

Measuring and Analyzing Trends in Recent Distributed Denial of Service Attacks

An Wang¹(✉), Aziz Mohaisen², Wentao Chang¹, and Songqing Chen¹

¹ George Mason University, Fairfax, USA
{awang10,wchang7,sqchen}@gmu.edu

² SUNY Buffalo, Buffalo, USA
mohaisen@buffalo.edu

Abstract. Internet DDoS attacks are prevalent but hard to defend against, partially due to the volatility of the attacking methods and patterns used by attackers. Understanding the latest of DDoS attacks can provide new insights for effective defense. But most of existing understandings are based on indirect traffic measures (e.g., backscatters) or traffic seen locally (e.g., in an ISP or from a botnet). In this study, we present an in-depth study based on 50,704 different Internet DDoS attacks directly observed in a seven-month period. These attacks were launched by 674 botnets from 23 different botnet families with a total of 9026 victim IPs belonging to 1074 organizations in 186 countries. In this study, we conduct some initial analysis mainly from the perspectives of these attacks' targets and sources. Our analysis reveals several interesting findings about today's Internet DDoS attacks. Some highlights include: (1) while 40% of the targets were attacked only once, 20% of the targets were attacked more than 100 times (2) most of the attacks are not massive in terms of number of participating nodes but they often last long, (3) most of these attacks are not widely distributed, but rather being highly regionalized. These findings add to the existing literature on the understanding of today's Internet DDoS attacks, and offer new insights for designing effective defense schemes at different levels.

Keywords: DDoS attack characteristics · Attack distribution and affinity

1 Introduction

That nature of Distributed Denial of Services (DDoS) attacks on the Internet has evolved in the last ten years due to their increasing complexity. Today's attacks are more prevalent due to the rise of botnets, large pools of infected machines that are well incentivized to pursue persistent criminal activities. Based on a recent report [1], an average DDoS attack is not detected until 4.5 h after its commencement and mitigation efforts do not start until 4.9 h after that. Furthermore, the operational impact, size, and consequences of DDoS attacks on large services on the Internet are widely reported. Recently, 3,000 open domain name

service (DNS) resolvers were capable of generating 300 Gbps DDoS traffic [11], and taking down Spamhaus, a popular spam tracking service. More recently, an amplification attack utilizing 4,529 network time protocol (NTP) servers was capable of generating a 325–400 Gbps of persistent attack traffic [31].

Efforts have been made continuously from both academia and industry to understand the DDoS attacks and defend against them. With ever-improving defenses, the attack strategies are constantly changing. Understanding the latest attack strategies is a key to successful defenses. The most recent literature on the problem is outdated, and utilizes measurements and analyses on DDoS attacks by means of inference from indirect traffic, such as backscatters, or from traffic collected locally, such as in a single Internet service provider (ISP) network or a university, or by infiltration into a botnet. While of a very high interest, a timely and large scale view of today’s Internet DDoS attacks is missing from the literature.

We present a timely measurement study of recent DDoS attacks launched by botnets. Our measurement is based on directly observed attack artifacts through anchor points deployed at a large number of major ISPs in a seven-month period. The attack workloads are collected by the monitoring and attribution unit in a commercial DDoS mitigation company located in the United States with global operational footprint. In this seven-month period, a total of 50,704 different DDoS attacks were observed, which were launched by 674 different botnets coming from 23 different botnet families. These attacks targeted 9026 different IPs that belong to 1074 organizations in 186 countries.

Our detailed analyses reveal several interesting observations about today’s Internet botnet DDoS attacks. While details are provided in the paper, some highlights include: (1) while 40% of the targets were attacked only once, 20% of the targets were attacked more than 100 times. This clearly highlights the inefficiency of defenses deployed by targets; (2) most of the attacks are not massive in terms of number of participating nodes but they often last long. This attacking strategy makes attacks stealthier and more difficult to detect; (3) most of these attacks are not widely distributed, but rather being highly regionalized. This may motivate some more effective DDoS defense development.

While there have been various studies on this topic [5, 36, 37] that admittedly compete with our study in the size of the utilized data, and the nature of the findings, we believe that our study is distinguished from the prior studies in two aspects. First, our study revisits the topic with a timely dataset obtained from operational mitigation and defense efforts. To our knowledge, the most recent operational look at the problem is based on a dataset that is at least five years older than ours [36], and many are based on datasets that are even more than ten years old [5, 37]. Second, as the trends of attacks evolve over time, the new dataset and findings obtained upon analyzing it offer new and unique insights that can be utilized for designing effective and customized defenses.

The rest of the paper is organized as follows. In Sect. 2, we describe our dataset including the overall data statistics and the data fields we utilized to do our analysis. In Sect. 3, we present some basic characterizations of DDoS attacks.

We discuss related work in Sect. 4 and conclude with a concise summary of our analyses and their implications in Sect. 5.

2 Data Collection

Our dataset is provided by the monitoring and attribution unit in a very large DDoS detection and mitigation company that is located in the United States, with partnerships of monitoring with a large number of ISPs for the sole purpose of attack detection and mitigation. The unit constantly monitors attack traffic to aid the mitigation efforts of its clients, using both active and passive measurement techniques [6, 34, 35]. For active measurements and attribution, malware families used in launching the various attacks are reverse engineered, and labeled to a known malware family using best practices [21]. A honeypot is then created to emulate the operation of the reverse-engineered malware sample and to enumerate all bots across the globe participating in the particular botnet. As each botnet evolves over time, new generations are marked by their unique hashes. The enumerated list of bots is then vetted on the participants in the active attacks.

Traces of traffic associated with various DDoS campaigns are then collected at various anchor points located at the aforementioned ISPs across the globe: North and South America, Asia, Europe, and Africa. The traces are then analyzed remotely to attribute and characterize attacks on various targets of interest. The collection of traffic is guided by two general principles: (1) that the source of the traffic is an infected host participating in a DDoS campaign, and (2) the destination of the traffic is a targeted client, as concluded from eavesdropping on C&C of the campaign using a live sample, or where the end-host is a customer of the said DDoS mitigation company.

2.1 High-Level Characteristics

The analysis is high level in nature to cope with the high volume of ingest traffic at peak attack times—on average there were 243 simultaneous verified DDoS attacks launched by the different botnets studied in this work. High level statistics associated with the various botnets and DDoS attacks are recorded every one hour. The workload we obtained ranges from August 28, 2012 to March 24, 2013, a total of 209 days (about seven months of valid and marked attack logs). In the log, a DDoS attack is labeled with a unique DDoS identifier, corresponding to an attack by given DDoS malware family on a given target. Other attributes and statistics of the dataset are shown in Table 1. We cannot reveal the capability of the capturing facility because attackers would learn such information, which is also critical to the business of the data source.

An interesting feature in Table 1 is the attack category, which refers to the nature of the DDoS attack by classifying it into different types based on the protocol utilized for launching it, including HTTP, TCP, UDP, Undetermined,

Table 1. Information of workload entries

| Field | Description |
|-----------|--|
| ddos_id | A global unique identifier for the specified DDoS attack |
| botnet_id | Unique identification of each botnet |
| category | Description of the nature of the attack |
| target_ip | IP address of the victim host |
| timestamp | The time when the attack started |
| end_time | The time when the attack ended |
| botnet_ip | The IP address of botnets involved in the attacks |
| asn | Autonomous system number |
| cc | Country in which the target resides (ISO3166-1 alpha-2) |
| city | City and/or state in which the target resides |
| latitude | Latitude of target |
| longitude | Longitude of target |

ICMP, Unknown, and SYN. Different from *Unknown*, *Undetermine* means that the attack type could not be determined based on the available information.

Among the fields listed in Table 1, the *longitude* and *latitude* of each IP address are obtained using a highly-accurate industrial geo-mapping service during trace collection. The mapping of the IP addresses happens in real time, making it resistive to IP dynamics. Beside the longitude and latitude, we also generate the individual *city* and *organization* of each IP address involved in an attack using a highly-accurate commercial grade geo-mapping dataset by Digital Envoy (Digital Element services [2]). We use such information for geographical analysis as presented in later sections.

Table 2 sums up some statistics of our dataset, including information from both the attacker and the target sides. Target statistics are illuminating. Over a

Table 2. Summary of the workload information

| Summary of attackers | | Summary of victims | |
|----------------------|--------|--------------------|-------|
| Description | Count | Description | Count |
| # of bot_ips | 310950 | # of target_ip | 9026 |
| # of cities | 2897 | # of cities | 616 |
| # of countries | 186 | # of countries | 84 |
| # of organizations | 3498 | # of organizations | 1074 |
| # of asn | 3973 | # of asn | 1260 |
| # of ddos_id | 50704 | | |
| # of botnet_id | 674 | | |
| # of traffic types | 7 | | |

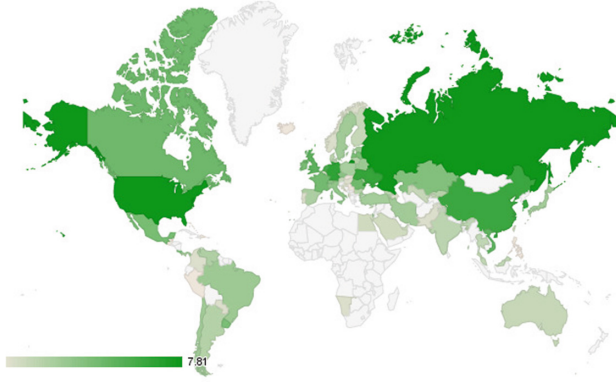


Fig. 1. Geographical distribution (The density is the log of the number of victims in each country).

period of 28 weeks, 50,704 different DDoS attacks were observed. These attacks are launched by 674 different botnets. These attacks targeted victims located in 84 different countries, 616 cities, involving 1074 organizations, residing in 1260 different ASes.

Based on the geographical information, Fig. 1 shows the popularity of DDoS targets at the country level. The color density indicates the number of attacks on each country. Clearly, and as widely believed, the US and Russia are the two most popular targets.

2.2 Discussions of the Dataset

Various related works are focused on radiation and port scanning [4, 5, 19, 24, 36] concerned with a single network (e.g., Tier-1 ISP [19] and sinkhole traffic [4, 36]), rather than DDoS attacks characterizations in a similar context to ours, thus preventing us from making a fair comparison with our data size and methodology for data collection. However, some of the recent studies show the relevance of our dataset in size. In comparison with [19], our study characterizes more than 50,000 verified attacks over 7 months observation period (compared to 31,612 alarms over a 4 weeks period in [19]). Note the fundamental difference between attacks and alarms, since a large number of triggered alarms in anomaly detection could be false alarms, while attacks are verified.

Nonetheless, the limitation of our dataset is that it does not cover all ISPs on the Internet, suggesting various forms of bias and incompleteness. We note that, however, our data collection also incorporates at-destination data collection, thus all statistics of interest and relevance are gathered in the process. We note that our data collection method is not subject to the shortcoming of locality bias highlighted in [5]: all malware families used for launching attacks that we study are well-understood and reversed engineered, and traffic sources utilized for launching the attacks are enumerated by active participation. To that end,

we believe that our data collection is representative to the characterized events, and that the length of the observation period is sufficient to draw conclusions on today’s DDoS attacks.

3 Attack Analysis

In this section, we present our analysis results across various DDoS attacks observed in our dataset. Through the analysis, we aim to understand the new trends of these attacks. We focus our analysis on the attack targets and sources in this study. We notice that not all of the 23 botnets logged in our dataset are active, and only 10 of them exhibit patterns and trends; we focus our study on those 10 active botnets. Namely, we study the DDoS attacks by the following families: Aldibot, BlackEnergy, Colddeath, Darkshell, DDoSer, DirtJumper, Nitol, Optima, Pandora, and YZF.

3.1 Botnet and Target Affinity

To take down a victim’s site (target), DDoS attacks could be launched continuously. To avoid being detected, some attacks could be split into multiple stages, and individual staged attacks could be launched periodically. Therefore, we first study how many attacks a victim received in our dataset. Along this line, we can identify those long-term targets and short-term targets for some DDoS malware families.

Figure 2 shows the popularity of targets distributions for all active families representing in Cumulative Distribution Function (CDF). There are 9026 unique targets in total. This figure shows that while 40% of the targets were attacked only once, 20% of the targets are attacked more than 106 times. The most

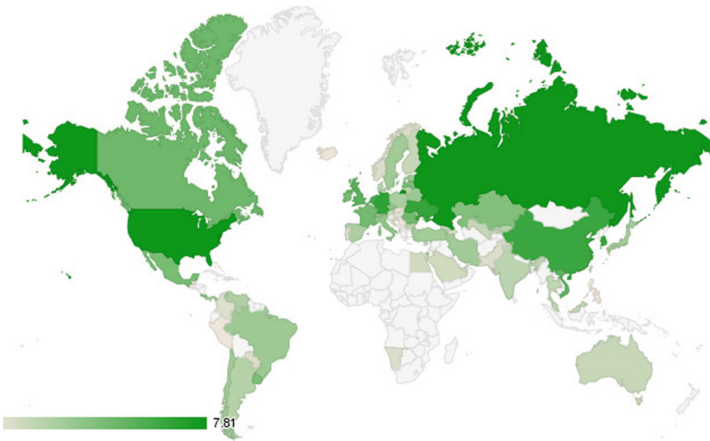


Fig. 2. Target popularity: the number of attacks a victim received over the period of 28 weeks.

popular target was attacked 940 times over the same period. After looking up the most popular target IP address, we found that this IP address belongs to the domain of HostGator, which is a Houston-based web hosting service, indicating that the real target could be an organization hosted by this service.

Table 3 summarizes some statistics of these targets by each botnet family. The second row shows the number of targets that were attacked multiple times by each family, the third row shows the number of unique targets attacked by each family, and the last row shows the percentage of repeatedly attacked targets by each family. This table shows that YZF and Ddosер often focus their attacks on some selected targets while the targets of Aldibot and Blackenergy are more distributed. This may indicate different attack patterns that *Aldibot* and *Blackenergy* focus on short-term targets while *YZF* and *Ddosер* may focus on long-term targets. Note that the short and long term patterns highlighted here also outline the different economical aspects of the attacks: while long-term persistent attacks highlight higher incentives of an adversary whereas short-term attacks may highlight a “hitman-like” strategy for service abruption.

Table 3. Number of targets attacked multiple times by each botnet family

| Family | Aldibot | Blackenergy | Colddeath | Darkshell | Ddosер | Dirtjumper | Nitol | Optima | Pandora | YZF |
|------------------------------------|---------|-------------|-----------|-----------|--------|------------|-------|--------|---------|------|
| #targets w/ multiple attacks | 12 | 105 | 127 | 486 | 89 | 3587 | 99 | 53 | 737 | 72 |
| #unique targets | 48 | 355 | 265 | 1029 | 131 | 5823 | 213 | 131 | 1775 | 72 |
| % | 25% | 30% | 48% | 47% | 68% | 62% | 46% | 40% | 42% | 100% |

Table 4. Target interest duration (h)

| Family | Aldibot | Blackenergy | Colddeath | Darkshell | Ddosер | Dirtjumper | Nitol | Optima | Pandora | YZF |
|------------|---------|-------------|-----------|-----------|--------|------------|-------|--------|---------|-----|
| Period (h) | 144 | 303 | 63 | 82 | 12.9 | 406 | 70.6 | 295 | 362 | 22 |

Table 4 further shows the average target interest duration of each family. From previous observations, we know that some of the targets will be attacked multiple times periodically. But this period may not last very long since we know that most DDoS attacks are money-driven. The table indicates that the average target interest also varies significantly across different families, ranging from half a day to half a month. The longest period we found is 200 days, of an attack by Dirtjumper (note that Table 4 only shows the average). It almost spans the whole observation period. The target is a Russian web hosting service company, which is known for hosting various malicious websites, indicating that the consistent attack on it is perhaps a form of retaliation to take those sites down.

Insights and Takeaways: The attack target analysis reveals that while about 40% of the targets were attacked only once, more than half of the targets were attacked more than once. This clearly indicates (1) such DDoS attacks are most likely driven by profit since most retaliatory and political attacks are temporary, and (2) the current defense mechanisms in practice fail to respond to such attacks promptly for efficient protections. On the other hand, each botnet family always has some long-term targets, which hints us to develop specific botnet based defenses. In the long run, the effective and practical defense may be to detect and defend various botnets based on their characteristics, rather than detect and defend DDoS since there is no clear pattern along time (last subsection) or on target selection.

3.2 Attack Size and Distribution

Beside the attack distribution, another factor for assessing attacks’ strength is the attack magnitude, which refers to how much DDoS traffic was seen towards a target. However, we did not log the raw packets carrying DDoS traffic—for contractual reasons. Rather, we use the number of unique IP addresses used for launching the attack as an estimator of its magnitude. As suggested in [19], to evade being detected quickly, IP spoofing in DDoS attacks is not very common. Thus, we use the number of unique IPs involved in an attack to estimate one aspect of the corresponding attack magnitude.




Fig. 3. Attack magnitude of each family

3.2.1 Attack Size – Contrary to the Recent Reports, Most Attacks Are Not Massive

Figure 3 shows the magnitude distributions of the DDoS attack launched by each family. Note that the x -axis is in log scale. Several interesting observations

can be found in this figure. First, most of the attacks are not massive in terms of magnitude. As shown in the figure, except some attacks by *Dirtjumper* and *Blackenergy*, attacks launched by other botnets typically involve less than 1000 unique IPs. Some botnets, such as *Aldibot* and *YZF*, launched a lot of small magnitude DDoS attacks and there are less than 40 unique IPs involved. However, massive DDoS attacks could involve several thousand or tens of thousands unique IPs. For example, the figure shows that the number of unique IPs involved in attacks of *Dirtjumper* and *Blackenergy* could be massive, with the maximum number of unique IPs used in a single *Blackenergy* attack being 2,365 and 37,584 for *Dirtjumper*, respectively.

Taking into consideration the magnitude and the duration of attacks, we aim to infer strategies utilized by botnets. We aggregate the 80% non-massive (small) attacks to calculate the average attack duration and compare it with that of the other 20% attacks—results are shown in Table 5. This table only includes the families that actually launched massive DDoS attacks. We found that 80% of the attacks lasted less than 13882s in our analysis. But from this table, we can see that, except for family *Dirtjumper*, the 80% non-massive DDoS attacks last longer than 13882s. This shows that most small attacks are not short-duration attacks.

For those families that didn't launch massive attacks, this is still true: attacks launched by *YZF* involved at most 6 bots; while the average duration of the attacks is 34053s. The above analysis indicates that most current DDoS attacks utilize a small number of bots with longer duration instead of launching massive DDoS attacks.

Table 5. Attack duration of small and massive attacks

| Family | Small attacks | Massive attacks |
|-------------|---------------|-----------------|
| Blackenergy | 5945 | 7430 |
| Colddeath | 16702 | 38831 |
| Dirtjumper | 10222 | 33568 |
| Optima | 31713 | 27410 |
| Pandora | 19121 | 46428 |

Figure 3 also indicates that the magnitude of attacks varies across different families, as well as within the same family. As shown in the figure, the attack magnitude of *Aldibot* and *Nitol* only slightly changed in 28 weeks, while *Blackenergy* changed attack magnitude dramatically for different attacks. Notice that on the left corner, there are two families, which are *Nitol* and *Ddoser*. The magnitude of their attacks is constantly small.

3.2.2 Attack Distribution – Most DDoS Attacks Are Not Widely Distributed, but Rather Being Regionalized

The attacker’s IP enables us to study their distribution as well. According to our data, it is obvious that each botnet family has its own geolocation preferences. Among the four most active families, *Optima* is mostly in Vietnam, India, Indonesia and Egypt; *Dirtjumper* is mostly in India, Brazil, Botswana and Pakistan; *Blackenergy* is mostly in Vietnam, Singapore, Thailand and United States; *Pandora* is mostly in Pakistan, India, Thailand and Indonesia. Generally, the geolocation preference could reflect the current global distribution of botnets. Most of them reside in the countries that lack proper protections like Southeast Asia area [30].

Among all the families, *Dirtjumper* covers the largest number of countries: 164. A comparable coverage is *Optima*’s: 153. Even though these families have very broad country coverages, the average number of country coverage for each attack is small, as shown in Table 6. From this table, we see that only *Pandora* has a very large average number; other families only launched attacks from a few countries. Figure 4 shows the CDF of country coverage of each single DDoS attack for each family. This figure conforms to the result in Table 6.

Table 6. Average number of country coverage for each botnet family

| Family | Aldibot | Blackenergy | Colddeath | Darkshell | Ddoser | Dirtjumper | Nitol | Optima | Pandora | YZF |
|------------------|---------|-------------|-----------|-----------|--------|------------|-------|--------|---------|-----|
| Country coverage | 4 | 3 | 2 | 1 | 1 | 3 | 1 | 9 | 2 | 1 |




Fig. 4. CDF of country coverage for each family

Insights and Takeaways: The attack size and distribution analysis indicates the change of the attack strategies (resources exhaustion vs. bandwidth attrition),

making it challenging to detect such attacks. On the other hand, the highly regionalized nature of most attacks can motivate some new defense strategies: disinfection, if infected hosts are in reach, or early stage filtering at the nearby ISPs connecting those hosts to the Internet [26].

4 Related Work

DDoS attacks have been intensively investigated and numerous countermeasures have been proposed to defend against them. As many DDoS attacks are launched by botnets, a popular approach is to disrupt the C&C channel of the botnet that launches DDoS attacks. However, most of current take-down methodologies are ad-hoc and their effectiveness are limited by the depth of knowledge about the specific botnet family involved in the attack. Nadji et al. [23] proposed a take-down analysis and recommendation system. As a proactive solution to DDoS attacks, several filtering schemes [12, 17, 25, 38, 39] are proposed to prevent flooding packets from reaching target victims. Chen et al. [7] proposed a defense that can detect DDoS attacks over multiple network domains. Overlay-based protection systems [15, 28] and statistical approaches [9, 13, 16, 18] offer another attractive alternative. Walfish et al. [33] advocated DDoS defense by offense, where attacked servers encourages all clients to send higher traffic volumes to attackers.

Historically, most defense systems such as Cisco IDSM-2 [8] and LADS [27] are deployed at the destination since it suffers most of the impact. Mirkovic et al. [20] proposed D-WARD, a DDoS defense system deployed at source-end networks that autonomously detects and stops attacks originating from these networks. Another detection mechanism [3] is proposed to be deployed at the source to help ISP network to detect attacks. Both studies can benefit from our source characterization to enable their operation. Huang et al. [10] addressed the lack of motivations for organizations to adopt cooperative solutions to defeat DDoS attacks by fixing the incentive chain. Our analysis of the shared fate, and diversity of targets, is a plausible incentive for enabling collaborative defense.

The continuous improvement on detection and defense has caused attackers to be adaptive as well. Thus, it is essential to understand the latest changes of DDoS attacks. In 2006, Mao et al. [19] presented their findings from measurement study of DDoS attacks relying on both direct measurements of flow-level information and more traditional indirect measurements using backscatter analysis, while Moore et al. [22] presented a backscatter analysis for quantitatively estimating DoS activity in the Internet based on a three-week dataset. Due to the growth of network address translation and firewall techniques, much of the Internet was precluded from the study by the traditional network measurement techniques. Thus, in the early days, work [5] proposed an opportunistic measurement approach that leverages sources of spurious traffic, such as worms and DDoS backscatter, to unveil unseen portion of Internet. The monitoring of packets destined for unused Internet addresses, termed as "background radiation", proved to be another useful technique to measure Internet phenomenon.

In 2004, Pang et al. [24] conducted an initial study of broad characteristics of Internet background radiation by measuring traffic from four large unused subnets. In 2010, a more recent study [36] revisited the same topic and characterized the current state of background radiation specifically highlighting those which exhibit significant differences. Our work serves as a revisit to those studies with new insights. Bailey et al. [4] designed and implemented the Internet Motion Sensors (IMS), a globally scoped Internet monitoring system to detect Internet threats, which includes a distributed blackhole network with a lightweight responder and a novel payload signature and caching mechanism. Xu et al. [37] presented a general methodology to build behavior profiles of Internet backbone traffic in terms of communication patterns of end-hosts and services.

In [32], a honey farm system architecture was proposed to improve honeypot scalability by up to six orders of magnitude while still offering quantitatively similar fidelity. Another Internet monitoring system, which primarily targets at early detection of worms, was presented in [40], using a non-threshold based “trend detection” methodology to detect presence of worms by using Kalman filter and worm propagation models.

Finally, theoretical attacks using botnets for “taking down the Internet” are studied in Crossfire [14], CXPST [26], and Coremelt [29]. The size, distribution, and coordination of attacks studied in this work highlight the feasibility of those theoretical attacks.

5 Conclusion

DDoS attacks are the most popular large scale attacks frequently launched on the Internet. While most of the existing studies have mainly focused on designing various defense schemes, the measurement and analysis of large scale Internet DDoS attacks are not very common, mainly due to the data un-availability, although understanding the DDoS attacks patterns is the key to defend against them. In this study, with the access to the dataset on such a large scale, we are able to collectively characterize today’s Internet DDoS attacks from different perspectives. Our investigation of these DDoS attacks reveals several interesting findings about today’s botnet based DDoS attacks. These results provide new insights for understanding and defending modern DDoS attack at different levels. While this study focuses on the DDoS attack characterization, in the future, we plan to leverage these findings to design more effective defense schemes. A direction is to combine the geolocation affinity information based on the botnet behavior pattern.

Acknowledgement. We would like to thank anonymous reviewers for their comments. This work was supported in part by an ARO grant W911NF-15-1-0262, NSF grant CNS-1524462, and the Global Research Lab. (GRL) Program of the National Research Foundation (NRF) funded by Ministry of Science, ICT (Information and Communication Technologies) and Future Planning (NRF-2016K1A1A2912757).

References

1. A ddos attack could cost \$1 million before mitigation even starts, October 2013. <http://bit.ly/MUXadv>
2. NetAcuity and NetAcuity Edge IP Location Technology, February 2014. <http://www.digitalelement.com/>
3. Akella, A., Bharambe, A., Reiter, M., Seshan, S.: Detecting DDoS Attacks on ISP Networks. In: ACM SIGMOD/PODS MPDS (2003)
4. Bailey, M., Cooke, E., Jahanian, F., Nazario, J., Watson, D., et al.: The internet motion sensor-a distributed blackhole monitoring system. In: Proceeding of NDSS (2005)
5. Casado, M., Garfinkel, T., Cui, W., Paxson, V., Savage, S.: Opportunistic measurement: extracting insight from spurious traffic. In: Proceeding of ACM Hotnets (2005)
6. Chang, W., Mohaisen, A., Wang, A., Chen, S.: Measuring botnets in the wild: some new trends. In: Proceeding of ACM ASIA CCS (2015)
7. Chen, Y., Hwang, K., Ku, W.S.: Collaborative detection of DDoS attacks over multiple network domains. IEEE TPDS (2007)
8. Cisco: Cisco Catalyst 6500 Series Intrusion Detection System, February 2014. <http://bit.ly/1hsppy9>
9. Feinstein, L., Schnackenberg, D., Balupari, R., Kindred, D.: Statistical approaches to DDoS attack detection and response. In: DARPA Information Survivability Conference and Exposition (2003)
10. Huang, Y., Geng, X., Whinston, A.B.: Defeating DDoS attacks by fixing the incentive chain. ACM ToIT 1 (2007)
11. Info Security Magazine: Spamhaus suffers largest ddos attack in history - entire internet affected, March 2013. <http://bit.ly/1bf3ZHZ>
12. Ioannidis, J., Bellovin, S.M.: Implementing pushback: router-based defense against DDoS attacks. In: Proceeding of NDSS (2002). <https://www.cs.columbia.edu/~smb/papers/pushback-impl.pdf>
13. Jin, S., Yeung, D.: A covariance analysis model for DDoS attack detection. IEEE ICC (2004)
14. Kang, M.S., Lee, S.B., Gligor, V.D.: The crossfire attack. In: Proceeding of IEEE S&P (2013)
15. Keromytis, A.D., Misra, A.D., Rubenstein, D.: SOS: an architecture for mitigating DDoS attacks. IEEE JSAC (2004)
16. Lee, K., Kim, J., Kwon, K.H., Han, Y., Kim, S.: DDoS attack detection method using cluster analysis. Expert systems with applications (2008)
17. Li, J., Mirkovic, J., Wang, M., Reiher, P., Zhang, L.: Save: source address validity enforcement protocol. In: Proceeding of IEEE ICC (2002)
18. Li, M.: Change trend of averaged Hurst parameter of traffic under DDOS flood attacks. Computers and Security (2006)
19. Mao, Z.M., Sekar, V., Spatscheck, O., van der Merwe, J., Vasudevan, R.: Analyzing large DDoS attacks using multiple data sources. In: Proceeding of ACM SIGCOMM LSAD (2006)
20. Mirkovic, J., Prier, G., Reiher, P.: Attacking DDoS at the source. In: Proceeding of IEEE ICNP, November 2002
21. Mohaisen, A., Alrawi, O., Larson, M., McPherson, D.: Towards a methodical evaluation of antivirus scans and labels. In: Information Security Applications (2014)

22. Moore, D., Shannon, C., Brown, D.J., Voelker, G.M., Savage, S.: Inferring internet denial-of-service activity. *ACM TOCS* **24**(2), 115–139 (2006)
23. Nadji, Y., Antonakakis, M., Perdisci, R., Dagon, D., Lee, W.: Beheading hydras: performing effective botnet takedowns. In: *Proceeding of ACM SIGSAC*, November 2013
24. Pang, R., Yegneswaran, V., Barford, P., Paxson, V., Peterson, L.: Characteristics of internet background radiation. In: *Proceeding of ACM IMC* (2004)
25. Park, K., Lee, H.: On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law Internets. In: *Proceeding of ACM SIGCOMM* (2001)
26. Schuchard, M., Mohaisen, A., Kune, D.F., Hopper, N., Kim, Y., Vasserman, E.Y.: Losing control of the internet: using the data plane to attack the control plane. In: *Proceeding of NDSS* (2011)
27. Sekar, V., Duffield, N., Spatscheck, O., van der Merwe, J., Zhang, H.: Lads: large-scale automated DDoS detection system. In: *Proceeding of USENIX ATC* (2006)
28. Stavrou, A., Keromytis, A.D.: Countering DoS attacks with stateless multipath overlays. In: *Proceeding of ACM CCS* (2005)
29. Studer, A., Perrig, A.: The coremelt attack. In: Backes, M., Ning, P. (eds.) *ESORICS 2009*. LNCS, vol. 5789, pp. 37–52. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-04444-1_3](https://doi.org/10.1007/978-3-642-04444-1_3)
30. Thomas, N.: Cyber security in East Asia: governing anarchy. *Asian Secur.* **5**(1), 3–23 (2009)
31. Vaughan-Nichols, S.J.: Worst DDoS attack of all time hits french site, February 2014. <http://zd.net/1kFDurZ>
32. Vrabie, M., Ma, J., Chen, J., Moore, D., Vandekieft, E., Snoeren, A.C., Voelker, G.M., Savage, S.: Scalability, fidelity, and containment in the potemkin virtual honeyfarm. *ACM SIGOPS* **5**, 148–162 (2005)
33. Walfish, M., Vutukuru, M., Balakrishnan, H., Karger, D., Shenke, S.: DDoS defense by offense. In: *Proceeding of SIGCOMM* (2006)
34. Wang, A., Mohaisen, A., Chang, W., Chen, S.: Capturing DDoS attack dynamics behind the scenes. In: Almgren, M., Gulisano, V., Maggi, F. (eds.) *DIMVA 2015*. LNCS, vol. 9148, pp. 205–215. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-20550-2_11](https://doi.org/10.1007/978-3-319-20550-2_11)
35. Wang, A., Mohaisen, A., Chang, W., Chen, S.: Delving into internet DDoS attacks by botnets: characterization and analysis. In: *Proceeding of IEEE DSN* (2015)
36. Wustrow, E., Karir, M., Bailey, M., Jahanian, F., Huston, G.: Internet background radiation revisited. In: *Proceeding of ACM IMC* (2010)
37. Xu, K., Zhang, Z.L., Bhattacharyya, S.: Profiling internet backbone traffic: behavior models and applications. In: *ACM SIGCOMM CCR*. No. 4 (2005)
38. Yaar, A., Perrig, A., Song, D.: SIFF: a stateless internet flow filter to mitigate DDoS flooding attacks. In: *Proceeding of IEEE S&P* (2004)
39. Yaar, A., Perrig, A., Song, D.: StackPi: new packet marking and filtering mechanisms for DDoS and IP spoofing defense. *IEEE JSAC* (2006)
40. Zou, C.C., Gong, W., Towsley, D., Gao, L.: The monitoring and early detection of internet worms. *IEEE/ACM TON* **5**, 961–974 (2005)

Information Security Applications

17th International Workshop, WISA 2016, Jeju Island,
Korea, August 25-27, 2016, Revised Selected Papers

Choi, D.; Guilley, S. (Eds.)

2017, XI, 398 p. 157 illus., Softcover

ISBN: 978-3-319-56548-4