An Evaluation of the Effectiveness of Measurement-based Anomaly Detection Techniques

Seong Soo Kim and A. L. Narasimha Reddy
Department of Electrical and Computer Engineering
Texas A&M University, College Station, TX 77843-3128, USA
{skim, reddy}@ece.tamu.edu

Abstract

A number of recent studies have proposed measurement based approaches to network traffic analysis. These techniques treat traffic volume and traffic header data as signals or images in order to make analysis feasible. We use trace-driven experiments and compare the performance of different strategies. Our evaluations on real traces reveal differences in the effectiveness of different traffic header data as potential signals for traffic analysis in terms of their detection rates and false alarm rates. Our results show that address distributions and number of flows are better signals than traffic volume for anomaly detection.

1. Introduction

A number of recent studies have pointed to the need for fast detection of malicious traffic for any mechanisms to be effective in thwarting network attacks [10, 2, 19]. If the spread of malicious traffic can be detected in real-time, we can alleviate and possibly contain such malicious traffic and improve availability of servers and network infrastructure.

Traditionally, Intrusion detection system (IDS) tools such as Snort [9] and Bro that rely upon operating system logs, process behaviors and firewall logs have been employed to monitor the network traffic. However, the frequent appearance of novel attacks compromises such analysis making detection of unknown malicious traffic difficult. Measurement-based IDS tools and network traffic analysis have recently started attracting attention as a potential complementary approach [1, 3, 4, 5].

Measurement-based tools analyze network traffic to observe statistical properties of traffic. Based on such measurements and some acceptable thresholds on normal network behavior, these tools try to classify traffic as normal or anomalous. These studies have considered traffic volume [3, 7, 4, 14], number of flows [5], address and port number distributions [13] as potential signals that can be analyzed in order to detect anomalies in network traffic.

While all these signals have been shown to be useful for analyzing network traffic, so far, no comprehensive study has been carried out about the relative usefulness of these traffic signals. Which signals are more effective for detecting anomalies? Which signals provide low false alarm rates? Different studies have employed different analysis techniques and different traces making such a comparison difficult. In this paper, we employ trace-driven evaluation to study the relative inherent effectiveness of various pieces of information in traffic headers for detecting traffic anomalies. We employ statistical analysis of the traffic data to compare the effectiveness of different pieces of information.

The rest of the paper is organized as follows. In section 2, we introduce the various traffic signals that have been proposed for analysis and anomaly detection. In section 3, we outline our approach for evaluating the effectiveness of these various signals. Our analysis is based on a number of real-world traces. In section 4, we present the results from our study. Section 5 concludes the paper.

This paper makes the following significant contributions: (a) provides a comprehensive evaluation of effectiveness of a number of signals derived from network traffic headers, (b) provides data from a number of real-world and simulated traces to compare the different traffic signals and (c) shows that distribution-based signals can be more effective and robust than volume-based signals in detecting traffic anomalies.

2. TRAFFIC SIGNALS

We broadly categorize anomaly detection schemes into two groups. The two groups are based on the property of information maintained or kept per sample. The volume-based signals keep track of an amount variable (such as traffic volume) as a time series. The distribution-based signals keep track of a distribution variable of the amount values over a domain of traffic header data (such as addresses, protocol numbers etc.). An exponential weighted moving average (EWMA) of the traffic signals is employed to consider the diurnal differences in traffic behavior. We briefly outline the various traffic signals below before we outline our approach for evaluation.

2.1 Volume-based Signals

The volume-based signals typically employ the traffic volume such as byte counts, packet counts and the number of flows. Many of the commonly exploited malicious attacks are based on high-bandwidth floods, or other repetitive streams of packets. Work in [3, 4] has analyzed traffic volume, measured in byte counts, packet counts and flow counts, using wavelets to detect anomalies in network traffic. For reasons of scalability, we look at the aggregate traffic volumes at the observation point (*not per-flow measurements*).

2.1.1 **Byte Counting.** This approach simply counts the number of bytes of traffic seen at the observation point during each sample. The traffic volume in bytes *b(t)* at each sample t is constructed as a time varying signal that can be analyzed to detect anomalies. Sudden variations in b(t) or traffic beyond normal statistical bounds can be considered as anomalies. Byte counting has O(1) processing cost per packet and O(1) storage cost per sample.

2.1.2 **Packet Counting.** In packet counting, traffic volume in packets *p(t)*, at each sample, is the time-varying signal that is analyzed to detect anomalies. Packet counts can be more useful than byte counts when links are carrying traffic to the capacity most of the time. Most of automated self-propagating codes use constant size packets and hence it is possible the packet counts can change as a result of an attack compared to normal traffic. Packet counting has *O(1)* processing cost per packet and *O(1)* storage cost per sample.

2.1.3 **Flow counting.** This approach counts the number of flows at the observation point. A flow can be classified by the 5-tuple (or by another definition) of source address, source port, destination address, destination port, protocol number. During each sampling period, the number of such distinct tuples are counted to generate the flow signal *f(t).* Hashing and other such techniques would be required to reduce the 5-tuple space to a single flow count number. The costs of such scheme would depend on the hashing techniques employed, the number of tuples chosen to define a flow and the number of flows likely to be seen at the observation point. Hashing can be accomplished in *O(1)* time per packet, individual flow observations can be marked in *O(1)* assuming no hash collisions and the number of flows can be counted in *O(n)* time, where *n* is the size of the flow space. More sophisticated approaches based on bloom filters can be employed to reduce the cost of such approaches to *log(n)* or *loglog(n)* [11, 12].

The number of flows could vary from the norm due to DDoS attacks, wide-scale worm propagation etc. It is expected through monitoring changes in the number of flows, it would be feasible to recognize such anomalies. FlowScan analyzes, visualizes and reports Internet traffic flow profiling on flow-centric measurements [5,6].

2.2 Distribution-based Signals

Distribution-based signals maintain the distribution of the amount data for each sample. This requires more space per sample, but allows more sophisticated analysis of traffic header data. We consider signals based on the protocol distribution and address distributions.

2.2.1 **Protocol composition.** This approach is based on the observation that during the attacks, the protocol employed by the attack traffic should see considerably more traffic than during normal traffic in general. In this approach, the amount of ICMP traffic ***i(t)***, TCP traffic ***t(t)***, UDP traffic ***u(t)*** and ETC. traffic ***e(t)*** is monitored as a fraction of the total traffic volume at each sampling interval. Because the proportion of each protocol in traffic is closely interrelated to each other, the increase of a proportion of one protocol makes the proportions of other protocols to decrease. We observe a considerable decrease of proportion of traffic other than TCP due to a moderate increase in the proportion of the already large fraction of TCP traffic. Protocol composition has *O(1)* cost per packet and *O(n)* storage cost per sample, where *n* is the number of protocols monitored.

2.2.2 **Image-based Signals.** In this approach, traffic distribution in a domain is used as a signal that can be analyzed. First, the traffic volume, such as normalized byte/packet counts and the number of flows, is measured along the packet header domain, such as IP addresses and port numbers. For example, traffic volume can be counted based on destination port numbers. Since the IP port number field is 16 bits, we would obtain 2^{16} = 64K values for each sample indicating the traffic distribution in the port number domain. For reducing the storage and computation complexity, we employ domain folding techniques. A datum at x.y.z.w address is represented in multiple subdomains at x, y, z and w. Specifically, the traffic header domain is processed in byte-segments which separate out each byte of the IP address (or the port number) [13]. This technique reduces the space complexity from 2^n to 256* (n/8), when 8-bit subdomains are considered. Each resultant traffic datum is converted to corresponding pixel intensity in image representation of traffic in the chosen domain [13]. The image-based signals then originate from the distribution of pixel intensity in each byte of the chosen domain.

According to the kinds of employed traffic data and the observed header domain, we categorize the image-based signals into address-based, flow-based and port-based signals. Address-based signal employs traffic volume distribution over address domain (either source address alone, or destination address alone, or a 2-dimensional source and destination address domain). Flow-based signal employs flow number distribution over address domain(s). Port-based signal employs traffic volume distribution over port number domain.

Image-based signals require two samples of packet header data 2**P*, where *P* is the size of the sample data. We also maintain summary information (pixel intensities) over a larger number of samples *S*, for statistical evaluation of the current data sample. So, the total space requirement is *O(P+S)*. In our example of source address domain analysis, *P* is originally 2^{32}, reduced to *4*256 (4 bytes of IP address * 256 values for each byte) = 1024*. *S* is originally 32*32, reduced to 1024 in real-time.

3. EVALUATION METHODOLOGY

3.1 Traces

In order to evaluate the different signals mentioned above, we employ three real traffic traces. We name these traces "*Access Link*", "*ISP*" and "*Campus*" based on where these traces are collected from.

KREONet2 (Korea Research Environment Open NETwork) trace used as the "*Access Link"* partially contained 2 weeks of network traffic data from Oct. 12, 2003 to Oct. 26, 2003, and contained actual worm attacks. Currently *Access Link* member institutions are over 230 organizations, which include 50 government research institutes, 72 universities, 15 industrial research laboratories, etc [16]. *Access Link* trace is a collection of NetFlow trace files collected by the 155Mbps international ATM link. In the trace employed, there are 5 major attacks and a few instantaneous probe attacks. Additionally we examined the signals on traces from a regional *ISP*[7] and live TAMU network. Due to space constraints, we are unable to present the results from these two later traces, but the trends are similar.

3.2 Measurement Criteria

We evaluate the performance of various signals based on the detection rates, false alarm rates and likelihood ratios. Measurement-based anomaly detection techniques have to contend with two types of errors. The true positive (sensitivity or detection β) is the probability that a statistical test will be positive for a true statistic. A type I error (false positive error α) occurs if a difference is declared when the null hypothesis is true. In other words, a false attack alarm is declared when the traffic is normal.

$$\alpha = Pr(\textbf{anounce } \mathbf{H_1} \,|\, \mathbf{H_0} \textbf{ is true}) = Pr(\textbf{detect anomaly} \,|\, \textbf{traffic is normal}) \quad \textbf{(1-1)}$$

$$\beta = Pr(\textbf{anounce } \mathbf{H_1} \,|\, \mathbf{H_1} \textbf{ is true}) = Pr(\textbf{detect anomaly} \,|\, \textbf{traffic is anomaly}) \quad \textbf{(1-2)}$$

The true negative (specificity 1-α) is the probability that a statistical test will be negative for a negative statistic. On the other hand, a type II error (false negative error 1-β) occurs if no difference is declared when the null hypothesis is false. To test non-nested complementary hypotheses, the LR (likelihood ratio) and NLR (negative likelihood ratio) are used as follows [15].

$$\textbf{LR} = \frac{\textbf{true positive rate}}{\textbf{false positive rate}} = \frac{\textbf{sensitivity}}{\textbf{1 - specificity}} = \frac{\beta}{\alpha} \qquad \textbf{(2)}$$

$$\textbf{NLR} = \frac{\textbf{false negative rate}}{\textbf{true negative rate}} = \frac{\textbf{1 - sensitivity}}{\textbf{specificity}} = \frac{\textbf{1 - }\beta}{\textbf{1 - }\alpha}$$

LR and NLR help to estimate the trade-off between the power of detection and false alarm. Ideally, LR is infinity and NLR is zero.

3.3 Statistical analysis based on 3σ

We employ statistical analysis of the network traffic data. Statistical analysis of traffic data requires only a model of normal traffic and hence possibly can distinguish new forms of attacks. We developed a theoretical basis for deriving thresholds for analyzing traffic signals and anomaly detection. For 3σ-based statistical analysis, we set 2 kinds of thresholds, a high threshold T_H and a low threshold T_L. When we respectively set the T_H and T_L thresholds to $\pm 3.0\sigma$ of aforementioned traffic signal distributions in ambient traffic, attacks can be detected with an error rate of 0.3% (if the signal is normally distributed) which can be expected as target false alarm rate by (3-1). We can judge the current traffic status by calculating the standard intensity deviation of signals, denoted by S_σ, in each sampling instant by (3-2).

$$X \sim N(\mu, \sigma^2) \rightarrow Pr\,(\mu - 3.0\sigma < X \leq \mu + 3.0\sigma) \approx 99.7\% \qquad (3\text{-}1)$$

$$\textit{traffic status} \begin{cases} \textit{normal}, & \text{if } T_L < S_\sigma < T_H \\ \textit{attack}, & \text{if } S_\sigma \leq T_L \text{ or } T_H \leq S_\sigma \end{cases} \qquad (3\text{-}2)$$

4. TRACE-DRIVEN EVALUATION

4.1 Volume-based Signals

4.1.1 **Byte counting and packet counting.** Table 1 shows the results of employing volume-based signals on the *Access Link* trace for anomaly detection. The 2nd and 3rd column of the Table 1 demonstrate the detection strength of the byte/packet counting signals.

The results show that byte count and packet count signals with statistical thresholds achieved detection rates of 11.0% and 18.3% respectively. The byte count signal is slightly better; however, the difference between the two measurements is marginal. These results show that the byte/packet counting signals may not be adequate for providing reliable signals for anomaly detection.

4.1.2 Flow counting. Flow counting shows significantly better performance in detecting anomalous traffic as shown in the 4th column of the Table 1. Flow counting with statistical thresholds achieved a detection rate of 95.1% at a false alarm rate of 0.73%. Flow counting is clearly superior to both byte counting and packet counting signals.

Most of the attacks in real traces consist of many flows with randomized addresses/ports, but with one or few attack packets per flow. This results in a large number of flows, but low attack traffic volume. These types of attacks can be effectively detected by flow counting rather than byte/packet counting.

Table 1. Results of volume-based signals

Signals	T.P.	F.P.	LR	NLR
Byte count	11.0%	0.11%	98.0	0.89
Packet count	18.3%	0.25%	72.4	0.82
Flow number	95.1%	0.73%	130.4	0.05

4.2 Protocol Composition

We employ 2 kinds of thresholds, a high threshold T_H that indicates that the fraction of volume of one of the network protocols increases abnormally and a low threshold T_L indicating that the fraction of the traffic volume of the network protocol decreases inordinately. When each protocol proportion in current input traffic is larger (or lower) than the 3σ of normal distribution for individual protocol, the detector declares anomalies.

Table 2. Results of protocol composition signals

Signals	T.P. β	F.P. α	LR	NLR
ICMP	72.9%	1.94%	37.6	0.28
TCP	81.0%	0.42%	192.3	0.19
UDP	77.5%	0.39%	197.2	0.23
ETC.	31.7%	0.00%	∞	0.68
Composite signal	80.8%	0.28%	288.0	0.19

Table 2 shows the measurement results for protocol composition. Protocol composition could achieve a detection rate of 80.8% at a false alarm rate of 0.28% as shown in Table 2. This shows that protocol composition could be a useful signal. Protocol composition signal has lower complexity than the flow counting signal.

4.3 Image-Based Signals

Real-time analysis may rely on less sophisticated analysis because of the resource demands and imminence of attacks. In real-time, we employ the light-weight approach simply using the variance of pixel intensities in the image for analysis and anomaly detection, which is denoted by distribution signal S_σ. Using the variance of these image signals as (4) for deriving thresholds, we can obtain an approximation of the energy distribution of the normalized packet counts within observation domain. Because it exploits the distribution property within a specific time interval, the analysis is not impacted by the diurnal variation of traffic.

$$S_\sigma = \left[\frac{1}{N} \sum_{k=1}^{N} (x_k - \bar{x})^2 \right]^{\frac{1}{2}} \qquad (4)$$

$$\text{, where} \begin{cases} x_k \text{ are pixel intensities, N=1024 in real-time} \\ x_k \text{ are DCT coefficients, N=16 in postmortem} \end{cases}$$

$$\text{and } \bar{x} = \frac{1}{N} \sum_{k=1}^{N} x_k$$

The detection signal is calculated instantaneously upon sampling instants for real-time analysis.

4.3.1 **Packet Distribution in Address Domain.** Usually the packet counts at the source addresses of the outbound aggregate traffic can illustrate the distributed properties of the network traffic usage. Meanwhile, the analysis at the destination address of the outgoing traffic can show the concentration of the flow target.

Table 3. Results of Address-based signals

Domain	TP β	FP α	LR	NLR

Source Addr.	81.5%	0.06%	1451.2	0.19
Dest. Addr	87.1%	0.42%	206.9	0.13
(Source, Dest)	94.2%	0.48%	197.5	0.06

The results of the address-based signals are shown in Table 3. The data illustrates that the true positive rate is 87.1 % (681 truly detected out of 782 anomalies) and the false positive rate is 0.42% (15 false alarms out of 3563 normal samples) for real-time detection based on destination address.

The results indicate that the different signals exhibit different strengths in anomaly detection. The destination address based signal performed better when compared with each other. In order to complete the analysis with lower detection latency with light-weight techniques, real-time analysis can only focus on small recent datasets. By default, the retained data are the latest 2 hour samples in real-time.

In our analyses, the (source, destination) based signal performed significantly better than the single dimensional signals. However, this signal has a higher storage and processing cost.

4.3.2 **Flow Distribution in Address Domain.** An analysis of the flow-based image could be effective for revealing flood types of attacks. When a flow is defined as the triple of source address / destination address / destination port, the flood-based attacks spread flows over the destination IP addresses (or ports) in random or dictionary mode style attacks. The distribution of the number of flows in address space would then be expected to be much different from its normal and historical distribution.

Results of analysis of flow-based images are shown in Table 4. As shown in the Table 4, the source address based images/signals generally exhibit higher confidence than the destination address based images/signal for detecting traffic anomalies due to bimodality of destination addresses. If source and destination address signals are jointly adopted, we can expect higher confidence in detection rate with a consequent deterioration of the false alarm rate.

The results from Tables 1 and Table 4 can be compared to understand the relative strengths of volume-based flow counting signal and the flow-based image signal. It is observed that flow-based images could reduce the false alarm rates. However, it is observed that flow-based images did not improve the detection rates when compared to the volume-based flow counting signal. From these results, the significant additional storage and processing cost entailed in flow-based images may not be warranted unless the reduction of the false alarm rate is paramount.

Table 4. Results of Flow-based signals

Domain	TP β	FP α	LR	NLR
Source Addr	90.3%	0.22%	402.1	0.10
Dest. Addr	56.5%	0.25%	223.8	0.44
(Source, Dest)	92.8%	0.42%	220.5	0.07

In our experiments, the destination flow-based signals failed to identify the 3rd attack in the *Access link* trace. This attack consists of a concurrent host scan aimed at specific destinations (high threshold), and the SQL Slammer worm which targeted random machines (low threshold). These two simultaneous conflicting attacks complicate the detection by offsetting the address distribution characteristics of each other. As a result, it shows that composite attacks may require multiple signals for analysis. The multidimensional signal is motivated from these observations.

4.3.3 **Packet Distribution in Port Domain.** Besides address domain, we could analyze and visualize the packet header information in port number domain. An analysis of the port number-based image can reveal portscan types of attacks. When a machine is the target of a portscan, the distribution of the exploited port numbers would deviate from its normal distribution.

Table 5. Results of Port distribution signals

Domain	TP β	FP α	LR	NLR
Source Port	83.4%	0.14%	594.1	0.17
Dest. Port	96.2%	0.17%	571.1	0.04
(Source, Dest)	96.8%	0.25%	383.2	0.03

The results in Table 5 indicate that port-based signal could be a powerful signal for anomaly detection achieving detection rates of up to 96% with very low false alarm rates. The simultaneous improvement in both the measures is a result of the nature of attacks which probe accessible ports in random or dictionary fashion for infiltration.

4.4 Sensitivity of Signals to Thresholds

In order to evaluate the effectiveness of employing different thresholds, we compare the detection results of schemes employing the image-based analysis with a volume-based signal, especially the number of flows. The anomaly detection measurements in combination of source and destination domains are shown in Fig 1. At medium confidence levels (3σ), the four kinds of image-based analyses (group 2 to group 5) do not offer significant advantage in detection rates over a volume-

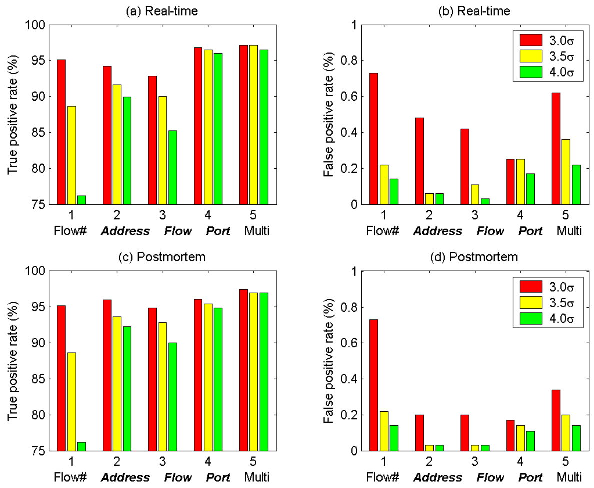

Figure 1. The impact of thresholds on results.

based signal of flow counting. However, the image-based signals offer significant advantage in false alarm rates. When higher confidence levels (3.5σ ~ 4.0σ) are considered (for decreasing the false alarm rates further), the image-based signals provide significantly better detection results than the flow counting approach. This clearly shows that the distribution-based signal offers more significant improvement in the detection of anomalies than volume-based signal.

5. CONCLUSION

In this paper, we have evaluated a number of signals proposed for detecting traffic anomalies. We evaluated the signals on three different real traces. We employed statistical techniques for unknown attack detection. Our evaluations indicate that the distribution-based signals, which track variations of distributions of the underlying traffic header domains, provide more reliable signals than the volume-based signals of traffic volume.

6. REFERENCES

[1] C. Estan, S. Savage and G. Varghese, "Automatically Inferring Patterns of Resource Consumption in Network Traffic", in Proc. of ACM SIGCOMM, August 2003.

[2] S. Staniford, V. Paxson and N. Weaver, “How to own the Internet in your spare time”, in Proc. Of 11th USENIX Security Symposium, Aug. 2002.

[3] P. Barford, et al, “A Signal Analysis of Network Traffic Anomalies”, Proc. of IMC workshop, Nov. 2002.

[4] Seong Soo Kim, A. L. N. Reddy and M. Vannucci, “Detecting traffic anomalies through aggregate analysis of packet header data”, in Proc. of Networking 2004, pp.1047-1059, May 2004.

[5] Dave Plonka, “FlowScan: A Network Traffic Flow Reporting and Visualization Tool”, Proc. of USENIX LISA Dec. 2000.

[6] P. Barford and D. Plonka, “Characteristics of Network Traffic Flow Anomalies”, in Proc. of ACM Internet Measurement Workshop (IMW), October, 2001.

[7] A. Hussein, J. Heidemann, and C. Papadopoulus, “A framework for classifying denial of service attacks”, ACM SIGCOMM, Aug. 2003.

[8] Robert V. Hogg and Allen T. Craig, Introduction to mathematical statistics, Macmillan Company, 1965.

[9] Martin Roesch, “Snort-Lightweight Intrusion Detection for Networks”, in Proc. of USENIX LISA ’99, November 1999.

[10] David Moore, et al, “Internet Quarantine: Requirements for Containing Self-Propagating Code”, Proc. INFOCOM, Apr. 2003

[11] C. Estan and G. Varghese, “New Directions in Traffic Measurement and Accounting”, in ACM SIGCOMM, Aug. 2002.

[12] M. Durand and P. Flajolet, “Loglog Counting of Large Cardinalities”, Proc. *11th Annual European Symposium on Algorithms* (ESA03)., Sept. 2003.

[13] S. Kim and A. L. N. Reddy, "A Study of Analyzing Network traffic as Images in Real-Time",Proc. INFOCOM, Mar. 2005.

[14] A. Lakhina, M. Crovella and C. Diot “Diagnosing network-wide traffic anomalies”, in ACM SIGCOMM, Sept. 2004.

[15] MathWorld, Web extensive mathematics resource, http://mathworld.wolfram.com/LikelihoodRatio.html.

[16] KREONet2 (Korea Research Environment Open NETwork2). Available: http://www.kreonet2.net.

[17] National Laboratory for Applied Network Research, “NLANR Network Traffic Packet Header Traces”, August, 2002.

[18] N. Weaver, S. Staniford and V. Paxson, “Very fast containment of scanning worms”, Proc. of Usenix Security Symp, Aug. 2004.