

Witnessing Distributed Denial-of-Service Traffic from an Attacker's Network

Sin-seok Seo

Dept. of Computer Science and Engineering
Pohang University of Science and Technology
(POSTECH), Korea
Email: sesise@postech.ac.kr

Young J. Won

IJ Research Lab.
Japan
Email: young@ij.ad.jp

James Won-Ki Hong

Division of IT Convergence Engineering
Pohang University of Science and Technology
(POSTECH), Korea
Email: jwkhong@postech.ac.kr

Abstract—In July 2009, surprising large-scale Distributed Denial-of-Service (DDoS) attacks simultaneously targeted US and South Korean government, military, and commercial websites. Initial speculation was that this was well-designed cyber warfare from North Korea, but the truth is still unknown. What was even more surprising was how these critical infrastructures were still vulnerable after a decade of research on DDoS attacks. These particular incidents, the so-called 7.7 (July 7th) DDoS attacks, were highlighted not just because of their success but also because of their well-coordinated strategy. The 3.3 (March 3rd, 2011) DDoS attacks had similar characteristics to the 7.7 DDoS attacks, but they were not as successful because of the rapid vaccination of the zombie hosts. In this paper, we suggest that it is worthwhile to take a step back from the target side of the DDoS attacks and look at the problem in terms of network traffic from the attacker's side. We collected a unique large-scale sample of DDoS attack traffic from the two real-world incidents (not simulated), and we provide an analysis of traffic patterns from the perspective of the attacker's hosting network.

Index Terms—DDoS, Monitoring, Traffic Analysis

I. INTRODUCTION

A typical Distributed Denial-of-Service (DDoS) attack is an attempt to make resources or services unavailable to the intended users by sending a large volume of traffic through the simultaneous cooperation of zombie hosts. A zombie host is infected by malware and attacks target machines upon commands from hackers or according to predefined schedules. Successful DDoS attacks prevent authorized users from accessing their legitimate resources or services. The detection and prevention of DDoS attacks is complicated because of their distributed nature.

DDoS countermeasures can be divided into two categories with respect to the deployment location: (1) target-side countermeasures and (2) attacker-side countermeasures. Most countermeasures against DDoS attacks try to tackle the problem from the target side of the network [1]–[5]. However, the DDoS attacks in July 2009 [6] proved that such approaches are not sufficient for the complete detection and prevention of

consistently evolving large-scale DDoS attacks. Accordingly, we propose that it is worthwhile to study the characteristics of DDoS attack traffic from the viewpoint of the attacker's network.

Two large-scale DDoS attacks occurred recently: one in July 2009 [6], and another in March 2011 [7]. The main targets were US and South Korean websites, and most of the zombie hosts in action were also located in South Korea (see Section III). We assumed that our campus network also had zombie hosts for these two attacks, and we captured traffic traces at the times of the attacks. We also captured normal traffic traces to compare them with the DDoS traffic traces. We analyzed the captured traffic traces to identify the distinctive characteristics of the DDoS attacks from the perspective of an attacker network. The analysis metrics included the number of packets, number of flows, traffic volume, protocol ratio, IP interaction graphlets, flow duration, and average packet size.

II. RELATED WORK

Most of the related work focused on detection and prevention methods for DDoS attacks from the target side of the networks [1]–[5]. However, there is still no perfect countermeasure against DDoS attacks.

Mirkovic *et al.* [8] proposed a series of DDoS impact metrics considering the end-user QoS requirements. They demonstrated that the proposed QoS metrics capture the impact of DDoS attacks more precisely than the legacy metrics. Mao *et al.* [9] analyzed DDoS attacks using multiple data sources obtained from both direct and indirect measurements. The analysis results showed the following: (1) 50 or less Autonomous Systems (ASes) are involved in 70% of DDoS attacks and (2) a small number of ASes produce about 72% of the total attack traffic. Xie *et al.* [10] proposed a method applying principle component analysis and independent component analysis to detect new application-layer-based DDoS attacks that utilize legitimate HTTP requests. These previous studies focused on finding the characteristics of DDoS attacks and detecting them from the target side.

A few studies have proposed the deployment of DDoS defense mechanisms at the attacker networks. Mirkovic *et al.* [11] constantly monitored two-way traffic flows between the network and the rest of the Internet, periodically comparing the

This research was supported by the KCC(Korea Communications Commission), Korea, under the “Novel Study on Highly Manageable Network and Service Architecture for New Generation” support program supervised by the KCA(Korea Communications Agency) (KCA-2011-10921-05003) and WCU (World Class University) program through the National Research Foundation of Korea funded by the Ministry of Education, Science and Technology (R31-2010-000-10100-0).

monitored results with normal flow models. Malliga *et al.* [12] proposed a system that drops suspicious traffic at the source network. Their scheme distinguishes between suspicious traffic and normal traffic using information entropy. However, the two DDoS defense schemes were validated only with simulated traffic.

In this paper, we are not concerned with detection or defense mechanisms from the target network. Rather, we deal with the traffic characteristics of DDoS attacks from the perspective of an attacker network.

III. DDOS DESCRIPTION

Recently, two large-scale DDoS attacks targeted the US and South Korea. The first attacks, the so-called 7.7 (July 7th) DDoS attacks, started at 02:00 on July 5th, 2009¹ and lasted until 18:00 on July 10th, 2009 [6]. The second attacks started at 17:00 on March 3rd, 2011 and the termination time is not specified. This means that it continued until the infected zombie hosts were cured. This event is referred to as the 3.3 (March 3rd) DDoS attacks [7]. In this section, we introduce these two DDoS attacks and describe several distinctive characteristics.

A. 7.7 DDoS Attacks in 2009

The malware that was designed for the 7.7 DDoS attacks was propagated through South Korean file-sharing websites. Hence, many of the IP addresses of the zombie hosts were identified as being located in South Korea. The 7.7 DDoS attacks consisted of five different attack phases. The first phase started at 02:00 on July 5th and lasted 12 h. It targeted eight US government websites including “www.whitehouse.gov.” The second phase started at 22:00 on the same day and ended at 18:00 on the next day (July 6th). The number of targets increased to 21, and included government, military, and commercial websites in the US. Although these two initial-phase attacks were very powerful, the US websites were no longer severely affected after a short while because the US government decided to block all traffic originating from South Korea.

The third phase started at 18:00 on July 7th and lasted 24 h, and it attacked 13 US and 13 South Korean websites. The targets were the South Korean government, military, and financial sectors and popular Internet portal websites, in addition to the US targets from the first and second attacks. At 18:00 on July 8th, the targets changed to 14 South Korean websites, and this fourth phase lasted 24 h. Note that the targets here included the websites of computer security companies, with the aim of hindering the update of computer virus vaccines. The third and fourth phase attacks were very successful, and most of the target websites were put out-of-service. Finally, the fifth phase started at 18:00 on July 9th and lasted 24 h attacking seven South Korean websites. However, this attack had relatively little impact on the target websites because the attack schedule was revealed, and defense mechanisms

¹We use Korea Standard Time (KST, UTC+09:00) as the default time zone in this paper.

were prepared in advance by the governments and security companies.

The 7.7 DDoS attacks have seven distinctive characteristics compared to usual DDoS attacks:

- *Uncertain attack objectives*: Typically, DDoS attacks have clear objectives (e.g., monetary or political), but the 7.7 DDoS attacks did not reveal their objectives. The initial speculation was that this was well-designed cyber warfare from North Korea, but the true objectives are still unknown.
- *Simultaneous attacks against multiple target websites*: Normally, DDoS attacks are limited to a small number of specific websites, but the 7.7 DDoS attacks targeted a vast range of websites simultaneously, including government, military, and commercial sites in the US and South Korea.
- *Autonomous attacks*: Zombie hosts are often controlled by Command and Control (C&C) servers to start or stop DDoS attacks. Accordingly, it is relatively easy to defend against such attacks once the IP addresses of the C&C servers are identified; we can simply block all communications from the C&C servers. Unlike the usual DDoS attacks, however, the zombie hosts in the 7.7 DDoS attacks had a predefined attack schedule and target website list, which meant that these zombie hosts attacked the target websites autonomously.
- *Large-scale zombie hosts*: The total number of infected zombie hosts is difficult to estimate, but the South Korean government and security companies reported that the number of zombies ranged from approximately 78,000 to 200,000. This huge number of zombie hosts made the 7.7 DDoS attacks very successful without relying on the IP spoofing technique.
- *Low-rate attacks*: Each zombie host of the 7.7 DDoS attacks generated 54.2 kbps of attack traffic. The amount of attack traffic per zombie host was very small, so it did not bother the zombie host users, and the attacks were not detected in the early stages by the DDoS defense systems.
- *Multiple attack types*: The 7.7 DDoS attacks exploited four different DDoS attack types, including *TCP Syn Flooding*, *UDP 80 Flooding*, *ICMP Flooding*, and *HTTP Get/POST Flooding*. By doing this, the 7.7 DDoS attacks exhausted both the target servers’ resources and the intermediate network bandwidth.
- *Corruption of zombie hosts*: The malware used for the 7.7 DDoS attacks was designed to erase certain document files and destroy hard disk drives by overwriting the Master Boot Record (MBR) of an infected zombie host after finishing the DDoS attacks. It was intended to leave no clue for traceback from the host log.

B. 3.3 DDoS Attacks in 2011

The 3.3 DDoS attacks occurred in of two phases. The first phase started at 17:00 on March 3rd, 2011, and 29 South Korean websites were targeted. The second phase started at 10:00 on March 4th, targeting 40 South Korean websites. The

TABLE I
ONE HOUR TRAFFIC TRACES AT 17:00.

Year	Date	DDoS Attack	Trace Size	Total Subnet IPs	Suspicious Subnet IPs
2009	03/31	No	30.7 GB	65,251	545
	07/08	Yes	27.3 GB	47,228	304
	07/09	Yes	25.9 GB	53,110	299
	08/12	No	32.0 GB	52,230	448
2011	03/04	Yes	12.8 GB	29,586	329
	03/14	No	15.7 GB	51,573	419


websites targeted in the 3.3 DDoS attacks were similar to those of the 7.7 DDoS attacks. The estimated number of zombie hosts was about 50,000, which was smaller than the 7.7 DDoS attacks. The 3.3 DDoS attacks had very similar characteristics to the 7.7 DDoS attacks in terms of the malware propagation path, uncertain objectives, simultaneous attacks, autonomous attacks, low-rate attacks, multiple attack types, and corruption of zombie hosts. However, the 3.3 DDoS attacks were not as successful as the 7.7 DDoS attacks because of the rapid counter actions by the government.

Although the 3.3 DDoS attacks were similar to the 7.7 DDoS attacks, there are several notable differences. First, the 3.3 DDoS attacks used C&C servers. In this case, however, the role of the C&C servers was different: they were used to provide additional malware program codes, and not for commanding the start or end of the attacks. Second, the malware of the 3.3 DDoS attacks modified the “hosts” file of Windows OS to prevent the update of computer virus vaccines, whereas the 7.7 DDoS attacks tried to achieve this by attacking the update servers of computer security companies. Third, the target list and the communication between the zombie hosts and C&C servers were encrypted to make analysis difficult. Finally, the termination time of the attacks was not specified, so the attacks continued until the infected zombie hosts were cured.

IV. TRAFFIC ANALYSIS


We speculated that there were zombie hosts in our campus network during the 7.7 and 3.3 DDoS attacks. We captured traffic traces at the POSTECH campus network, in order to identify the distinguishable characteristics of the DDoS attack traffic from the perspective of an attacker network. The campus network had two core switches and four routers that were internally connected using 1-Gbps links, and it was connected to the external Internet via two 1-Gbps links. We captured the traffic traces through an optical tap that was attached to one of these two external links.

The traffic traces were captured over 1 h, starting at 17:00 each day. We stored the first 96 bytes of a packet, which was enough to analyze the transport-layer and partial application-layer information. A summary of the traffic traces is shown in Table I. We captured traffic traces on March 31st (before the 7.7 DDoS attacks), July 8th and 9th (during the 7.7 DDoS attacks), and August 12th (after the 7.7 DDoS attacks) in 2009, and on March 4th (during the 3.3 DDoS attacks) and



(a) Number of Packets (b) Number of Flows (c) Traffic Volume

Fig. 1. Total traffic summary of four 2009 traces. Filled bars: DDoS-free; unfilled bars: DDoS.



(a) Number of Packets (b) Number of Flows (c) Traffic Volume

Fig. 2. Suspicious traffic summary of four 2009 traces. Filled bars: DDoS-free; unfilled bars: DDoS.

14th (after the 3.3 DDoS attacks) in 2011. Note that the 2011 traces are smaller than the 2009 traces; this is because our campus network traffic was divided into three links with the recent addition of an extra 1-Gbps link. “Total Subnet IPs” means the number of IP addresses that belonged to our subnet. “Suspicious Subnet IPs” means the number of internal IP addresses that had at least one connection with the target websites of the DDoS attacks. In this work, we regarded a host that had a connection with the target websites as a suspicious zombie.

A. Traffic Summary

Fig. 2 shows the number of packets, the number of flows, and the traffic volume of the four traffic traces in 2009. We used unidirectional flows with a time-out of 300 s (TCP, UDP, and ICMP flows only). The number of packets on 07/08 and 07/09 (unfilled bars; during the 7.7 DDoS attacks) was generally larger than on 03/31 and 08/12 (filled bars; DDoS-free). On the contrary, the number of flows and the traffic volume of the DDoS traffic traces tended to be slightly smaller than for the normal traffic traces. However, these gaps are insignificant, and it is difficult to say whether they have any statistical meaning. This implies that the zombie hosts of the 7.7 DDoS attacks tried using low-rate attacks, so they did not have a significant impact on the total traffic of our campus network.

These characteristics change when we derive the three metrics using only the suspicious traffic traces (see Fig. 3). We regarded a POSTECH host as an infected zombie if it communicated with one or more of the target websites of the DDoS attacks. Fig. 3 was derived using traffic traces coming from or going to a suspicious zombie. The numbers of packets and flows on 07/08, the day of the most severe DDoS attacks, were considerably larger than the others in Fig. 3; meanwhile, the traffic volume was lower. This implies that the suspicious zombie hosts sent a huge number of DDoS attack packets that did not carry meaningful payload data to target

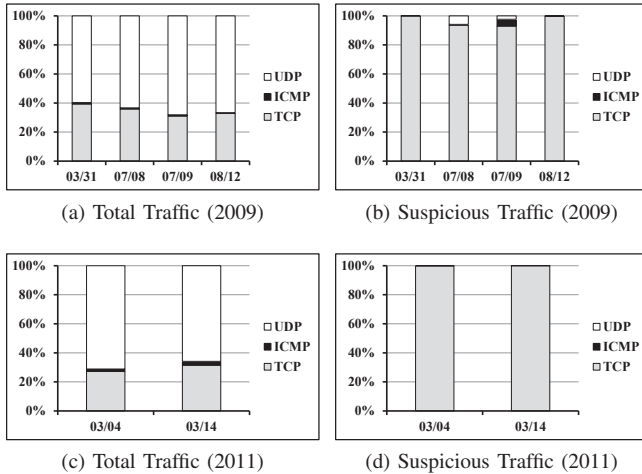


Fig. 3. Ratio of flow numbers per protocol (TCP, UDP, ICMP) at 17:00.

websites. Note that the 07/09 traffic trace does not show the same pattern as any of the other traces: it shows the lowest number of packets, lowest number of flows, and lowest traffic volume. We consider the following two explanations. First, many zombie hosts were cured at that time, so the DDoS traffic was reduced. Second, the majority of the target websites were still unavailable, so only a very small amount of normal traffic existed.

We carried out the same analysis for the two 2011 traffic traces. There was no notable difference between these traces except that the number of packets, number of flows, and traffic volume of the 03/04 traffic traces were lower than those of the 03/14 traces. This tendency arises from the different weekly pattern of the people at POSTECH; 03/04 was a Friday and 03/14 was a Monday. The impact of the 3.3 DDoS attacks on our traffic traces was insignificant, implying that there were few zombie hosts in our campus network.

B. Protocol Ratio

Fig. 4 illustrates the ratios of protocols (TCP, UDP, and ICMP) in terms of the number of flows. Fig. 4a is derived using all the traffic traces, and no significant differences are observed between the DDoS-free (03/31 and 08/12) and DDoS (07/08 and 07/09) traffic traces. UDP accounts for more than 60%, TCP accounts for 30–40%, and ICMP makes up the remainder. However, this is not true for Fig. 4b, which was derived using only the suspicious traffic traces. TCP accounts for the majority of traffic in the 03/31 and 08/12 traffic traces, whereas UDP and ICMP are also present in the 07/08 and 07/09 traffic traces together with TCP. TCP traffic was dominant because the DDoS targets were websites that used HTTP over TCP. On the contrary, there was a relatively large amount of UDP and ICMP traffic on 07/08 and 07/09 because the 7.7 DDoS attacks used these protocols to make the target websites unavailable.

Regarding the 2011 traffic traces, they show similar patterns to the DDoS-free traffic traces in 2009: UDP accounts for more than 60%, TCP for 30–40%, and ICMP for the remainder in

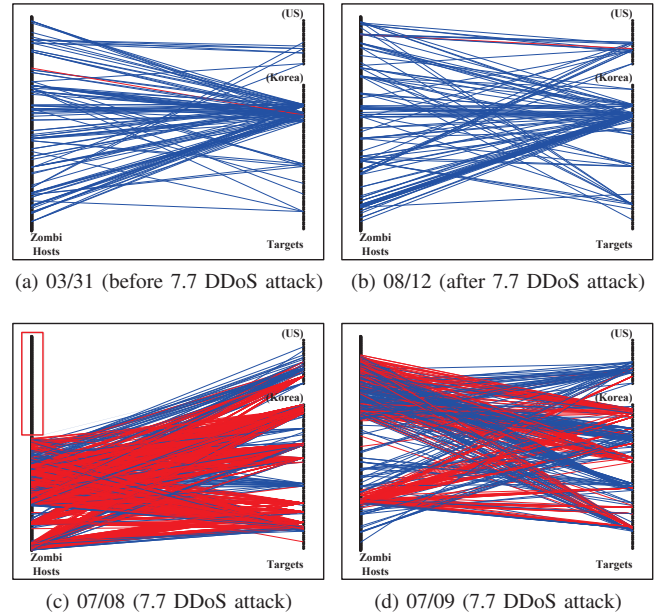


Fig. 4. IP interaction graphlets between suspicious zombie hosts and target websites in the USA and South Korea during the 7.7 DDoS attacks in 2009. One-hour traffic traces starting at 17:00 are used.

the total traffic traces (Fig. 4c). TCP accounts for almost 100% in the suspicious traffic traces (Fig. 4d).

According to these figures (Figs. 2–4), we can say that the primitive statistical information on the total traffic from the attacker’s side of the network is inappropriate for the detection of DDoS attacks that exploit low-rate and multiple attacks. We need to adopt a method that monitors the connections between intranet hosts and *potential* DDoS attack target websites or per-IP behavior [13]. On the basis of this rationale, the following analyses consider only suspicious traffic traces.

C. IP Interaction Graphlets

Figs. 5–8 represent the interaction behavior between suspicious zombie hosts and DDoS targets using social-level graphlets [14]. In this paper, the graphlet shows the interaction between the IP addresses of zombie hosts and DDoS targets. The suspicious zombie hosts are represented as nodes on the left, and the target websites are on the right. In the 7.7 DDoS attacks, the target website nodes are divided into US and South Korean ones. A line connects a zombie node and a target node if they have communicated. If the average packet size of the flow is larger than 64 bytes, the line is blue. Otherwise, the line is red.

Fig. 5 shows interaction graphlets from four traces in 2009. These four graphlets show clear differences between DDoS and DDoS-free traffic. Most of the lines are colored blue, and they are concentrated on several target websites when there are no DDoS attacks (see Figs. 5a and 5b). On the contrary, in Figs. 5c and 5d, a considerable number of connection lines are colored red, and the flows are distributed over all the target websites. Note that there are no flows originating from zombie hosts that are positioned in the upper part of the 07/08

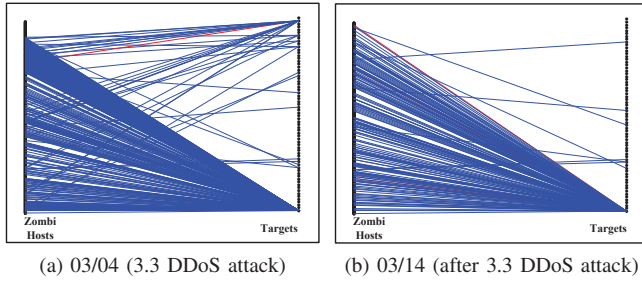


Fig. 5. IP interaction graphlets between suspicious zombie hosts and target websites in South Korea during the 3.3 DDoS attacks in 2011. One-hour traffic traces starting at 17:00 are used.




Fig. 6. IP interaction graphlets between suspicious zombie hosts and target websites in the USA and South Korea from 07:00 to 11:00 on July 9th, 2009.

graphlet (red box in Fig. 5c). Of all our traffic traces, this phenomenon is observed only in the 17:00 and 18:00 traffic traces on 07/08, 2009. We suspect that this was caused by the temporal malfunction or performance degradation of the traffic switch.

The graphlets from two traffic traces in 2011 are shown in Fig. 6. These two graphlets are very similar: the average packet sizes of most flows are larger than 64 bytes, and most connections are concentrated on several target websites. These patterns are similar to the two DDoS-free traces from 2009 (Figs. 5a and 5b). Thus, we can say that our campus network had few zombie hosts during the 3.3 DDoS attacks.

Figs. 7 and 8 show graphlets using one-hour traffic traces from 07:00 to 11:00 and from 16:00 to 20:00 on 07/09. From these graphlets, we can recognize the changing DDoS attack phase according to the malware's schedule and/or life patterns of the POSTECH faculty and students. In Fig. 7, we see that only a few connections existed between 07:00 and 08:00, but the number of connections increased after 08:00 (the start




Fig. 7. IP interaction graphlets between suspicious zombie hosts and target websites in the USA and South Korea from 16:00 to 20:00 on July 9th, 2009.

of the working day). This woke up the DDoS malware, and the zombie hosts started to attack the target websites. The same attack patterns lasted until 17:00 (Fig. 8a). Suddenly, the number of connections of the suspicious zombie hosts in the upper part increased between 17:00 and 18:00 (see Fig. 8b). This was because variants of the DDoS malware were scheduled to start another attack phase at this time. The number of connections decreased significantly after 18:00 (after working hours).

Analyzing IP interaction graphlets proved that a considerable number of zombie hosts existed in our network during the 7.7 DDoS attacks, whereas fewer existed during the 3.3 DDoS attacks. In addition, we could identify the changing DDoS attack phases according to the life patterns of the zombie host users and the malware schedule. The main features of the 7.7 DDoS attacks that are derived from the IP interaction graphlets are twofold: (1) the average packet sizes of most connections were less than 64 bytes, and (2) the connections were distributed over all the target websites. These features are more or less restricted to the 7.7 DDoS attacks, but we show that the patterns of the IP interaction graphlets during DDoS attacks are very different from the patterns with DDoS-free traffic. Consequently, the visualized IP interaction graphlet could be a very useful method for the detection of various kinds of evolving DDoS attacks from the side of the attacker network.

D. Flow Duration

Fig. 9 presents the flow duration cumulative distribution functions (CDFs) of the six traces. In the 2009 traffic traces, we can clearly identify different patterns between the DDoS (07/08 and 07/09) and DDoS-free traces (03/31 and 08/12).




Fig. 8. CDFs of flow duration.

The proportion of flows that lasted less than 0.01 s was small in the DDoS-free traces, whereas about 10% of the flows lasted less than 0.01 s during the 7.7 DDoS attacks. This implies that the zombie hosts sent a series of DDoS packets (such as *TCP Syn Flooding*) in a very short time period, and changed their target websites. In addition, the proportion of flows that lasted more than 10 s during the 7.7 DDoS attacks was higher than when there were no DDoS attacks. The zombie hosts consistently sent DDoS attack packets to the target websites for more than 10 s. These results indicate that the 7.7 DDoS attacks consisted of multiple attack types. Note that the number of flows lasting more than 10 s on 07/08 was considerably larger than on 07/09, and the 7.7 DDoS attacks were more severe on 07/08 than 07/09. The two 2011 flow-duration CDFs show very similar patterns. Accordingly, we conclude that there were very few zombie hosts in our campus network during the 3.3 DDoS attacks.

E. Average Packet Size per Suspicious Zombie Host

Fig. 10 shows the CDFs of average packet size per suspicious zombie host in the six traffic traces. We show that they are similar to those in section IV-D. Regarding the four average-packet-size CDFs of 2009, the average packet sizes of the DDoS traffic traces tends to be lower than those of the DDoS-free traffic traces. This tendency is more evident on 07/08 than on 07/09, implying that the 7.7 DDoS attacks were more severe on 07/08 than on 07/09. In the average-packet-size CDFs of 2011, almost the same pattern is seen. Again, we can say that there were very few zombie hosts in our campus network during the 3.3 DDoS attacks.

V. CONCLUDING REMARKS

DDoS attacks inflict severe loss on target servers by draining the resources of servers and/or networks. Various detection and prevention methods have been proposed, but the successful 7.7 DDoS attacks revealed the vulnerabilities of existing DDoS

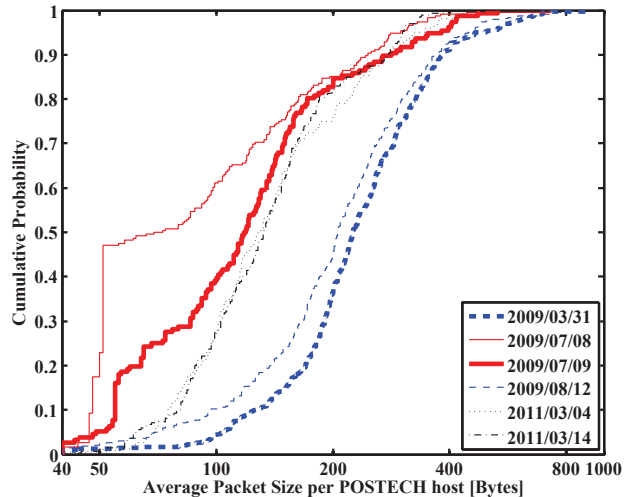


Fig. 9. CDFs of average packet size per suspicious zombie host.

countermeasures. Most existing DDoS countermeasures concentrate on handling DDoS attacks from the victim's network. In this paper, we suggest that it is worthwhile to take a step back from the victim's side, and look at the problem of DDoS from the attacker's viewpoint.

To identify the distinguishable characteristics of DDoS traffic from the perspective of an attacker network, we captured traffic traces from two real-world DDoS attacks (the 7.7 DDoS attacks in 2009 and the 3.3 DDoS attacks in 2011) from networks hosting the zombies. Through analysis of the data, we were able to make the following observations. (1) The primitive statistical information from all the traffic in the attacker side is inappropriate for the detection of DDoS attacks. Rather, we need a method that can monitor the connections between attacker network hosts and potential DDoS attack targets or per-IP behavior [13]. (2) A considerable number of zombie hosts existed in our campus network during the 7.7 DDoS attacks in 2009, whereas there were only a few zombie hosts during the 3.3 DDoS attacks in 2011. (3) The various metrics that have been analyzed in this paper could be used to detect and mitigate DDoS attacks on the attacker side of the network.

For future work, we are planning to deploy a system using the analysis metrics described in this paper on our campus network to monitor and detect DDoS attacks from the attacker side of the network, before the attack traffic reaches the targets. In addition, we aim to derive a formal method that will quantify the IP interaction graphlets to allow the automatic detection of DDoS attacks.

REFERENCES

- [1] S. S. Kim and A. L. N. Reddy, "Statistical techniques for detecting traffic anomalies through packet header data," *IEEE/ACM Transactions on Networking*, vol. 16, no. 3, pp. 562–575, Jun. 2008.
- [2] P. Du and S. Abe, "Detecting DoS attacks using packet size distribution," in *Proc. 2nd International Conference on Bio-Inspired Models of Network, Information, and Computing Systems (BIONETICS '07)*, Budapest, Hungary, Dec. 10–13, 2007, pp. 93–96.

- [3] F. Al-Haidari, M. Sqalli, K. Salah, and J. Hamodi, "An entropy-based countermeasure against intelligent DoS attacks targeting firewalls," in *Proc. 2009 IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY '09)*, Imperial College London, UK, Jul. 20–22, 2009, pp. 41–44.
- [4] R. K. C. Chang, "Defending against flooding-based Distributed Denial-of-Service attacks: A tutorial," *IEEE Communications Magazine*, vol. 40, no. 10, pp. 42–51, Oct. 2002.
- [5] X. Yang, D. Wetherall, and T. Anderson, "TVA: A DoS-limiting network architecture," *IEEE/ACM Transactions on Networking*, vol. 16, no. 6, pp. 1267–1280, Dec. 2008.
- [6] "Large-scale DDoS attacks in the United States and South Korea," White Paper, Internet Initiative Japan Inc., Nov. 16, 2009.
- [7] S. Jang, "3.3 DDoS report: In South Korea," Hauri, Tech. Rep., Mar. 5, 2011.
- [8] J. Mirković, A. Hussain, B. Wilson, S. Fahmy, P. Reiher, R. Thomas, W.-M. Yao, and S. Schwab, "Towards user-centric metrics for Denial-Of-Service measurement," in *Proc. ACM Workshop on Experimental Computer Science (ExpCS '07)*, San Diego, CA, USA, Jun. 13–14, 2007, pp. 1–14.
- [9] Z. M. Mao, V. Sekar, O. Spatscheck, J. van der Merwe, and R. Vasudevan, "Analyzing large DDoS attacks using multiple data sources," in *Proc. ACM SIGCOMM Workshop on Large-scale Attack Defense (LSAD '06)*, Pisa, Italy, Sep. 11, 2006, pp. 161–168.
- [10] Y. Xie and S.-Z. Yu, "Monitoring the application-layer DDoS attacks for popular websites," *IEEE/ACM Transactions on Networking*, vol. 17, no. 1, pp. 15–25, Feb. 2009.
- [11] J. Mirković, G. Prier, and P. Reiher, "Attacking DDoS at the source," in *Proc. IEEE 10th International Conference on Network Protocols (ICNP '02)*, Paris, France, Nov. 12–15, 2002, pp. 312–321.
- [12] S. Malliga, A. Tamilarasi, and M. Janani, "Filtering spoofed traffic at source end for defending against DoS / DDoS attacks," in *Proc. IEEE 17th International Conference on Computer, Communication and Networks (ICCCN '08)*, St. Thomas, Virgin Islands, USA, Aug. 3–7, 2008, pp. 1–5.
- [13] Y. Zhang, Q. Liu, and G. Zhao, "A real-time DDoS attack detection and prevention system based on per-IP traffic behavioral analysis," in *Proc. IEEE 3rd International Conference on Computer Science and Information Technology (ICCSIT '10)*, Rome, Italy, Apr. 28–30, 2010, pp. 163–167.
- [14] T. Karagiannis, K. Papagiannaki, and M. Faloutsos, "BLINC: Multi-level traffic classification in the dark," in *Proc. ACM SIGCOMM '05*, Philadelphia, PA, USA, Aug. 22–26, 2005, pp. 229–240.