

Delving into Internet DDoS Attacks by Botnets: Characterization and Analysis

An Wang
George Mason University

Aziz Mohaisen
Verisign Labs

Wentao Chang
George Mason University

Songqing Chen
George Mason University

Abstract—Internet Distributed Denial of Service (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 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. In this study, we present an in-depth analysis 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 9,026 victim IPs belonging to 1,074 organizations in 186 countries. Our analysis reveals several interesting findings about today’s Internet DDoS attacks. Some highlights include: (1) geolocation analysis shows that the geospatial distribution of the attacking sources follows certain patterns, which enables very accurate source prediction of future attacks for most active botnet families; (2) from the target perspective, multiple attacks to the same target also exhibit strong patterns of inter-attack time interval, allowing accurate start time prediction of the next anticipated attacks from certain botnet families; (3) there is a trend for different botnets to launch DDoS attacks targeting the same victim, simultaneously or in turn. These findings add to the existing literature on the understanding of today’s Internet DDoS attacks, and offer new insights for designing new defense schemes at different levels.

I. INTRODUCTION

Today, Internet Distributed Denial of Services (DDoS) attacks are prevalent with the ease of access to large numbers of infected machines, collectively called botnets. According to a recent report [1], the duration, intensity, and diversity of attacks are on the rise: a year-over-year analysis shows that the average DDoS attack size has increased by 245% in the fourth quarter of 2014, compared to the same quarter of 2013, and by 14% from the previous quarter of the same year, with an average attack of 7.39 Gbps. Furthermore, the same report shows that all industry verticals are targeted by attacks. Another report reveals a clear increase in the average duration of DDoS attacks from 60 minutes in the first quarter of 2014 to 72 minutes in second quarter of the same year, which translates to 20% increase [2]. Additionally, recent DDoS attacks have witnessed an uptrend in operational impact, size, and consequences [3], [4], with the largest reported attacks exceeding 500 Gbps [5]. Today’s malicious actors are not limited to sophisticated machines, like servers and personal computers; recent DDoS attacks were reportedly utilizing fridges [6], and other massive scanning activities were done using embedded devices, including monitoring cameras and security doors [7].

Security researchers in academia and industry devoted enormous efforts to understanding DDoS attacks and defending against them. As defenses are deployed, attacks evolved and became more sophisticated to circumvent those defenses. Understanding the current trends in today’s DDoS attacks

and their attack vectors is an important phase in devising effective defenses. Existing studies in this regard are based on indirect traffic analyses and artifacts, such as backscatters, or traffic collected locally, or by infiltrating into a botnet. A large scale view of today’s Internet DDoS attacks is missing in the literature and calls for further investigation.

In this paper, we present our study of DDoS attacks analysis. As most of the DDoS attacks nowadays are launched by botnets, the dataset utilized in this study focuses on DDoS attacks launched by various botnet families across the Internet. In a seven-month period captured in our dataset, 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 9,026 different IPs that belong to 1,074 organizations in 186 countries.

Our detailed analyses revealed several interesting observations about today’s Internet botnet DDoS attacks. While details are provided in the paper, some highlights include:

- Geolocation analysis shows that the geospatial distribution of the attacking sources follows certain patterns, which enables very accurate source prediction of future attacks for most active botnet families.
- From the target perspective, multiple attacks to the same target also exhibit strong patterns of inter-attack time interval, allowing accurate start time prediction of the next anticipated attacks from certain botnet families.
- There is a trend for different botnets to launch DDoS attacks targeting the same victim, simultaneously or in turn.

These findings offer new insights for designing effective and/or customized defense schemes at different levels.

Organization. In Section II, we describe our dataset including the overall data statistics and the data fields we utilized to do our analysis. In Section III, we present an overview of these DDoS attacks. In Section IV, we analyze the geolocation affinity of attacking sources and their targets. In Section V, we present in depth collaboration analyses between different botnets in a family or across families. We discuss related work in Section VI and conclude with a concise summary of our analyses and their implications in Section VII.

II. DATASET COLLECTION AND METHODOLOGY

Dataset. Our dataset is provided by third-party through monitoring Internet infrastructures, using both active and passive measurement techniques. For active measurements and attribution, malware families used in launching various attacks are

TABLE I. 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

reverse engineered, and labeled to a known malware family using best practices. Hosts participating in the given botnet, by communicating with pieces of infrastructure infected by that malware family (e.g. the command and control) are then enumerated and monitored over time, and their activities are logged and analyzed.

Collection methodology. As each botnet evolves over time, new generations are marked by their unique (MD5 and SHA-1) hashes. Traces of traffic associated with various botnets are collected at various points on the Internet, in cooperation with various ISPs. Traffic logs are then analyzed to attribute and characterize attacks. The collection and analysis are guided by two general principles: 1) the source of the traffic is an infected host participating in a botnet attack, and 2) that the destination of the traffic is a targeted client, as concluded from eavesdropping on command and control of the campaign using a live malware samples.

By tracking temporal activities of 23 different known botnet families, the dataset captures a snapshot of each family every hour from 08/29/2012 to 03/24/2013, a total of 207 days, or about seven months. There are 24 hourly reports per day for each botnet family. The set of bots or controllers listed in each report are cumulative over the past 24 hours. The 24-hour time span is measured using the timestamp of the last known bot activity and the time of logged snapshot.

The analysis is high level in nature to cope with the high volume of ingest traffic at peak attack times. As shown later, on average, there was 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 29, 2012 to March 24, 2013, a total of 207 days (about seven months of valid and marked attack logs). In the log, each 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 I.

Features and statistics. An interesting feature in Table I is the attack category, which refers to the nature of the DDoS attacks by classifying them into various types based on the protocol utilized for launching them; HTTP, TCP, UDP, Undetermined, ICMP, Unknown, and SYN. Different from *Unknown*, *Undetermined* means that the attack type could not be determined based on the available information.

The *longitude* and *latitude* of each IP address in Table I are

TABLE II. 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		

obtained using a highly-accurate geo-mapping service during the trace collection. The mapping of the IP addresses is a real-time process, 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 [8]). We use such information for geographical analysis as presented later.

Table II sums up some statistics of our dataset, including information from both the attacker and the target sides. Target statistics are illuminating. Over a period of 28 weeks, 50,704 different DDoS attacks were observed. These attacks were launched by 674 different botnets. These attacks targeted victims located in 84 different countries, 616 cities, involving 1,074 organizations, and residing in 1,260 different autonomous systems (ASes).

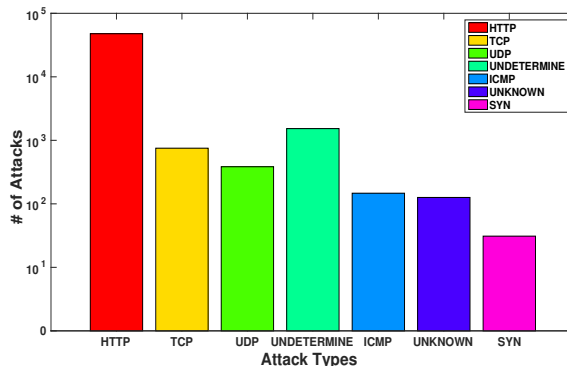


Fig. 1. Popularity of attack types

Attack mechanisms. Based on the traffic type information, Figure 1 shows the statistic of different protocols. Clearly, the dominant protocol used in these attacks is HTTP, followed by UDP and TCP. Table III further shows the breakdown of transport types used by different botnet families. The last column in the table shows the number of attacks belonging to each type. Note that a botnet could utilize multiple attack types. For example, Blackenergy supports different transport mechanisms of attack traffic, including HTTP, TCP, UDP, ICMP and SYN. The variety of transport mechanisms explains the family's popularity. Furthermore, the large share of HTTP as means of transport in this family highlights the preferred target of attacks, namely web services.

Comparison and limitations. A large body of the related work [9], [10], [11], [12], [13] are on radiation and port scanning measurements. However, most of that work is concerned

TABLE III. PROTOCOL PREFERENCES OF EACH BOTNET FAMILY

Protocol	botnet family	# of attacks
HTTP	colddeath	826
	darkshell	999
	dirtjumper	34620
	blackenergy	3048
	nitol	591
	optima	567
	pandora	6906
yzf	177	
TCP	blackenergy	199
	nitol	345
	yzf	182
	aldibot	26
UDP	blackenergy	71
	ddoser	126
	yzf	187
UNDETERMINED	darkshell	1530
ICMP	blackenergy	147
UNKNOWN	optima	126
SYN	blackenergy	31

with a single network (Tier-1 ISP [11], sinkhole traffic [10], [13]). On the other hand, our work is on DDoS attack characterization at a larger scale. This fundamental difference between our work and the prior literature makes it difficult to make direct comparisons between our work and the prior literature.

Towards the limitations of our data collection, one may argue that not covering all ISPs on the Internet for data collection may bias our data, and thus our findings. We note that; however, our data collection also incorporates at-destination data collection, thus all statistics of interest are gathered in the process. For the data size, and in comparison to [11], our study characterizes more than 50,000 attacks over seven months observation period (compared to 31,612 alarms over a period of four weeks in the prior work). Note, the fundamental difference between attacks and alarms is that a large number of triggered alarms in anomaly detection systems could be false alarms, while attacks are verified alarms

Note that our data collection method is not subject to the shortcoming of locality bias highlighted in [12]: all malware families used for launching attacks that we study are well-understood at the time of the data collection and reversed engineered, and traffic sources utilized for launching the attacks are enumerated by active measurement. 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 some conclusions on DDoS attacks on the Internet today.

III. OVERVIEW OF DDoS ATTACKS

In this section, we present an overview of DDoS attacks logged in our dataset. We recognize that not all of the 23 botnets logged in our dataset are active all the time. Among them, 10 families are more active than others – a complete analysis of all 23 botnet families can be found in [14]. To this end, in this section we focus on analyzing and characterizing attacks launched by those 10 active families. Namely, we study the DDoS attacks launched by Aldibot, BlackEnergy, Colddeath, Darkshell, DDoSer, DirtJumper, Nitol, Optima, Pandora, and YZF.

A. Attack Distribution

In the 28 weeks covered in our dataset, we observed over 50,000 DDoS attacks launched using bots that belong to the 10 active botnets. First we study the attack distribution along time. To do this, we extract the beginning time of each attack and plotted the aggregate number of attacks over the period of 28 weeks in Figure 2. In this figure, the y -axis represents the number of different DDoS attacks, and the x -axis represents the time (date). We find that on average there are 243 DDoS attacks launched by the 10 botnet families every day. The maximum number of simultaneous DDoS attacks per day was 983 attacks, which happened on August 30, 2012. All of these attacks were launched by Dirtjumper and the targets were located in the same subnet in Russia, suggesting a strong relationship between the different attacks.

Although we observe fluctuations in the number of attacks over time, we did not find any obvious daily, weekly, or monthly patterns in Figure 2 that are common in other Internet activities (e.g., diurnal patterns in web access). This is, however, anticipated since DDoS attacks typically are not user-driven, thus lack recurring patterns.




Fig. 3. Attack interval: all attacks and family-based attacks

Simultaneous Attacks. We further extract the intervals between DDoS attacks. We define the intervals between two DDoS attacks similar to that of the inter-arrival time: the time interval between any two consecutive attacks launched by the same botnet family (or on the same target; across multiple families). Figure 3 shows the CDF of the attack intervals across all attacks and attacks launched by each family. Note that x -axis is in log scale.

Attack intervals observed from all attacks and family-based attacks show consistent patterns. Clearly, more than half of the attacks are launched simultaneously, which is less likely to be a coincidence—we will investigate that later. For family based attacks, we found that the longest attack interval was 59 days, almost two months. Also, 80% of the attack intervals lasted less than 1081 seconds, which is roughly 18 minutes. The average DDoS attack interval was 3060 seconds and the standard deviation was 39140 seconds. Those numbers, and by observing the CDF in Figure 3, tell that the attack intervals follow two extremes: except for 15% of the attack falling in the [1,000, 10,000] seconds interval, the majority of the attacks (about 50%) are concurrent, with less than 1% of the attacks at least one order of magnitude larger than the rest of attack intervals.




Fig. 2. Daily Attack Distribution

As these concurrent attacks are very interesting, we take a closer look at them. We find that they can be classified into two categories: attacks launched by a single botnet family and attacks launched by multiple families. Attacks in the first category happened 3692 times and attacks in the second category happened 956 times.

For the first category, we found that seven out of the 10 botnet families exhibit such behavior. Among all families, Dirtjumper is the most active in launching simultaneous attacks; 10% of the attacks launched by Dirtjumper are simultaneous. For the second category, we found that most common combinations were Dirtjumper with Blackenergy and Dirtjumper with Pandora, which happened 391 and 338 times respectively. This finding is very interesting, and further investigation is dedicated to understand it in §V.

From a family’s perspective, Figure 4 further shows the intervals of all attacks by *Dirtjumper* in the order of their occurrence; the x -axis represents the attack number and the y -axis represents the corresponding interval in seconds. From this figure, we observe that the attack intervals are random. While *Dirtjumper* is used as an example, other families exhibit the same pattern of random interval distribution.




Fig. 4. Dirtjumper attack intervals

Figure 5 shows the attack interval CDF for each family, where the x -axis represents the attack intervals in seconds and each color represents a single family. Note that the x -axis is in log scale (base 2) to highlight the trend and pattern in the intervals for the various families. From this figure we observe that Blackenergy, Aldibot and Optima launch 40%-50% of attacks simultaneously or within a short time frame. We also observe that both Aldibot and Optima have no attacks with intervals that are less than 60 seconds. This could be a strategy utilized to evade detections. Finally, from the same




Fig. 5. DDoS attack intervals for each botnet family

figure, we observe that the activeness of botnets differ by an order of magnitude, with Nitel and Aldibot being the least active ones.

B. Attack Duration

The duration of an attack is one aspect that measures its strength and longevity. In our dataset, the measurement of duration is in a way aggregate and doesn’t differentiate between providers and their capability. Figure 6 depicts the durations of all DDoS attacks, where the x -axis represents the attacks along time while the y -axis represents the attack duration in seconds. Simultaneous attacks are ordered based on IP addresses. As shown, the attack duration varies significantly: while the average duration is 10,308 seconds, the median is only 1,766 seconds, with a standard deviation of 18,475 seconds (which indicates wide-spread).

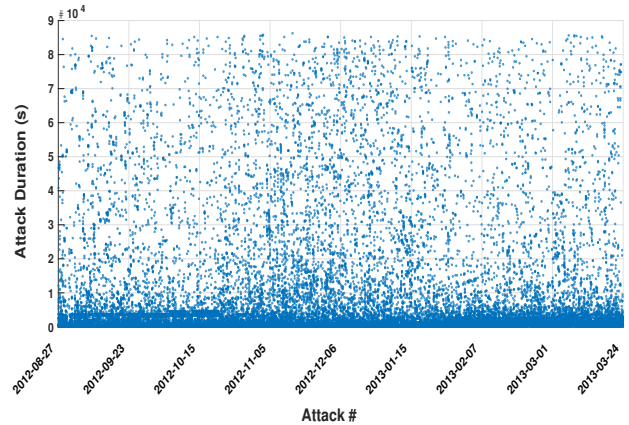


Fig. 6. Attack Duration

Figure 7 further shows the corresponding CDF of the attack duration. As shown, 80% of the attacks last for less than 13,882 seconds (\sim four hours). Choosing four hours as the cut-off for the majority of attacks duration is perhaps not arbitrary. This value indicates that four hours might be a reasonable duration for DDoS attacks to be detected and mitigated. An adaptive attacker using such a strategy would evade detection for the longest possible time for most attacks. That is, the longer the attack lasts, the higher its chances are of being detected. By limiting attack to four hours, the attacker can successfully reduce the detection rate, and thus can repetitively launch more attacks later without risking being blacklisted. Compared with the literature [11], where it was shown that 80% of attacks in a comparable study last for less than 1.25 hours, this finding is interesting in itself: DDoS attacks are becoming more persistent by lasting longer; however, their duration is still smaller than the required time frame for detections.




Fig. 7. Duration CDF

Summary. 80% of the attacks have a duration less than four hours, where targets are constantly attacked. This is more likely to be a strategy, rather than the effectiveness of defenses. This further demands *automatic* detection and defense instead of any *semi-automatic* or *manual* approaches. Only the former can effectively respond in such a short time frame. Without such an automatic system in place, the detection is not possible for one-time attack targets. For targets that are repetitively attacked, investigation of the attack intervals may be helpful.

IV. ANALYSIS AND PREDICTION OF DDoS TARGET AND SOURCE

Having analyzed the attacks distribution and duration, we now shift our attention to the geolocation of these attacks from the target and source perspectives, respectively. 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 log. Along this line, we can identify those long-term targets and short-term targets for some DDoS malware families.

A. Source Analysis

Geolocation affinity is a direct indicator of how an attacker is geo-spatially distributed. To further quantify the geolocation affinity, we extract all the bots involved in DDoS attacks for each family and aggregate the number of these bots per

week. Thus, we are able to observe the attack source and their migrations over weeks. We define such changes as a shift pattern. Figure 8 shows the dynamic per week as a shift pattern of *Optima*. The grey bars represent the shift among the same group of countries while the green bars represent the shift to new countries. From this figure we can see clearly that most of the attack sources will be limited to the same group of countries (other families not shown have the same pattern, confirming that most of these attacks are highly regionalized). Next, we explore how the geolocations of different bots participating in attacks change over time.




Fig. 8. Botnet shift pattern of Optima

In our dataset, each DDoS attack could be illustrated by a series of snapshots along time. In each snapshot, as discussed in §II, IP addresses of all bots evolved at the given time were recorded. Since every IP address corresponds to a single location (longitude and latitude pair), we are able to identify the locations of all the bots involved on a map. We use such information to characterize source locations. First, we find the geological center point of the various locations of IP addresses at any time. Then, we calculate the distance between each bot and this center point (using Haversine formula), and add the distances together. In our analysis, the distance has a sign to indicate direction: positive indicates east or north, and negative indicates west and south. For simplicity, we consider the absolute value of the sum of all distances; a sum of zero means that participating bots are geographically symmetric. We use these distances to represent the geolocation distribution of the bots. We calculate this value across all the families and plot the CDF of geolocation distributions in Figure 9.

In this figure, six families with at least 10 snapshots (with active attacks for more than 10 days) are reported. From Figure 9, we observe that not all the families follow the same distribution of location proximity. For the families *Optima* and *Blackenergy*, the distances exhibit a normal distribution, whereas other families have a skewed distribution. The families *Dirtjumper* and *Pandora* both have more than 40% distribution distances of zero, indicating complete geographical symmetry. Later, we will show that *Dirtjumper* and *Pandora* collaborate with each other closely, which may explain the similar distribution of their geolocation distances. Furthermore, the different distribution patterns suggest that geolocation distribution is less likely to be random, but rather part of the attack and infection strategy, which could be further confirmed later.

To further explore the dynamics behind the geolocation




Fig. 9. CDF of geolocation distribution

changes of each DDoS attack, we arrange all the geolocation distribution values of all the DDoS attacks launched by each family in time order. Then, we plot the geolocation distances along time. Figure 10 and Figure 11 show the result for *Pandora* and *Blackenergy*, respectively.




Fig. 10. *Pandora* geolocation distances




Fig. 11. *Blackenergy* geolocation distances

In these two figures, the x -axis represents the time when each snapshot was taken, and the y -axis represents the geolocation distance value. From the above figures, we see that periodic patterns exist in both cases, and the distance values appear in stationary states, meaning that the values vary around

a certain mean value. This indicates that these values are predictable or even stable.

To verify our conjecture, we next build a prediction model over this data. To build the model, we use the Autoregressive Integrated Moving Average (ARIMA) model, which is one of the popular linear models in time series forecasting. The popularity of the ARIMA model is because of its statistical properties in the model building process. In addition, ARIMA models are quite flexible in that they can present several different types of time series [15].

To evaluate the results of our prediction model, we split our data into two parts, the first half is for training and the other half is used for prediction and evaluation. For the prediction part, we use the last 2,700 values (2,700 is a randomly picked number. This value shouldn't affect our prediction results). Again, due to the space limit, we only present the results for the same two families *Pandora* and *Blackenergy*. The prediction results are shown in Figure 12 and Figure 13.




Fig. 12. *Pandora* geolocation distance prediction




Fig. 13. *Blackenergy* geolocation distance prediction

In these figures, the x -axis represents the predicted points while the y -axis represents the geolocation distance value. The predicted results are shown by the dotted red curve and the ground truth values are marked by the blue lines. From these figures, we can clearly observe that the predicted results are almost identical with the ground truth value. We further calculate the numerical statistics for all the families except for *Darkshell* since there are not enough data points for training the model. The results are listed in Table IV.

TABLE IV. STATISTICS FOR GEOLOCATION DISTANCE PREDICTION

Family	Group	Mean	Standard Deviation	Cosine Similarity
<i>Blackenergy</i>	prediction	3968.41327074	1955.53446352	0.96021535
	ground truth	3970.63145063	2294.35821194	
<i>Pandora</i>	prediction	562.621714849	1809.21755865	0.94645314
	ground truth	569.208453348	1842.49554775	
<i>Dirtjumper</i>	prediction	1203.90188001	925.844685329	0.8476428
	ground truth	1229.08641657	1033.65534376	
<i>Optima</i>	prediction	3526.62483465	1150.51297736	0.94071316
	ground truth	3545.75239893	1717.76781969	
<i>Colddeath</i>	prediction	356.474286966	753.241975174	0.80935886
	ground truth	341.60923214	933.823379971	

We compare two groups of artifacts in this table: the prediction and the ground truth values. We calculated the mean value and the standard deviation value of both groups. Further, we compared these two groups by calculating their cosine similarity with each other. From this table, we can see that for all the families, both the mean value and the standard deviation are close to those of the ground truth, except for family *Dirtjumper* and *Colddeath*; the predicted results represent more than 90% similarity to the ground truth.

Insight into defenses: These results reveal several insights including: (1) The geolocation dynamics of bots involved in DDoS attacks exhibit certain patterns for different botnet families. (2) Attack source geolocation changes can be accurately predicted by using a proper model. (3) Such information combined with changes of the attack volumes can be used for forecasting how DDoS attacks evolve over time, thus allowing one to deploy or adjust defenses accordingly.

B. Target Analysis

Country-level analysis. Now, we turn our attention to the country-level preference of families and their victims. The third column in Table V shows the top five popular targeted countries of each active family. Most families have a specific preference over specific areas or organizations. The top five most popular target countries are the United States of America (USA), targeted by 13, 738 attacks, Russia, targeted by 11, 451 attacks, Germany, targeted by 5, 048 attacks, Ukraine, targeted by 4, 078 attacks, and the Netherlands, targeted by 2, 816 attacks. The Aldibot and Dirtjumper families' preferred target country is the USA; Colddeath's is India; the Optima, Pandor and YZF families' is Russia; the Darkshell and Nitol families' is China and Ddoser's is Mexico.

Organization-level analysis. Similar to country-level analysis, we have also conducted organization-level analysis. Our results show that the targets were narrowly distributed within several organizations. Figure 14 shows the organization-level analysis in February 2013 for Pandora. In this figure, the size of the markers on the map represents the number of attacks toward a specific target. From this figure, we can easily identify some hotspots in Russia and the USA. Among all the families, Dirtjumper has a wider presence by attacking more organizations than any other family. Also, we found that most attacks were aimed towards web hosting services, large-scale cloud providers and data centers, Internet domain registers and backbone autonomous systems, where massive

TABLE V. COUNTRY-LEVEL DDoS target STATISTICS

Family	Countries	Top 5	Count
Aldibot	14	USA	32
		France	11
		Spain	8
		Venezuela	8
		Germany	4
Blackenergy	20	Netherlands	949
		USA	820
		Singapore	729
		Russian	262
		Germany	219
Colddeath	16	India	801
		Pakistan	345
		Botswana	125
		Thailand	117
		Indonesia	112
Darkshell	13	China	1880
		South Korea	1004
		USA	694
		Hong Kong	385
		Japan	86
Ddoser	19	Mexico	452
		Venezuela	191
		Uruguay	83
		Chile	66
		USA	48
Dirtjumper	71	USA	9674
		Russian	8391
		Germany	3750
		Ukraine	3412
		Netherlands	1626
Nitol	12	China	778
		USA	176
		Canada	15
		United Kingdom	10
		Netherlands	6
Optima	12	Russian	171
		Germany	155
		USA	123
		Ukraine	9
		Kyrgyzstan	7
Pandora	43	Russian	2115
		Germany	155
		USA	123
		Ukraine	9
		Kyrgyzstan	7
YZF	11	Russian	120
		Ukraine	105
		USA	65
		Germany	39
		Netherlands	19

network resources are possessed and play a critical function in the operations of other Internet services.

Attack interval. Besides the geolocation, we also conducted an analysis on the attack intervals of each target of each family. Similar to the analysis of the target geolocation change, we sorted the attacks with respect to their time and calculate

Targets in: 2013-02-01

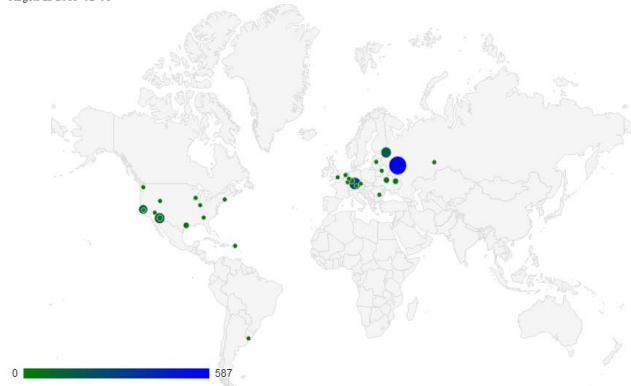


Fig. 14. Pandora target preference (organization-level)




Fig. 15. Blackenergy attack intervals

the attack intervals between consecutive attacks towards the same target. By doing that, we obtain a series of attack intervals for each target; this information is also time-related and might be utilized to characterize attack behaviors and patterns. Figure 15 displays two examples of targets by the family family *Blackenergy*.

In Figure 15, the x -axis represents the attack interval along the time and the y -axis represents the interval value in seconds. The figure shows some repeated patterns of peaks and dips of their attack interval series. Besides the periodic pattern, they also present stationary state concerning the mean value of the attack interval values.

Prediction. This characterization alludes to the possibility of predicting those series by modeling using an ARIMA model to forecast the next attack interval value, thus the start time of the next attack. To verify this possibility, we construct the model as described earlier. Figure 16 and Figure 17 show the prediction results. In both cases, we split the data into two equal halves, one for the training pool and the other for prediction and evaluation.

In both figures, the x -axis represents the time and the y -axis represents the predicted values and ground truth values. The ground truth values are marked by the dotted curves while the predicted values are marked by the solid blue lines. From those figures, it is clear that the predicted values match the ground truth consistently.




Fig. 16. *Blackenergy* attacking interval prediction for target 103.1.x.x




Fig. 17. *Blackenergy* attacking interval prediction for target 198.154.x.x

We calculated various statistics for the ground truth and predicted series for two instances as shown in Table VI. We note that the model is capable of predicting the original series with more than 90% accuracy, which confirms our initial conjecture, and highlights the potential of predicting attack intervals accurately in many cases. In this case, accuracy is denoted by cosine similarities between the prediction results and the ground truth. By looking further into the nature of the two instances of attacks captured by the series, we unveil several interesting findings. Most importantly, both instances are common and a recurring target for the families *Blackenergy* and *Dirtjumper*. However, those patterns are not common among or shared with other families, indicating that they are sufficient of identifying those families in particular as a result of the prediction.

Insight into defenses: The country and organization level target analyses provide insights for defenses. For example, findings concerning the country-level characterization can set some guidelines on country-level prioritization of disinfection and botnet takedowns. Organization-level characterization and findings associated with that can hint on the possible role provisioning can play in maximizing protection capabilities.

Understanding the attack interval pattern guides preparation for the attacks beforehand by allocating needed resources. This guidance is even more educated when the evolution of attack is predicted as shown in this section. Prediction with a high accuracy facilitates cost-effective provisioning of resources and

TABLE VI. STATISTICS FOR ATTACKING INTERVAL PREDICTION

Target	Group	Mean	Standard Deviation	Cosine Similarity
103.1.x.x	prediction	3579.07596949	1435.59818583	0.92620508
	ground truth	3534.82825485	1901.6606734	
198.154.x.x	prediction	4019.27512744	1473.30582867	0.93394388
	ground truth	4040.81449893	2187.42513432	

minimizes damages caused by DDoS attacks.

V. ANALYSIS OF COLLABORATIVE ATTACKS

So far, DDoS attacks were analyzed individually. Based on the target analysis discussed earlier, we found that different botnets (in the same family corresponding to different generations, or from different families) may collaborate to attack the same target. They may launch attacks at the same time or alternate their attacks in a way that indicates collaboration. In the following, we elaborate on this collaboration.

Table VII shows the collaboration results using both intra-family and cross-family collaborations. Basically, if different botnets are targeting the same target, and their starting time is simultaneous (or within a 60 second timeframe from each other), and their duration difference is within half an hour, then they are regarded as collaborations. As shown in this table, 121 of the detected collaborations are between different families. Among these collaborations, we observe that two families, namely Dirtjumper and Darkshell, have the most intra-family collaborations. Next, we look into these intra-family collaborations (between different botnet IDs of the same family) and inter-family collaborations in details.




Fig. 18. Intra-family collaborations of Dirtjumper

A. Concurrent Attacks

Figure 18 shows the collaboration attack magnitude by the family Dirtjumper. For clarity with respect to the multiple variables, we plot a three dimensional (3D) figure characterizing Dirtjumper: the *x*-axis represents each unique botnet ID, the *y*-axis represents the date of collaboration, and the *z*-axis represents the attack volume. From this figure, we can see that for most collaborations, there are two botnets involved, where the average number of botnets involved in the collaboration is 2.19. Such collaborations may be due to a guided action by botmasters, or as instrumented by bots themselves (e.g., multiple entities behind various attacks coincided to utilize the same resources to attack the same target at random).

Looking into Figure 18, we also find that for most bars along the same timestamp, they have the same height. Such an observation reduces the likelihood of involvement of the previously mentioned entities in these collaborations. That is, for all the botnets involved in the collaboration, detailed instructions were perhaps given for the attack magnitude. While that being a random coincidence is possible, it is not plausible, and that further highlights the potential of close collaborations between different botnets.

In addition to the collaborative attacks launched by botnets from the same family, we found that there are attacks launched by botnets from different botnet families. From Table VII, we can see that all families involved in inter-family collaborations had collaborated with Dirtjumper. Among these collaborations, Dirtjumper and Pandora collaborated with each other the most. Our next analysis will focus on those two families.

The collaborations between Dirtjumper and Pandora involved 96 unique targets, which were located in 16 countries, 58 organizations and 61 ASes. Among the 16 countries, the most popular three countries were Russia, the USA and Germany; with 31, 26 and 14 attacks per country, respectively. On the other hand, for Pandora, the average duration of an attack was 6,420 seconds (107 minutes), while the duration was 5,083 seconds (87.7 minutes) per attack for Dirtjumper.




Fig. 19. Inter-family collaborations between Dirtjumper and Pandora

Figure 19 shows the duration and attack magnitude of collaborations between Dirtjumper and Pandora as they change over time. Note that the left *y*-axis represents the attack duration while the right *y*-axis represents the attack magnitude. Both of the *y*-axes are in log scale. The histogram shows the attack magnitude and the curve shows the attack durations. From this figure, we observe that the attack magnitude for these two families are almost equal for most of the attacks, and the duration of these two families are almost identical. Another observation we make is that the attack magnitudes are not very high for both families except for an outlier. Finally, we observe that the time span of collaboration lasted from October 2012 until January 2013, covering nearly 16 weeks.

TABLE VII. BOTNETS COLLABORATION STATISTICS

Collaboration Type	Blackenergy	Colddeath	Darkshell	Ddoser	Dirtjumper	Nitol	Optima	Pandora	YZF
Intra-Family	0	0	253	134	756	17	1	10	66
Inter-Family	1	1	0	0	121	0	1	118	0

This long-term collaboration between Dirtjumper and Pandora highlights a close tie between the two families.

B. Multistage Attacks

Thus far, we consider the collaboration as multiple individual DDoS attacks are launched at the same time. Besides this kind of collaboration, another form of collaboration could be multiple DDoS attacks happening continuously one after another. Next, we investigated this type of collaboration among botnets. For this purpose, we extract the DDoS attacks on a given target that happen consecutively (i.e., the second attack happens at the end of the first attack, or within 60 second margin over overlap). For this type of attack, the results show that only intra-family collaborations were involved. Furthermore, we found that four families had this type of collaboration; Darkshell, Ddoser, Dirtjumper and Nitol.

Among all the families and collaborations, Ddoser has the longest consecutive DDoS attack involving 22 continuous attacks that lasted for more than 18 minutes on August 30, 2012. On average, the mean interval between two consecutive attacks was 0.11 seconds (a median of three seconds) with a standard deviation of 23 seconds (bursty period)




Fig. 20. Consecutive attack interval

Figure 20 displays the CDF of the intervals between two consecutive attacks. The figure shows that nearly 80% of the consecutive attacks happened within 30 seconds. In practice, this anticipated, and highlights the potential intelligence behind those coordinated attacks: a longer interval would potentially allow targets to deploy various defense mechanisms, and are not likely to be logged in our dataset.

Figure 21 shows the attack magnitude of *all* consecutive DDoS attacks. In this figure, the x -axis represents the 28 week timespan of our dataset, and the y -axis represents all the targets attacked by these consecutive DDoS attacks. Each dot represents a single DDoS attack. In this figure, the dots displayed consecutively in a row indicate that the attacks were happened consecutively. Finally, the size of each marker represents the attack magnitude of each DDoS attack and




Fig. 21. Consecutive attack magnitude

the different colors represent different families. We observe that the attack magnitudes of different collaborating families are relatively stable during the consecutive attacks, except for Dirtjumper that has several attacks of a very large magnitude.

Summary. Intra- and inter-family collaborations could be due to an underlying ecosystem, the evolution of a botnet family, or the evolution of defense mechanisms, which all make defending against them daunting tasks. Devising defenses that employ this insight for attack attribution with an in-depth understanding of the participating hosts in each family is imperative.

VI. RELATED WORK

Research on prevention and mitigation of DDoS attacks remains one of the hottest topics in the security community. DDoS attacks have been intensively investigated and numerous counter measures have been proposed to defend against them. Huang et al. [16] addressed the issue of a lack of motivation for organizations to adopt existing cooperative solutions to defeat DDoS attacks by fixing the incentive chain. As many DDoS attacks are launched by botnets, another popular approach to defend against attacks is to disrupt the C&C channel of the botnet that launches the DDoS attack. However, most current take-down methodologies are often ad-hoc and their effectiveness are limited by the depth of knowledge about the specific botnet family involved in the attack. A comprehensive measurement and analysis of behaviors of different botnet families are provided in [14]. To look closer to the botnet take-down problem, Nadji et al. [17] proposed a take-down analysis and recommendation system called rza, which not only allows a postmortem analysis of past take-downs but also provides recommendations for future take-down actions. As a proactive solution to DDoS attacks, several filtering schemes [18], [19], [20], [21], [22], which must execute on IP routers, have been proposed to prevent flooding packets from reaching target victims. Chen et al. [23] proposed a new defense system that can detect DDoS attacks over multiple network domains. Overlay-based protection systems such as Secure Overlay Services [24] offer another attractive alternative, as it requires

no changes to existing network routing infrastructure and minimal collaboration from ISPs. In their follow-up work [25] Stavrou and Keromytis proposed a novel, multiple-path overlay network that adopts a spread-spectrum-like communication paradigm to address the limitations in existing overlay-based approaches. Statistical approaches [26], [27], [28], [29] are also applied to perform anomaly detection of DDoS attacks. Work [26] identified DDoS attacks by computing entropy and frequency-sorted distributions of selected packet attributes of live traffic. Lee et al. [29] proposed an efficient method for proactive detection of DDoS attacks using cluster analysis in which cubic clustering criterion (CCC) is used on selected variables for clustering. Another research work [30] advocated DDoS defense by offense. Authors designed and implemented an application-level defense named Speak-Up, in which victimized servers encourage all clients to automatically send higher volumes of traffic to attackers. Defense mechanisms can be classified into two broad types depending on their deployment location: either on the destination of attacks (victim) or on the source of attacks. Historically, most defense systems such as Cisco IDSM-2 [31] and Large-scale Automated DDoS detection System [32] are deployed at the destination since it suffers most of the impact. Mirkovic et al. [33] proposed D-WARD, a DDoS defense system deployed at source-end networks that autonomously detects and stops attacks originating from these networks.

There have been several works on understanding unique characteristic of DDoS attacks, such as their types, durations and patterns. Wood and Stankovic extracted distinct features of DDoS attacks that are unique to sensor networks [34], while Geng et al. [35] focused on unique aspects of DDoS attacks to ad hoc networks. Mirkovic et al. [36] and Specht and Lee [37] proposed taxonomies of DDoS attacks and defenses. Work [38] highlighted important features of each attack and defenses and outlined pros and cons of the various defenses. In another study, Peng et al. [39] presented a comprehensive survey of the causes of Denial of Service (DoS) attacks and the state-of-art mechanisms for detecting and mitigating those attacks.

To understand the nature of DDoS attacks and develop effective counter measures, measurement studies are conducted on the DDoS data collected from network traffic of real attacks. Mao et al. [11] presented findings from a measurement study of DDoS attacks relying on both direct measurements of flow-level information and more traditional indirect measurements using backscatter analysis. Compared to their work, the data source in our study is purely direct measurement of Internet traffic. Moore et al. [40] presented a backscatter analysis for quantitatively estimating DoS attack activity on the Internet. They applied their approach to a three-week long dataset to study DoS attacks. Widespread DoS attacks were observed in their study. Similar to our work, they discussed the attack size, length and other characteristics followed by a victim classification including their geographical distribution. Their study was performed in 2006. Several new trends in Internet DDoS attacks are revealed in our work over their findings. In a very recent work, Rossow [41] revisited other UDP-based network protocols and identified protocols that are susceptible to amplification attacks. 14 protocols of various services including network services such as Network Time Protocol, Simple Network Management Protocol, legacy services, p2p file sharing network and so on were shown to be

vulnerable and can be abused by distributed reflective denial-of-service (DRDoS) attacks.

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, work [12] proposed an opportunistic measurement approach that leverages sources of spurious traffics such as worms, DDoS backscatter, etc. to unveil unseen portions of the Internet. The monitoring of packets destined for unused Internet addresses, termed as "background radiation," proved to be another useful technique to measure the Internet phenomenon. In 2004, Pang et al. [9] conducted an initial study of broad characteristics of Internet background radiation by measuring traffics from four large unused subnets. Both filtering techniques and active transponders are used to perform passive analysis and activity analysis of internet traffic. In 2010, a more recent study [10] revisited the same topic and characterized the current state of background radiation specifically highlighting those that exhibit significant differences. Some new trends are exposed including rapid growth outpacing the growth in productive network traffic, trends toward increasing SYN and decreasing SYN-ACK traffic, etc. In another research work, Bailey et al. [13] designed and implemented the Internet Motion Sensor (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. Another Internet monitoring system, which primarily targets early detection of worms, was presented in [42]. It used a non-threshold based "trend detection" methodology to detect presence of worms by using Kalman filter and worm propagation models.

Some other studies focus on traffic characterization on IP backbones in an attempt to build a general profile in terms of behaviors. Xu et al. [43] presented a general methodology to build behavior profiles of Internet backbone traffic in terms of communication patterns of end-hosts and services. Their work used data mining techniques to automatically discover patterns from link-level traffic data and provided plausible interpretations of those patterns.

VII. CONCLUSION

DDoS attacks are 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, although understanding DDoS attacks patterns is the key to defending against them. In this study, with the access to a large scale dataset, we were able to collectively characterize today's Internet DDoS attacks from different perspectives. Our in-depth 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 against modern DDoS attacks at different levels (e.g., organization and country). While this study focuses on DDoS characterization, in the future, we plan to leverage these findings to design more effective defense schemes.

VIII. ACKNOWLEDGEMENT

This work is partially supported by National Science Foundation (NSF) under grant CNS-1117300. The views and

opinions expressed in this paper are the views of the authors, and do not necessarily represent the policy or position of VeriSign, Inc.

REFERENCES

- [1] —, “Verisign distributed denial of service trends report,” http://www.verisigninc.com/en_US/cyber-security/ddos-protection/ddos-report/index.xhtml, February 2015.
- [2] —, <http://news.softpedia.com/news/Volumetric-DDoS-Attacks-Decrease-in-Q2-2014-Compared-to-Q1-451160.shtml>, July 2014.
- [3] Info Security Magazine, “Spamhaus suffers largest DDoS attack in history – entire internet affected,” March 2013. [Online]. Available: <http://www.infosecurity-magazine.com/news/spamhaus-suffers-largest-ddos-attack-in-history/>
- [4] S. J. Vaughan-Nichols, “Worst DDoS attack of all time hits french site,” February 2014. [Online]. Available: <http://www.zdnet.com/article/worst-ddos-attack-of-all-time-hits-french-site/>
- [5] P. Olson, “The largest cyber attack in history has been hitting hong kong sites,” *Forbes*, November 2014.
- [6] M. Starr, “Fridge caught sending spam emails in botnet attack,” <http://bit.ly/1j5Jacl>, Jan 2014.
- [7] Wikipedia, “Carna botnet,” <http://bit.ly/1slx1E6>, 2014.
- [8] —, “NetAcuity and NetAcuity Edge IP Location Technology,” <http://www.digitalelement.com/>, Feb 2014.
- [9] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson, “Characteristics of internet background radiation,” in *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*. ACM, 2004, pp. 27–40.
- [10] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston, “Internet background radiation revisited,” in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 2010, pp. 62–74.
- [11] Z. M. Mao, V. Sekar, O. Spatscheck, J. van der Merwe, and R. Vasudevan, “Analyzing Large DDoS Attacks using Multiple Data Sources,” in *Proceedings of ACM SIGCOMM Workshop on Large-Scale Attack Defense*, 2006.
- [12] M. Casado, T. Garfinkel, W. Cui, V. Paxson, and S. Savage, “Opportunistic measurement: Extracting insight from spurious traffic,” in *Proc. 4th ACM Workshop on Hot Topics in Networks (Hotnets-IV)*, 2005.
- [13] M. Bailey, E. Cooke, F. Jahanian, J. Nazario, D. Watson *et al.*, “The internet motion sensor—a distributed blackhole monitoring system,” in *NDSS*, 2005.
- [14] W. Chang, A. Mohaisen, A. Wang, and S. Chen, “Measuring botnets in the wild: Some new trends,” in *ACM ASIACCS*, 2015.
- [15] G. P. Zhang, “Time series forecasting using a hybrid arima and neural network model,” in *Neurocomputing*, 2003, pp. 159–175.
- [16] Y. Huang, X. Geng, and A. B. Whinston, “Defeating DDoS attacks by fixing the incentive chain,” *ACM Transactions on Internet Technology*, vol. 7, no. 1, 2007.
- [17] Y. Nadji, M. Antonakakis, R. Perdisci, D. Dagon, and W. Lee, “Beheading hydras: performing effective botnet takedowns,” in *Proceedings of the 2013 ACM SIGSAC conference on Computer and Communications Security*, pp. 121–132, Nov. 2013.
- [18] K. Park and H. Lee, “On the effectiveness of route-based packet filtering for distributed dos attack prevention in power-law Internets,” in *Proceedings of ACM SIGCOMM*, 2001.
- [19] J. Li, J. Mirkovic, M. Wang, P. Reiher, and L. Zhang, “Save: Source address validity enforcement protocol,” in *Proc. of IEEE International Conference on Computer Communications*, 2002.
- [20] J. Ioannidis and S. M. Bellovin, “Implementing pushback: Router-based defense against DDoS attacks,” in *Proc. of Internet Society Symposium on Network and Distributed System Security*, 2002. [Online]. Available: <https://www.cs.columbia.edu/~smb/papers/pushback-impl.pdf>
- [21] A. Yaar, A. Perrig, and D. Song, “Siff: a stateless internet flow filter to mitigate DDoS flooding attacks,” *IEEE Symposium on Security and Privacy*, 2004.
- [22] —, “Stackpi: New packet marking and filtering mechanisms for DDoS and ip spoofing defense,” *IEEE Journal on Selected Areas in Communications*, 2006.
- [23] Y. Chen, K. Hwang, and W.-S. Ku, “Collaborative Detection of DDoS Attacks over Multiple Network Domains,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 18, pp. 1649–1662, 2007.
- [24] A. D. Keromytis, A. D. Misra, and D. Rubenstein, “SOS: An Architecture For Mitigating DDoS Attacks,” *IEEE Journal on Selected Areas of Communications*, 2004.
- [25] A. Stavrou and A. D. Keromytis, “Countering DoS Attacks With Stateless Multipath Overlays,” in *Proceedings of the 12th ACM Conference on Computer and Communications Security*, pp. 249–259, 2005.
- [26] L. Feinstein, D. Schnackenberg, R. Balupari, and D. Kindred, “Statistical approaches to DDoS attack detection and response,” in *DARPA Information Survivability Conference and Exposition*, 2003.
- [27] S. Jin and D. Yeung, “A covariance analysis model for DDoS attack detection,” *IEEE International Conference on Communications*, 2004.
- [28] M. Li, “Change trend of averaged hurst parameter of traffic under ddos flood attacks,” *Computers and Security*, 2006.
- [29] K. Lee, J. Kim, K. H. Kwon, Y. Han, and S. Kim, “DDoS attack detection method using cluster analysis,” *Expert Systems with Applications*, vol. 34, pp. 1659–1665, 2008.
- [30] M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenke, “DDoS defense by offense,” in *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 303–314, 2006.
- [31] Cisco, “Cisco Catalyst 6500 Series Intrusion Detection System,” Feb 2014. [Online]. Available: <http://www.cisco.com/c/en/us/products/interfaces-modules/catalyst-6500-series-intrusion-detection-system-idsm-2-services-module/index.html>
- [32] V. Sekar, N. Duffield, O. Spatscheck, J. van der Merwe, and H. Zhang, “Lads: Large-scale automated DDoS detection system,” in *Proc. of USENIX Annual Technical Conference*, 2006, pp. 171–184.
- [33] J. Mirkovic, G. Prier, and P. Reiher, “Attacking DDoS at the Source,” in *Proceedings of 10th IEEE International Conference on Network Protocols*, pp. 312–321, Nov. 2002.
- [34] A. D. Wood and J. A. Stankovic, “A Taxonomy for Denial-of-Service Attacks in Wireless Sensor Networks,” *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, 2004.
- [35] X. Geng, Y. Huang, and A. B. Whinston, “Defending wireless infrastructure against the challenge of DDoS attacks,” *Mobile Networks and Applications*, vol. 7, no. 3, pp. 213 – 223, 2002.
- [36] J. Mirkovic and P. Reiher, “A Taxonomy of DDoS Attack and DDoS Defense Mechanisms,” *ACM SIGCOMM Computer Communications Review*, vol. 34, pp. 39–54, Apr. 2004.
- [37] S. M. Specht and R. B. Lee, “Distributed Denial of Service: Taxonomies of Attacks, Tools and Countermeasures,” *International Workshop on Security in Parallel and Distributed Systems*, pp. 543–550, Sep. 2004.
- [38] C. Douligeris and A. Mitrokotsa, “DoS attacks and defense mechanisms: classification and state-of-the-art,” *Computer Networks*, 2004.
- [39] T. Peng, C. Leckie, and K. Ramamohanarao, “Survey of network-based defense mechanisms countering the dos and DDoS problems,” *ACM Comput. Surv.*, vol. 39, 2007.
- [40] D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage, “Inferring internet denial-of-service activity,” *ACM Transactions on Computer Systems (TOCS)*, vol. 24, no. 2, pp. 115–139, 2006.
- [41] C. Rossow, “Amplification hell: Revisiting network protocols for DDoS abuse,” in *NDSS Symposium 2014*, 2014.
- [42] C. C. Zou, W. Gong, D. Towsley, and L. Gao, “The monitoring and early detection of internet worms,” *IEEE-ACM Transactions on Networking*, vol. 13, no. 5, pp. 961–974, 2005.
- [43] K. Xu, Z.-L. Zhang, and S. Bhattacharyya, “Profiling internet backbone traffic: behavior models and applications,” in *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 4. ACM, 2005, pp. 169–180.