

Detecting Distributed Attacks using Network-Wide Flow Traffic

Anukool Lakhina
Dept. of Computer Science,
Boston University
anukool@cs.bu.edu

Mark Crovella
Dept. of Computer Science,
Boston University
crovella@cs.bu.edu

Christophe Diot
Intel Research
Cambridge, UK
christophe.diot@intel.com

1. INTRODUCTION

Distributed denial of service attacks have become both prevalent and sophisticated. Botnet-driven attacks can be launched from thousands of worm-infected and compromised machines with relative ease and impunity today. The damage caused by such attacks is considerable: the 2004 CSI/FBI computer crime and security survey found that DDOS attacks are the second largest contributor to all financial losses due to cybercrime [3]. Further, distributed attacks are expected to increase both in sophistication and damage [1]. Containing distributed attacks is therefore a crucial problem, one that has not been adequately addressed.

One reason why distributed attacks are difficult to contain is because defenses against these attacks are typically deployed at edge networks, near the victim. Deploying defenses at the edge makes detecting attacks easier, since one simply needs to monitor incoming traffic volume for an unusually large burst. However, containing and mitigating such attacks from the edge is ineffective for two reasons. First, filtering the malicious attack traffic requires identifying the (potentially thousands of) attackers, which is complicated, especially if the source addresses are spoofed. Second, even if accurate filtering was feasible at the edge, it cannot prevent attackers from consuming the victim's bandwidth, and denying service to legitimate users. Thus edge-based defenses against distributed attacks have limited value.

On the other hand, defending against distributed attacks at the backbone (*i.e.*, carrier networks) overcomes the hurdles of edge-based defenses. In principle, backbone networks can detect and identify the origins of malicious sources involved in a distributed attack that traverses the backbone. Thus backbone networks are well-suited to mitigate distributed attacks, before they cause harm to the victim at the edge. However, distributed attacks are challenging to detect in the backbone because they do not cause a visible, easily detectable change in traffic volume on individual backbone links. To effectively detect distributed attacks in the backbone, one therefore needs to simultaneously analyze all traffic across the network.

In this work, we present our methods to detect distributed attacks in backbone networks using sampled flow traffic data. Distributed attacks are traditionally viewed to be fundamentally more difficult to detect than single-source attacks. In contrast, we demonstrate that the more distributed an attack is, the better our methods are at detecting it. This is because our methods analyze correlations across all *network-wide* traffic simultaneously, instead of inspecting traffic on individual links in isolation. In addition, our methods are highly sensitive to the attack intensity; we show that attacks rates of less than 1% of the underlying traffic can be detected successfully by our methods.

The rest of this paper is organized as follows. In the next section

we show how network-wide traffic summaries can be assembled, and present the data we have processed from the Abilene Internet2 backbone network. Then, in Section 3, we describe the multiway subspace method for detecting attacks in network-wide flow data. We evaluate our methods on actual DDOS attack traces in a series of experiments and present results in Section 4. Finally, we conclude in Section 5.

2. NETWORK-WIDE FLOW DATA

Our methods analyze all traffic that traverses the network. To obtain such network-wide flow traffic, we must collect the ensemble of origin-destination flows (OD flows) from a network. The traffic in an Origin-Destination flow is the set of traffic that enters the network at specific point (the origin) and exits the network at the destination. For this study, we assembled the set of OD flows for the Abilene network.

Abilene is the Internet2 backbone network, connecting over 200 US universities and peering with research networks in Europe and Asia. It consists of 11 Points of Presence (PoPs), spanning the continental US. We collected three weeks of sampled IP-level traffic flow data from every PoP in Abilene for the duration of December 8, 2003 to December 28, 2003. Sampling is periodic, at a rate of 1 out of 100 packets, and flow statistics are reported every 5 minutes; this allows us to construct traffic timeseries with bins of size 5 minutes.

To aggregate the IP flow data at the OD flow level, we must resolve the egress PoP for each flow record measured at a given ingress PoP. This egress PoP resolution is accomplished by using BGP and ISIS routing tables, as detailed in [2]. After this procedure is completed, there are 121 OD flows in Abilene.

Our final post-processing step constructs timeseries at 5 minute bins for traffic summaries of each OD flow. The traffic summary we use is the sample entropy of the four main traffic features (source IP, destination IP, source port and destination port). Sample entropy captures the *distribution* of each traffic feature in a manner that reveals unusual changes in the distribution. An analysis of the merits of distributional-based analysis of traffic features for anomaly diagnosis can be found in [6].

To summarize, the network-wide flow traffic we study is the multivariate timeseries of sample entropy of traffic features for the ensemble of Abilene's OD flows.

3. THE MULTIWAY SUBSPACE METHOD

To detect distributed attacks, it is necessary to examine network-wide traffic - as captured by the set of OD flows - simultaneously. The multiway subspace method accomplishes this task and is described in [6]; we review the main ideas here.

Thinning Rate	0	10	100	1000	10000
Attack Intensity (pps)	2.75e4	2.75e3	275	27.5	25.9
Attack Intensity (%)	93%	57%	12%	1.3%	0.13%

Table 1: Intensity of injected attack, in # pkts/sec (pps) and percent of (single) OD flow traffic.




Figure 1: Detection results from injecting multi-OD flow attacks (across 2 to 11 OD flows).

The multiway subspace method separates the ensemble of OD flow timeseries into normal and anomalous attributes. Normal traffic behavior is determined directly from the data, as those temporal patterns that are most common to the set of OD flows. This extraction of common trends is achieved by Principal Component Analysis (PCA). As shown in [7], PCA can be used to decompose the set of OD flows into their constituent common temporal patterns.

A key result of [7] was that a handful of dominant temporal patterns are common to the hundreds of OD flows. The multiway subspace method exploits this result by designating these dominant trends as normal, and the remaining temporal patterns as anomalous. As a result, each OD flow can be reconstructed as a sum of normal and anomalous components. In particular, we can write, $\mathbf{x} = \hat{\mathbf{x}} + \tilde{\mathbf{x}}$, where \mathbf{x} denotes the traffic of all the OD flows at a specific point in time, $\hat{\mathbf{x}}$ is the reconstruction of \mathbf{x} using only the dominant temporal patterns, and $\tilde{\mathbf{x}}$ contains the residual traffic.

Once this separation is completed, detection of unusual events requires monitoring the size (ℓ_2 norm) of $\tilde{\mathbf{x}}$. The size of $\tilde{\mathbf{x}}$ measures the degree to which \mathbf{x} is anomalous. Statistical tests can then be formulated to test for unusual large $\|\tilde{\mathbf{x}}\|$, based on a desired false alarm rate [5].

As demonstrated in [6], the multiway subspace method can detect a broad spectrum of anomalous events, at a low false alarm rate. Further, these anomalies can span multiple traffic features, and also occur in multiple OD flows. Our focus in this work is to specifically evaluate the power of network-wide traffic analysis, via the multiway subspace method, to detect distributed attacks.

4. DETECTION RESULTS

In this section we specifically study how effective our methods are at detecting distributed attacks. We first describe our experimental setup, where we inject traces of a known distributed denial of service attack in the Abilene network-wide flow data. Next, we present results from applying the multiway subspace method to detect these injected attacks.

4.1 Injecting Distributed Attacks

To evaluate our detection method, we decided to use an actual distributed denial of service attack packet trace and superimpose it onto our Abilene flow data in a manner that is as realistic as possible. This involved a number of steps which we describe below.

We use the distributed denial of service attack trace collected and described in [4]. This 5-minute trace consists of packet headers without any sampling. It was collected at a Los Nettos regional ISP in 2003, and so exemplifies an attack on an edge network. We extracted the attack traffic from this attack trace by identifying all packets directed to the victim. We then mapped header fields in the extracted packets to appropriate values for the Abilene network.

Then, to construct representative distributed attacks, we divided the attack trace into k smaller traces, based on uniquely mapping the set of source IPs in the attack trace onto k different origin PoPs of Abilene. The splitting was performed so that each of the k groups has roughly the same amount of traffic. Next, we injected this k -partitioned trace into k OD flows sharing the same destination PoP (the victim of the DDOS attack). For each destination PoP, we injected the k OD flow attack into all possible combinations of k sources, *i.e.*, $\binom{p}{k}$ combinations where $p = 11$ is the number of PoPs in the Abilene network. We repeated this set of experiments for every destination PoP in the network; thus for a given choice of $2 \leq k \leq p$, we performed $\binom{p}{k} \cdot p$ total experiments. For each multi-OD flow injection, we recorded if the multiway subspace method detected the attack.

Finally, we repeated the entire set of experiments at different thinning rates to measure the sensitivity of the detection methods to lower intensity DDOS attacks. We thinned the original attack trace by selecting 1 out of every N packets, then extracted the attack and injected it into the Abilene OD flows, as described above. The resulting attack intensity for the various thinning rates is shown in Table 1. The table also shows the percent of all packets in the resulting OD flow that was due to the injected anomaly.

While these multi-OD flow experiments are designed to span a number of origin PoPs sharing a common destination PoP, our

detection methods do not assume any fixed topological arrangement on the malicious OD flows. The results from these experiments give us insight into the performance of the multiway subspace method in detecting attacks that are dwarfed in individual OD flows, but are only visible network-wide, across multiple OD flows.

4.2 Results

We now present results on using the multiway subspace method to detect multi-OD flow attacks. The detection rates (averaged over the entire set of experiments) from injecting DDOS attacks spanning $2 \leq k \leq 11$ OD flows are shown in Figure 1. Figure 1(a) and (b) present results when the detection threshold is $\alpha=0.999$ (equivalent to asking for a false alarm rate of $1-\alpha$, or 1 in 1000) and $\alpha=0.995$ (equivalent to a false alarm rate of 5 in 1000).

Both figures show that we can effectively detect attacks spanning a large number of OD flows. In fact, the detection rates are generally higher for larger k , *i.e.*, for attacks that span a larger number of OD flows. For example, in Figure 1(a) we detect 100% of DDOS attacks that are split evenly across all the 11 possible origins PoPs in the Abilene network, even at a thinning rate of 1000. From Table 1, the average intensity of the DDOS attack trace in each of the 11 OD flows at a thinning rate of 1000 is $\frac{27.5}{11} = 2.5$ packets/sec.

In Figure 1(b), we relaxed the detection threshold to $\alpha = 0.995$. In this setting, we detect about 82% of all DDOS attack traffic spanning 10 OD flows, at thinning rates of even 10000, which corresponds to an attack with intensity of 0.259 packets/sec in each of the 10 participating OD flows individually. Such low-rate attacks are a tiny component of any single OD flow, and so are only detectable when analyzing multiple OD flows simultaneously.

Thus, the results here underscore the power of network-wide analysis via the multiway subspace method.

5. CONCLUSIONS

Distributed attacks are an important problem facing networks today. We argue that distributed attacks are best mitigated at the backbone. Detecting distributed attacks at the backbone requires departing from traditional single-link traffic analysis and adopting a *network-wide* view to traffic monitoring. In this work, we applied the multiway subspace method on network-wide flow data to detect distributed attacks in the Abilene backbone network. Through a series of controlled experiments, we demonstrated that the multiway subspace method is well suited to detect massively distributed attacks, even those with low attack intensity.

6. ACKNOWLEDGEMENTS

We thank Alefiya Hussain for providing the multi-source denial of service attack trace.

This work was supported in part by NSF grants ANI-9986397 and CCR-0325701, and by Intel.

7. REFERENCES

- [1] CERT Research. 2004 Annual Report. At www.cert.org/archive/pdf/cert_rsrch_annual_rpt_2004.pdf.
- [2] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True. Deriving traffic demands for operational IP networks: Methodology and experience. In *IEEE/ACM Transactions on Networking*, pages 265–279, June 2001.
- [3] L. A. Gordon, M. P. Loeb, W. Lucyshyn, and R. Richardson. 2004 CSI/FBI Computer Crime and Security Survey.

Available at www.gocsi.com/forms/fbi/csi_fbi_survey.jhtml, 2004.

- [4] A. Hussain, J. Heidemann, and C. Papadopoulos. A Framework for Classifying Denial of Service Attacks. In *ACM SIGCOMM*, Karlsruhe, August 2003.
- [5] A. Lakhina, M. Crovella, and C. Diot. Diagnosing Network-Wide Traffic Anomalies. In *ACM SIGCOMM*, Portland, August 2004.
- [6] A. Lakhina, M. Crovella, and C. Diot. Mining Anomalies Using Traffic Feature Distributions. In *ACM SIGCOMM*, Philadelphia, August 2005.
- [7] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. D. Kolaczyk, and N. Taft. Structural Analysis of Network Traffic Flows. In *ACM SIGMETRICS*, New York, June 2004.