TCP Flow Analysis for Defense against Shrew DDoS Attacks

Yu Chen and Kai Hwang*

University of Southern California, Los Angeles, CA 90089, USA

Abstract - The shrew or RoS attacks are low-rate DDoS attacks that degrade the QoS to end systems slowly but not to deny the services completely. These attacks are more difficult to detect than the flooding type of DDoS attacks. In this paper, we explore the energy distributions of Internet traffic flows in frequency domain. Normal TCP traffic flows present some form of periodicity because of TCP protocol behavior. Our results reveal that normal TCP flows can be segregated from malicious flows using some energy distribution properties. We discover the spectral shifting of attack flows from that of normal flows. Combining flow-level spectral analysis with sequential hypothesis testing, we propose a novel defense scheme against shrew DDoS or RoQ (reduction-of-service) attacks. Our detection and filtering scheme can effectively rescue 99% legitimate TCP flows under the RoS attacks.

Keywords: Network security, Internet traffic spectrum, low-rate DDoS attacks, reduction-of-quality attacks, digital signal processing, spectral analysis.

1. Introduction

Understanding Internet traffic patterns is undeniably a critical task to secure network management and resource optimization. In order to detect network attacks and respond swiftly, the defense system must be designed to distinguish anomalies embedded in legitimate traffic. This requires solid understanding of the Internet traffic patterns of variant flows under different protocols such as TCP and UDP. More than 85% of Internet traffic and most DDoS attacks apply the TCP protocol [15]. However, it is infeasible to segregate legitimate TCP flows from malicious attack traffic flows just using protocol information in the packet header.

Typical DDoS flooding attacks are characterized by sustained high-rate or high volume. Recently, a variant of DDoS attacks has been identified that is even more difficult to detect. They are called *shrew* attacks [5], [21] or *Reduction-of-Service* (RoS) attacks. These attacks do not denial the clients from services completely [9]. Instead, they throttle the TCP throughput heavily and reduce the quality of service gradually. These new attacks are also known as low-rate TCP targeted DDoS attacks [21] or *reduction of quality* (RoQ) attacks [13] by other researchers.

Due to different protocols or applications, the periodicity of traffic could be used as a reliable signature for traffic monitoring or anomaly detection [7]. Indeed, when the powerful DDoS attack tools can generate packets and manipulate the header information, it is non-trivial for attackers to control the statistic properties of attack traffic to mimic the behaviors of legitimate flows.

Figure 1 shows a network attacking scenario. The RoS attacks are launched in a similar fashion as the DDoS attacks using handlers and zombies. Both legitimate and attack flows are initiated from edge networks. The victim systems are also at the edge network. The core network of routers used by the *Internet service provider* (ISP) is shown in the middle of Fig.1. We simulate this ISP core network environment to test the RoS defense scheme proposed.

This research is part of our efforts to integrate DSP core techniques with reconfigurable hardware for Internet traffic analysis. Our goal is to enable routers to detect the network anomalies quickly and respond swiftly. We need to push the traffic flow analysis work to a lower layer of the packet processing procedure. Then spectral analysis is carried out in parallel with regular router functionalities.

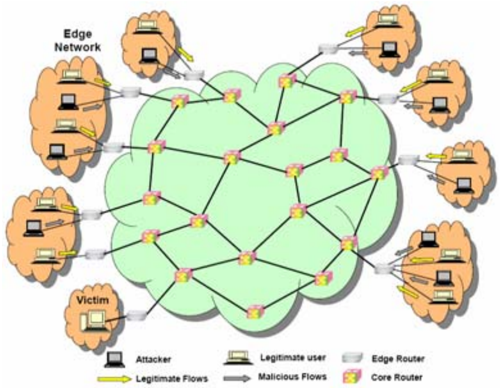

Figure 1 Reduction-of-service (RoS) attack scenario on end systems at Internet edge networks

In this paper, we use spectral analysis to establish a *traffic spectrum* that describes the behavior of Internet traffic flows using frequency-domain characteristics. In contrast to earlier reports, we reveal frequency-domain characteristics of Internet traffic at the flow level. We identify the frequency components that lead to network anomalies. As a case study, we verify the effectiveness of spectral analysis for defense against shrew or RoS attacks.

The rest of the paper is organized as follows: Section 2 gives a brief review of related works. Section 3 discusses the rationale, methodology and results of the quantification of traffic spectrum. A spectral hypothesis testing scheme is specified in Section 4 to detect and cut off RoS attacks. Section 5 presents the NS2 simulation results using our detection and filtering algorithms. Finally, we summarize the contributions and conclude this paper.

***** Manuscript accepted for presentation at *IEEE International Conference on Communications* (ICC 2007), Glasgow, Scotland, UK, June 24-28, 2007. The research work reported here was supported by a NSF ITR Grant 0325409. Corresponding author: Kai Hwang, USC Internet and Grid Research Lab, EEB 212, Los Angeles, CA 90089. E-mail: kaihwang@usc.edu, Tel.: (213) 740-4470.

2. Related Works

The Internet research community has recognized the efficacy of using *digital signal processing* (DSP) techniques to detect DDoS attacks [5], [6], [7], [12], and [17]. These technologies sample the number of packet arrivals periodically and treated it as a time domain signal series. Using traffic spectrum information, we propose a novel approach that can distinguish attack flows from legitimate flow. Enlightened by sequential hypothesis testing theory [4], [19], we check the traffic energy distribution over frequency bands instead of using time series.

The motivation of spectral analysis is that normal TCP flows must present some form of periodicity in packet transmission and the periods are related to the *round-trip times* (RTTs). Fourier transform then manifests as an effective tool that allows researchers to examine traffic properties in frequency domain. In general, signal processing methodologies applied in network traffic analysis could be divided into three categories: (1) *spectral analysis,* (2) *wavelet analysis*, and (3) *statistical anomaly analysis.*

He *et al.* [15] indicated that dominant link frequency is independent of the number of flows. Instead, it depends on the link bandwidth and packet size distribution. Giles *et al.* [12] suggested detect DoS attacks using spectral analysis techniques on the backscatter packets. Considering that TCP traffic exhibits a periodicity on its *Power Spectrum Density* (PSD), the lack of periodicity could indicate that DoS attacks are raging on [7]. Furthermore, the PSD of multi-sourced DDoS attacks are distributed in lower frequency band comparing to single-sourced DoS attacks [17].

Previously, we proposed to detect the low-rate TCP-targeted DDoS attacks on aggregated traffic level based on the energy distribution shifting and template matching technique [6]. We suggested to filter shrew attack flows by focusing on the extraordinary high energy allocated at the low frequency band [5]. Sun *et al.* [24] also proposed to extract the signature of low-rate attack streams through analyzing the PSD of the autocorrelation sequence.

Wavelet has been adopted to analyze Internet traffic properties since late 1990s [1]. The fundamental idea behind wavelets is to analyze traffic data using different scales or resolutions. Wavelet system can effectively isolates both short and long-lived traffic anomalies. Barford *et al.* [3] used wavelet to decompose the signal into low-frequency part, mid-frequency part and high-frequency part. Thus variant resolutions of anomalies detection are achieved by adjust wavelet parameters.

Huang *et al.* [16] developed a prototype tool based on wavelet to detect network performance problems. Kim *et al.* [20] proposed to use *wavelet de-noising* method to separate queuing delay caused by network congestion from other delay variations. Luo and Chang [23] have studied the characteristics of low-rate DDoS attack using a wavelet approach. They observed anomalies in fluctuation of incoming traffic rate and declining of outgoing TCP ACKs incurred by pulsing streams.

Statistical anomaly analysis is another approach that is insensitive to the header spoofing. He and Hou [14] present the influence of different sampling techniques with an in-depth study of three sampling techniques, and proposed a novel static systematic sampling method.

Feinstein *et al.* [11] suggested identifying DDoS attacks by computing the entropy of certain packet attributes. Li *et al.* [22] studied the dynamic behaviors of statistical-based filtering technique against DDoS attacks, their result presents that such a filtering scheme may perform bad if the attacker is more dynamic than perceived. Thottan and Ji [25] described a statistical signal processing technique based on abrupt change detection.

This paper explores much deeper in Internet traffic flow properties in the frequency domain. Beyond the low frequency band, we extend the study to cover the characteristics of real Internet traffic both with periodicity (TCP) and without periodicity (UDP/ICMP). We propose a novel concept *traffic spectrum* to describe the quantified energy distribution over the entire frequency band. Enlightened by the sequential multi-hypotheses testing theory [4], we design a *spectral hypothesis testing* (SHT) algorithm to characterize Internet traffic spectrum.

3. Internet Traffic Flow Analysis

This section introduces the Internet trace dataset used and analyzes the traffic energy distribution. Then we specify the RoS attack detection and flow filtering scheme in Section 4.

A. *Abilene-III Internet Traffic Traces*

The traffic dataset used in our experiments is the Abilene-III Internet trace data [26]. This dataset is the first publicly available trace of 10 Gigabit Internet backbone traffic. It was collected on June 01, 2004 at the OC192c Packet-over-SONET link from Internet2's Indianapolis Abilene router node towards Kansas City. The OC192MON hardware supports both time stamping and global CDMA/GPS synchronization [27].

Although the source IP addresses are often spoofed in attack packets, for our purpose, it is legitimate to use the 5-tuple {*Source IP, Source Port, Destination IP, Destination Port, Protocol*} as a flow identifier. Using the Abilene-III trace data, we identified more than 20,000 TCP flows, more than 2,000 UDP flows and several flows on top of other protocols. For each flow, we obtained its time series of packet arrival sequence numbers using the globally synchronized time stamp of each packet.

B Energy Spectral Distribution

Because of the traffic/congestion control mechanisms of protocols on which applications are implemented and RTT, we observe the periodicity in packet arrivals on traffic flows. Periodic signals consisting of different frequency ingredients present different properties. The energy distribution patterns over frequencies may vary from protocol to protocol if the traffic flows are viewed in frequency domain. These variants could be detected conveniently using digital signal-processing techniques.

Therefore, we take the number of received packets of each flow as the signal and sample it every 1 *ms*. Nyquist sampling theorem indicates that the highest frequency of our analysis is 500 Hz. The packet arrivals are modeled by a random process: $\{X(t), t = n\Delta, n \in N\}$, where *Δ* is a constant time interval chosen as 1 *ms*, *N* is the set of positive integers, and for each *t*. *X*(*t*) is a random variable that represents the total number of packet arrivals at one router in (*t-Δ, t*]. This random process is referred to as *packet process* [7]. We assume a *Wide Sense Stationary* random process. The autocorrelation function of the discrete random signal *X*(*t*) is defined by:

$$R_{xx}[m] = \frac{1}{N-m} \sum_{n=0}^{N-m+1} (x[n]x[n+m]) \qquad \textbf{(1)}$$

$R_{xx}[m]$ is the correlation of the packet process and itself at interval *m*. Autocorrelation function is capable of enforcing periodicity. Then we compute the *Power Spectral Density* (PSD) of the flow by converting the autocorrelation into frequency domain using *discrete Fourier transform* (DFT):

$$DFT(R_{xx}(m), K) = \frac{1}{N} \sum_{n=0}^{N-1} R_{xx}(m) \times e^{-j2\pi kn/N} \qquad \textbf{(2)}$$

where K=0,1,2…N-1.

Figure 2(a) is the PSD of a TCP flow randomly drawn from the Abilene-III trace dataset. TCP flow presents a clear periodicity, and the positions of peaks are related to the RTT of the communication.

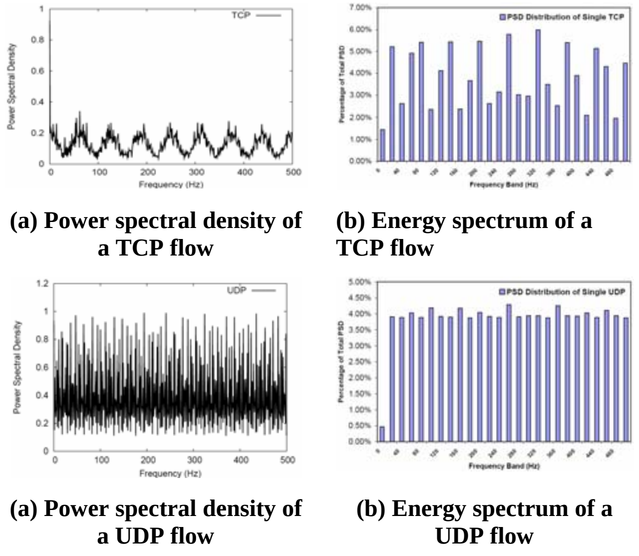

(a) Power spectral density of a TCP flow

(b) Energy spectrum of a TCP flow

(a) Power spectral density of a UDP flow

(b) Energy spectrum of a UDP flow

Figure 2. Comparison of Power spectrum density and energy distribution between TCP and UDP flows.

The width of each frequency band is set to 20 Hz. That is DC (0 Hz), (0, 20 Hz], (20 Hz, 40 Hz], ···, (480 Hz, 500 Hz]. The convention (*a, b*] implies the range greater than *a* and smaller than or equal to *b*. Figure 2(b) compares the variant energy percentages over different frequency bands of the same TCP flow.

The strength of PSD of TCP flows is varied over frequencies. Without considering the DC band ingredient, the *standard deviation* (STD) of the UDP flow (about 0.1%) is only approximately 1/10 of the STD of the TCP flow (about 1.2%). Thus, beside the strength of energy over certain frequencies, the standard deviation of energy variance could be used as the second signature to segregate TCP flows from UDP flows.

C. Traffic Spectrum Characteristics

Since such PSD pattern resembles the visible light spectrum that different colors are presented in different frequencies, we call it *traffic spectrum*. In this paper, we established a framework in which Internet traffic can be categorized in frequency domain based on the analysis of:

- Pattern of energy over the frequency bands;
- Standard deviation of such energy distribution.

The first metric reveals the energy distribution of the flows over frequencies. The second metric describes statistically whether one type of flow may biased to certain special frequencies. As shown in Fig.2 (a, b), TCP flows locate more energy to some frequency band than others. This spectrum describes the flow level energy distribution.

Although the exact probability distribution of Internet flows' energy on different frequency bands are not available, *Central Limit Theorem* indicates that if the sample space is large, the sampling distribution approaches a *Gaussian* (*Normal*) distribution with mean *μ* and variance σ^2. Thus, we can describe energy distribution using normalized amplitude spectrum at different frequency bands using Gaussian distribution model:

$$G(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left\{-\frac{(x-\mu)^2}{2\sigma^2}\right\} \qquad \textbf{(3)}$$

In the Abilene-III trace dataset, TCP packets accounts 88%, UDP packets for 9%, ICMP packets for 2%. The remaining 1% packets are others. We have analyzed the energy spectra of both TPC and UDP flows. Therefore, our traffic spectra cover majority of traffic on the network. Figure 3 plots the TCP traffic spectrum we obtained from the Abilene-III trace dataset, which contains the energy distribution patterns of more than 20,000 TCP flows.

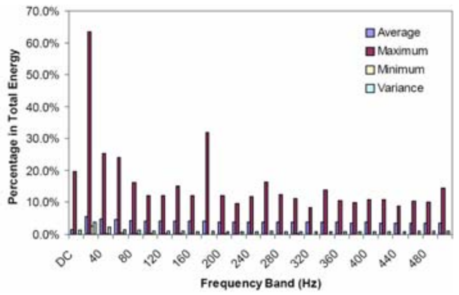

Figure 3. Traffic spectrum over 20,000 TCP flows

Table 1 presents low frequency part of the statistical values of the traffic spectrum we obtained. Beside the mean energy distribution in frequency bands, we also consider ±*3σ* errors representing a confidence level of 99.7%. This is considered a high-precision detection process.

Table 1. Internet spectrum in low frequency band

Flow Type	DC Band		(0, 20] Hz		(20, 40] Hz	
	AVE	±3σ	AVE	±3σ	AVE	±3σ
TCP	1.454%	±0.037	5.66%	±0.120	4.85%	±0.067
UDP	0.447%	±0.021	4.36%	±0.041	4.22%	±0.031
Flow Type	(40, 60] Hz		(60, 80] Hz		(80, 100] Hz	
	AVE	±3σ	AVE	±3σ	AVE	±3σ
TCP	4.49%	±0.050	4.30%	±0.037	4.15%	±0.031
UDP	4.13%	±0.026	4.08%	±0.024	4.03%	±0.021

In general, sample series from UDP flows do not show periodicity as found in TCP flows. UDP flows present an evenly distributed spectrum. Although the average strength of TCP and UDP look similar in Table 1, they are of entirely different. For UDP flows, the energy does distribute evenly over all frequencies. For TCP flows, the average values become similar because peaks of flows are located at different points determined by the RTTs.

It is feasible to segregate normal TCP flows from those aperiodic ones using the distribution of the *standard deviation* (STD) of traffic spectrum. Figure 4 presents the STD's Gaussian distribution and the detection probabilities with given threshold. The *False Negative Rate* (*Fn*) is defined as the probability that TCP flows are categorized as UDP flows; the *False Positive Rate* (*Fp*) is the probability that UDP flows are recognized as TCP flows; the *TCP Detection Rate* (*Fd*) is the probability that a TCP flow is recognized correctly.

Among the flows in Abilene-III trace, we discovered some periodic UDP flows (< 5 %) and some aperiodic TCP flows (< 8 %). This phenomenon leads to the overlapping between the STD curves of TCP and UDP. Higher false positive rate is experienced with higher TCP detection rate.

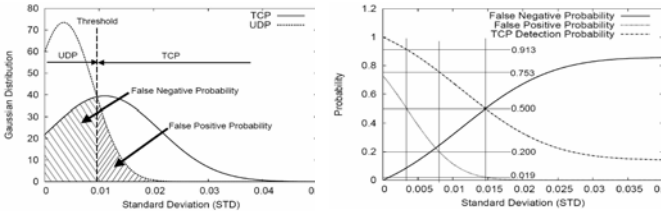

(a) Comparing Gaussian distribution of STD

(b) Segregation probabilities using STD

Figure 4. STD distribution of energy distribution of flows from Abilene-III trace dataset.

4. Defense Scheme Against RoS Attacks

RoS attackers often launch attacks through multiple *zombies* and spoof packet header information to escape from being caught by traceback techniques. However, it is non-trivial to manipulate the frequency domain characteristics of attacking flows. The attacking period has to be close to the *Retransmission Time Out* (RTO) to throttle TCP flows effectively [21], [23]. Even if the source IP addresses carried in packet header are falsified, the energy distribution pattern will betray such malicious flows to detection mechanisms using traffic spectrum.

As a case study, this section applies the traffic spectrum to defend against RoS attacks. Starting with an introduction of the spectrum of RoS attack flows, we present a detection procedure based on the sequential hypothesis testing theory [4] and a filtering algorithm to cut off the RoS flows.

A. Spectrum of RoS Attack Flows

Identified in 2001 [9], the RoS attacks are categorized as the stealthy, harder-to-detect DDoS attacks. Instead of constantly injecting traffic flows with huge rates into the network, attackers send burst pulses periodically. Such low-rate attacks have high peak rate while maintaining a low average rate to exhibit "stealthy" behaviors.

As shown in Fig. 5, a single source RoS attack is modeled as a packet flow with an attack period of *T*, length of the burst *L*, and the burst rate *R*. The period *T* is calculated by the estimated TCP RTO timer of legitimate sources. During the burst with a peak rate *R*, the low-rate pulses create a bursty but severe congestion on the links to the victim. In fact, recent research carried by Chertov *et al.* [8] indicates that variants of RoS attacks can still be effective even if the attack period is not tuned very precise.

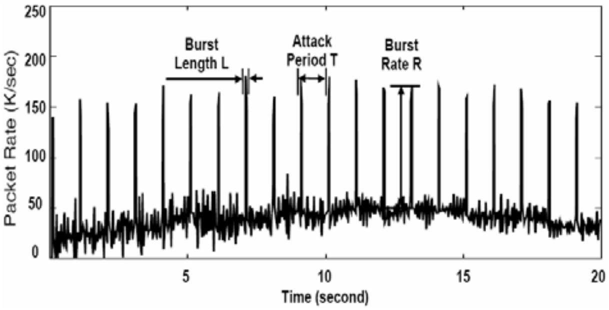

Figure 5. An illustration of typical TCP traffic flows mingled with low-rate RoS attacks.

In distributed scenarios, attacks launched by multiple zombies could lower their individual traffic rates further, thereby making detection much harder. The distributed attack sources could decrease its average traffic rate either by lowering the peak rate or using longer attack periods. Detecting the signs of such attacks using traffic series in time domain is ineffective. However, periodic pulse streams present characteristics totally different from normal TCP flows in frequency domain. The energy spectrum of shrew attack flows Fig. 6 presents a low-frequency distribution.

The major part of the total energy (80% in average) is located in the frequency from 0 Hz to 60 Hz. The frequency bands higher than 200 Hz are not shown in this figure, as less energy will be found in them. The pattern of Fig.6 is calculated over a sample space with 8,000 samples from the NS-2 simulation experiments. We mingled RoS attacks on real-life Abilene-III Internet trace data as shown in Fig.5. Due to page limit, we concentrate on TCP flows. This method can be extended to detect and fence off shrew DDoS attacks embedded in UDP flows as well.

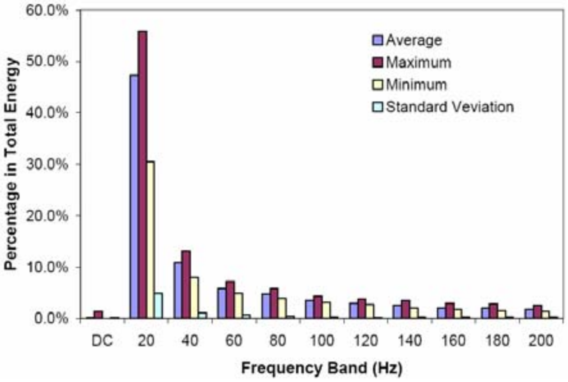

Figure 6. Energy Spectrum of low-rate DDoS attacks

B. Sequential Hypothesis Testing

Recently, sequential hypothesis testing has been introduced into network security research area. In particular, the *Threshold Random Walk* algorithm was proposed to detect malicious scanners by monitoring the sequence of succeeded and failed remote connection request [19]. In contrast to fixed length sampling hypothesis testing, the sequential hypothesis testing reaches a conclusion as soon as available data is enough. It saves computing time since it does not need to wait for the whole sample series.

As shown by Fig.3 and Fig.6, we can distinguish a normal TCP flow from a RoS attack flow using spectral analysis. To make the filtering decision, we propose a *spectral hypothesis testing* (SHT) method. The general sequential hypothesis testing decreases computing overhead by making decision without requiring the whole time series. Our SHT does not need to scan through the entire spectrum. It saves processing time, since it stops scanning as soon as the decision can be made with a desired accuracy.

We consider two hypotheses H_0 and H_1. H_0 is the hypothesis that the flow under investigation is a normal TCP flow. H_1 is the hypothesis that the flow is a RoS attack flow. The input vector Y is the spectrum of the flow. Each item y_i in the vector Y corresponds to the energy percentage of the i-th frequency band.

Given the two hypotheses, there are 4 possible results when decision is made. For our purpose, we define detection rate P_d as the probability that the algorithm selects H_1 when H_1 is in fact true. Also, we define the false positive rate P_f as the probability that the algorithm selects H_0 when H_1 is in fact true. For user specified the desired false positive rate α and detection rate β, we'd like to have:

$$P_f \le \alpha \,,\ P_d \ge \beta \tag{4}$$

In defense against RoS attacks, our goal is to make the decision as quickly as possible. Starting from the DC band, we scan through the spectrum vector Y and calculate a *likelihood ratio function* as follows:

$$L(Y) \equiv \frac{\Pr[Y \mid H_1]}{\Pr[Y \mid H_0]} = \prod_{i=1}^{n} \frac{\Pr[y_i \mid H_1]}{\Pr[y_i \mid H_0]} \tag{5}$$

This function $L(Y)$ indicates to what extent the spectrum is associated with a RoS attack flow. The larger is $L(Y)$, the more likely the flow is a RoS attack. The likelihood ratio is compared with an upper threshold t_l and a lower threshold t_0. The *sequential hypothesis testing* is defined by the following rules:

$$\begin{cases} L(Y) \le t_0 : & \textit{Hypothesis}\ \ H_0\ \ \textit{is}\ \ \textit{True} \\ t_0 < L(Y) < t_1 : & \textit{Say}\ \ "\textit{Need}\ \ \textit{more}\ \ \textit{data}" \\ L(Y) \ge t_1 : & \textit{Hypothesis}\ \ H_1\ \ \textit{is}\ \ \textit{True} \end{cases} \tag{6}$$

Using sequential hypothesis testing, it is not necessary to scan through the entire spectrum. Starting from DC band, we calculate $L(Y)$ and compare it with the thresholds t_0 and t_1. If $L(Y) \le t_0$, the flow is identified as a normal TCP flow. If $L(Y) \ge t_1$, the flow is identified as a RoS attack flow. If the value of $L(Y)$ lies in between, we continue scanning until a conclusion can be drawn. Figure 7 illustrates the sequential decision-making process.

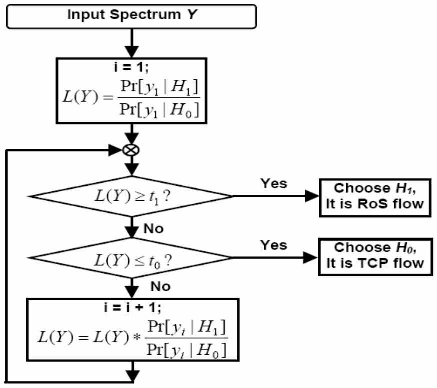
Input Spectrum Y
i = 1;
$L(Y) = \frac{\Pr[y_1 \mid H_1]}{\Pr[y_1 \mid H_0]}$
$L(Y) \ge t_1$?
Yes
Choose H_1, It is RoS flow
No
$L(Y) \le t_0$?
Yes
Choose H_0, It is TCP flow
No
i = i + 1;
$L(Y) = L(Y) * \frac{\Pr[y_i \mid H_1]}{\Pr[y_i \mid H_0]}$

Figure 7. Illustration of spectral hypothesis testing (SHT) process to detect RoS attacks amid legitimate TCP flows.

The thresholds t_1 and t_0 should be chosen such that false positive rate and detection rate satisfies the constraint of Eq.(4). As specified by Eq. (5) and Eq. (6), when hypothesis H_1 is selected, we have:

$$\frac{\Pr[y_1, \cdots, y_n \mid H_1]}{\Pr[y_1, \cdots, y_n \mid H_0]} \ge t_1 \tag{7}$$

The probability $\Pr[y_1,\ldots y_n|H_1]$ is at least t_1 time as big as $\Pr[y_1,\ldots y_n|H_0]$ when H_1 is selected regardless of when the test terminated [19]. The probability measure of where H_1 selected when H_1 is true is at least t_1 times the probability where H_1 is selected when H_0 is true. The first of these probability measures is actually the detection rate P_d, and the second is the false positive rate P_f. Thus, we have the upper bound on threshold t_1:

$$t_1 \le P_d / P_f \tag{8}$$

Similarly, we can deduct the lower bound on threshold t_0:

$$t_0 \ge \frac{1-P_d}{1-P_f} \tag{9}$$

If we choose the thresholds equal to these bounds, then the P_d and P_f are replaced with the user-chosen false positive rate α and detection rate β.

$$\begin{cases} t_1 = \beta / \alpha \\ t_0 = \dfrac{1-\beta}{1-\alpha} \end{cases} \tag{10}$$

Equation (10) indicates that we obtain the detection thresholds t_0 and t_1 with the user-specified detection rate β and false positive rate α.

C. Detection of RoS Attacks

Figure 8 presents the procedure to identify malicious flows containing RoS attacks. Our algorithm starts by checking the (0, 20 Hz] area, which is shown in Fig.8(a) with the energy distribution. In this band, distribution of the RoS attack flows is overlapped with that of both UDP and TCP flows. It would lead to a large false rate if we made a decision based only on this band.

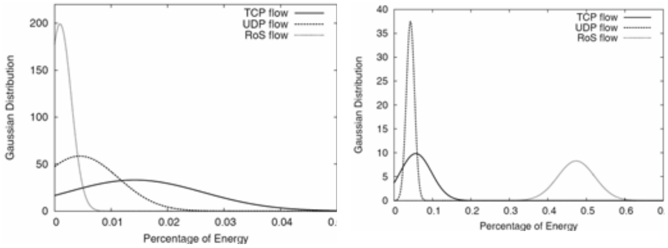

(a) Comparison of energy spectrum in DC band

(b) Comparison of Energy spectrum in (0 Hz, 20Hz] band

Figure 8. Two step in identifying low-rate TCP-targeted DDoS attack streams

The major advantage of sequential hypothesis testing is that it can reach decision promptly with limited sample points. As soon as current samples are sufficiently large to draw a conclusion, it needs not to wait for future arriving samples. In other words, a shorter temporal sequence makes a quicker testing. It is desired in defense against RoS attacks that the scheme can detect and response quickly before extended damages are done.

To determine whether a flow is malicious using traffic spectrum, we apply the sequential hypothesis testing on a frequency sequence instead of a temporal sequence. We compare the energy distribution patterns and calculate the STD along the frequency axis. The RoS attack streams present an obvious low-frequency biased energy distribution pattern. Using the SHT filtering method, we identify the malicious flows by comparing only the percentage of total energy located in the low frequency (0, 20 Hz] band.

D. Filtering of RoS Attack Flows

We propose to cut off the malicious flows identified with RoS attacks using the SHT algorithm. Although the source IP addresses are generally spoofed in attack packets, it is safe to use the 5-tuple {*Source IP, Source Port, Destination IP, Destination Port, protocol*} as the flow identifier.

To minimize the storage overhead incurred by the extra data structures needed, we store only the output of a hash function with the label as the input instead of the label itself. Our filtering algorithm manages the packet labels using three data structures: *Malicious Flow Table* (MFT), *Suspicious Flow Table* (SFT) and *Legitimate Flow Table* (LFT) as shown in Fig. 9.

If an incoming packet label is in LFT, this packet is routed normally. If it is in the MFT, this packet is dropped. If it is in SFT, we continue sampling until enough for meaningful spectral analysis. If there is no matching in any table, this packet belongs to a new flow and it would be added into SFT, then sampling begins.

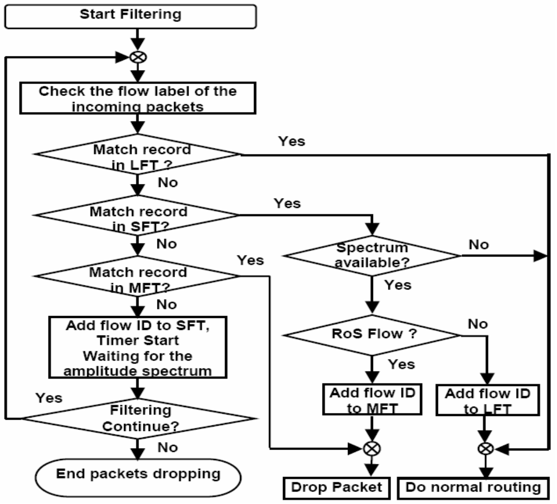

Fig. 9. The filtering algorithm for cutting off malicious attack flows and rescue legitimate TCP flows. (LFT: Legitimate Flow Table, SFT: Suspicious Flow Table, MFT: Malicious Flow Table)

Once a sample series available for a suspicious flow, we convert it into frequency domain. Then we identify whether it is a RoS attack flow using the SHT algorithm as illustrated by Fig. 7. If it is a legitimate TCP flow, we move its label into LFT. All further incoming packets of this flow will be routed normally. Otherwise, we move the label into MFT and start cutting off the flow.

5. Simulation Results and Performance Analysis

We have verified the effectiveness of using the Internet spectrum analysis scheme in defense against RoS attacks using the NS-2 network simulator, which is a widely used packet-level discrete event simulator.

A. Simulation Setup

Our simulations were carried out with several topologies generated by the GT-ITM toolkit from Georgia Tech. [28]. One of the topology we simulated was shown in Fig. 1. The background traffic in our simulation is generated using parameters from the real-life Abilene-III Internet trace dataset [26]. We mingle them with RoS attack flows with a period *T* between 0.5s and 3.0s, the burst period *L* is in the range (30 – 90 ms). For single-source attacks, the burse rate *R* varies in 0.5 - 2 MB/s. In distributed attacks from multiple sources, *R* varies in 0.1 - 2 MB/s.

We set the link capacity of the last hop to the victim as 2 Mbps. Since all TCP variants are equally vulnerable to RoS stream [21], we use TCP-Reno in our experiments. Their delays to the Internet are randomly distributed from 60 ms to 120 ms uniformly. In all experiments, we set $\alpha = 0.01$ and $\beta = 0.99$.

B. Experimental Results

We report the performance results from the NS-2 simulation experiments in Fig.10. Figure 10(a) plots the ROC curve showing the detection rate of RoS attacks versus the false positive rate. As shown in Fig.8, the RoS attack flows present a very different spectral pattern at the low frequency band. Our SHT scheme achieves a 99.7% detection rate by scanning through the low frequency band.

The false positive rate is resulted from a few TCP/UDP flows (< 3%) with periods close to that of RoS attack flows. They are identified as RoS attack flows. Although they are not launched purposely to hurt other legitimate flows, such type of flow does throttle the throughput of innocent TCP flows sharing link with them [21]. Thus, it is desired to suppress them in some scenarios. The ROC curve shows that we can tolerate 3% false alarms and still yield 99.7% detection accuracy.

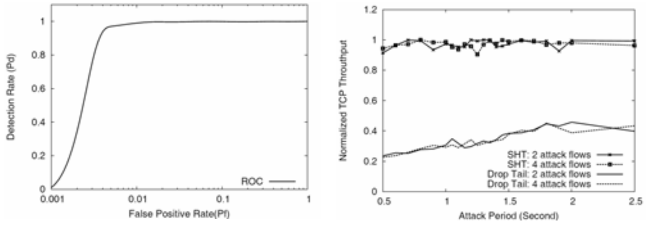

(a) ROC curve of flow detection performance

(b) Normalized TCP throughput

Figure 10. Detection rate of shrew attacks and normalized TCP throughput of SHT scheme compared with the Drop-tail filtering performance

We also study the sustained TCP throughput achieved by SHT filtering algorithm, compared with that of the *Drop Tail* algorithm. The metric *normalized throughput, ρ* is defined by the ratio of average throughput achieved by the TCP flows under RoS attacks to the maximum throughput achievable over the data links.

Figure 10(b) presents the scenario of two TCP flows under RoS attacks from 2 and 4 zombies. The x-axis is the RoS attack period and the y-axis is the normalized TCP throughput achieved. The normalized throughput indicates the percentage of legitimate TCP flows preserved or rescued from deleting the malicious flows containing RoS attacks.

The lower is the normalized throughput, the greater is the impact from RoS attacks against legitimate flows. It is clear that under the Drop Tail algorithm, the TCP throughput of legitimate flows is maintained only 30% to 40% from the peak throughput. Using our SHT scheme as described in Section 4 and Section 5, we can save 90% to 99% of the legitimate TCP flows, quite an impressive result.

6. Conclusions and Discussions

Analyzing Internet traffic spectrum in frequency domain enables us to solve some network anomaly problems that cannot be solved effectively by volume-based traffic monitoring in real time. We have performed flow-level spectral analysis over Internet traffic spectrum. Our detection and filtering methods were tested with traffic dataset from the real-life Abilene-III Internet traces. The shrew DDoS or RoS attacks were artificially generated and mixed up the Abilene-III background traffic.

Our scheme effectively rescues legitimate TCP flows from RoS attacks, which are very hard to detect in time domain for their stealthy properties [13], [21]. Our results verify the spectrum shifting in multi-source RoS attacks. We identify the following distinct advantages of using the flow-level spectral analysis for defense against shrew DDoS or RoS attacks.

- The scheme distinguishes normal TCP flows from others by observing the energy distribution and its standard deviation. The detailed ingredients are revealed by a PSD of the entire traffic spectrum.
- The scheme detects malicious RoS attack flows accurately and swiftly. Legitimate TCP applications are saved from the attack flows.
- The scheme segregates legitimate TCP flows from flooding DDoS attack flows. This property is very helpful to minimize the collateral damage to legitimate flows while packet-dropping mechanism is adopted.

Attackers often use source spoofing to hide themselves from being caught by traceback techniques. This is not a problem to our spectral analysis. The limit of the spectral analysis scheme lies in the difficulty in handling the dynamically spoofed source IP. Consider the extreme case, if the source IP address varies for each packet, we cannot collect a sample series to perform spectral analysis. It is still an open problem to identify malicious flows that exhibit transient behaviors.

To deploy distributed security scheme in an ISP core network, the system scalability is desired over high-speed data links. The major obstacle lies in the conflict between limited computing power in routers to perform the spectral analysis tasks for real-time detection and filtering of malicious flows over the high-speed links [17]. Programmable network processors have been suggested to provide flexible network services in high-speed networks [18]. Field Programmable Gate Arrays (FPGAs) are also suggested to security information processing. Recently, an *intrusion detection system* (IDS) was implemented on a FPGA platform, which can process 32,768 complex rules at a data rate as high as 10 Gbps [2].

The DSP methods are suitable for implementation using network processors or FPGAs. In on-going efforts, developing our spectral analysis scheme on FPGA. We are testing the spectrum analysis scheme and its hardware implementation on the DETER testbed [10]. That will allow us to evaluate the scheme in an environment closer to the Internet reality. We are extending the SHT algorithm and sequential multi-hypothesis test framework to address other types of DDoS attacks that present variant patterns in frequency domain.

References

[1] P. Abry and D. Veitch, "Wavelet analysis of long-range-dependent traffic", *IEEE Trans. on Information Theory*, 44, 1 (1998), 2–15

[2] M. Attig and J. Lockwood, “A Framework For Rule Processing in Reconfigurable Network Systems,” *Proc. of IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM)*, Napa, CA, April 17, 2005.

[3] P. Barford, J. Kline, D. Plonka, and A. Ron, “A signal analysis of network traffic anomalies”, *ACM Proc. Internet Measurement Workshop,* Marseille, France, Nov. 6-8, 2002

[4] C. Baum and V. Veeravalli, “A Sequential Procedure for Multihypothesis Testing”, *IEEE Transactions on Information Theory*, Vol. 40, No. 6. Nov. 1994.

[5] Y. Chen, K. Hwang, and Y.-K. Kwok, "Filtering of Shrew DDoS Attacks in Frequency Domain," *the First IEEE LCN Workshop on Network Security (WoNS 2005),*Sydney, Australia, Nov. 15-17, 2005

[6] Y. Chen and K. Hwang, "Collaborative Detection and Filtering of Shrew DDoS Attacks using Spectral Analysis," *Journal of Parallel and Distributed Computing,* 2006.

[7] C-M Cheng, H. Kung and K. Tan, “Use of Spectral Analysis in Defense Against DoS Attacks”, *Proceedings of IEEE GLOBECOM 2002*, Taipei, Taiwan

[8] R. Chertov, S. Fahmy, and N. Shroff, “Emulation versus Simulation: A Case Study of TCP-Targeted Denial of Service Attack,” *Proc. of 2nd IEEE CreateNet Conference on Testbeds and Research Infrastructures,* March 2006.

[9] M. Delio, “New Breed of Attack Zombies Lurk”, http://www.wired.com/news/technology/0,1282,43697,00.html, as of May. 20, 2006.

[10] DETER and EMIST Team Members, “Cyber Defense Technology Networking and Evaluation”, *Comm. ACM*, vol. 47, no. 3, Mar. 2004, pp. 58–61.

[11] L. Feinstein, D. Schnackenberg, R. Balupari, and D. Kindred, “Statistical Approaches to DDoS Attack Detection and Response”, *Proc. of DARPA Information Survivability Conference and Exposition (DISCEX03)*, Wash. D.C. 2003

[12] K. Giles, D. Marchette, and C. Priebe, "On the Spectral Analysis of Backscatter Data", *Proc. of Hawaii Int’l Conf. on Statistics, Mathematics, and Related Fields*, 2004

[13] M. Guirguis, A. Bestavros, I. Matta, and Y. Zhang, "Reduction of Quality (RoQ) Attacks on Internet End Systems," *IEEE INFOCOM,* 2005

[14] G. He and J. Hou, “An In-Depth, Analytical Study of Sampling Techniques For Self-Similar Internet Traffic”, the *25th Int’l Conf. on Distributed Computing Systems (ICDCS 2005)*, Columbus, OH.

[15] X. He, C. Papadopoulos, J. Heidemann, and A. Hussain, “Spectral Characteristics of Saturated Links”, *Technical Report USC-CSD-TR-827, University of Southern California Computer Science Department*, June, 2004

[16] P. Huang, A. Feldmann, and W. Willinger, "A non-intrusive, wavelet-based approach to detecting network performance problems", *Proc. ACM SIGCOMM Internet Measurement Workshop*, 2001

[17] A. Hussain, J. Heidemann, and C. Papadopoulos, “A Framework for Classifying Denial-of-Service Attacks,” *Proceedings of the ACM SIGCOMM Conf. on Network Architectures and Protocols*, 2003

[18] Intel, “Case Study: IDT™ PAX.port™ 2500 content inspection engine (CIE) and Intel® IXP2400 network processor”, www.intel.com/design/network/casestudies/idt_04.pdf, 2004

[19] J. Jung, V. Paxson, A. Berger, and H. Balakrishnan, “Fast Portscan Detection Using Sequential Hypothesis Testing”, *IEEE Symp. on Security and Privacy 2004*, Oakland, CA

[20] M. Kim, T. Kim, Y. Shin, S. Lam, and E. Powers, "A Wavelet-Based Approach to Detect Shared Congestion", *SIGCOMM'04*, Aug.30 - Sept. 3, 2004, Portland, OR.

[21] A. Kuzmanovic and E. Knightly, “Low-Rate TCP-Targeted Denial of Service Attacks—The Shrew vs. the Mice and Elephants,” *ACM SIGCOMM 2003*, Aug. 2003.

[22] Q. Li, E. Chang, and M. Chan, "On the Effectiveness of DDoS Attacks on Statistical Filtering," *IEEE INFOCOM,* 2005

[23] X. Luo and R. Chang, "On a New Class of Pulsing Denial-of-Service Attacks and the Defense", *Network and Distributed System Security Symposium 2005* (NDSS'05)

[24] H. Sun, J. Lui, and D. Yau, “Defending Against Low-rate TCP Attacks: Dynamic Detection and Protection,” *Proc. IEEE International Conference on Network Protocols* (ICNP), Berlin, Germany, October 5-8, 2004

[25] M. Thottan and C. Ji, “Anomaly Detection in IP Networks”, *IEEE Trans. On Signal Processing*, August 2003

[26] Abilene-III Trace Data - Illustrated, http://pma.nlanr.net/Special/ipls3.html

[27] 2004 Internet2 IPLS Abilene backbone instrumentation, http://pma.nlanr.net/Sites/ipls-200406/

[28] “GT-ITM: Georgia Tech Internet Topology Models,” http://www.cc.gatech.edu/projects/gtitm/, Nov. 03, 2005