Detecting Traffic Anomalies using Discrete Wavelet Transform

Seong Soo Kim[1], A. L. Narasimha Reddy[1], and Marina Vannucci[2]

[1] Department of Electrical Engineering, Texas A&M University,
College Station, TX 77843-3128, USA,
{skim, reddy}@ee.tamu.edu,
[2] Department of Statistics, Texas A&M University,
College Station, TX 77843-3143, USA,
mvannucci@stat.tamu.edu

Abstract. We propose a traffic anomaly detector operated in postmortem and real-time by passively monitoring packet headers of traffic. We analyze the correlation of destination IP addresses of outgoing traffic at an egress router. Based on statistical bounds on normal traffic patterns of the correlation signal of destination addresses, sudden changes can be used to detect anomalies in traffic behavior. For more computational efficiency, we suggest a correlation calculation using a simple data structure. These correlation data are processed through coefficient-selective discrete wavelet transform for effective and high-confidence detection. We present two kinds of mechanisms for postmortem and real-time detection modes. We evaluate the effectiveness of those two mechanisms by employing network traffic traces. [3]

1 Introduction

The frequent attacks on network infrastructure, using various forms of bandwidth attacks, have led to an increased interest for developing techniques for analyzing and monitoring network traffic. If efficient analysis tools were available, it could become possible to detect the attacks, anomalies and to appropriately take action to suppress the attacks before they have had much time to propagate across the network. In this paper, we study the possibilities of traffic-analysis based mechanisms for attack and anomaly detection.

Our approach monitors a packet header of network traffic at regular intervals and analyzes it to find if any abnormalities are observed in the traffic. By observing the traffic and correlating it to previous states of traffic, it may be possible to see whether the current traffic is behaving in an ordinary manner. In the case of bandwidth anomalies such as flash crowds and denial of service (DoS) attacks, the usage of network may be increased and abnormalities may show up in traffic pattern. Abrupt increase or decrease of traffic access pattern could signify the onset of an anomaly such as worm propagation. Our methodology relies on analyzing packet header data in order to provide indications of possible abnormalities in the traffic. Our approach to detecting anomalies envisions two kinds of detection mechanisms: off-line and on-line modes.

[3] This work is supported by an NSF grant ANI-0087372, Texas Higher Education Board, Texas Information Technology and Telecommunications Taskforce and Intel Corp.

2 Related Work

In terms of bandwidth consumption, long term high-rate flow can be identified using partial state information [8]. Using a sample and hold approach, once a suspicious entry is detected; every subsequent packet belonging to the flow is kept monitoring and will be updated [11].

Many rule-based approaches, such as intrusion detection systems, try to match the established rules to the potential DoS attack from external incoming traffic near the victims. Moreover, the pushback is a cooperative defending mechanism through which the information exchanges amongst core routers [9, 10]. In contrast, some approaches proactively seek a method that suppresses the overflowing of traffic in the source [5]. It usually uses a rate-limited control for reducing the monopolistic consumption of available bandwidth to diminish the effect of attack [5, 7, 10].

Traditionally, various forms of signature have been utilized for representing the whole contents or certain identities. By expanding utilization of signature, network securities use the signature-based algorithm prevalently. The disproportion of bidirectional flows can be used as the signature of anomalistic traffic [4]. The changing ratios (i.e., the rate of decrease) between the flow numbers of neighboring specific bit-prefix aggregate flows can be calculated and used for peculiarities [6]. Besides, there are some distinguished transform approaches to emphasize the anomaly of the traffic [1, 3, 12]. Using traffic volume such as byte counts as signal, a wavelet system shows performance to expose the variance of the anomalies.

3 Our Approach

3.1 Traffic analysis at the source

The (spoofed) source address of outbound traffic from an AD (administrative domain) could be filtered because router can know internal address range unlike destination addresses. On the other hand, destination address is likely to have a strong correlation with itself over time since the individual accesses have strong correlation over time. Recent studies have shown that the traffic can have strong patterns of behavior over several timescales [3]. It is possible to infer that some correlation exists on their weekly or daily consumption patterns. We hypothesize that the destination addresses will have a high degree of correlation over longer timescales. If this is the case, sudden changes in correlation of outgoing addresses can be used to detect anomalies in traffic behavior.

3.2 General Mechanism of the detector

Our detection mechanisms can be organized in three steps. The first step is traffic parser, in which the correlation coefficient signal is generated from packer header traces or NetFlow records as input, in section 4. The second step processes wavelet transforms to study the address correlation over several timescales. We selectively reconstruct decomposed signal across specific timescales based on the nature of attacks and network administrator's focus, in section 5. The final stage is detection, in which attacks and

anomalies are detected using historical thresholds of traffic to see whether the traffic's characteristics are out of regular norms. Sudden changes in the analyzed signal will lead to some indication that could be used to alert the network administrator of the potential anomalies in the network traffic. In this paper, we consider some statistical summary measures as explained in section 6.

3.3 Traces

To verify the validity of our approach, we run our algorithm on two kinds of traffic traces. First, we examine the detector on traces from the University of Southern California which contains real worm attacks. Additionally to inspect the sensitivity of our detector over attack of various configurations, we employ simulated virtual attacks on the University of Auckland traces of addresses collected over a campus access link. These traces range in length from 3 days to several weeks.

3.4 Attacks

We consider nine kinds of virtual attacks as shown in Table 1. These simulated attacks cover many kinds of behaviors and allow us to deterministically test diverse attacks.

Table 1. The Nine kinds of simulated Attacks

	1 (2,I,SD)	2 (2,I,SR)	3 (2,I,R)	4 (2,P,SD)	5 (2,P,SR)	6 (2,P,R)	7 (1,P,SD)	8 (1,P,SR)	9 (1,P,R)
$Duration$	2 hours	2 h.	2 h.	2 h.	2 h.	2 h.	1 hour	1 h.	1 h.
$Persistent$	intermittent	int.	int.	persistent	per.	per.	per.	per.	per.
IP	single destination	semi-random	random	single dest.	semi-rand.	rand.	single dest.	semi-rand.	rand.

– **Persistency.** The first 3 attacks send malicious packets in on-off type. More sophisticated attackers attempt to conceal their intentions through repeating attack and pause periods. So, it is intended to model intelligent and crafty attackers that attempt to dilute their trails. The other remnant attacks continue to assault throughout the attack.
– **IP address.** The 1^{st} attack among every 3 attacks targets for a (semi) single destination IP address. This target may be really one host in case of 32-bit prefix, occasionally aggregated neighboring hosts in case of *x*-bit prefix. The 2^{nd} attack style composes the IP address in which specific portion of the address structure preserves the identical value and the rest of the address is generated randomly for the infiltration efficiency. The 3^{rd} type is randomly generated.

Our attacks can be described by a 3-tuple (duration, persistency and IP address). We superimpose these attacks on ambient traces, which are the University of Auckland traces [2]. The mixture ratios of normal traffic and attack traffic range 2:1 to 10:1 in packet count. Replacement of normal traffic with attack traffic is easier to detect and hence not considered here.

4 Signal Generation

4.1 The correlation coefficient

Our approach collects packet header data at an AD's edge over a sampling period. Individual fields in the packet header are then analyzed to observe anomalies in the traffic. To study the correlation embedded in the IP address, we use its correlation coefficient which is a normalized measure of the linear relationship in random variable.

For each address, a_m, in the traffic, we count the number of packets, p_{mn}, sent in the sampling instant, s_n. In order to compute address correlation coefficient signal, we consider two adjacent sampling instants *n-1* and *n* . We can define IP address correlation coefficient signal in sampling point *n* as

$$\rho(n) = \frac{\sum_m (p_{mn-1} - \overline{p_{n-1}}) * (p_{mn} - \overline{p_n})}{\sqrt{\sum_m (p_{mn-1} - \overline{p_{n-1}})^2} \sqrt{\sum_m (p_{mn} - \overline{p_n})^2}} \quad (1)$$

When correlation coefficient signal varies between -1 and 1, if an address a_m spans the two sampling points, we will obtain a positive contribution to $\rho(n)$. A popular destination address a_m contributes more to $\rho(n)$ than an infrequently accessed destination.

4.2 Data structure for computing correlation coefficient

In order to minimize storage and compute efficiently, we employ a simple but powerful data structure. This data structure consists of 4 arrays *"p[4]"* . Each array expresses one of the 4 bytes in an IP address. Within each array, we have byte-sized 256 locations, for a total of 4*256 bytes = 1024 bytes. A location *p[i][j]* is used to record the packet count for the address *j* in i^{th} field of the IP address through scaling. This provides a concise description of the address instead of unique 2^{32} locations. We generate this approximate signal by computing a correlation coefficient over the address in two success samples, i.e., by computing

$$\rho_{in} = \frac{\sum_{j=0}^{255} (p[i][j][n-1] - \overline{p[i][n-1]}) * (p[i][j][n] - \overline{p[i][n]})}{\sqrt{\sum_{j=0}^{255} (p[i][j][n-1] - \overline{p[i][n-1]})^2} \sqrt{\sum_{j=0}^{255} (p[i][j][n] - \overline{p[i][n]})^2}} \quad (2)$$
$$, where \;\; i = 1, 2, 3, 4$$

We present simple example to illustrate the information can be stored this data structure as shown in Fig. 1. The packet counts of each flow are recorded to the corresponding position of each IP address segment.

Our approach could introduce errors when the addresses segments match even if addresses themselves don't match. From comparison between correlation coefficient signal of the full-32 bit address and our data structure, we see that the difference are negligible i.e., our approach does not add significant noise. Moreover, we examine the similarity of above signals with cross-correlation coefficient, which has actually $\rho_{XY} \approx 0.74$. We think that these signals have a close positive correlation interchangeably.

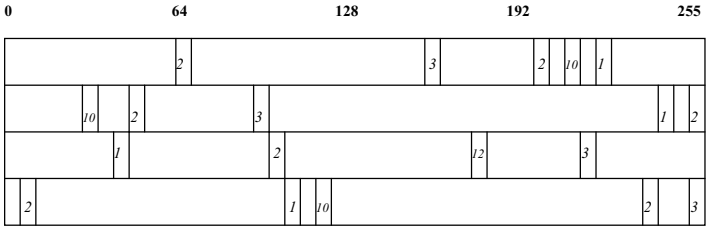

Fig. 1. Data structure for computing correlation coefficient. Suppose that only five flows exist, their destination IP addresses and packet counts are as follows. IP of F1 = 165. 91.212.255, P1 = 3; IP of F2 = 64. 58.179.230, P2 = 2; IP of F3 = 216.239. 51.100, P3 = 1; IP of F4 = 211. 40.179.102, P4 = 10; IP of F5 = 203.255. 98. 2, P5 = 9

5 Discrete Wavelet Transform

5.1 DWT (Discrete Wavelet Transform)

Wavelet technique is one of the most up-to-date modeling tools to exploit both non-stationary and long-range dependence. In real situations, we encounter signals which are characterized by abrupt changes and it becomes essential to relate to the occurrence of an event in time. Wavelet analysis can reveal scaling properties of the temporal and frequency dynamics simultaneously unlike Fourier Transform used in [12]. Through signal can be detected in certain timescales and in certain position of the timescales, we can induce the frequency and temporal components respectively. We compute a wavelet transform of this correlation signal over a given time. A multilevel one-dimensional DWT consists of decomposition (or analysis) and reconstruction (or synthesis) [13].

For decomposition, starting from a signal *s* , the first step of the transform decomposes *s* into two sets of coefficients, namely approximation coefficients cA_1, and detail coefficients cD_1. The input *s* is convolved with the low-pass filter Lo_D to yield the approximation coefficients. The detail coefficients are obtained by convolving *s* with the high-pass filter Hi_D. This procedure is followed by down sampling by 2. The second step decomposes the approximation coefficient cA_1 into two sets of coefficients using the same method, substituting *s* by cA_1, and producing cA_2 and cD_2, and so forth. At level *j* , the wavelet analysis of the signal *s* has the following coefficients, [cA_j, cD_j, cD_{j-1}, cD_{j-2}, ... , cD_2, cD_1].

For reconstruction, starting from two sets of coefficients at level *j* , that is cA_j and cD_j, the inverse DWT synthesizes cA_{j-1}, up-samples by inserting zeros and convolves the up-sampled result with the reconstruction filters Lo_R and Hi_R. For a discrete signal of length *n* , DWT can consist of $\log_2 n$ levels at most.

5.2 Our specification: coefficient-selective reconstruction

Our specification is daubechies-6 two-band filter. The filtered signal is down-sampled by 2 at each level of the analysis procedure; the signal of each level has an effect that

sampling period extends 2 times. Consequently it means that the wavelet transform can identify the changes in the signal over several timescales. When we use *t* seconds as sampling period, the time range at level *j* extend $t * 2^j$ seconds. And a 1-minute sampling interval and 30-second sampling duration are used to reduce the amount of data and computational complexity.

The network operators can select reconstructed levels that they wish to be captured. We assume that the network administrators are interested in detecting shorter anomalies of sufficient intensity and anomalies of more than 30-minute duration. In order to detect these attacks, we extract only the 1^{st}, 5^{th}, 6^{th} and 7^{th} levels in decomposition and reconstruct the signal based only on coefficients at these levels.

6 Detection

6.1 Thresholds setting through statistical analysis

We develop a theoretical basis for deriving thresholds for anomaly detection. We first investigate the distribution of the wavelet reconstructed signal in the University of Auckland free of attacks. To examine random variable density, we select only some levels of the DWT decomposition of the ambient trace without attacks and reconstruct the signal based on those levels. We then look at some statistical properties. The Fig. 2(b) and 2(e) show the distribution and histogram of the reconstructed signal of the ambient traces in postmortem mode. We verify normality through the Lilliefors test for goodness of fit to normal distribution with unspecified mean and variance. The postmortem transformed data have a normal distribution at 5% significance level, namely X ˜ N(0, 0.026^2). By selecting some of the levels through selective reconstruction, we have removed some of the features from the signal that were responsible for the non-normality in the original signal.

We gather the 4-week traces and analyze their statistical summary measures. The statistical parameters of network traffic, such as mean, variance and autocorrelation function, are stationary distributed given different days. From the viewpoint of communications, the ambient trace could be considered as wide-sense stationary Gaussian white noise, on the other hands, the attack and anomaly could be considered as random signal.

When we set the thresholds to -0.078 and 0.078 respectively, these figures are equivalent to $\pm 3.0\sigma$ confidence interval for random variable X.

$$P(\mu - 3.0\sigma < X \leq \mu + 3.0\sigma) \approx 99.7\% \tag{3}$$

This interval corresponds to 99.7% confidence level by (3). With such thresholds, we can detect attacks with error rate of 0.3%.

We also analyze the reconstructed signal at our selected levels of ambient trace in real-time mode, which shows approximately normal distribution. The Fig. 2(c) and 2(f) show the distribution and histogram of the reconstructed signal of the ambient traces in real-time mode.

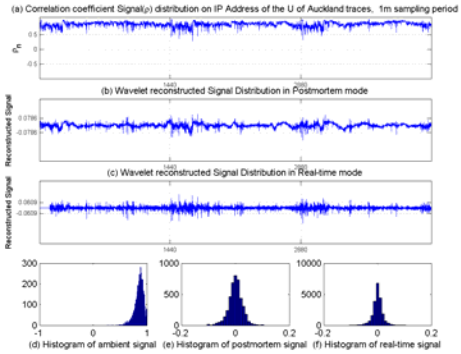

Fig. 2. Various distribution of the ambient traces in the Univ. of Auckland

Fig. 3. The Univ. of Auckland detection results in postmortem and real-time mode

6.2 Detection anomalies using the real attack trace

Our postmortem and real-time approaches are applied to the USC traces which contain real network attacks. Detection results are shown in Fig. 4. The Fig. 4(a) illustrates a correlation coefficient signal of IP addresses that is used for wavelet transform. The Fig. 4(b) is the wavelet-transformed and reconstructed signal in postmortem and its detection results. The detection signal is shown with dots at the bottom of the each sub-picture. The Fig. 4(c) shows the wavelet-transformed and reconstructed signal in real-time and its accumulated results. Through traffic engineering, we can identify the attack in which a specific internal compromised machine continued to attack a few external destinations.

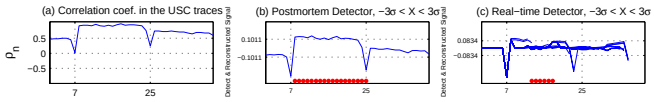

Fig. 4. The USC trace detection results in postmortem and real-time mode

6.3 Detection anomalies using the simulated attack trace

Detection results on the University of Auckland traces in 3 days are shown in Fig. 3. The Fig. 3(a) illustrates a correlation coefficient signal of IP addresses and the Fig. 3(b) is the wavelet-transformed and reconstructed signal in postmortem and its detection results and the Fig. 3(c) shows the accumulated results in real-time.

The simulated nine attacks are staged between the vertical lines, shown in the figure. Overall, our results show that our approach may provide a good detector of attacks in both modes. Moreover, the detections in the early points of every day, namely sampling

points are near at the 1450 and 2900, are turned out the regular flash crowds included in the original traces such as file backup.

Discussion of Postmortem analysis. The postmortem analysis and detectors can rely on datasets over long periods of time and we use whole 3-day correlation data all at once. The reconstructed signals of first 3 attacks (*,I,*) show an oscillatory fashion because of their intermittent attack patterns, while the remaining six attacks, namely (*,P,*), give a shape of hill and dale at attack times due to persistence.

The attacks on a single machine, the 1st attack among every 3 attacks described in (*,*,SD), reveal the high valued correlation which means the current traffic is concentrated on (aggregated) a single destination. Detection signals in the form of dots show that these typed attacks can be detected effectively. On the other hand, the semi-random typed attacks, that is (*,*,SR), and random styled attacks, namely (*,*,R), illustrate low correlations which means traffic is behaving in irregular pattern. Consecutive detection signals indicate the length of attacks and also imply the strength of anomalies.

In order to evaluate the effectiveness of employing DWT, we compare the detection results of our scheme employing DWT with a scheme that directly employs statistical analysis of the IP address weighted correlation signal. When confidence levels of most interest (90% ˜ 99.7%) are considered, DWT provides significantly better detection results than the simpler statistical analysis without applying of DWT. This shows that DWT offers significant improvement in the detection of anomalies.

Discussion of Real-time analysis. The real-time detection requires that the analysis and the detection mechanism rely on small datasets in order to keep such on-line analysis feasible. If we want to investigate a specific level j , it requires 2^j samples for reconstruction at least. So, the most recent 2-hour correlation data make use of detecting an attack of less than 2-hour duration. We take the moving window and majority to accommodate faster detection while reducing the false alarms. As the Fig. 3(c) shows, our detector achieves acceptable attack detection performance in on-line analysis as well as in off-line analysis.

Table 2 shows the overall timing relationship between detection latency and the setting of the confidence level of our simulated attacks in real-time mode. As we expect, the higher the confidence level, the higher the detection latency. The detection against the (*,*,SD) typed attacks, (aggregated) single destination, can achieve a prompt response. And the (*,*,R) typed attacks, randomly generated destination, can generally be detected more quickly than the semi-random type attacks described in (*,*,SR) because of the resulting lower correlation with random attacks.

6.4 Adaptive filtering in real-time analysis

In order to detect anomalies of different unknown durations, we considered adaptive filtering of the traffic signal. Adaptive filtering continues to search for the proper levels of timescales suitable to the nature of the attack. At normal times, the detector monitors only the reconstructed signal based on lower level coefficients (say, aggregated 1, 2 and 3 levels), which can identify the most detailed and instantaneous change in the traffic

Table 2. The Relation between Latencies and Confidence levels in Nine attacks in Real-time mode

	c	1	2	3	4	5	6	7	8	9	f	f
	l[a]	(2,I,SD)	(2,I,SR)	(2,I,R)	(2,P,SD)	(2,P,SR)	(2,P,R)	(1,P,SD)	(1,P,SR)	(1,P,R)	p[b]	n[c]
1.0σ	68	0[d]	1	2	0	0	0	0	0	0	11	0
1.5σ	86	0	1	2	0	0	0	0	2	0	7	0
2.0σ	95.5	1	2	5	0	0	0	0	5	2	5	0
2.5σ	98.5	1	2	5	0	3	0	0	7	2	3	0
3.0σ	99.7	1	3	5	0	10	2	0	8	6	2	0
3.5σ	99.95	1	5	10	0	36	2	0	11	6	2	0
4.0σ	99.99	2	15	13	0	X[e]	6	0	X	8	1	2

[a] confidence level in percentage
[b] false positive is counted a series of relevant signal as 1
[c] false negative
[d] latency is measured by minute unit
[e] X means non-detection

signal. Once a possible detection of anomaly is identified at these levels, the detector expands its investigative scope into higher levels gradually, for example at levels 2, 3 and 4, for improving the identification accuracy or robustness. It may help to reveal the substance of attack and to diminish the false alarm under unknown conditions. False alarms can be reduced by not declaring the detection of an anomaly until consecutive alarms are raised at multiple levels. The traffic signal at higher levels is considered as the identification progresses. On the other hand, when any anomaly is not detected, the reconstructed signals return to lower levels gradually. The Table 3 shows the results of such an approach. It may induce more false positives but achieve faster detection campared to non adaptive method. The results indicate that it may be feasible to detect traffic anomalies with low latency even when we consider attacks of unknown length.

Table 3. The Latencies in Nine kinds of attacks in Adaptive Filtering

	c	1	2	3	4	5	6	7	8	9	f	f
	l	(2,I,SD)	(2,I,SR)	(2,I,R)	(2,P,SD)	(2,P,SR)	(2,P,R)	(1,P,SD)	(1,P,SR)	(1,P,R)	p	n
3.0σ	99.7%	1	2	2	0	1	1	4	1	1	5	0

7 Future work and Conclusion

The duration of the samples and the number of samples have a strong impact on the accuracy of the results and the latency for detecting an attack. The samples with a smaller sampling period cause to more false positives but lead to faster identification of the attack. It is a pivotal factor for the real-time approach we discussed. Thus, as a

further research, the relation between sampling rate and latency should be investigated from statistical point of view.

We plan to study containment approaches along the multiple dimensions of addresses, port numbers, protocols and other such header data based on a detection tool. We also plan to study the effectiveness of the analysis of traffic at various points in the network, at the destination network and within the network core.

We studied the feasibility of analyzing packet header data through wavelet analysis for detecting traffic anomalies. Specifically, we proposed the use of correlation coefficient of destination IP addresses in the outgoing traffic at an egress router. Our results show that statistical analysis of aggregate traffic header data may provide an effective mechanism for the detection of anomalies within a campus or edge network. We studied the effectiveness of our approach in postmortem and real-time analysis of network traffic. The results of our analysis are encouraging and point to a number of interesting directions for future research.

Acknowledgement We are very grateful to Deukwoo Kwon for his comments and reviews on statistical analysis, to Alefiya Hussain at USC for her help in accessing traces.

References

1. Ramanathan, A.: "WADeS: A Tool for Distributed Denial of Service Attack Detection", *TAMU-ECE-2002-02, Master of Science Thesis*, August 2002
2. National Laboratory for Applied Network Research (NLANR), measurement and operations analysis team: "NLANR network traffic packet header traces", accessed in August 2002
3. Barford, P., Kline, J., Plonka, D., Ron A.: "A Signal Analysis of Network Traffic Anomalies", in *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, November 2002
4. Gil, T., Poletto, M.: "MULTOPS: A Data-Structure for Bandwidth Attack Detection", in *Proceedings of the 10th USENIX Security Symposium*, Washington, D.C., USA, August 2001
5. Mirkovic, J., Prier, G., Reiher P.: "Attacking DDoS at the Source", in *10th IEEE International Conference on Network Protocols*, Paris, France, November 2002
6. Kohler, E., Li, J., Paxson, V., Shenker, S.: "Observed Structure of Addresses in IP Traffic", in *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, November 2002
7. Garg, A., Reddy, A.: "Mitigation of DoS attacks through QoS regulation", in *Proc. of IWQOS workshop*, May 2002
8. Smitha, Kim, I., Reddy, A.: "Identifying long term high rate flows at a router", in *Proc. of High Performance Computing*, December 2001
9. Mahajan, R., Bellovin, S., Floyd, S., Ioannidis, J., Paxson, V., Shenker, S.: "Controlling High Bandwidth Aggregates in the Network (Extended Version)", in *ACM SIGCOMM Computer Communication Revies, Volume 32, Issue 3*, July 2002
10. Ioannidis, J., Bellovin, S.: "Implementing Pushback: Router-Based Defense Against DDoS Attacks", in *Proceedings of Network and Distributed System Security Symposium*, San Diego, California, February 2002
11. Estan, C., Varghese, G.: "New Directions in Traffic Measurement and Accounting", in *ACM SIGCOMM 2002*, Pittsburgh, PA, USA, August 2002
12. Cheng, C., Kung, H., Tan, K.: "Use of spectral analysis in defense against DoS attacks", in *Proc. of IEEE Globecom*, 2002
13. The MathWorks. Inc.: MatLab software, ver 6.1.0.450 Release 12.1, May 2001