Network Traffic Behavior Analysis by Decomposition into Control and Data Planes

Basil AsSadhan, Hyong Kim, José M. F. Moura, Xiaohui Wang
Carnegie Mellon University
Electrical and Computer Engineering Department
5000 Forbes Ave, Pittsburgh, PA, USA
bsadhan, kim, moura, xiaohuiw@ece.cmu.edu

Abstract

In this paper, we analyze network traffic behavior by decomposing header traffic into control and data planes to study the relationship between the two planes. By computing the cross-correlation between the control and data traffics, we observe a general 'similar' behavior between the two planes during normal behavior, and that this similarity is affected during abnormal behaviors. This allows us to focus on abnormal changes in network traffic behavior.

We test our approach on the Network Intrusion Dataset provided by the Information Exploration Shootout (IES) project and the 1999 DARPA Intrusion detection Evaluation Dataset from the MIT Lincoln Lab. We find that TCP control and data traffic have high correlation levels during benign normal applications. This correlation is reduced when attacks that affect the aggregate traffic are present in the two datasets.

Keywords— Network traffic analysis, cross-correlation function, abnormal behavior, anomaly detection, long-range dependence.

1 Introduction

It is common in today's computer networks to have traffic abnormal behaviors such as attacks, failures, and malfunctions. Such abnormal behaviors negatively impact the network's operation and security, and cost financial losses[1]. Network traffic analysis is important, as it helps to detect these abnormal behaviors in a timely manner. However, the continuous emergence of diverse set of applications and the large amount of network traffic hinder the process. This requires continuous characterization/modeling of normal and/or abnormal behaviors, which is difficult to achieve, specifically for normal behavior.

Detection schemes based on the continuous characterization of *abnormal* behavior resulting from attacks are typically known as signature based detectors [6]. Detection schemes based on characterizing normal behavior are typically known as anomaly based detectors [6]. We use a less general approach than network anomaly based detection. We look for violations of specific behaviors that affect the aggregate traffic behavior as opposed to violations that are carried in the content of one or few packets, where signature based detection techniques are used to detect the latter type. A description of a bad behavior affecting the aggregate traffic behavior (e.g., number of TCP SYN packets, data traffic rate) is predefined on a heuristic basis at the detection system. Such approaches that look for specific behaviors are referred to by some (e.g., [3]) as *network based behavioral approaches*.

We propose to detect such network violations through network traffic analysis. One of the main characteristics of our work is that we decompose network header traffic into control and data planes. This enables us to reduce the amount of traffic to look at as data traffic generation is based on control traffic generation. Hence, we anticipate that the two traffic groups should have 'similar' behaviors during benign normal applications. In addition, since certain network abnormal behaviors appear mainly at the control plane, we expect the behavior of the control and data traffic to differ during these abnormal behaviors. We compute the cross-correlation function between the control and data traffics to study the similarity between different groups of traffic during normal behavior and their dissimilarities during abnormal behavior.

We report our results on the Network Intrusion Dataset provided by the Information Exploration Shootout (IES) project [2] and the 1999 DARPA Intrusion detection Evaluation Dataset from the MIT Lincoln Lab [4]. We found

[1] The CSI/FBI annual Computer Crime and Security Survey addresses the amount of financial losses [1].

that certain types of network attacks affecting the aggregate traffic decrease the correlation between the control and data traffics in the bidirectional traffic between the enterprise LAN and the external network.

Previous work in network and host anomaly based detection have looked into analyzing the observed data (e.g., network traffic, system calls, etc). However no available study, to the best of our knowledge, has looked at the control and data planes separately. The closest one to us that we are aware of, is the work done by Kompella et. al in [3]. They introduce a new data structure, *Partial Completion Filters* (PCF), to detect partial completion[2] and scanning attacks. They look at the aggregate traffic behavior, count the number of SYN packets, and compare it to the count of FIN packets for a given destination or source host. PCF differs from our approach: it is more specific since it looks only at the relation within the control traffic (i.e., the SYN and FIN packets) as opposed to the relation between the control and data traffic. This implies its detection would be faster since it processes a much smaller number of traffic packets; but its detection capability and scope is much less than ours, since it only looks at control packets.

The rest of the paper is organized as follows: Section 2 explains our methodology in decomposing network traffic into control and data and the techniques we use to detect abnormal behavior. The datasets used in the study are described in Section 3 along with the needed analysis and preprocessing. We examine the effect of abnormal behaviors on the cross-correlation function between the control and data traffic sequences in Section 4. Conclusions are presented in Section 5.

2 Methodology: Decomposing Network Traffic into Control and Data

We decompose the network header traffic in a given time interval into control and data planes to be able to observe any dissimilarities between the traffic in these two planes. Such dissimilarities are an indication of abnormal behavior in that time interval. We decompose the header traffic into control and data because we assume that data traffic generation is based on control traffic generation. In fact we focus on enterprise LANs, where control traffic are the packets that set, maintain, or tear down a connection, and data traffic are those packets that are concerned with the actual transmission of data. Thus, during normal behavior the two traffic groups should be correlated. This correlation is expected to be affected by network abnormal behaviors that manifest themselves mainly at one of the two planes (typically the control plane) and not the other.

We only consider TCP traffic as it constitutes the majority of Internet traffic, and it can be easily decomposed into control and data planes using the flag and sequence numbers fields in the TCP packet's header. Packets having any of the following flags are treated as control packets: SYN, FIN, or RST. In addition, pure acknowledgment packets that have empty sequence number fields are also treated as control packets. All other TCP packets are treated as data packets.

To observe the effect of abnormal behaviors on the similarity between the control and data traffic, we develop an anomaly based detector that compares the two traffic groups and measures their dissimilarity. We first extract from the packets arriving at the network traffic monitor a set of discrete time sequences. This is done by aggregating different packet header-related features over a suitable aggregation interval. These discrete time sequences will enable us to apply our method of comparison. We list four issues that need to be identified to develop the anomaly based detector. They are:

1. The aggregation interval from which features are extracted.
2. The set of extracted features to be used in the comparison.
3. The size of the time-window over which the comparison takes place.
4. The method of comparison.

The aggregation interval to extract features from the packets header is selected based on the packet rate of the traffic, such that the packet count variability (variance) is high. This implies that the aggregation interval should not be too small nor too large, as either will result in low variability.

Features are extracted by counting a set of packet header-related fields that are originating or destined to a certain host(s) over the selected aggregation interval to produce several time sequences. We refer to these features as count-features. The count-features that we have used include: the number of packets, the number of bytes, and the number of different addresses. We select these features for the simplicity and their usage by others.

The size of the time-window over which the comparison takes place depends on the method of comparison. However, there are generally two bounds:

1. A lower bound to have a large enough time-window to notice the similarity between the control and data during normal benign traffic.

[2] Partial completion attacks are also known as claim-and-hold attacks. The attack basically attempts is to claim a resource and hold it by not releasing it, hence denying service to legitimate user (e.g., TCP SYN flooding attacks) [3].

2. An upper bound in order to notice an abnormal behavior with a short duration in the traffic and not miss it due to being suppressed by the remaining background traffic.

These two bounds are related to the traffic rate, the higher the traffic rate is; the tighter the two bounds would be. In selecting this window, there is a trade-off between reducing the false positive rate by selecting a large window and reducing the false negative rate by selecting a small window. In addition, selecting a small window is useful in narrowing down the time interval where the abnormal behavior took place. A sliding period (e.g., 1% of the window's size) is used to slide the window to have faster detection of abnormal behavior rather than waiting for the next whole window to pass.

Next we describe one of the methods of comparison that we have conducted, which is anomaly detection through observing the time variation of the cross-correlation function between the control and data traffic sequences.

2.1 Anomaly detection: Time Variation of Cross-Correlation

We observe the effect of abnormal behaviors on the similarity between the control and data traffic sequences, for different count features, by estimating the cross-correlation function between the two sequences. We use the correlation function as they are used in many fields such as pattern recognition to compare between different objects or signals.

Let X_i and Y_i represent a count feature sequence for the control and data traffic, respectively. Define $\mathbf{X}_{N,n_0}$ and $\mathbf{Y}_{N,n_0}$, where each one is a vector of size N starting at element n_0 and ending at element $n_0 + N - 1$. The cross-correlation function ρ_{XY} is then computed using a correlation window of size N and as a function of the starting element n_0 by:

$$\rho_{XY}(N, n_0) = \frac{< \mathbf{X}_{N,n_0} - \mu_{\mathbf{X}_{N,n_0}}, \mathbf{Y}_{N,n_0} - \mu_{\mathbf{Y}_{N,n_0}} >}{||\mathbf{X}_{N,n_0} - \mu_{\mathbf{X}_{N,n_0}}|| \cdot ||\mathbf{Y}_{N,n_0} - \mu_{\mathbf{Y}_{N,n_0}}||}$$

$$= \frac{\Sigma_{i=0}^{N-1}(X_{i+n_0} - \mu_{\mathbf{X}_{N,n_0}})(Y_{i+n_0} - \mu_{\mathbf{Y}_{N,n_0}})}{\sqrt{\Sigma_{i=0}^{N-1}(X_{i+n_0} - \mu_{\mathbf{X}_{N,n_0}})^2 \Sigma_{i=0}^{N-1}(Y_{i+n_0} - \mu_{\mathbf{Y}_{N,n_0}})^2}},$$

where $\mu_{\mathbf{X}_{N,n_0}}$ and $\mu_{\mathbf{Y}_{N,n_0}}$, are the estimated means of $\mathbf{X}_{N,n_0}$ and $\mathbf{Y}_{N,n_0}$, respectively.

The cross-correlation function ρ_{XY} ranges between -1 and 1, where a high value near 1 indicates high correlation, a low positive value near 0 indicates low correlation, and a value equal or less than 0 indicates no correlation in our case study.

3 Datasets

3.1 The Network Intrusion Dataset

We obtained the network intrusion dataset provided by the Information Exploration Shootout (IES) project [2]. The network traffic is collected at the gateway connecting an enterprise LAN with the external network (Internet) using tcpdump. Only *header* information of the packet that passed by the network interface of the gateway was captured by tcpdump. This included the communication between the enterprise LAN and the external network, and the traffic communication within the LAN. The filters of tcpdump were specified to only collect TCP and UDP packets but we only consider TCP packets.

To protect the identity of the hosts that were communicating with each other while the network traffic was collected, the IP address of each *external* host is *anonymized*. All *internal* hosts, however, are anonymized to a *single* IP address to ensure the privacy of the network by not revealing the network topology. This makes the analysis of the traffic communication *within* the LAN part very limited.

The dataset consists of four files[3], the first is believed to be free of attacks and the rest contain attacks that were *simulated* and stored, each file containing instances of a single different attack behavior. The three attacks stored in the files are: TCP SYN flooding Denial of Service (DoS), Password Guessing, and Port Scanning. Each file has 16-20 minutes of traffic, and the average total bit rate in the four files, based on our calculations, is on the order of 1 Mbps. Although the attacks are inserted by the providers of the dataset, their times and targets are not provided.

3.2 The 1999 DARPA Intrusion detection Evaluation Dataset

We obtained the 1999 DARPA Intrusion detection Evaluation Dataset from the MIT Lincoln Lab [4]. The network traffic in this dataset is synthetically generated on an isolated testbed by simulation. To provide realistic background traffic, a traffic model is developed based on actual network activity observed at a number of locations, including sites monitored by the Air Force Information Warfare Center, data collected directly from an active Air Force Base network and information gathered by Lincoln Labs staff from selected location. The traffic model included a range of application, such as FTP, HTTP, telnet, POP3 and SMTP Email, SQL queries, ICMP and finger traffic, SSH sessions and IRC connections. Based on the observed traffic patterns, each of the application-level protocols is modeled independently.

[3]There is actually an additional fifth file, but it is surprisingly almost identical to the second file, therefore, we do not consider it.

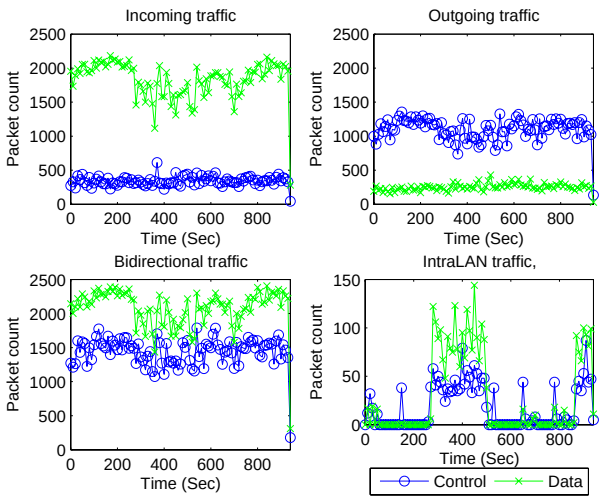

Figure 1. Packet counts of incoming, outgoing, bidirectional, and intraLAN of Base file of the IES dataset. Aggregation interval: 10 seconds.

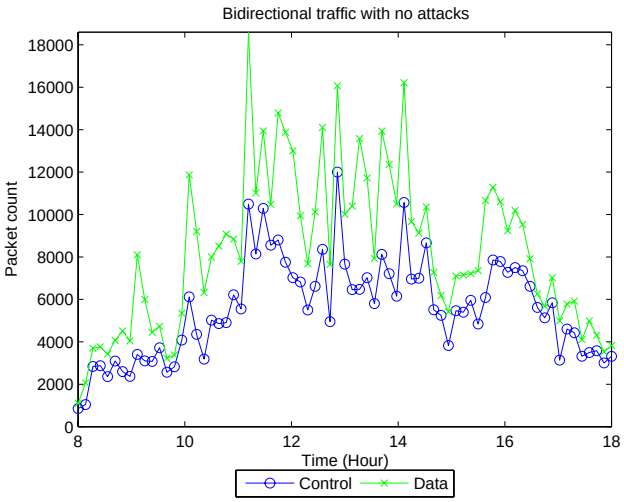

Figure 2. Packet count of bidirectional traffic of Tuesday of week 1 of the 1999 DARPA dataset. Aggregation interval: 500 seconds.

The traffic generation process was automated by executing scripts on selected hosts that initiated connections and provided interaction among those hosts and other network nodes. The traffic model was used to determine when connections would be established. For each protocol or service the model consisted of a set of mean arrival rates, collected in 15 minute intervals, and those rates were used in each interval to generate Poisson connection arrivals. The traffic generation system was run each weekday for five weeks. Three weeks of training data were collected from March 1 to March 20, 1999 and two weeks of test data were collected from March 29 to April 10, 1999. The average total bit rate, based on our calculations, is on the order of 50 Kbps

The network traces generated during the first and third weeks do not include any attack. The second week contains several attacks for training purpose and the last two weeks of the testing data contain 201 instances of about 56 types of attacks. These attacks are listed in a separate file with related information, such as start time, duration, destined host, attack name, etc.

There are four main categories of attacks as described in the attack files: Denial of Service (DoS), User to Root (U2R), Remote to Local (R2L) and Probes. More than 10 different attacks are simulated among each of these four categories. We note that the categorization is based on the attacker's intent, i.e., the effect that the attack leaves on the system rather than the way that the attack manifested itself in the observed system [5].

The dataset consists of many sets, and we look only at the tcpdump data that was collected by a sniffer located at the external side of the LAN. The tcpdump data contains all of the traffic generated inside the LAN destined to the outside world and all of the traffic generated outside the LAN destined to it. As in the IES dataset, we only consider here TCP header traffic.

A main contrast between this dataset and the IES dataset is that it is synthetically generated on an isolated testbed by simulation. This implies that its normal traffic is guaranteed to be free of attacks, however, it does not represent real network traffic accurately.

3.3 Analysis of Datasets

In this section, we plot the time sequences of the packet and different address counts as a visual way to observe the level of similarity between the control and data traffic in different directions of the traffic.

For this purpose, we use large aggregation intervals, 10 seconds in the IES dataset, and 500 seconds in the 1999 DARPA dataset, to make the graphs easier to visual. The two intervals are selected based on the traffic durations in the files that are considered in the two datasets, 20 minutes in the IES dataset, and 10 hours in the 1999 DARPA dataset.

Figure 1 shows the time sequences for the packet count for the base file[4]. We notice that the control and data traffic sequences at the incoming traffic (top-left) and the outgoing traffic (top-right) do not vary similarly. However, when

[4]The base file has only instances of normal traffic, no attack traffic is simulated.

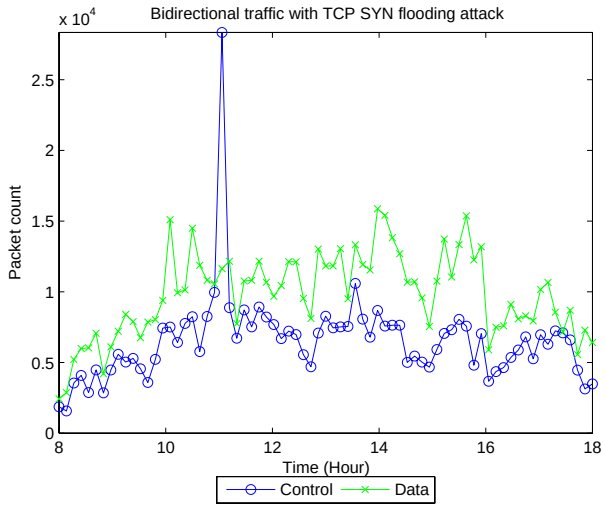

Figure 3. Packet count of bidirectional traffic of Thursday of week 2 of the 1999 DARPA dataset. Aggregation interval: 500 seconds.

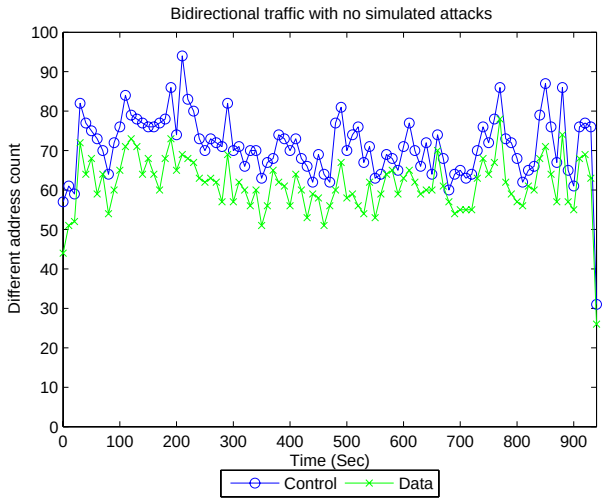

Figure 4. Different address counts of bidirectional of Base file of the IES dataset. Aggregation interval: 10 seconds.

the two traffic sequences are combined in the bidirectional traffic (bottom-left), the variation is similar. This should be expected, since an increase in the control outgoing traffic (e.g., visiting a website) should cause an increase in the incoming data and not in the outgoing data traffic (e.g., downloading website's content), and vice versa. The variation of the control and data traffic sequences is also similar in the intraLAN traffic (bottom-right). These similarities indicate that the criteria used in decomposing the traffic into control and data traffic is successful. In the incoming traffic (top-left), the data traffic rate is higher than the control traffic rate, whereas in the outgoing traffic (top-right), the control traffic rate is higher than the incoming traffic rate. This is also expected, since at an end user most of the downstream traffic is data traffic, whereas the upstream traffic has more control traffic than data traffic. Finally, in the intraLAN traffic (bottom-right), the packet count has zero values sometimes. This is because there is no internal communication over the corresponding 10-second aggregation interval.

Figure 2 shows the bidirectional traffic of Tuesday of week 1 in the 1999 DARPA dataset. The traffic of this day is free of any attacks, hence, we also notice the similar variation between the control and data traffic sequences. The effect of a TCP SYN flooding attack on the similarity between the control and data traffic sequences is shown Figure 3. The figure shows the time sequences of the packet count for the bidirectional traffic for Thursday of week 2 in the 1999 DARPA dataset, which has instances of normal traffic along with attack traffic. We notice that the control and data traffic sequences vary similarly most of the time except for certain points (e.g., around 11:00 am) where an attack has occurred.

Figure 4 shows the time sequences for the different address count for the bidirectional traffic on the base file in the IES dataset. We remind the reader that the whole enterprise LAN is represented in the dataset by one anonymized address. Therefore, the different address count here really represents the number of external hosts communicating with the LAN plus the single internal address. We notice a similar variation between the control and data traffic sequences, however, not as the one shown in the packet count case. We also examine the byte count behavior, and found it to have a similar behavior to the packet count. Except that for the byte count, the data traffic is about one order of magnitude higher than the control traffic. This is because data packets have much more bytes than control packets.

4 Anomaly detection: Time Variation of Cross-Correlation

We report the effect of several attacks in the two datasets on the similarity between the control and data traffic sequences by computing the cross-correlation function. We only consider the bidirectional traffic, since the control and data traffic do not vary similarly at the incoming and outgoing traffic. In addition, we do not consider the intraLAN traffic as it has long periods of non activity, which limit the analysis.

We first select a count aggregation interval of 1 sec-

ond for both the IES and 1999 DARPA datasets. This selection is based on the average traffic rate in each of the two datasets, 448 packets/second and 32 packets/second, respectively. Second, since the attack durations in the IES dataset are not provided, we use an arbitrary correlation window of 100 seconds with a sliding period of 1 second to compute the cross-correlation of the different counts of the bidirectional control and data. In the 1999 DARPA dataset, the attacks' execution durations range from very few seconds to about half an hour. However, the effect of the attack on the network traffic is typically longer. Therefore, we use three correlation windows to compare how they detect attacks with different durations. We use a 500, 1000, and 2000 second correlation windows with a sliding period of 1% of the window's size (i.e., 5, 10, and 20 seconds). We report our results on the IES dataset first, then the 1999 DARPA dataset.

Figure 5 shows the time variation of the cross-correlation function between the bidirectional control and data packet, byte, and different address counts of the base file in the IES dataset. The x axis represents the start of the correlation window. One can observe from the figure, that the cross-correlation functions of the packet count (top) and byte count (middle) vary similarly. This similar variation is expected due to the direct relation between the two counts. The cross-correlation function of the packet count has higher values since the packet count is less dependable on the application used. The range of the cross-correlation function is between $0.2 - 0.8$ for the packet count and $0.1 - 0.7$ for the byte count. The range for the different address count is higher and is between $0.7 - 0.9$. The positive correlation levels in Figure 5 agrees with the similarity between the packet, byte, and different address counts of the control and data traffic sequences shown in Figure 1 (bottom-left), and Figure 4. These observations will be useful in detecting attacks that manifest themselves mainly at the control traffic plane.

The same graphs are produced for the three attack files and are shown in Figures 6-8. One can see that the byte and packet counts vary similarly as well. The effect of the attacks can be seen in the parts of the graph where the cross-correlation function goes below the threshold, and below 0 in some cases. These drops can be interpreted as an attack in that interval and helps in zooming into where the attack has most probably taken place. We are not able to fully verify the accuracy of the detection method since the attacks are not labeled in the IES dataset.

The level of drop at the three counts varies depending on the type of attack. It can be seen from Figures 7 and 8, which represent the effect of the password guessing and scanning attack, respectively, that the correlation of the different address count is not affected when compared with the effect of correlation between the control and data traffic at

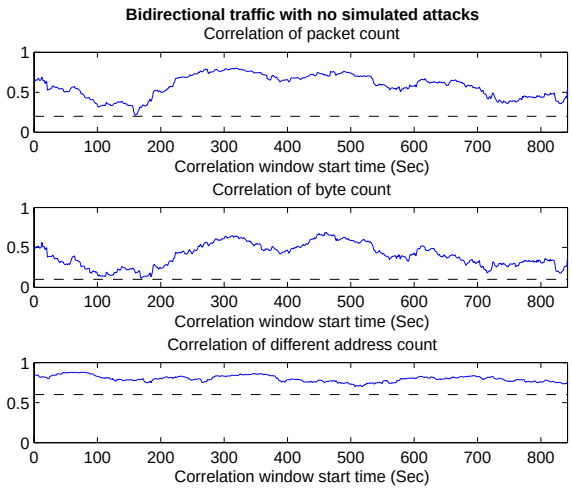

Figure 5. Time variation of the cross-correlation function between the bidirectional control and data packet, byte, and different address counts for Base file in the IES dataset. Aggregation interval: 1 second, Correlation window: 100 seconds, Sliding period: 1 second.

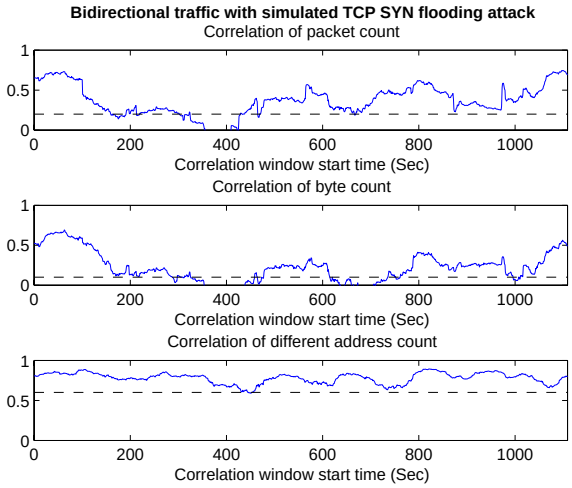

Figure 6. Time variation of the cross-correlation function between the bidirectional control and data packet, byte, and different address counts for Attack 1 file in the IES dataset. Aggregation interval: 1 second, Correlation window: 100 seconds, Sliding period: 1 second.

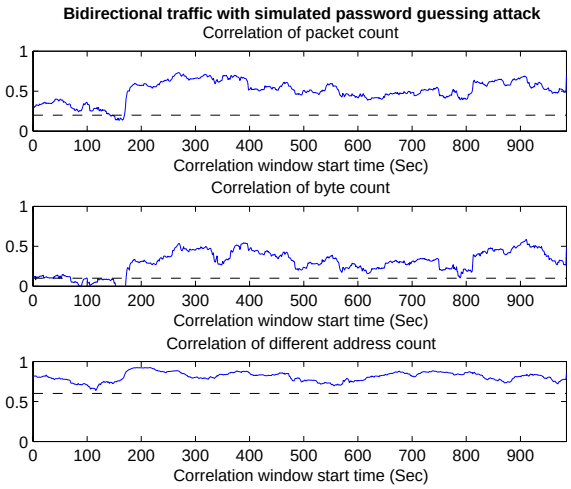

Figure 7. Time variation of the cross-correlation function between the bidirectional control and data packet, byte, and different address counts for Attack 2 file in the IES dataset. Aggregation interval: 1 second, Correlation window: 100 seconds, Sliding period: 1 second.

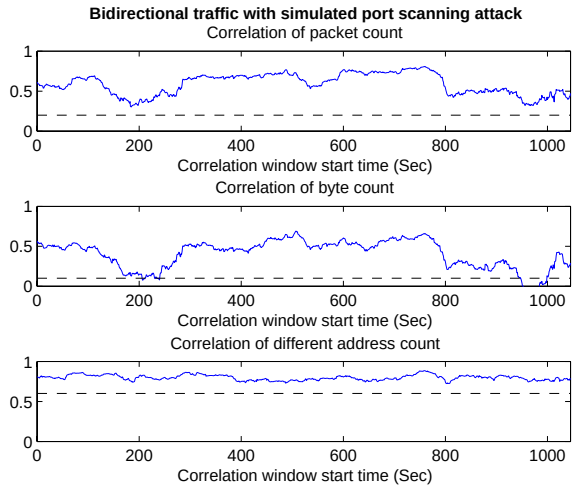

Figure 8. Time variation of the cross-correlation function between the bidirectional control and data packet, byte, and different address counts for Attack 3 file in the IES dataset. Aggregation interval: 1 second, Correlation window: 100 seconds, Sliding period: 1 second.

the byte count. This because the two attacks were only carried by one host, as described in [2]. The correlation at the packet count is less affected by the attacks when compared with the byte count, specifically at the scanning attack. This indicates the importance of looking at both counts to detect abnormal behavior. The correlation of the different address count at Figure 6, which represents the TCP SYN DoS attack, was affected a bet by the attack since the attack involved spoofing IP addresses. Using other additional counts, e.g., the count of different destination ports, to compare the effect of attacks on the correlation between the control and data would be useful in detecting certain types of attacks, e.g., port scanning attacks.

For our results on the 1999 DARPA dataset, we only discuss the packet count and one attack behavior due to the limited space. Figure 9 shows the time variation of the cross-correlation function between the bidirectional control and data packet counts for Tuesday of week 1 in the 1999 DARPA dataset, which has no attacks, using three different correlation windows (500, 1000, and 2000 seconds). As it can be seen from the figure, the correlation levels are high over the three correlations windows. The correlation levels are higher as the size of the correlation window gets larger. This is because using longer periods in the comparison between the control and data traffic sequences will suppress the effect of short-duration dissimilarities. The high correlation levels in the figure agree with the similarity between the packet count of the control and data traffic sequences shown in Figure 2.

Figure 10 shows the time variation of the cross-correlation function between the bidirectional control and data packet count for Thursday of week 2 in the 1999 DARPA dataset. As it can be seen from the figure around 11:00 am, the correlation level has a sudden drop where it reaches low values in all of the correlation window sizes. However, the smaller the window, the narrower the interval of detection is. Looking at the attacks action times in the description of the 1999 DARPA dataset, we find that a TCP SYN flooding attack took place at 11:04 am, which can also be seen from the control and data packet counts in Figure 3. Since TCP SYN packets are treated in this study as control packets, the correlation level between the control and data traffic sequences is low during the time of the attack.

5 Conclusions

We analyze network traffic behavior by decomposing header traffic into control and data planes. We focus on network abnormal behaviors that affect the aggregate traffic behavior. We explain that during normal benign applications, the generation of data traffic is based on the genera-

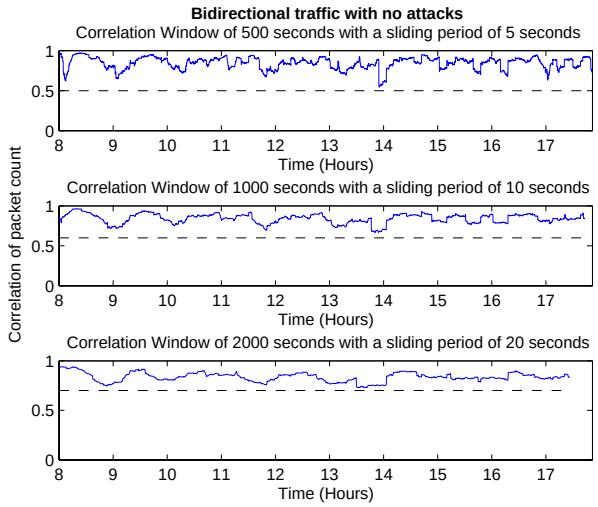

Figure 9. Time variation of the cross-correlation function between the bidirectional control and data packet count for Tuesday of week1 of the 1999 DARPA dataset. Aggregation interval: 1 second.

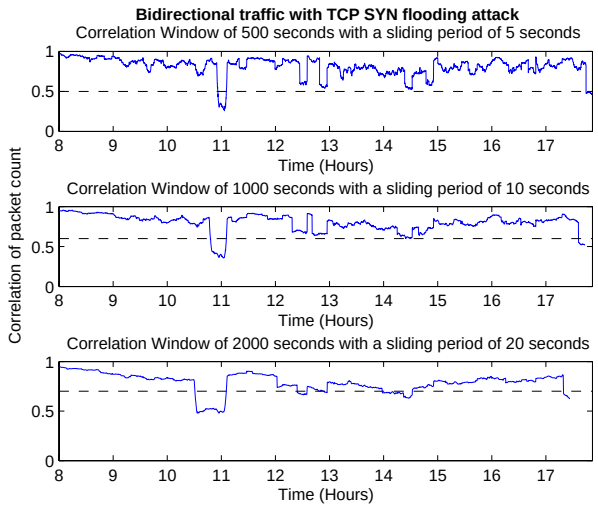

Figure 10. Time variation of the cross-correlation function between the bidirectional control and data packet count for Thursday of week2 of the 1999 DARPA dataset. Aggregation interval: 1 second.

tion of control traffic. Hence, both traffic sequences have similar time variations. We verify this with the TCP traffic of both the IES dataset and the 1999 DARPA dataset. The similarity in time variation between the bidirectional control and data traffic sequences is affected by certain simulated attacks in the two datasets. We confirm this effect by repeatedly computing the cross-correlation function between the two sequences over a sliding correlation time-window, where sudden drops in the correlation levels are caused by these attacks. This decomposition helps in detecting network abnormal behaviors that manifest themselves mainly at either the control or data traffic. In our future work, we will perform additional advanced analysis to take into account the rate of change in the cross-correlation variation on real recently collected traffic and examine the aggregate traffic per destination host.

In parallel work, we also examine the Long-Range Dependence (LRD) behavior of the traffic. We observe that the attacks in the IES dataset cause the incoming control traffic and/or the outgoing data traffic to fail to exhibit LRD behavior, while the traffic as a whole still exhibits LRD behavior. These two traffic subgroups, generally, have lower volume than the incoming data and outgoing control traffics. Hence, this will lead to significantly reducing the amount of network traffic to consider in detecting network abnormal behavior. Our analysis with the 1999 DARPA dataset does not confirm these observations, as the two traffic subgroups still exhibit LRD behavior in the presence of attacks. As a result, more traffic analysis of real collected network traffic is needed to reach more accurate conclusions regarding the LRD behavior of the control and data behavior.

References

[1] CSI/FBI survey. http://www.gocsi.com/forms/fbi/csi_fbi_survey.jhtml, 2007.

[2] Information exploration shootout project. http://ivpr.cs.uml.edu/shootout/about.html. Accessed January 2007.

[3] R. R. Kompella, S. Singh, and G. Varghese. On scalable attack detection in the network. In *USENIX/ACM Internet Measurement Conference*, Taormina, Sicily, Italy, Oct. 2004.

[4] R. Lippmann, J. Haines, D. Fried, J. Korba, and K. Das. The 1999 DARPA off-line intrusion detection evaluation. In *the 3rd International Workshop on Recent Advances in Intrusion Detection (RAID 2000)*, New York, NY, USA, Oct. 2000.

[5] J. McHugh. Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by lincoln laboratory. *ACM Transactions on Information and System Security*, 3(4):262–294, Nov. 2000.

[6] J. McHugh. Intrusion and intrusion detection. *International Journal of Information Security*, 1(1):14–35, Aug. 2001.