Detecting DoS Attacks Using Packet Size Distribution

Ping Du
National Institute of Informatics,
Tokyo, Japan
duping@nii.ac.jp

Shunji Abe
National Institute of Informatics,
Tokyo, Japan
abe@nii.ac.jp

ABSTRACT

Enabling early detection of Denial of service (DoS) attacks in network traffic is an important and challenging task because DoS attacks have become one of the most serious threats to the Internet. In this paper, we develop an IP packet size entropy (IPSE)-based DoS detection scheme in which the entropy is markedly changed when traffic is affected by an attack. Through our analysis, we find that the IPSE-based scheme is capable of detecting not only long-term attacks but also short-term attacks that are beyond the volume-based schemes' ability to detect.

Keywords

Denial of service attack, Network security, Attack detection

1. INTRODUCTION

In recent years, denial of service (DoS) [1] attacks have caused significant financial loss and have become one of the most serious security threats to the Internet. In a DoS attack, a malicious user often cripples a victim by simply flooding the target with many legitimate-looking requests. Launching a DoS attack is very easy by using some attack tools, but detection is still an open issue because of the complex nature of network traffic.

Many DoS detection schemes have been proposed and basically they can be divided into two types: volume-based and feature-based. In a volume-based detection scheme, attacks are detected by identifying abrupt changes in traffic volume. Although volume-based detection schemes [2, 4, 6] have been successful in isolating large traffic changes, a large class of short-term DoS attacks do not cause detectable disruptions in traffic volume because they have only minor effects on the traffic volume. On the other hand, a feature-based detection scheme [3, 5] detects attacks by inspecting changes in the distributional aspects of packet header fields. Although feature-based detection schemes can detect even small-volume attack traffic, inspecting the header fields of every packet to collect and analyze the features is too exhausting a method to detect attacks in real time.

In this paper, considering that different application traffic has different packet size distributions and that this distribution changes during DoS attacks, we propose an IP packet size entropy (IPSE)-based scheme in which the attacks are detected by observing time series of packet size entropy. A spike in the time series indicates that a possible DoS attack is under way. We distinguish DoS traffic from legitimate traffic at the detected possible attack points. Different from existing volume-based and feature-based methods, our research studies DoS traffic characteristics from the perspective of the IP packet size distribution, which has not been used in attack detection yet.

The rest of the paper is organized as follows. Section 2 elaborates on the utility of the IPSE-based scheme for detecting DoS attacks and introduces a method for distinguishing attack traffic from legitimate traffic. Section 3 discusses the performance of our proposed IPSE-based attack detection scheme in experiments using real DARPA traffic-trace data sets. Section 4 concludes our proposal and outlines our future work.

2. PACKET SIZE ENTROPY-BASED DETECTION SCHEME

This section describes the IP packet size entropy (IPSE)-based detection scheme. We first show how IP packet size entropy can be used to detect a potential DoS attack.

2.1 Using entropy to detect traffic anomalies

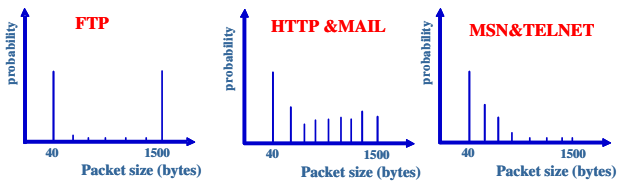

Figure 1: Illustration of packet size distribution for different applications

As shown in Fig.1, many applications have typical packet sizes with respect to requests and responses or data and acknowledgments. For long-term TCP sessions such as FTP applications, traffic mostly consists of simple acknowledgment packets with 40 bytes and full data packets with 1500 bytes. For short-term TCP sessions such as MSN&TELNET, each data packet only contains simple text messages with

Bionetics '07, December 10-13, 2007, Budapest, Hungary

small sizes. In an HTTP application, each object on a web page corresponds to a short-term TCP session. Hence, HTTP packets have a wide range of byte sizes because of the different object sizes.

On the other hand, attacks usually produce packets independent of the response from the victim. Moreover, flooding-based attack traffic often consists of packets with identical sizes. For example, a SYN flooding attack traffic consists of SYN packets with 40 bytes and an ICMP flooding attack traffic consists of ICMP packets with 1500 bytes. Hence, we believe that the distribution of packet size is changed under attacks and that analysis of the packet size distribution can identify attacks on some degree especially when some special IP packet size distribution appears.

How to effectively describe the packet size distribution in a manner that provides necessary information for attack detection is the key question. After conducting observations, we find that entropy, which describes the degree of dispersal or concentration of a distribution, is an effective metric for extracting the properties of the packet size distribution in a manner that is appropriate for attack detection. By observing the time series of the entropy of packet size, we can expose the changes in packet size distribution and detect attack points.

Suppose an observation window contains S packets at time t; the entropy of the packet size at time t is defined as

$$H(t) = -\sum_{l}(\frac{n_l}{S})log(\frac{n_l}{S}), \quad (1)$$

where n_l is the number of times packets with size l in the observation window. The time series of entropy consists of the entropies calculated within a sliding observation window of specified size S. The length of this sliding observation window should depend on the duration of the attack traffic that we wish to capture. If we denote the duration of the attack traffic by S_0 packets, we need, in the ideal situation, to have $q = S_0/S \approx 1$. If the quotient q is too small, the anomaly may be blurred and lost. If the quotient is too large, we may be overwhelmed by "attacks" that are of very little interest to the network operator. Our current experiment focuses on anomalies with durations of at least 200 packets. The entropy takes on a small value when size distribution of observed packets are concentrated (i.e, all packets are of the same size) and takes on a large value when the size distribution is dispersed.

We observed the time series of packet size entropy for the traffic from the DARPA/MIT Lincoln Laboratory off-line intrusion detection evaluation data set [7], which has been widely used for testing intrusion detection systems [8]. As indicated in Fig.2, a short-term ICMP flooding attack and a long-term SYN flooding attack happened at 09:18:15 and 11:20:15 on 03/11/1999, respectively. The ICMP flooding attack lasted for 0.3s and the SYN flooding attack lasted for 120s. The top plot of Fig.2 shows that not only the long-term SYN flooding attack causes a spike in the graph of the time series of packet size entropy; the short-term ICMP flooding attack does as well. On the other hand, as shown in the bottom plot of Fig.2, the ICMP flooding attack does not cause a detectable change in traffic volume. These analysis results show that the entropy of the IP packet size is a more suitable metric than volume because it successfully captures both long-term and short-term attacks.

The threshold of entropy H_{th} for reporting an alarm can

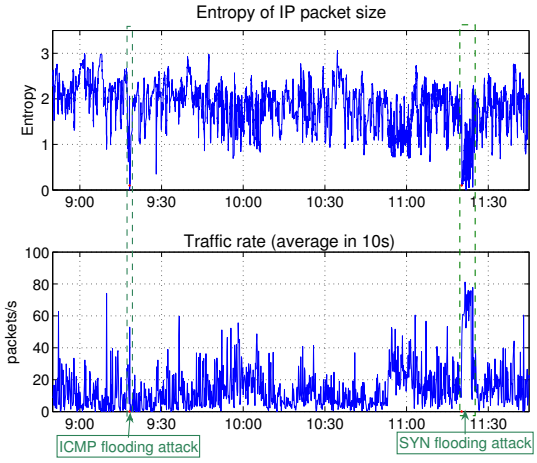

Figure 2: Example of DoS attacks viewed in terms of entropy (top plot) and volume (bottom plot)

be gotten by self-learning of legitimate traffic data for a certain period. After checking clean legitimate DARPA data sets for two weeks, we found that then entropies are mostly distributed in [0.6, 3.5] and events with entropies of less than 0.5 only happened two or three times per day. This is a very low false alarm rate. (As described in [7], a system with 10 false alarms per day is preferred.) Here, two consecutive alarms are calculated as one alarm when their interval is less than one second since the network operator has already been alerted by the first alarm and too many alarms are not meaningful. Hence, we set 0.5 as the default threshold of entropy H_{th} for studying the DARPA traffic data. When the entropy is less than H_{th}, it indicates that a possible denial of service is under way.

2.2 Discrepancy between DoS attack and legitimate traffic

As introduced above, one potential problem for the IPSE-based detection scheme is false alarms when many legitimate packets arrive simultaneously. Here, we will try to solve the problem by analyzing the different packet arrival processes of legitimate applications and attacks. Here, we would like to make the same assumption as in [1]; an attacker will do his best to cripple the victim by sending data with the maximum rate possible and will consistently make requests for higher rates than legitimate clients. Because any computer and network interface has a maximum possible transmission rate due to hardware or operating system limits, the attacker's sending rate will be usually at a constant rate.

As shown in Fig.2, besides the ICMP flooding attack that happened on 09:18:15 which causes a small spike on the time series of packet size entropy, there is another spike at 09:28:06 which is caused by an FTP DATA session. We compare the packet arrival processes in Fig.3 by counting the number of arrival packets in each of 10ms. The packets of the ICMP flooding attack traffic arrive at a constant rate, whereas the packet arrival process of the FTP DATA session is burstier.

Figure 4 shows the packet arrival process of the SYN flooding attack traffic of one observation window in which its packet size entropy is less than 0.5. In the figure, the num-

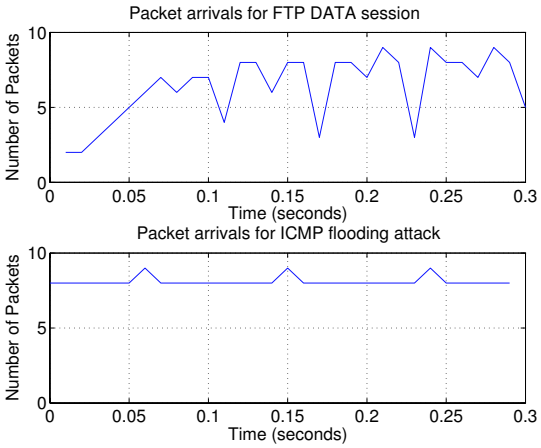

Figure 3: Packet arrival processes for FTP DATA session (top plot) and ICMP flooding attack traffic (bottom plot)

ber of arriving packets is counted in 100ms. The results show that the packet arrival rate also approximates constant.

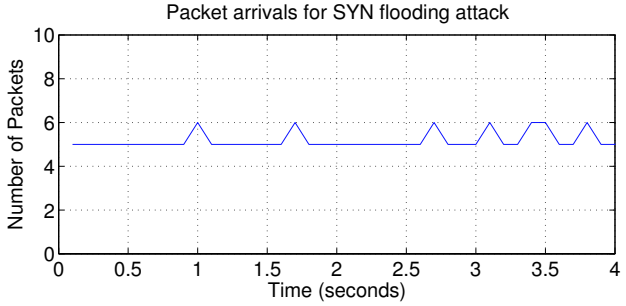

Figure 4: Packet arrival process for SYN flooding attack traffic

While Figs.3 and 4 provide some intuition to judge whether the packets arrive at a constant rate by visual inspection, it is difficult to automate and quantify this idea in an implementation. Whether the packet arrival rate is constant or not can also be judged by calculating the variance of the number of arrival packets in a time unit. Suppose X_n is the number of packets in the nth time interval and $\sum_n X_n = S$. We assume $\{X_n\}$ to be a wide-sense stationary discrete stochastic process, with mean $\mu = E[X_n]$. The variance is defined as $Var[X_n] = E[(X_n - \mu)^2]$. Without loss of generality, we can use a new parameter, *deviation* D, which is defined as $D = Var[\frac{X_n}{\mu}] = E[(\frac{X_n}{\mu} - 1)^2]$, to represent the variance of the packet arrival process. The calculation results for the packet arrival processes in Figs.3 and 4 are shown as Table 1.

According to the calculation results, we can see that D is a very suitable metric to judge whether the packets arrive at a constant rate or not. The D of legitimate traffic $D_{legitimate} \gg 0.01$, whereas the D of DoS attack traffic $D_{DoS} \ll 0.01$. Therefore, we can distinguish the DoS attack traffic from legitimate traffic by calculating D. When $D < 0.01$, the traffic can be judged to be DoS attack traffic.

Table 1: Statistical comparison of different types of traffic

Traffic Type	Mean (μ)	Variance ($Var[X_n]$)	Deviation ($D = Var[X_n]/\mu^2$)
FTP DATA	6.36	4.7	0.12
ICMP flooding	8.1	0.093	0.0014
SYN flooding	5.18	0.148	0.0055

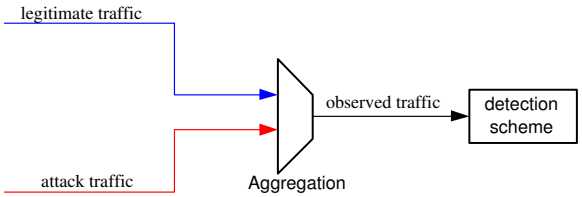

Figure 5: The environment of DoS attack detection experiments

3. EXPERIMENTAL EVALUATION

In this section, we investigate the performance of the IPSE-based algorithm presented in the previous section. The performance metrics are considered as follows: (1) detection probability: $DP = \frac{\text{number of successful detections}}{\text{number of attacks}}$; and (2) detection time: the detection delay after the detection starts.

The environment of DoS/DDoS experiments is shown in Fig.5, in which we inject the attack traffic of different rates and durations into the real traffic data sets and apply IPSE-based detection scheme to diagnose these known attacks. We perform this repeatedly for each kind of attacks so as to check how does the attack traffic rate and duration affect the method's performance.

Our experiments use real network traffic taken from the MIT Lincoln Laboratory. The data set taken on 03/08/2000 contains 11 hours of collected packets (08:00-19:00) and has a mean rate of about 13 packets/s measured in 10 second intervals.

3.1 Detection of short-term high-rate attacks

Our first step is to detect short-term attacks with a high rate, as in the ICMP flooding attack in Fig.2. In each experiment, we generate multiple attacks and inject them into the DARPA traffic for detection. The inter-arrival time between consecutive attacks is exponentially distributed with mean value 10 minutes. The attacks in the same experiment are modelled with the same duration and constant rate. For different experiments, one attack traffic consists of 200, 400, and 600 packets respectively and the attack traffic rate varies from 100 packets/s to 1000 packets/s. Without loss of generality, all packets are of 1500 bytes, which is the same as that of ICMP flooding attack traffic. Table 2 shows that the IPSE-based detection scheme has excellent performance against high rate attacks, since it yields very high detection probabilities.

3.2 Detection of long-term low-rate attacks

An important issue in detecting attack traffic is when it is aggregated with a large amount of additional traffic. Intuition would say that the attacks with higher rates can be

Table 2: Detection of short-term high-rate attacks

Attack Rate (packets/s)	200 packets	400 packets	600 packets
100	90%	100%	100%
200	90%	100%	100%
300	95%	100%	100%
400	100%	100%	100%
500	100%	100%	100%
600	100%	100%	100%

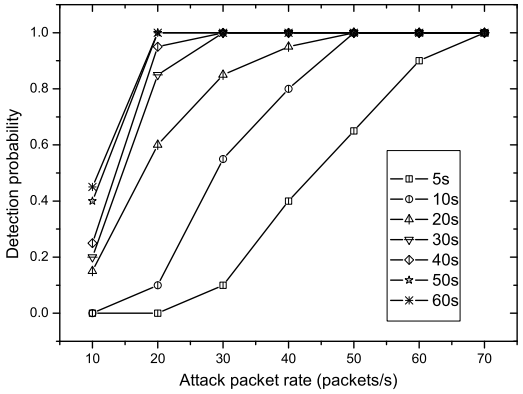

Figure 6: Detection probability for different attack durations

detected with higher certainty. Our experiments consider the effect of the packet rate of attacks on detection performance. For different experiments, the duration of one attack is set from 5 seconds to 60 seconds and the attack rate varies from 10 packets/s to 70 packets/s. The packet size is set to 40 bytes, which is the same as that of SYN flooding attack traffic.

Figure 6 shows the impact of attack rate and attack duration on the detection probability. In the figure, the curve "αs" denotes the detection probability when the attack traffic duration is set to α seconds. These results indicate that the detection probability increases as attack rate and attack duration increase.

Figure 7 shows the tradeoff between the detection time and detection probability for different attack packet rates. Note that a curve "βpps" corresponds to the attack traffic with a rate of β packets/s. The analysis shows that it takes a long time to ensure a high detection probability.

4. CONCLUSION AND FUTURE WORK

In this paper, we described an IP packet size entropy (IPSE)-based DoS detection scheme, which was capable of detect not only long-term attacks but also short-term attacks which did not cause abrupt changes in traffic volume by observing the time series of the entropy of packet size.

Although our proposal can not detect all attacks completely (actually no scheme can), it is a novel approach with a simple implementation for DoS detection. If a wily attacker knows our detection scheme and modifies his strategy,

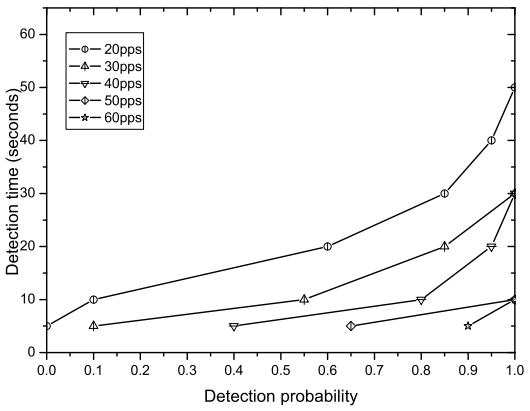

Figure 7: Tradeoff between detection time and detection probability for different attack packet rates

there is still a possibility to find new differences between the packet size distributions of the human-participating legitimate traffic and the machine-automating attack traffic. For example, when the attacker generate packets with randomized sizes, the packet size entropy of the attack traffic will be expected much larger than that of legitimate traffic. Our future work is to detect more stealthy attacks.

5. REFERENCES

[1] A. Hussain, J. Heidemann, and C. Papadopoulos, "A Framework for Classifying Denial of Service Attacks", Proc. ACM SIGCOMM 2003.

[2] J.Haggerty, T.Berry, Q.shi, and M.Merabik, "DiDDeM: a system for early detection of TCP SYN flood attacks," Proc. IEEE GLOBECOM 2004, pp.2037–2042.

[3] A. Lakhina, M. Crovella, and C. Diot, "Mining Anomalies Using Traffic Feature Distributions," Proc. ACM SIGCOMM 2005.

[4] P. Barford, J. Kline, D. Plonka, and A. Ron, "A signal analysis of network traffic anomalies," Proc. ACM SIGCOMM InternetMeasurement Workshop 2002.

[5] H. Wang, D. Zhang, and K. Shin, "Detecting SYN Flooding Attacks," Proc. IEEE Infocom 2002, no. 1, pp. 1530–1539, June 2002.

[6] A. Dainotti, A. Pescape, and G. Ventre, "Wavelet-based Detection of DoS Attacks," Proc. Globecom 2006, vol. 25, no. 1, pp. 1452–1457, Nov. 2006.

[7] R. Lippmann, et al., "The 1999 DARPA Off-line Intrusion Detection Evaluation", Computer Networks 34(4) 579–595, 2000. Data is available at http://www.ll.mit.edu/IST/ideval/

[8] R. Sekar, A. Gupta, J. Frullo, T. Shanbhag, S. Zhou, A. Tiwari and H. Yang, "Specification Based Anomaly Detection: A New Approach for Detecting Network Intrusions", Proc. ACM CCS, 2002.