would create too many alerts if not aggregated. We can alert now with more meaningful attribution of attacks, such as targets, ports, duration, and packet count, which are analyzed in this study.

4.3 Aggregation

The aggregation algorithm used in this study performs comprehensive aggregation at the network block level without prior examination or filtering of packets. This lack of filtering causes other communication packets observed in the same time period to be aggregated together. This, to some extent, affects the accuracy of the analysis results. To address this issue, scrutinizing the payload and TTL of packets is considered feasible. Since the effectiveness of DRDoS attacks heavily relies on amplification rates, attackers need to capture devices that can serve as reflectors through network scans beforehand. Attackers are likely to pay attention to which payload yields the most powerful results during this capture process. Furthermore, it is unlikely that the payload used in DRDoS attacks has the same content as the payload of packets with different purposes. Therefore, even if attacks other than carpet bombing are accidentally observed in the same time period, examining the payload content should reveal some differences. Additionally, scrutinizing the TTL values of packets can help eliminate mistakenly aggregated packets. TTL, in network

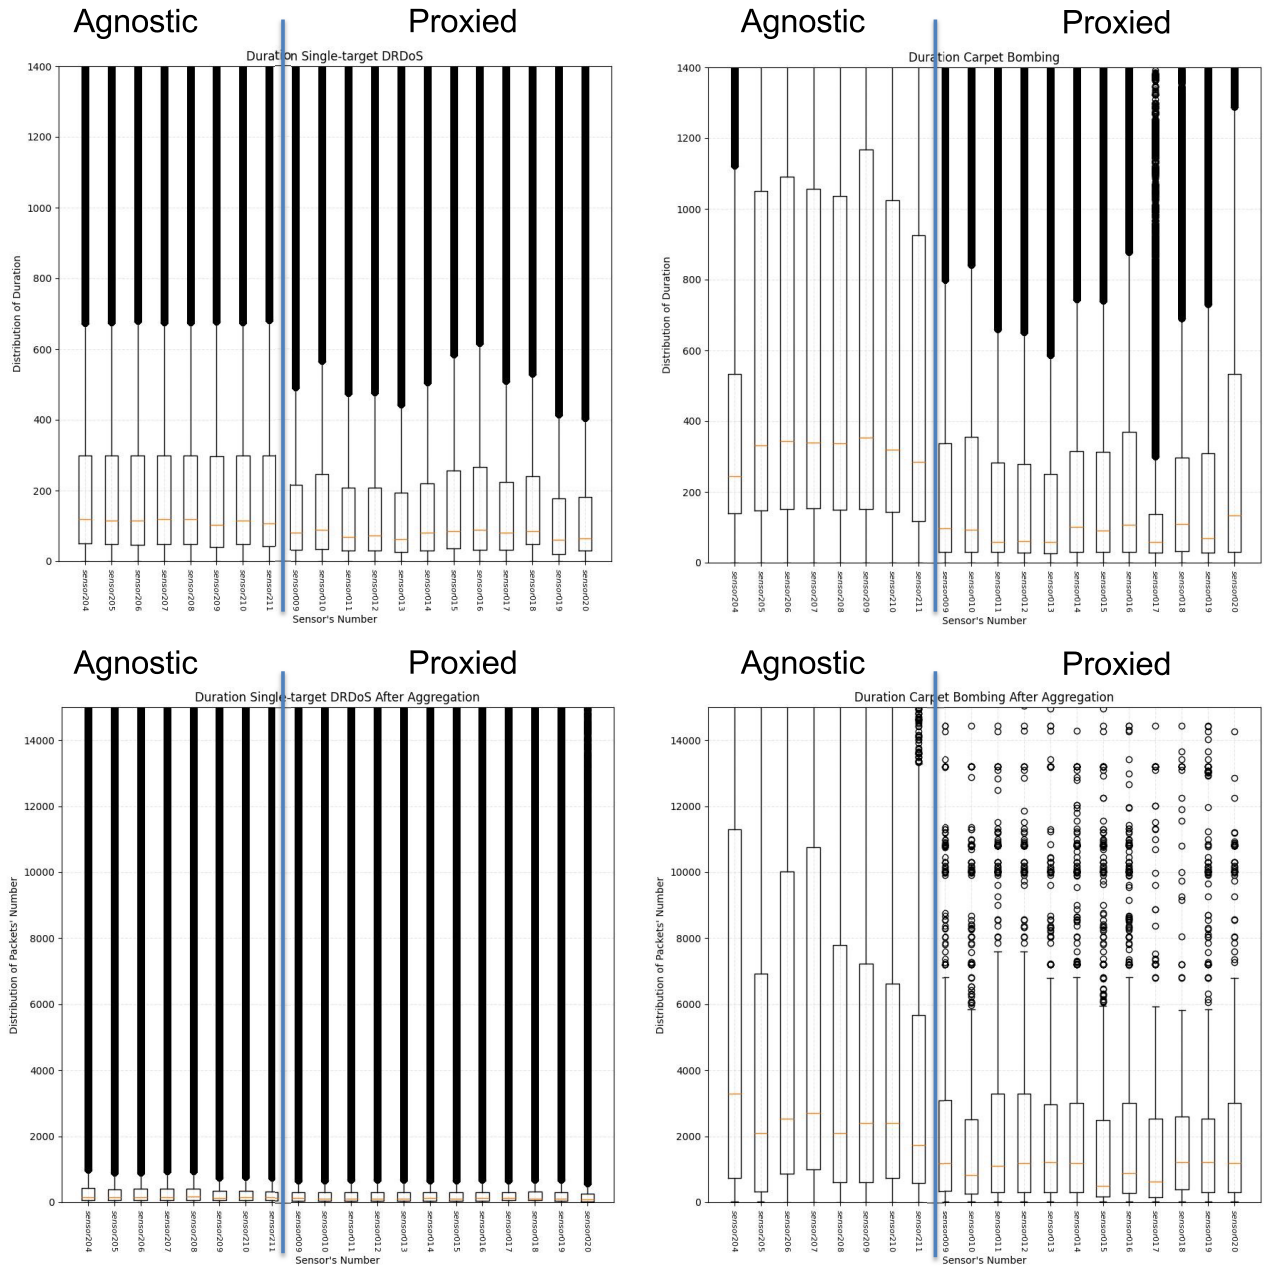

Fig. 16 The distribution of attack duration before and after aggregation.

communication, indicates the time or "hops" a packet can survive inside a network during transmission. In other words, TTL can be seen as a measure of the distance between the attacker and the observation system. While this method may not be effective for carpet bombing from multiple sources, it can be reasonably effective for carpet bombing from a single source.

While we have outlined methods to enhance the accuracy of aggregation, it is essential to note that our honeypots do not observe all attacks occurring globally; they only capture a subset of attacks. Nevertheless, we believe that these results are highly meaningful.

5. Related Work

While DDoS attacks continue to significantly impact the Internet, countermeasures against such attacks have progressed. Amid these efforts, a new attack method called carpet bombing in DDoS attacks has emerged in recent years, presenting a new challenge. Unlike targeting a single IP address, carpet bombing aims at large network blocks owned by major organizations, companies, or Internet Service Providers (ISPs), posing a substantial threat. A recent report indicates a 300% increase in carpet bomb DDoS attacks in 2022 [31]. Carpet bombing is not only utilized in DDoS attacks but is also employed in a specific type of DDoS attack known as DRDoS attacks.

Existing research has provided limited insights into carpet bombing. With Tiago et al. have designed and implemented a honeypot that emulates reflectors for nine protocols (Chargen, DNS, NTP, Memcached, QOTD, SSDP, CoAP, CLDAP, and Steam) that are exploited in DRDoS attacks. They described several features of multiprotocol attacks and compared them to monoprotocol attacks that occurred in the same period, and characterized the carpet bombing attacks. They defined a carpet

Table 3 The distribution of TTL types.

Agnostic Carpet Bombing				**Agnostic Single-target DRDoS Attacks**		
Types of TTL	Count	Percentage		Types of TTL	Count	Percentage
1	3,352	42.323%		1	10,893,676	55.723%
Over 100	1,242	15.682%		2	3,939,947	20.154%
2	1,022	12.904%		Over 100	2,793,243	14.288%
4	710	8.965%		3	1,146,474	5.864%
3	592	7.475%		4	416,899	2.133%
5	564	7.121%		10	168,070	0.860%
6	53	0.669%		5	110,926	0.567%
10	40	0.505%		6	43,225	0.221%
16	36	0.455%		7	22,639	0.116%
7	27	0.341%		8	9,467	0.048%
Proxied Carpet Bombing				**Proxied Single-target DRDoS Attacks**		
Types of TTL	Count	Percentage		Types of TTL	Count	Percentage
6	708	10.884%		1	8,895,407	70.277%
Over 100	697	10.715%		2	2,209,832	17.459%
10	673	10.346%		Over 100	555,989	4.393%
11	589	9.055%		3	498,372	3.937%
9	566	8.701%		10	192,620	1.522%
2	563	8.655%		4	189,597	1.498%
4	531	8.163%		5	55,250	0.436%
5	450	6.918%		6	24,850	0.196%
7	447	6.872%		9	14,306	0.113%
8	391	6.011%		7	11,909	0.094%

Table 4 The overall port abuse situation.

Agnostic Carpet Bombing				**Agnostic Single-target DRDoS Attacks**		
Port Number	Count	Percentage		Port Number	Count	Percentage
389	27,356,814	38.659%		389	40,695,464	33.682%
443	2,962,647	4.187%		53	17,346,993	14.357%
53	2,193,519	3.100%		123	11,131,308	9.213%
3702	2,009,941	2.840%		3283	6,894,138	5.706%
37810	1,970,564	2.785%		111	5,218,515	4.319%
3283	1,600,392	2.262%		3702	4,920,915	4.073%
123	1,258,800	1.779%		37810	4,056,924	3.358%
161	1,221,397	1.726%		1434	4,007,279	3.317%
11211	279,059	0.394%		161	3,445,195	2.851%
10074	242,509	0.343%		137	2,789,850	2.309%
Proxied Carpet Bombing				**Proxied Single-target DRDoS Attacks**		
Port Number	Count	Percentage		Port Number	Count	Percentage
123	5,799,413	48.543%		123	57,145,019	57.482%
161	2,641,138	22.107%		53	13,549,753	13.630%
53	1,833,632	15.348%		11211	9,820,489	9.878%
11211	1,264,930	10.588%		161	9,272,469	9.327%
19	301,046	2.520%		1900	4,960,942	4.990%
1900	88,268	0.739%		19	4,275,713	4.301%
17	18,509	0.155%		17	388,545	0.391%

bombing attack as an attack targeting multiple IP addresses from the same CIDR block. The results from 731 days of data collected by their honeypot showed more than 3.7% of all attacks employed carpet bombing, affecting 21.8% of the victims observed. Also, they showed that when attackers target a larger fraction of a CIDR block, the number of requests per host tends to be smaller [30].

Their study analyzed traditional DRDoS attacks and carpet bombing from the perspective of exploited protocol types, but did not compare these two types of attacks from the viewpoint of targeted IP addresses. The definition of carpet bombing has been ambiguous thus far, with specific attack targets, volumes, duration, and other characteristics yet to be clarified. We defined carpet bombing as attacks characterized by extreme concentration, aiming to minimize ambiguity and analyze the features of carpet bombing, comparing it with traditional DRDoS attacks. We believe that this is the first study to analyze the differences between carpet bombing and single-target DRDoS attacks, which could facilitate more effective mitigation against these attacks.

In our study, it was observed that pre-aggregated carpet bombing attacks exhibited dispersed attacks with relatively fewer packets per IP address and shorter attack durations, whereas post-aggregated carpet bombing attacks showed relatively higher packet counts and longer attack durations. This suggests that pre-aggregated carpet bombing attacks might not have been recognized as attacks. Moreover, traditional single IP address-based attack detection methods might incorrectly identify a single carpet bombing attack as multiple DRDoS attack events. Understanding the differences in attack volume and duration can facilitate updating traditional detection methods to accurately recognize communications that are indeed attacks but are currently unrecognized as such, enabling the correct identification of attack events without dispersion.

Table 5 The distribution of port types.

Agnostic Carpet Bombing				**Agnostic Single-target DRDoS Attacks**		
Types of Port	Count	Percentage		Types of Port	Count	Percentage
1	5,484	69.242%		1	17,562,271	89.834%
2	1,074	13.561%		2	1,062,221	5.433%
3	579	7.311%		3	256,674	1.313%
4	521	6.578%		8	133,729	0.684%
Over 100	107	1.351%		9	126,500	0.647%
5	102	1.288%		7	113,491	0.581%
6	39	0.492%		4	95,749	0.490%
7	8	0.101%		6	74,322	0.380%
10	2	0.025%		5	72,585	0.371%
15	1	0.013%		10	33,520	0.171%
Proxied Carpet Bombing				**Proxied Single-target DRDoS Attacks**		
Types of Port	Count	Percentage		Types of Port	Count	Percentage
1	4,798	73.759%		1	11,631,672	91.895%
2	1,327	20.400%		2	843,667	6.665%
3	361	5.550%		3	169,959	1.343%
4	18	0.277%		4	12,173	0.096%
5	1	0.015%		6	77	0.001%
				5	29	0.000%

We have been providing semi-real-time alerts to Japanese ISPs through ICT Information Sharing and Analysis Center (ICT-ISAC), Japan. Also, we have been issuing DDoS alerts to the nonprofit international security organization Shadowserver Foundation, who share these provided alerts with 132 National CSIRTs covering 173 countries and territories and over 6,500 organizations worldwide [32], [33]. These alerts were used by the network operators and National CSIRTs to mitigate the attacks and understand their trends. We believe that our study is the first step towards more accurate alerting with better understanding of the nature of two types of attacks: carpet bombing and (traditional) single-target attacks.

6. Conclusion & Future Work

Distributed Denial of Service attacks continue to pose a significant threat to the Internet. While countermeasures against these attacks are advancing, so too are the techniques employed by attackers. This study aimed to analyze and compare the characteristics of carpet bombing and single-target Distributed Reflection Denial of Service attacks, shedding light on their differences. Despite the ongoing challenges and the ambiguity surrounding the definition of carpet bombing, this research represents a crucial step forward in addressing the issues obscured by the previously unseen fog. Although the method used to aggregate attacks in this study is imperfect, recognizing highly concentrated attack events as carpet bombing has minimized the impact of potential packet mixing and the uncertainty of the attack scope. Furthermore, the comparison of features revealed distinct differences between carpet bombing and single-target-focused DRDoS attacks, contributing significantly to the development of effective countermeasures against carpet bombing.

Looking ahead, improving the accuracy of aggregation through attack packet fingerprints could enhance the understanding of carpet bombing. Separating carpet bombing from the multitude of existing attacks could simplify the development of countermeasures. Additionally, considering aggregation on actual network ranges allocated according to demand in the real world may contribute to the precision of aggregation.

Acknowledgments This research was partly conducted under a contract of “Research and development on IoT malware removal / make it non-functional technologies for effective use of the radio spectrum” among “Research and Development for Expansion of Radio Wave Resources (JPJ000254)”, which was supported by the Ministry of Internal Affairs and Communications, Japan. A part of these research results were obtained from the commissioned research (No.JPJ012368C05201) by National Institute of Information and Communications Technology (NICT), Japan. This work was supported by JSPS KAKENHI Grant Number 22H03588. This work was supported by JSPS KAKENHI Grant Number 21KK0178 for international collaboration.

References

[1] FORTINET: What Is DDOS Attack, available from ⟨https://www.fortinet.com/resources/cyberglossary/ddos-attack⟩ (accessed 2023-12-05).

[2] arturai: DRDOS ATTACKS, available from ⟨https://www.arturai.com/en/support/faqs/drdosattacks#::text=DrDoS%20stands%20for%20Distributed%20Reflection,victim%20hosts%20to%20the%20target⟩ (accessed 2023-12-05).

[3] Makita, D., Nishizoe, T., Koide, T., Tsutsumi, T., Kanei, F., Mori, H., Yoshioka, K., Matsumoto, T., Inoue, D. and Nakao, K.: Development of Integrated DRDoS Attack Observation System toward Early Response, *SCIS2015* (2015).

[4] Nishizoe, T., Makita, D., Yoshioka, K. and Matsumoto, T.: Observing DRDoS Attacks with Protocol-noncompliant Honeypot, *The 32nd Symposium on Cryptography and Information Security SCIS2015* (2015).

[5] Kramer, L., Krupp, J., Makita, D., Nishizoe, T., Koide, T., Yoshioka, K. and Rossow, C.: AmpPot: Monitoring and Defending Amplification DDoS Attacks, *Proc. Research in Attacks, Intrusions, and Defenses* (*RAID15*), Lecture Notes in Computer Science, Vol.9404, pp.615–636 (2015).

[6] Makita, D., Nishizoe, T., Yoshioka, K., Matsumoto, T., Inoue, D. and Nakao, K.: DRDoS Attack Alert System for Early Incident Response, *Information Processing Society of Japan*, Vol.57, No.9, pp.1974–1985 (2016).

[7] Yoshioka, K., Makita, D., Nishizoe, T., Matsumoto, T., Inoue, D. and Nakao, K.: Real-Time Detection and Alerting of DRDoS Attacks, *The Institute of Electronics, Information and Communication Engineers* (2016).

[8] Makita, D., Yoshioka, K., Matsumoto, T., Nakazato, J., Shimamura, J. and Inoue, D.: Correlation Analysis between DNS Honeypot and Darknet toward Proactive Countermeasures against DNS Amplification Attacks, *Infornatiom Processing Society of Japan*, Vol.56, No.3, pp.921–931 (2015).

[9] Koidem, T., Makita, D., Yoshioka, K. and Matsumoto, T.: Observation

and Analysis of TCP-based Reflection DDoS Attacks Using Honeypot, *Research in Attacks, Intrusions, and Defenses* (*RAID15*), Poster Session (2015).
[10] Gao, Y., Feng, Y., Kawamoto, J. and Sakurai, K.: A Machine Learning Based Approach for Detecting DRDoS Attacks and Its Performance Evaluation, *Proc. 11th Asia Joint Conference on Information Security* (2016).
[11] Tsunoda, H., Ohta, K., Yamamoto, A., Ansari, N., Waizumi, Y. and Nemoto, Y.: Detecting DRDoS attacks by a simple response packet confirmation mechanism, *Computer Communications*, Vol.31, pp.3299–3306 (2008).
[12] Choi, H., Park, H. and Lee, H.: A Study on Amplification DRDoS Attacks and Defense, *The Journal of Korea Institute of Information, Electronics, and Communication Technology*, Vol.8, No.5, pp.429–437 (2015).
[13] Ryba, F.J., Orlinski, M., Wahlisch, M., Rossow, C. and Schmidt, T.C.: Amplification and DRDoS Attack Defense – A Survey and New Perspectives, arXiv:1505.07892 (2016).
[14] Nuiaa, R.R., Manickam, S., Alsaeedi, A.H. and Alomari, E.S.: A new proactive feature selection model based on the enhanced optimization algorithms to detect DRDoS attacks, *International Journal of Electrical and Computer Engineering* (*IJECE*), Vol.12, No.2, pp.1869–1880 (2022).
[15] Subramani, K., Perdisci, R. and Konte, M.: Detecting and Measuring In-The-Wild DRDoS Attacks at IXPs, Bilge, L. et al. (Eds.), *DIMVA 2021*, LNCS 12756, pp.42–67 (2021).
[16] Wei, W., Chen, F., Xia, Y. and Jin, G.: A Rank Correlation Based Detection against Distributed Reflection DoS Attacks, *IEEE Communications Letters*, Vol.17, No.1 (2013).
[17] Nawrocki, M. Kristoff, J. Hiesgen, R. Kanich, C. Schmidt, T.C. and Wählisch, M.: SoK: A Data-driven View on Methods to Detect Reflective Amplification DDoS Attacks Using Honeypots, *Proc. Euro S&P '23*, IEEE (2023).
[18] Bekeneva, Y. and Shorov, A.: Development of protection mechanisms against DRDoS-attacks and combined DRDoS-attacks, *Vibroengineering Procedia*, Vol.12, Jve International Ltd. (2017).
[19] Raviv, I.: DDoS Carpet-Bombing – Coming In Fast And Brutal, available from ⟨https://www.radware.com/blog/ddos-protection/2023/07/ddos-carpet-bombing-coming-in-fast-and-brutal/⟩ (accessed 2023-12-05).
[20] NETSCOUT: Defending Against Carpet Bombing DDoS Attacks, available from ⟨https://www.netscout.com/use-case/carpet-bombing-attacks⟩ (accessed 2023-12-05).
[21] Nishizoe, T., Makita, D., Yoshioka, K. and Matsumoto, T.: Early Detection of New DRDoS Attack with Protocol-Agnostic Honeypot, *The Institute of Electronics, Information and Commuication Engineers Information and Communication System Security* (*ICSS*) (2017).
[22] Nmap, available from ⟨http://nmap.org/⟩.
[23] Durumeric, Z., Wustrow, E. and Halderman, J.A.: ZMap: Fast Internet-wide Scanning and Its Security Applications, *USENIX Security Symposium* (2013).
[24] Mao, Q., Makita, D., Yoshioka, K. and Matsumoto, T.: Analysis of DRDoS Honeypot Observation Data to Understand the Actual Situation of Carpet bombing DRDoS Attacks, *IPSJ Journal*, Vol.64, No.9 (2023).
[25] Python, available from ⟨https://www.python.org/⟩.
[26] MaxMind GeoIP Databases, available from ⟨https://www.maxmind.com/en/geoip-databases⟩.
[27] Ziv, M., Izhikevich, L., Ruth, K., Izhikevich, K. and Durumeric, Z.: ASdb: A System for Classifying Owners of Autonomous Systems, *ACM Internet Measurement Conference* (*IMC*) (2021).
[28] ASrank, available from ⟨https://asrank.caida.org/⟩.
[29] MaxMind GeoIP2 Connection Type Databases, available from ⟨https://dev.maxmind.com/geoip/docs/databases/connection-type⟩.
[30] Heinrich, T., Obelheiro, R.R. and Maziero, C.A.: New Kids on the DRDoS Block: Characterizing Multiprotocol and Carpet Bombing Attacks, *PAM 2021*, pp.269–283 (2021).
[31] Juniper 2023 DDOS THREAT INTELLIGENCE REPORT, available from ⟨https://www.juniper.net/content/dam/www/assets/analyst-reports/us/en/2023/corero-ddos-threat-intelligence-report.pdf⟩.
[32] Amppot: Honeypot for Monitoring Amplification DDoS Attacks, available from ⟨https://sec.ynu.codes/dos⟩.
[33] Shadowserver, available from ⟨https://www.shadowserver.org/what-we-do/⟩.

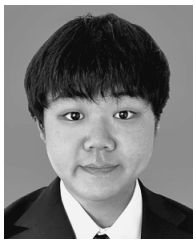

Qingxin Mao completed his Master's program (M.Sc. in Informatics) in the Department of Information Environment, Graduate School of Environment and Information Sciences, Yokohama National University in September 2022. Entered his Doctoral program (Ph.D. in Informatics) in the Department of Information Environment, Graduate School of Environment and Information Sciences, Yokohama National University in April 2023.

Daisuke Makita received his Ph.D. in Engineering from Yokohama National University, Japan, in 2017. Since 2014, he has been working as a researcher at the National Institute of Information and Communications Technology (NICT) in Japan. His research focuses on cybersecurity, with a particular emphasis on the observation and analysis of cyber-attacks.

Michel van Eeten is a professor of cybersecurity at Delft University of Technology. His team analyses large-scale Internet measurement and incident data to identify how the markets for Internet services deal with security risks. He has conducted empirical studies for the ITU, OECD, European Commission and the Dutch government on the economics of malware, the impact of cybercrime and the role of ISPs in mitigating botnets and bad hosting. He is also a member of the Dutch Cyber Security Council.

Katsunari Yoshioka is a Professor at Yokohama National University. His research interests cover a wide area of system security and network security, including malware analysis and IoT security. He received the commendation for science and technology from the minister of MEXT, Japan, in 2009, the award for contribution to Industry-Academia-Government Collaboration by the minister of MIC, Japan, in 2016, and the Culture of Information Security Award in 2017.

Tsutomu Matsumoto received a Doctor of Engineering degree from the University of Tokyo in 1986. He is a Distinguished YNU Professor of Yokohama National University and a Fellow of the National Institute of Advanced Industrial Science and Technology (AIST) and the Director of the Cyber Physical Security Research Center at AIST. He has been interested in the research and education of Embedded Security Systems such as IoT devices, cryptographic hardware, in-vehicle networks, instrumentation and control security, tamper resistance, biometrics, artifactmetrics, and countermeasures against cyber-physical attacks. He serves as the chair of the Japanese National Body for ISO/TC68 (Financial Services) and the Cryptography Research and Evaluation Committees (CRYPTREC) and as an associate member of the Science Council of Japan (SCJ). He received the IEICE Achievement Award, the DoCoMo Mobile Science Award, the Culture of Information Security Award, the MEXT Prize for Science and Technology, and the Fuji Sankei Business Eye Award.