Journal of Southeast University (English Edition) Vol.20 No.1 Mar. 2004 ISSN 1003—7985

Anomaly detection for network traffic flow

Shan Rongsheng Li Jianhua Wang Mingzheng

(Department of Electronic Engineering, Shanghai Jiaotong University, Shanghai 200030, China)

Abstract: This paper presents a novel mechanism for detecting flooding-attacks. The simplicity of the mechanism lies in its statelessness and low computation overhead, which makes the detection mechanism itself immune to flooding-attacks. In this paper, SYN-flooding, as an instance of flooding-attack, is used to illustrate the anomaly detection mechanism. The mechanism applies an exponentially weighted moving average (EWMA) method to detect the abrupt net flow and applies a symmetry analysis method to detect the anomaly activity of the network flow. Experiment shows that the mechanism has high detection accuracy and low detection latency.

Key words: anomaly detection; intrusion detection; denial of service; port scan

The Internet has undergone phenomenal growth in the past decades and become one of the indispensable infrastructures of our society. However, the Internet is prone to network attacks because of the inherent vulnerability of TCP/IP protocol suite. The popular web sites often suffer from denial of service (DoS) attacks. It was shown that more than 90% of the DoS attacks used TCP[1]. Detecting the abnormal variations of network traffic is a necessary way to detect and defend DoS attacks because they always lead to variation of characteristics of network traffic. Currently, many network administrators identify these attacks through analyzing net-flow anomalies. Visual analysis of traffic flow anomalies has grouped these anomalies into two general categories:

- **Network operation anomalies** They include network abnormal behavior caused by configuration changes and caused by traffic reaching environmental limits.
- **Network abuse anomalies** Two types of network abuse that can be identified are DoS flooding-attacks and fast port-scans. DoS attack is a malicious behavior intended to cripple an online service. Fast port-scan is also a prevalent attack to glean useful information of the target. Our research focuses on network abuse anomalies.

Several mechanisms have been proposed to counter SYN-flooding, such as SYN cache[2], SYN cookies[3], SYN defender[4], SYN proxying[5], and SYN-kill[6]. However, these mechanisms are dependent on states, i.e., states are maintained for each TCP connection and state computation is required. Experiments with SYN attacks on commercial platforms show that the flooding rate to overwhelm an unprotected server is 14 000 SYN packets/s. If the detection system based on state maintaining is integrated into these servers, it will intensely drop the end-to-end performance of TCP and increase the delay of creating a TCP connection.

Available methods to detect fast SYN-scan are also state maintaining or computation consuming, such as snort[7] and packet header detection (PHD)[8]. A snort port-scan preprocessor must maintain the state information of each TCP session, and the detection system itself is subjected to DoS attack. PHD has low detection rate since it doesn't take into account the TCP flag and our experiment shows that it also has a high negative alarm rate.

The above solutions based on the TCP state or complex computation make the defense mechanism itself subject to flooding-attacks. This paper presents a simple and efficient method to detect flooding-attacks. In our method, we use time series analysis, which is based on an exponentially weighted moving average (EWMA), and traffic symmetry analysis. It can accurately detect those attacks that produce abrupt net-flows, such as SYN-flooding, and fast SYN-scan, etc. Experiments based on Defense Advanced Research Projects Agency (DARPA) evaluation data shows that the proposed detection mechanism has high accuracy and low latency.

1 Flooding Detection Based on Flow Anomaly

According to the attributes of network abuse anomalies, we select two kinds of measures that are effective in detecting flooding-attacks.

- **Intensity measures** These measures track the number of attribute records that occur in a fixed-time

Received 2003-07-15.

Foundation item: The National High Technology Research and Development Program of China(863 Program) (No.2002AA145090).

Biographies: Shan Rongsheng (1971—), male, graduate; Li Jianhua (corresponding author), male, professor, lijh888@sjtu.edu.cn.

interval, and can detect abnormal bursts of activity.

• **Symmetry measures** These measures describe the symmetry between two attributes in a fixed-time interval, and can be used to detect the anomaly **activity with broken symmetry.**

1.1 Intensity anomaly detection

To detect a network traffic anomaly in real-time, we use an intensity measure of network behaviors to describe the degree of busyness of the network. The higher the intensity is, the busier the network is. Under normal conditions, the number of network events that occur in a fixed-time interval varies with. If a network abuse anomalous event occurs, the intensity measure will burst.

Recently, there has been considerable work on the modeling of the arrival process of TCP connection requests. It is reported in Ref.[9] that the statistics of TCP connection request arrivals have shown significant changes in the past few years, along with Internet traffic itself: in the early 1990's, the dominant TCP connections were FTP and Telnet sessions, and the arrival process was Poisson[10]. However, after the web became the predominant source of TCP connections, the arrival process displays heavy tails in its inter-arrival time[11]. Furthermore, recent Internet traffic analyses have shown that arrival process is not even stationary and dependent on the average arrival rate[12]. For such a dynamic and complicated entity like the Internet, it may not be possible to model the total number of TCP connections at all times with a simple parametric model.

EWMA techniques can be univariate or multivariate, and detect changes in process mean (mean shifts), process variance (variance changes), and relationships among multiple variables (counter-relationship)[13]. This paper focuses on detecting significant changes of event intensity for intrusion detection. The event intensity is a single variable measuring the characteristics of events in an information system. Hence, this paper considers only univariate techniques to detect anomalies or possible intrusions.

So we apply time series analysis based on EWMA to illustrate the theory of intensity anomaly detection. We use the SYN sequence in the DARPA IDS evaluation dataset as an example to study the EWMA detection method. The DARPA dataset is provided by MIT Lincoln Lab, and widely used in evaluating IDS algorithms and systems[14].

Let x_t denote the rate of SYNs that occur in a fixed time interval. Through analyzing the normal data of the first and the third week of the DARPA 1999 intrusion detection evaluation data, we find x_t is not a stationary process. Fig.1 describes the statistic data of SYN on the first day of the first week. The N axis describes the sampling interval sequence. Fig. 2 describes the SYN sequence after differencing.

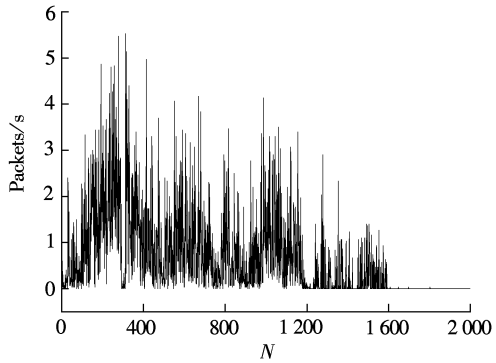

Fig.1 SYN sequence

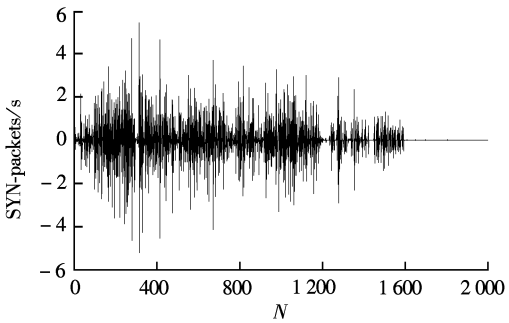

Fig.2 SYN sequence after differencing

Through differencing the SYN sequence, we get a stationary process z_t, namely

$$(1-B)x_t = z_t \qquad t>1,\ z_t \sim N(0,\ \sigma_t^2) \tag{1}$$

where B is the backshift operator, z_t is a normal process with zero mean and σ_t variance. The upper and lower control limit of z_t, UCL_t and LCL_t, are individually as follows:

$$\mathrm{UCL}_t = 3\hat{\sigma}_t,\ \mathrm{LCL}_t = -3\hat{\sigma}_t \tag{2}$$

In Eq.(1), we evaluate a variance σ_t from the history data of z_t, and usually use the N history data to evaluate the value of σ_t:

$$\tilde{\sigma}_t = \sqrt{\frac{1}{N}\sum_{j=t-N-1}^{t} z_j^2} \tag{3}$$

This method of evaluating standard deviation is simple moving average. The method gives an equal weight ($1/N$) to each observation, and the evaluating value is apparently dependent on the data length N.

But the practical observed data, such as SYN, SYNACK measures, changed largely with time shift. Each observation has a close relationship with the adjacent observations, and otherwise has little relationship with the distant observations. We use

EWMA to evaluate the value of σ_t, namely:

$$\hat{\sigma}_t = \sqrt{(1-\lambda)\sum_{j=0}^{\infty}\lambda^j z_{t-j}^2} \tag{4}$$

where $\lambda\ (0<\lambda<1)$ is a decay factor whose value determines the weight of each observation along with the effective sample length. The method gives a different weight to each observation, giving more weight to recent observations and less weight to older observations. Another characteristic of EWMA is that the evaluation equation of the variance is transformed to a recursive equation, which is suitable for a computer dealing with mass data. The evaluation value of the parameter σ_i, is given by

$$\hat{\sigma}_t^2 = \lambda\hat{\sigma}_{t-1}^2 + (1-\lambda)z_t^2 \tag{5}$$

Because the scope of the value of λ is $0<\lambda<1$, the weight of z_{t-j}^2 is $(1-\lambda)\lambda^j \to 0$ (when $j\to\infty$). So, the tolerance L_K is defined as

$$L_K = (1-\lambda)\sum_{j=K}^{\infty}\lambda^j = \lambda^K \tag{6}$$

where K is the effective length of observations. Eq.(6) classifies the relations among L_K, K and λ. So, when L_K is given a fixed value, the burst judgment function of the abrupt intensity measure is defined as

$$f(z_t) = \begin{cases} 1 & z_t > 3\sqrt{(1-\lambda)\sum_{j=0}^{K-1}\lambda^j z_{t-j}^2} \\ 0 & |z_t| < 3\sqrt{(1-\lambda)\sum_{j=0}^{K-1}\lambda^j z_{t-j}^2} \\ -1 & z_t < -3\sqrt{(1-\lambda)\sum_{j=0}^{K-1}\lambda^j z_{t-j}^2} \end{cases} \tag{7}$$

During testing, if z_t value for an observation exceeds the upper control limit UCL, the $f(z_t)$ value is set to 1 and a rise-jump signal of a burst flow is generated; and if z_t value for an observation is under the lower control limit LCL, the $f(z_t)$ value is set to -1 and a drop-jump signal of the burst flow is generated; otherwise the $f(z_t)$ value is set to 0 and no burst occurs.

Many anomalous network behaviors can last for a period of time. To avoid frequent alerts, a definite beginning time and end time of the anomalous event must be given. The rise-jump and the drop-jump of the burst flow are helpful in confirming the beginning and the end time of the anomalous event.

1.2 Symmetry anomaly detection

Symmetry is an obvious phenomenon in many net-flows. For example, under normal conditions, the discrepancy between the collected number of SYNs and SYNACKs is very small, compared to the total number of TCP connection requests. Furthermore, the one-to-one match between SYN and SYNACK is independent of the sample time and network sites.

Let y_t denote the rate of SYNACKs that occur in a fixed time interval. Obviously, under normal network conditions, the value of x_t is equal to y_t. Fig.3 describes the statistic data of SYNs and SYNACKs of the first day of the DARPA 1999 evaluation data.

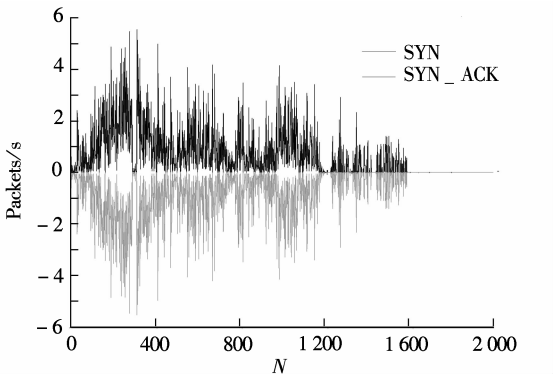

Fig.3 The symmetry of SYNs and SYNACKs

We define the symmetry of TCP SYN and SYN/ACK as

$$\varepsilon_t = \frac{|x_t - y_t|}{x_t} \tag{8}$$

Under a normal network environment, ε_t is close to 0. But a misconfigured system and incorrect operation will lead to a shift in the value of ε_t. Furthermore, when an anomalous event occurs, the symmetry of traffic will be badly broken and ε_t will be far from zero. In order to ensure that the value of ε_t can efficiently represent the symmetry balance capability of the observed measures, and distinguish between network operation anomalies and network abuse anomalies, we must give a practical definition of the symmetry threshold. Through analyzing the practical net-flow, we find that when the flow is small, anomalous events hardly ever happen. Furthermore, the heavier the flow is, the higher the probability of anomalous events happening is. So we design an adaptive symmetry detection algorithm, in which the threshold value of the symmetry can be adjusted according to the changes of the network traffic flow. The auto-adapting threshold is defined as

$$s_t = \frac{\beta\ln x_t}{x_t} \qquad s_t \in [0,\ 0.3] \tag{9}$$

where β is an adjustable parameter. In principle, we consider that the symmetry is obviously broken if ε is more than 0.3. Obviously, $\beta \leqslant \frac{0.3x_t}{\ln x_t}$. Let $\beta = \frac{0.3\alpha}{\ln\alpha}$, where α is the max security control limit of net flow,

and let 98% of total normal $x_t \in [0, \alpha]$. When the intensity of net-flow is less than the value of α, anomalous events hardly happen. The symmetry detection function is

$$f(y_t, x_t) = \frac{|\varepsilon_t|}{s_t} = \frac{|x_t - y_t|}{\beta \ln x_t} > 1 \tag{10}$$

If the above expression is true, the symmetry is broken.

In order to avoid the effect of the time interval, we must optimize the above algorithm and design a cumulative detection. When there is a rise-jump of a flow x in time t_0, and the event maintains N sampling interval, let

$$\varepsilon_i = \frac{\sum_{t=t_0}^{t_0+i} |x_t - y_t|}{\sum_{t=t_0}^{t_0+i} x_t}, \quad s_i = \frac{\beta \log \sum_{t=t_0}^{t_0+i} x_t}{\sum_{t=t_0}^{t_0+i} x_t}$$

where $0 \leqslant i \leqslant N$. So the judgment function is

$$f_i(y_t, x_t) = \frac{\sum_{t=t_0}^{t_0+i} |x_t - y_t|}{\beta \log \sum_{t=t_0}^{t_0+i} x_t} > 1 \tag{11}$$

If the above expression is true, the symmetry is broken. Obviously, Eq. (10) is a special case of Eq. (11) when $i = 0$.

2 Experiment

To evaluate and validate our method, we have conducted simulation experiments on DARPA 1999 evaluation data[14]. We record the number of SYN and SYNACK packets during every observation period t_0, which determine the detection resolution. As usual, the first retransmission interval of SYN is 6 s, the second retransmission interval of SYN is 24 s[15], so we set $t_0 = 30$ s. However, the parameter is tunable and our algorithm is not very sensitive to this choice.

After training the normal DRAPA 1999 evaluation data, we find the 98% value of total x belongs to $[0, 3]$. So let $\alpha = 3$, then $\beta \approx 0.81$. Furthermore, because there is a close relationship among the adjacent observations, let $\lambda = 0.85$ and the minimum of tolerance $L = 0.001$. We can get $K \approx 42$ from Eq.(6).

We have conducted detection experiment on the first day of the 5th week of the DARPA 1999 evaluation data. The experimental result is described in Fig.4. Fig.4(a) describes the intensity detection of the differencing SYN sequence; Fig.4(b) describes the symmetry detection between SYNs and SYNACKs. The detected attacks are given in Tab.1.

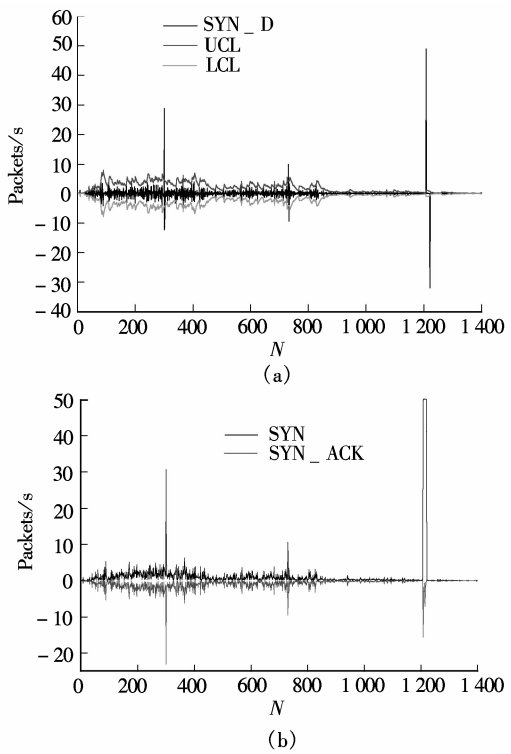

Fig. 4 The first day of DARPA evaluation data. (a) Intensity detection of SYN_D; (b) Symmetry detection between SYNs and SYNACKs

Tab.1 The detection results

Index	Attack name	Start time	End time
1	Apache2	10:29:51, Apr. 5, 1999	10:30:51, Apr. 5, 1999
2	Apache2	14:05:45, Apr. 5, 1999	14:06:15, Apr. 5, 1999
3	Neptune	18:04:26, Apr. 5, 1999	18:10:26, Apr. 5, 1999

In the experiment, our algorithm detects three flooding attacks, where the first Apache2 attack happened on 10:29:51, April 5, 1999, the second Apache2 attack happened on 14:05:45, April 5, 1999. A Neptune attack happened on 18:04:26, April 5, 1999. The Apache2[14] attack is a DoS attack against an apache web server where a client sends a request with many http headers. If the server receives many of these requests it will slow down, and may eventually crash. The Neptune[14] attack is a typical SYN flood DoS on one or more ports. But, the algorithm can't detect a scan attack which happened on 9: 43: 11, April 5, 1999. The attack lasts 4 min with 10 s interval between every attack packet. It isn't a flooding attack. In allusion to the slow and stealthy scan attack, we use the active detection technique which is described in another paper. There is a discrepancy between the beginning time of the attack detected by our algorithm and the real result, which is less than a sample interval.

3 Conclusions

1) The algorithm is stateless and requires low computation overhead, which makes it immune to flooding attacks;

2) The algorithm is insensitive to access pattern;

3) Cumulative detection is employed, which makes the detection robust;

4) Intensity and symmetry measures are applied simultaneously, which make the detection exact.

The experimental results show that both the accuracy and the real-time are very good. Furthermore, our method can be extended to detect other flooding attacks that can break the traffic symmetry, such as synKill, Ipsweep, NMAP fast scan, and SATAN, etc.

References

[1] Moore D, Voelker G, Savage S. Inferring Internet denial of service activity [A]. In: *Proceedings of the 10th USENIX Security Symposium* [C]. Washington DC, 2001. 9－22.

[2] Lemon J. Resisting SYN flooding DoS attacks with a SYN cache [A]. In: *Proceedings of USENIX BSDCon*'2002 [C]. San Francisco, 2002.89－97.

[3] Bernstein D J. SYN cookies [EB/OL]. http: //cr.yp.to/syncookies.html. 2000/2003-07-08.

[4] Check Point Software Technologies Ltd. SynDefender [EB/OL]. http: //www. checkpoint. com/products/protect/firewall-1.html. 2002/2003-07-08.

[5] Netscreen Technologies Ltd. Firewall appliance[EB/OL]. http: //www.netscreen.com/. 2002/2003-07-08.

[6] Schuba C L, Krsul I V, Kuhn M G, et al. Analysis of a denial of service attack on TCP[A]. In: *Proceedings of IEEE Symposium on Security and Privacy* [C]. Los Alamitos: IEEE Computer Society Press, 1997. 208－223.

[7] Roesch M. Snort-lightweight intrusion detection for networks[A]. In: *Proceedings of the* 13*th Conference on Systems Administration (LISA* '99) [C]. Seattle, Washington,1999. 229－238.

[8] Staniford S, Hoagland J, McAlerney J. Practical automated detection of stealthy portscans [A]. In: *ACM Computer and Communications Security IDS Workshop* [C]. Athers, Greece, 2000.1－7.

[9] Feldmann A. Characteristics of TCP connection arrivals [A]. In: Park K, Willinger W, eds. *Self-Similar Network Trac and Performance Evluation* [C]. John Wiley and Sons, 2000. 367－399.

[10] Caceres R, Danzig P B, Jamin S, et al. Characteristics of wide area TCP/IP conversations[A]. In: *Proceedings of ACM SIGCOMM* '91 [C]. Zurich, Switzerland, 1991. 101－112.

[11] Paxson V, Floyd S. Wide area traffic: the failure of poisson modeling [J]. *IEEE/ACM Transactions on Networking*, **1995, 3**(3) 226－244.

[12] Cleveland W S, Lin D, Sun D. IP packet generation: statistical models for start times on connection-rate superposition [A]. In: *Proceedings of ACM SIGMETRICS* [C]. California, 2000. 166－177.

[13] Bowerman B L, O'Connell R T. *Forecasting and time series: an applied approach*. 3rd ed. [M]. Thomson, 2003.

[14] Lincoln Laboratory, Massachusetts Institute of Technology. DARPA intrusion detection evaluation [EB/OL]. http: //www. ll. mit. edu/IST/ideval/index. html. 2001/2003-07-08.

[15] Stevens R W. *TCP/IP Illustrated Volume: The protocols* [M]. New York: Addison-Wesley, 1994.178－179.

网络流量异常检测

单蓉胜　李建华　王明政

(上海交通大学电子工程系, 上海 200030)

摘要: 提出了一种新颖的网络洪流攻击的异常检测机制.这种检测机制的无状态维护、低计算代价的特性保证了自身具有抗洪流攻击的能力.本文以检测 SYN 洪流攻击为实例详细阐述了异常检测机制.这个机制应用 EWMA 方法检测网络流的突变, 并运用对称性分析方法检测网络流的异常活动.测试结果表明本文所提出的检测机制具有很好的检测洪流攻击的准确度, 并具有低延时特性.

关键词: 异常检测; 入侵检测; 拒绝服务攻击; 端口扫描

中图分类号: TP393