

Analyzing Large DDoS Attacks Using Multiple Data Sources

Z. Morley Mao
University of Michigan
zmao@umich.edu

Vyas Sekar
Carnegie Mellon University
vyass@cs.cmu.edu

Oliver Spatscheck
AT&T Labs-Research
spatsch@research.att.com

Jacobus van der Merwe
AT&T Labs-Research
kobus@research.att.com

Rangarajan Vasudevan
University of Michigan
ranga@umich.edu

ABSTRACT

We present a measurement study analyzing DDoS attacks from multiple data sources, relying on both *direct* measurements of flow-level information, and more traditional *indirect* measurements using backscatter analysis. Understanding the nature of DDoS attacks is critically important to the development of effective counter measures to this pressing problem. While much of the community's current understanding of DDoS attacks result from indirect measurements, our analysis suggests that such studies do not give a comprehensive view of DDoS attacks witnessed in today's Internet. Specifically, our results suggest little use of address spoofing by attackers, which imply that such attacks will be invisible to indirect backscatter measurement techniques. Further, at the detailed packet-level characterization (e.g., attack destination ports), there are significant differences between direct and indirect measurements. Thus, there is tremendous value in moving towards direct observations to better understand DDoS attacks. Direct measurements additionally provide information inaccessible to indirect measurements, enabling us to better understand how to defend against attacks. We find that for 70% of the attacks fewer than 50 source ASes are involved and a relatively small number of ASes produce nearly 72% of the total attack volume. This suggests that network providers can reduce a substantial volume of malicious traffic with targeted deployment of DDoS defenses.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations—*Network monitoring*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'06 Workshops September 11-15, 2006, Pisa, Italy.
Copyright 2006 ACM 1-59593-417-0/06/0009 ...\$5.00.

General Terms

Measurement, Experimentation, Security

Keywords

Distributed denial of service attacks, DDoS attacks, IP spoofing, attack characterization

1. INTRODUCTION

Internet distributed denial of service (DDoS) attacks are becoming increasingly prevalent [1]. To prevent the discovery of attack sources, attackers have been known to spoof the source IP addresses of packets in DDoS attack traffic. These spoofed addresses were often chosen randomly from the IPv4 space, which allowed a technique called backscatter analysis [2] to be used to infer the prevalence of such spoofed DDoS attacks on the Internet. This technique works by measuring the amount of unwanted traffic sent to unused address blocks. Backscatter traffic originates from attack targets responding to the attack packets by replying to spoofed source addresses. Indeed, much of the current understanding of the nature of DDoS attacks is a result of analyzing such backscatter data by monitoring lightly used or unused address blocks [3, 2].

Relying on backscatter analysis does not provide a complete picture of all possible DDoS attacks. Ingress filtering [4] at interfaces to customer networks can block randomly spoofed traffic by only permitting traffic with addresses known to belong to the customers¹. Backscatter data might also be unavailable if the attacked IP address is not responding to the attack packets. This can be due to the fact that the attacker attacks a target IP address which is not assigned to any host; the attacked host is down; the Internet access link of the attack target is unavailable; or, the attack traffic is filtered. In addition, not all DDoS attacks employ address spoofing. For example, it is known that botnets [6] consisting of thousands of compromised machines can be used to launch DDoS attacks, frequently through public IRC channels, without source address spoofing. From the attacker's point of view, this is an attractive alternative,

¹A recent study [5] shows that a large number of networks still do not deploy ingress filtering, thus allowing spoofed attacks to originate from them.

as attackers do not need to be concerned with ingress filtering which thwart the use of random source address spoofing. Furthermore, attack sources are not owned by attackers and thus, do not reveal their identity.

Our first contribution in this paper is a first-of-its-kind analysis of DDoS attacks that uses two independent data sources namely, *indirectly* measured DDoS activity using backscatter data from a mostly unused /8 network, as well as *directly* measured DDoS activity in an ISP network. For the latter set, we use a combination of independently collected Netflow and alarms from a commercial DDoS detection system, as well as a recently developed large-scale automated DDoS detection system to indicate the presence of flow anomalies [7]. Using simple heuristics we are able to verify that most of the flow anomalies were indeed DDoS attacks. We compare the DDoS attack characteristics from these data sets and find that almost all the attacks in the directly measured data set are *not* present in the backscatter data. If indeed large spoofed attacks were present, a monitored address as large as ours (an unused /8 network) is highly likely to witness backscatter traffic. Also, such large attacks would register volume-based hits on at least the LADS detection system if not the commercial detection system as well. Therefore, the very minimal overlap suggests that most DDoS attacks in today’s networks are unlikely to be detected using the backscatter approach, and direct measurements are required.

Our second contribution is a characterization of DDoS attacks using the directly measured attack data. To our knowledge this study, in which we analyze DDoS packet-level properties and attack sources and targets, is the first study of its kind to be published. Such analysis is beneficial in designing more effective and practical defense mechanisms, understanding mitigation deployment in large networks, as well as developing defense strategies deployable at end-hosts.

The paper is organized as follows. In Section 2, we describe our data sets and the methodology we used to filter the data sets to allow a fair comparison. In Section 3 we present a characterization of DDoS attacks for all the data sets that we use. Also, we determine the extent to which attacks in our data set use randomly spoofed IP packets, and subsequently analyze the sources and targets of attacks in our data sets. We then conclude with a concise summary of contributions and their implications.

2. DATA AND METHODOLOGY

We now describe the three data sources used in our study, as well as our approach to validate the presence of DDoS attack traffic in the flow based data. To allow direct comparisons across the data sets, we preprocess them by focusing only on large attacks targeting address ranges advertised by the tier-1 ISP under study. We expect large randomly spoofed attacks to generate sufficient number of packets to be observed in the /8 backscatter data set.

2.1 DDoS Data Sources

Flow Data from commercial anomaly detection system: We used a commercial flow based DDoS detection system deployed within a tier-1 ISP and collected large DDoS alarms generated by this system over 4 weeks in March 2006. Since the algorithms used in this system are proprietary in nature, we used separately collected Netflow

data, pertaining to the same time duration as the alarms, to study and verify the alarms generated by the DDoS detection system. We use the following steps to derive the flow-based attack traces: (1) Collect alarms for significant flow anomalies from commercial flow based network anomaly detectors deployed in key locations in a large ISP. (2) Correlate the flow anomaly alarms into attack instances, where alarms are combined if they target the same set of destination prefixes and occur with no more than 15 minutes idle time in between. (We focus on targeted attacks.) (3) Retrieve all sampled Netflow records covering the entire network destined toward the attack destination during the attack period.

Flow Data from a custom anomaly detection system: The other data source we have for analyzing DDoS attacks is a home-grown DDoS detection system called LADS – Large-scale Automated DDoS detection System [7] deployed in the tier-1 provider. LADS is based on a triggered multi-stage architecture for scalable, accurate, and cost-effective large-scale attack detection. Conceptually, the initial stages consist of low-cost anomaly detection mechanisms that provide information to traffic collectors and analyzers to reduce the search space for further traffic analysis. Successive stages of the triggered framework, invoked on demand and therefore much less frequently, then operate on data streams of progressively increasing granularity (*e.g.*, flow or packet header traces), and perform more fine-grained analysis. Our system makes use of two data sources: SNMP and Netflow, both of which are readily available in commercial routers today. We adopt a two-stage approach in LADS. In the first stage, we detect volume anomalies using low-cost SNMP data feeds (packets per second counters). These anomalies are then used to trigger flow-collectors that obtain Netflow records for the appropriate routers, interfaces, and time periods. For this stage, we build a traffic prediction model, using packet rate as the metric, for each customer egress interface using historical traffic data over a 5-week period. This prediction model is then used to identify traffic anomalies over the observation period. In the second stage, we then perform automated analysis of the flow records, using uni-dimensional aggregation and clustering techniques, to generate alarm reports indicative of DDoS attacks targeted at customer networks. Here, flow-records are partitioned into 4-different categories (All, ICMP, SYN, RST), and we identify destinations receiving a large traffic flood within the duration specified by the coarse-grained volume anomaly. More details about the system design and implementation can be found in [7].

The bandwidth thresholds depend on the type of flow data and the capacity of the customer’s access interface. For high-capacity interfaces (>5Mbps) we set the threshold to be 10Mbps for aggregate attacks, and 2.5Mbps for specific attacks (SYN, RST, ICMP). For low-capacity interfaces, the thresholds are set to be equal to the access-link capacity. These thresholds are selected to focus on attacks that would impact customers. Our current deployment of LADS monitors in excess of 50,000 customer egress interfaces and collects flow data from over 500 routers in the provider’s backbone for analysis. For the dataset used in this paper, we collected 31612 alarms generated by LADS over a four-week period in March 2006.

Backscatter Data: The backscatter data consist of traffic logs detailing the timestamp, packet type, source and

destination IP address and port numbers obtained from a mostly unused /8 network over the same time period as the flow data. There are very few address blocks from the /8 actively used. Thus, traffic received at the unused addresses is most likely illegitimate – a result of replies to attack packets, measurement probing, worm scanning, misconfiguration *etc.* To exclude packets that are not replies of DDoS attack packets, we take the same set of steps as suggested by prior work on backscatter analysis [2] — that is, we only consider flows with more than 100 packets and lasting more than 60 seconds. We define flows to be consecutive packets with the same source IP address and protocol, based on a 5-minute time out to achieve resilience to temporary outages while not combining unrelated traffic flows [2]. This data set is called the *Backscatter* set in the remainder of this paper, and consists of 3491 events that originate from IP addresses belonging to the tier-1 ISP network. Note that we only focus on backscatter events originating from addresses of the tier-1 ISP and its customers, as we can observe the corresponding attacks targeted to such addresses using the above-mentioned anomaly detection systems.

2.2 Data Processing

Flow Data: To ensure a sufficient number of samples in our data, we therefore focus our analysis on large DDoS attacks which during the time period of the DDoS alarm, transmitted at least 10 million packets through the ISP guaranteeing us at least 25 sampled flow records per attack². Furthermore, we only consider attacks with an average packet rate of at least 6666 packets per second, ensuring that at least one sample per minute is received, given on average 20MB per sample flow record. Considering both attacks that target destinations within the ISP and attacks that originate from customers connected to the ISP gives us a filtered set comprising of 46 potential attacks. This set is referred to as the *LargeFlow* data set. We apply the same filtering to the flow data from the LADS system, giving the *LADS* data set comprising of in 536 attack instances.

Backscatter Data: To allow direct comparison between the Netflow based data (consisting of large DDoS attacks) and the backscatter based data, we filtered the backscatter data based on the same criteria used to classify large ISP DDoS attacks, *i.e.*, with at least 10 million packets total, and average packet rate of at least 6666 packets per second. Note that we scale the above packet rate and packet count by 256 given that we monitor a /8 block, roughly $\frac{1}{256}$ of the IPv4 address space. This filtered set, comprising of 248 events will be referred to as the *LargeBackscatter* data set.

2.3 Validation of Flow Anomalies of the Commercial System

Unlike the backscatter data, which are most likely a result of attacks or misconfigurations, the flow anomaly based DDoS data can contain false positives, especially given the proprietary nature of the commercial DDoS detection mechanism. Note that the flow anomalies from the LADS system, a public DDoS detection system, have already been validated independently [7]. We now describe how we vali-

²Given 10 million packets, assuming 100Byte average packet sizes, the probability of observing at least one sample based on the sampling algorithm [8, 9] is calculated to be very close to 1 using the formula of $1 - e^{(-1 * T / 20000000)}$ with T being the total number of bytes.

date, with high probability, that the data produced by the commercial system is indeed the result of DDoS attacks. Specifically, from the *LargeFlow* data set, we created a filtered data set, called the *SureFlow* set, for which we independently validated the flow records as being part of a DDoS attack by using the following heuristics. We assume a flow based attack trace is a real attack if, considering all flows associated with the trace, any of the following holds true:

- More than 95% of packets in the flows are UDP packets originating from a large number of source IPs (potential UDP flooding attacks).
- More than 95% of packets in the flows are ICMP packets (potential ICMP flooding attacks).
- More than 90% traffic is TCP and all TCP packets have only a single flag. (Most of these flags are SYN, RST, ACK, an indication of SYN flooding attacks or reflector attacks)³

The *SureFlow* set consists of 41 events, containing a little more than 89% of the original *LargeFlow* data set of 46 events. Using the independent Netflow data we therefore are able to verify that at least 89% of the large attack alarms are real DDoS attacks. Note that this does not imply that the remaining 11% of the attacks are not real DDoS attacks, it just means we could not verify them as attacks using our simple heuristics.

Although the *SureFlow* set provides an upper bound of 11% on the number of false positives (the case in which the commercial DDoS detection system marks non-DDoS traffic as a DDoS attack), it is substantially harder to provide a bound on false negatives (no alarms are generated while a DDoS attack was present). Although false negatives do not introduce false data in our analysis, they might bias the analysis of our results. This appears to be a fundamental problem with any DDoS detection mechanism. However, given that the characteristics of large DDoS attacks (compared to small DDoS attacks) are quite unusual and, therefore, easy to detect, we expect the number of false negatives of the commercial DDoS detection system to be low for large DDoS attacks. Furthermore, we also use LADS, a recently developed DDoS detection system with a publicly known algorithm that focuses on flooding attacks on customer interfaces to improve detected attack coverage.

3. DATA ANALYSIS RESULTS

We now detail the data analysis using the four data sets previously mentioned: *Backscatter (BS)*, the backscatter set corresponding to the 100 packet, 60 second rule [2]; *Large-Backscatter (LBS)*, the backscatter set filtered to correspond to large DDoS attacks; *SureFlow (SF)*, the flow-based set for which we were able to verify the existence of DDoS traffic in the flow records; and *LADS*, the flow data using custom volume-based anomaly detection. We compare the data sets whenever possible and correlate them to infer the prevalence of random address spoofing.

³It is possible that the complete failure of a busy server might cause an increase in the number of SYN packets being sent to a particular IP address as clients attempt to re-establish connectivity. This might cause the resulting traffic flows to be incorrectly classified as a SYN flooding attack. Since our study is limited to large DDoS attacks, we do not expect this to be a problem.

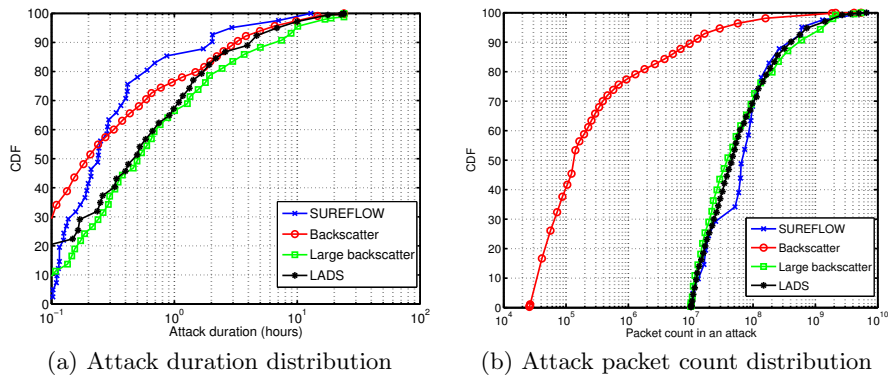


Figure 1: Attack duration and packet count properties

3.1 Attack Characterization

We present a detailed characterization of the attack traffic captured from the four data sets above. In the process, we also point out important network properties and traffic characteristics of the attack traffic that make them stand out from regular traffic.

Traffic properties: We first focus on general traffic properties in terms of attack durations, packet counts, and packet rate. Figure 1(a) depicts the distribution of attack duration for four data sets described above. About 86% of the SureFlow attacks last less than one hour, but some persist as long as 12 hours. The attack duration for both the Backscatter and LargeBackscatter data sets appear to be slightly longer, with about 67% of the LargeBackscatter attacks lasting less than one hour. This could be due to our conservative way of classifying attacks in the backscatter data: as long as there are no idle periods of longer than 5 minutes, data packets are grouped into the same attack flow. This may result in unrelated events grouped together. The attack duration distribution for LADS alarms matches very closely to that for LargeBackscatter attacks. Note that the overall distribution for all the data sets are quite similar.

We find strong correlation between attack duration and the amount of attack traffic in both packet and byte count for the SureFlow data set, and similarly strong correlation for the LADS data set. This means that long-lived attacks usually have more attack traffic. The correlation is much weaker for the Backscatter data set. This is probably an artifact of the 5 minute idle time used to define flows for backscatter data. This implies it is possible that several separate flows are classified as a single flow, resulting in overall lower attack rate for longer-lived attack. The distributions of the total number of packets in each attack are shown in Figure 1(b). Given that our detector uses large traffic volume as one of the criteria to generate alarms, it is not surprising that the attacks in the SureFlow, Largebackscatter, and LADS data set on average have at least 40 Mega packets, very likely aggregated from many traffic flows. The packet count distribution across these data sets match very well. The Backscatter data set has a wider distribution including many smaller attacks with an average of 130K packets. This shows that the majority of the backscatter events, likely caused by the use of spoofed IPs, are quite small in size.

Directly related to the traffic volume metric is the traffic rate. To have a first order estimate, we plot the distribu-

tion of the average packet rate for individual attacks. Figure 2(a) shows the distribution. Again the data points for the three data sets consisting of SureFlow, Largebackscatter, and LADS match reasonably well with medians ranging between 23K and 62K packets per second. The attacks in SureFlow have the highest rate followed by LADS and then Largebackscatter. These rates can be sustained without noticeable impact by today’s core ISP networks, but may impact servers or even firewalls [2]. Based on the Backscatter data, the largest attack has an estimated rate of 280K packets per second. The largest attack observed in the anomalous flow data set is close to 1 million packets per second. This difference might be due to the fact that very large attacks may not be randomly spoofed, thus not visible in the backscatter data. We confirm this conjecture later.

The traffic properties described above are fairly coarse-grained, and according to these properties there is a strong similarity between attacks observed from LargeBackscatter and those from the anomalous flow data sets: SureFlow and LADS.

Packet details: We now examine the packet header and the packet type to study properties such as distribution of port numbers, protocol types, and packet sizes. We do not have information on packet payload due to the aggregate nature of our traffic data, but we can speculate on the application type based on port numbers. Again, we highlight properties of attack traffic that make it stand out from regular traffic. Figure 2(b) displays *the distribution of the number of destination ports in attacks*. In the case of the backscatter data, if the data is a result of spoofed attack, the source port of the backscatter packets correspond to the destination port of the attack targets. Only about 27% SureFlow attacks target a single port, indicative of a single application under attack, the corresponding number is 43% for LADS, 57% for Backscatter, and close to 90% for LargeBackscatter data set. The distribution of the number of destination ports varies significantly across the four data sets, with the LargeBackscatter having the smallest number of destination ports and both LADS and Backscatter data having the largest number of destination ports on average. In fact, the actual values of the destination ports are also quite different across the data sets. The destination ports receiving most packets or highest average packet rate often include service ports for applications such as HTTP, SSH, DNS, IRC. Other popular ports are not well-known and are suspected to be used by peer to peer applications.

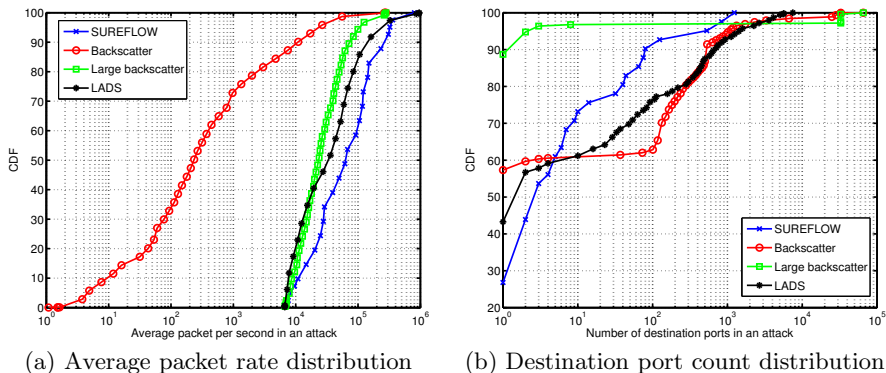


Figure 2: Attack packet rate and destination port count properties

Protocol	SF-PPS	SF-BPS	LADS-PPS	LADS-BPS	LBS-Pkts	LBS-Attacks
TCP	73.06	72.48	49.98	72.15	99.93	99.60
UDP	17.08	17.08	46.18	15.11	0.07	0.40
ICMP	9.86	10.44	2.45	0.52	0.00	0.00

Table 1: Pctg contribution of protocols to attack traffic (SF: SureFlow, PPS: Packets/sec, BPS: Bytes/sec, LBS: LargeBackscatter, Pkts: Packets count, Attacks: Attack count)

Next we characterize the *distribution of IP protocols* by examining the protocol field in the IP header of the traffic in Netflow data and by inference from the backscatter packet types using previously described techniques [2]. Table 1 shows the average traffic contribution in bytes per second and packets per second across all attacks for all the protocols found for the SureFlow and LADS data set. For the LargeBackscatter data set, the table shows the corresponding percentages for the number of packets and the number of attacks respectively. Not surprisingly, similar to regular traffic, TCP is the dominant protocol, UDP being the second highest contributor. Interestingly, the percentage of UDP attacks detected in the SureFlow and LADS data is significantly higher than in the LargeBackscatter data. A possible explanation is that UDP data on random ports are typically blocked by firewalls, which might reduce the potential amount of UDP traffic generating backscatter.

Protocol	BS-Pkts	BS-Attacks	BSAll-Pkts	BSAll-Attacks
TCP	98.99	95.76	98.68	25.45
UDP	0.25	0.78	0.40	59.50
ICMP	0.03	0.90	0.03	0.26

Table 2: Pctg contribution of protocols to attack traffic in backscatter data (BS: backscatter using 60 sec 100 packet filter, BSAll: all backscatter data without filtering, Pkts: Number of packets, Attacks: Number of attacks)

A recent work in progress presented by Nazario [10] also based on backscatter data from a mostly unused /8 network suggested that attacks are shifting from TCP to the UDP protocol. However, our analysis did not confirm this. One possible explanation for this discrepancy is explained by Table 2 where the data set is not filtered based on source addresses belonging to the customer of the tier-1 ISP. Here, *BSALL* refers to all backscatter instances without any filtering, and BS is the 60 second 100 packet filtered Backscatter

set. Surprisingly, for BSALL, we observe a large number of UDP based attacks contributing little traffic, *i.e.*, 59.5% of the attacks which contribute only 0.4% of the number of packets. (Recall that for this data set we did not filter out smaller attacks.) We believe that such events are unlikely caused by real DDoS attacks and more likely due to probing into the dark address space.

Given that TCP is the dominant protocol in attack traffic, we analyze the *distribution of TCP flags* as shown in Table 3. The table shows the average contribution of particular TCP flags only for the flags contributing to more than 1% of the attack traffic in all SureFlow and LADS attacks. To our surprise, we found significantly large amounts of packets in SureFlow and LADS are due to ACK packets, contributing to more than 60% of packet rate and data rate. This implies that these packets may result from reflector attacks targeting the host. Such traffic will not be visible in the backscatter data, as it only consists of packets destined to the unused address block. We also observe the prevalence of SYN-flooding based attacks, as SYN packets contribute to more than 20% of the packet and bit rate. The most popular packet types in the LargeBackscatter data confirms the presence of a large number of SYN flooding attacks, as the most popular packet type is SYNACK. We observe much fewer ICMP port unreachable and echo reply packets which is possibly resulting from UDP and ping flooding attacks.

Finally, we examine the distribution of *packet sizes* in attack traces for SureFlow. Note, it is impossible to understand this property by using backscatter data alone. We found many attacks purely consisting of packets smaller than 100 Bytes. In about 6% of attacks, no packets were smaller than 100 Bytes. However, in about 83% of attacks, all traffic consist of these small packets. Such properties clearly can be used to identify DDoS attacks.

3.2 Spoofing Analysis Results

Given the discrepancies of some of the packet and traffic properties between the backscatter data and the DDoS data from the ISP, we correlate them further to understand their common subset, which would indicate the use of source address spoofing.

Random spoofing: Given the 41 large attacks using the commercial detection system and 536 large attacks using the LADS system observed at the large ISP over our measurement period, we found only 4(0.7%) such alarms matching the backscatter data. Note that all these attack targets belong to customers of the ISP. We use the follow-

SureFlow			LADS			LargeBackscatter		
TCP flag	PPS	BPS	TCP flag	PPS	BPS	Packet type	attack freq	packet cnt
ACK	65.75	66.62	ACK	72.67	63.09	SYNACK	95.56	91.88
SYN	29.13	27.60	PSH,ACK	18.54	35.31	ICMP_UNREACH_PORT	1.21	0.07
PSH,ACK	2.64	3.16	SYN	5.83	0.70	ICMP_ECHOREPLY	0.40	0.00
RST,PSH,ACK	2.42	2.42	SYN,ACK	1.36	0.15			

Table 3: Average contribution by packets with particular TCP flags (flags contributing at least 1% of traffic)

ing simple method to perform the correlation: if there exists some backscatter packets coming from the detected attack target identified in the DDoS alarm during the duration of the alarm (with 5 minutes time window), then we consider this a match. Note that here we do not filter out backscatter attacks based on the size. In fact, among the four matching attack instances, in the backscatter data none of them generated sufficient number of packets and at high enough rate to be considered a large attack.

Another way to determine if an attack used random IP address spoofing is to check if it contains flow records with private nonroutable source IP addresses. Even though the flow records in the SureFlow set are sampled, some flow records with nonroutable source IPs should still be present in our data for attacks which use purely random IP address spoofing. This is due to the fact that a large fraction of the IP address space is not routable and our analysis is limited to attacks which produced at least 25 sampled flow records. Among the 41 large attacks in the SureFlow set, we found 4 attacks with at least one flow record using nonroutable source IPs. Interestingly, these attacks are not visible in the backscatter data. One explanation is that these attacks are indeed using randomly spoofed source addresses, however, the attack target either didn’t generate the backscatter packets or the packets were filtered. Combining these two results we therefore conclude, with high likelihood, that less than 1% of the attacks in the flow-based data set use random IP address spoofing.

Local spoofing: Note that the above discussion focuses on random address spoofing which is much more likely to be detected using backscatter data or nonroutable source IPs. To overcome ingress filtering, attack tools can perform local spoofing or spoof addresses from the local network using the knowledge of the routing address block. It is also possible that the attackers were in fact performing random spoofing but ingress filtering somewhere in the network dropped all but a subset of the attack packets. To understand the possible occurrence of local spoofing, we perform simple clustering at the /24 granularity and count the number of unique source IPs in each /24 for each attack in the SureFlow data set. We found some indication of possible local spoofing in 4 attacks (which are not the ones found in the backscatter data) where there are more than 200 IP addresses participating in the attack from the same /24 network. We believe these events very likely result from local spoofing, as it is unlikely for an attacker to own an entire /24.

Our results also indicate that even though it is possible for compromised machines in botnets to spoof source addresses, they do not appear to be doing so in large numbers.

3.3 Sources and Targets of Attacks

We now provide an initial characterization of the network elements involved in DDoS attacks as observed in our data sets. In particular, we analyze attack sources and targets, and their implications for effective network defense.

3.3.1 Source Analysis

To our knowledge, there has been no systematic study on understanding DDoS attack sources to date, as backscatter data inherently do not have such information. Some related work examined hosts infected by particular worms. For example, Kumar *et al.* [11] recently characterized Witty-infected hosts by exploiting the worm’s structures to determine properties such as the number of disks. Although worm-infected hosts can participate in DDoS attacks, there has been no study directly examining the DDoS attack sources. We fill this important gap by taking a first look at the network properties of DDoS attack sources. Given the disjoint nature between attacks observed in backscatter data and the attacks directly observed in the large ISP in our study, we are confident that most attack sources we discovered correspond to actual IP addresses of hosts that took part in DDoS attacks.

We first seek to understand “how distributed” these attacks are, *i.e.*, how many network entities are taking part in an attack? Second, we attempt to understand whether the same network entities are repeatedly involved in attacks. We perform this analysis at two levels of granularity. First we study the ASes from which attacks originate, indicating the ultimate attack sources. Second, we examine the network ingress interfaces where attacks enter the ISP under our observation, which is important from a network management and mitigation perspective.

Distributed nature of attack sources: In Figure 3, we show the distributed nature of attack sources from three data sets - the *SureFlow* data set; the *LADS* data set; and the *SureFlow_extended* data set which is an extended version of the *SureFlow* data set covering an 11-month period up to March 2006. From the figure, it is clear that a sizable number of attacks originate from few sources and ingress interfaces for all three data sets. For instance, for *SureFlow* data set, there are fewer than 100 source IPs and 50 ASes that were involved in about 70% of attacks, and similarly fewer than 0.1% of ingress interfaces were involved in all attacks. This indicates that DDoS attacks are much less distributed than their name implies.

Topological predictability of attack sources: Tables 4 and 5 capture the volume contributions of ASes and network interfaces to attacks — originating from a particular AS and entering the ISP network via a particular interface. The former table represents the *SureFlow* data set while the latter the *SureFlow_extended* data set. The tables show results for three different “bins” corresponding to the individual contribution of each entity (AS or ingress interface). As shown in the tables, attacks tend to originate from the same set of networks, and for the ISP under observation enter through the same set ingress interfaces. For example, for the *SureFlow* data set, just 2 ASes, each individually contributing at least 1% of total attack traffic, together contribute more than 72% of attack traffic observed. Similarly, 0.01% of the ingress interfaces will carry more than 90% of

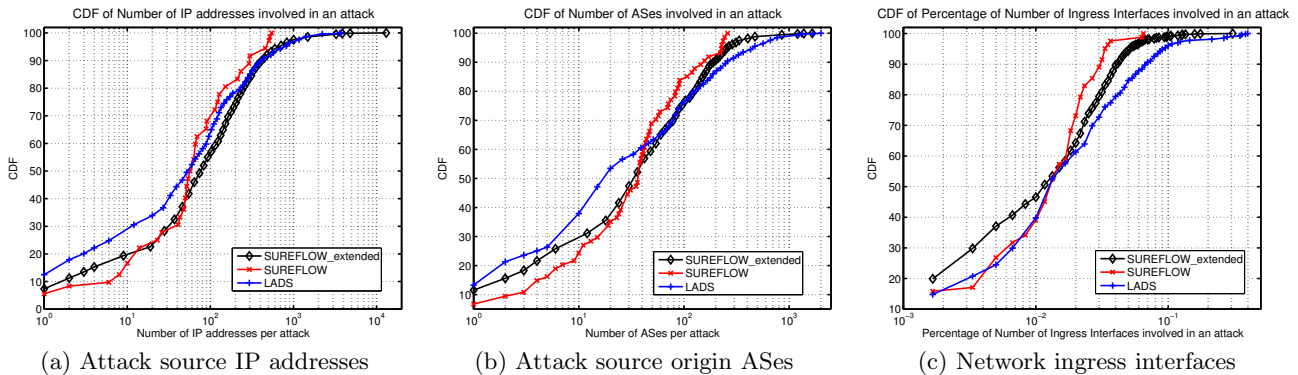


Figure 3: Distributed nature of attack sources

Individual volume contribution	Number of ASes	AS volume contribution	Percent of ingress interfaces	Interface volume contribution
> 1.0%	2	72.08	0.01	93.26
<= 1.0% and > 0.1%	54	13.56	0.027	6.38
<= 0.1%	1901	14.36	0.087	0.36

Table 4: Distribution of origin ASes and ingress interfaces for the SureFlow data set.

the attack traffic by volume, while each such interface carries at least 1% attack traffic. Even more interestingly, from the *SureFlow_extended*, we find that less than 1.1% of all ingress interfaces participated in any DDoS activity over a 11-month period. Note that, due to smart sampling of Netflow records, our source analysis considers a subset of the actual distinct number of sources involved in the attacks. However, the numbers stated above are not expected to be significantly different. Hence, such predictability is very useful for attack detection and mitigation purposes. For example, an ISP only has to deploy mitigation equipment at about 2% of its ingress interfaces to be able to mitigate all DDoS traffic within our observation period before it enters the ISP’s backbone.

3.3.2 Target Analysis

We analyzed the targets of attacks from all three data sets. Across all three data sets, we find that many attacks target customers of service providers. These include end-users of broadband Internet Service Providers (cable and DSL), network service providers for small businesses, web-hosting services, providers of network telecommunications like VoIP, and network customer care services. In fact, for the *SureFlow* data set, we found that more than 90% of all targets were likely end-users or small businesses, who leased network connectivity from lower tiered service providers. The numbers for the *SureFlow_extended* and LADS data sets were 80% and 73% respectively. There were very few university-based users in the targets attack data sets. Likewise, there were very few of the Fortune 500 corporations targeted.

In terms of frequency, most targets were victims of a single or a small number of attacks. For the *SureFlow* data set, about 99% of targets featured in 1 or 2 attacks; while for the *SureFlow_extended* and LADS data sets, these numbers were 91% and 83% respectively. Moreover, the more frequent targets in all three data sets were clustered fairly close together with no single target being the most favored over an extended period. Small businesses were the most favored amongst all targets across all 3 data sets.

3.4 Result Summary and Implications

Here we summarize the main observations from our study. For the analysis using both direct and indirect measurements we found that at a coarse-grained aggregate level, large attacks observed in backscatter data match well with those in direct measurements from Netflow data. These similarities do not hold when more fine-grained attack properties, such as the type of service under attack and IP protocols used, are being considered. We list some specific results: (i) Most attacks (at least 70%) last for less than an hour. (ii) Packet rates are in the tens of thousands per second, maximum close to 1 million packets per second. (iii) Most attacks use TCP. Here we saw some difference with LADS data showing much higher UDP involvement (46% by packets for LADS versus less than 1% for backscatter). (iv) Most TCP based attacks are ACK or SYN only floods. (v) Less than 1% of the directly measured attacks produced backscatter. Using the direct measurements we found that: (i) Most attacks (83%) consist only of packets smaller than 100 Bytes. (ii) We saw evidence of local spoofing in a small number (4) of the attacks. (iii) Attacks are only mildly distributed (fewer than 50 ASes were involved in about 70% of the attacks as attack sources, and less 0.1% of ingress interfaces were involved in all attacks). (iv) There is significant predictability in attacks both in terms of their originating AS as well as from which interface they enter a large ISP network. (v) Small businesses seem to be the most common targets of attacks.

These results have significant implications for attack defense. With respect to attack detection and understanding attacks, relying on indirect measurements is clearly not sufficient given current trends in attacks, since very few if any of attacks appear to be using spoofed source addresses. As a corollary, we also find that direct measurements can provide significantly more diagnostic capability that can better guide the design and deployment of attack defenses. In fact, there are positive implications for attack defense. We find that from the perspective of service providers DDoS attacks are really not as distributed as they are made out to be. Since the vast majority of malicious traffic arises from a

Individual volume contribution	Number of ASes	AS volume contribution	Percent of ingress interfaces	Interface volume contribution
> 1.0%	18	32.08	0.02	89.80
<= 1.0% and > 0.1%	126	30.58	0.042	8.53
<= 0.1%	15743	37.34	1.083	1.67

Table 5: Distribution of origin ASes and ingress interfaces of the SureFlow_extended data set.

small set of ASes and network ingress points, providers can ensure significant protection for their customers with even limited (but intelligently targeted) deployment of DDoS defense mechanisms (e.g. [12]).

4. CONCLUSION

Our work is a first study at combining multiple independent data sources to study large DDoS attacks. We examined backscatter data from a mostly unused /8 network along with flow anomaly based DDoS data from a tier-1 ISP network. The attack characterization indicates that most properties such as attack duration, packet count, packet rate, and dominant protocol type match fairly well in the two data sets. However, we do observe strong discrepancies in other properties such as the number of attack destination ports. One possible explanation for such differences is that they cover different types of DDoS attacks as shown by the very small overlap between them: one consists entirely of spoofed attacks, the other are mostly unspoofed. Using direct DDoS attack measurements, we performed a first analysis of several DDoS properties which is impossible using indirect measurements.

Acknowledgments

We are very grateful to Evan Cooke for providing us the backscatter data. We would like to thank Brian Rexroad for his insightful comments on an earlier version of this paper. The work is partially supported by NSF grants CNS-0430204, CNS-0433540, and ANI-0331653.

5. REFERENCES

- [1] R. Richmond, "Firms Join Forces Against Hackers," *Wall Street Journal*, March 28, 2005.
- [2] D. Moore, G. Voelker, and S. Savage, "Inferring Internet Denial of Service Activity," in *Proceedings of the 2001 USENIX Security Symposium*, August 2001.
- [3] J. Mirkovic, S. Dietrich, D. Dittrich, and P. Reiher, *Internet Denial of Service: Attack and Defense Mechanisms*. Prentice Hall, 2005.
- [4] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing." RFC 2267, January 1998.
- [5] R. Beverly, "The Spoofer Project: Inferring the Extent of Internet Source Address Filtering on the Internet." roceedings of USENIX Steps to Reducing Unwanted Traffic on the Internet Workshop (SRUTI 2005), July 2005.
- [6] The HoneyNet Project & Research Alliance, "Know your Enemy: Tracking Botnets." <http://www.honeynet.org/papers/bots/>, March 2005.
- [7] V. Sekar, N. Duffield, J. van der Merwe, O. Spatscheck, and H. Zhang, "LADS: Large-scale Automated DDoS Detection System," in *Proc. USENIX Annual Technical Conference*, 2006.
- [8] N. Duffield, C. Lund, and M. Thorup, "Charging from sampled network usage," in *ACM SIGCOMM Internet Measurement Workshop*, 2001.
- [9] N. Duffield and C. Lund, "Predicting Resource Usage and Estimation Accuracy in an IP Flow Measurement Collection Infrastructure," in *ACM SIGCOMM Internet Measurement Conference*, 2003.
- [10] J. Nazario, "Trends in Denial of Service Attacks." WIP of Usenix Security 2003, 2003.
- [11] A. Kumar, V. Paxson, and N. Weaver, "Exploiting Underlying Structure for Detailed Reconstruction of an Internet-scale Event," in *Proceedings of ACM IMC*, 2005.
- [12] "Cisco Guard." <http://www.cisco.com/en/US/products/ps5888/>.