Anomaly Detection Systems for Distributed Denial of Service Attacks

by

Assad Raza
B.Sc., University of Sindh, Pakistan, 2012

A Report Submitted in Partial Fulfillment of the Requirements for the Degree of

MASTER OF ENGINEERING

in the Department of Electrical and Computer Engineering

Anomaly Detection Systems for Distributed Denial of Service Attacks

by

Assad Raza
B.Sc., University of Sindh, Pakistan, 2012

Supervisory Committee

Dr. T. Aaron Gulliver, Supervisor
(Department of Electrical and Computer Engineering)

Dr. Samer Moein, Departmental Member
(Department of Electrical and Computer Engineering)

Supervisory Committee

Dr. T. Aaron Gulliver, Supervisor
(Department of Electrical and Computer Engineering)

Dr. Samer Moein, Departmental Member
(Department of Electrical and Computer Engineering)

ABSTRACT

Distributed Denial of Service (DDOS) attacks persist and are growing stronger. According to the latest data, 2016 has seen DDOS attacks which were large in both frequency and size [1]. DDOS attacks have been investigated extensively and various countermeasures have been proposed to protect networks from these attacks. However, DDOS is still considered to be the major threat to current networks and there is a need for Anomaly Detection Systems (ADSs) to accurately detect DDOS attacks. Furthermore, network traffic now has significant Peer to Peer (P2P) traffic. P2P traffic in Europe accounts for more than a quarter of all bandwidth, and 40 percent of all packets sent. Previous work has shown that P2P traffic can have a negative impact on the accuracy of ADSs. A P2P traffic preprocessor was proposed in [2] to compensate for the adverse impact of P2P traffic on ADSs. In this project, two well-known anomaly detectors, namely Network Traffic Anomaly Detector (NETAD) and Maximum Entropy Anomaly Detector (MaxEnt), are evaluated with and without this P2P traffic preprocessor for the detection of DDOS attacks. Performance of these ADSs has also been evaluated for the detection of TCP and UDP flood Denial of Service (DOS) attacks. Results are presented which show that using this P2P traffic preprocessor improves the ability of these ADSs to detect attacks.

Contents

List of Tables

List of Figures

Glossary

DOS ..... Denial of Service

DDOS ..... Distributed Denial of Service

ADS ..... Anomaly Detection System

IDS ..... Intrusion Detection System

P2P ..... Peer to Peer

TCP ..... Transmission Control Protocol

UDP ..... User Datagram Protocol

ICMP ..... Internet Control Message Protocol

ARP ..... Address Resolution Protocol

IP ..... Internet Protocol

MAC ..... Media Access Control

IT ..... Information Technology

DNS ..... Domain Name System

SDN ..... Software Defined Networking

IoT ..... Internet of Things

MaxEnt ..... Maximum Entropy Anomaly Detector

NETAD ..... Network Traffic Anomaly Detector

ROC ..... Receiver Operating Characteristics

ACKNOWLEDGEMENTS

I would like to thank:

Sardar Ali and Adil Sulehri, for supporting me in the low moments.

Dr. T. Aaron Gulliver, for mentoring, support, encouragement, and patience.

Dr. Samer Moein, for his time and guidance.

My Siblings, for their love and motivation.

Logic will get you from A to B. Imagination will take you everywhere.
Albert Einstein

DEDICATION

This work is dedicated to my parents for their endless love, support, and encouragement. Thank you for teaching me to believe in myself, and in my dreams.

Chapter 1

Introduction

Recent reports show a substantial increase in broadband internet connectivity and an exponential growth in IT infrastructure worldwide [3]. The internet provides the means of communication for business growth. This readily available network connectivity has resulted in operational benefits and increased efficiency. However, systems connected to the internet are vulnerable to network intrusion and attacks. These attacks have grown in frequency and potential to harm networks in recent years.

The most familiar threats to a network are from malware, spam, phishing, Denial of Service (DOS), and Distributed Denial of Service (DDOS) attacks. They cause financial losses in the billions of dollars to businesses and pose a serious threat to information confidentiality [4]. The associated service disruptions can result in long-term loss of credibility. Unlike other threats, DOS and DDOS attacks do not attempt to breach a security perimeter. Rather, they make network resources unavailable to users. DOS attacks can last for long periods of time, making them the most destructive type of attack to an organization.

It is estimated that Peer to Peer (P2P) traffic consumes about 40%-60% of the total bandwidth used by an average internet user. P2P file sharing traffic, excluding commercial video streaming applications, has increased substantially over the years and was reported to be 13,797 petabytes (PB) per month in 2015 [5], as shown in Figure 1.1. Efforts have been made to regulate P2P traffic by deploying traffic classifiers [5], [6] which has resulted in a decrease in the overall volume of P2P traffic. However, the proportion of P2P traffic in current internet traffic is forecast to remain unchanged over the next few years [5].

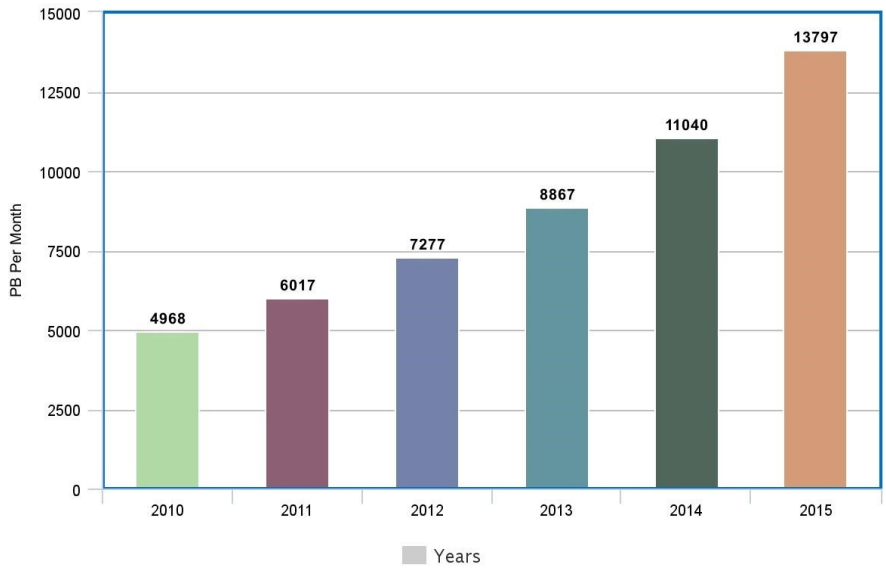

Figure 1.1: Global consumer P2P file sharing traffic over the last 5 years [5].

Although P2P traffic is a large proportion of current internet traffic, minimal work has been done to address its impact on network intrusion detection, particularly when anomaly detectors are used to detect DDOS attacks. Furthermore, P2P traffic has features which overlap with many of those used by anomaly detectors to identify malicious traffic. These features include:

- multiple connection attempts;
- large number of failed connections;
- heavy flow at connection startup.

These similarities between P2P and attack traffic makes it difficult for anomaly detectors to distinguish between legitimate and attack traffic, thus degrading the performance of Anomaly Detection Systems (ADSs). Performance measures commonly employed are the percentage of attack traffic detected (detection rate), and the percentage of legitimate traffic incorrectly marked as attack traffic (false positive rate). Efforts made in evaluating the performance degradation of ADSs in the presence of P2P traffic [2] reveal that:

- P2P traffic causes a significant degradation in detection accuracy for all attack types and classes;
- P2P traffic has a large impact on the detection of low rate attacks;
- training ADSs with a dataset containing P2P traffic does not improve detection accuracy because of the overlapping features of P2P and attack traffic.

1.1 Intrusion Detection Systems

Intrusion detection techniques can be classified as signature-based or behavior-based. Signature-based detection systems use known patterns of attacks to match and identify intrusion events. These patterns are referred to as signatures. On the other hand, anomaly-based detection systems detect deviations from normal or expected behavior to identify an attack on the network without any prior knowledge of the attack [7].

An ADS considers any traffic that is new or unusual as suspicious. They are created by modeling what is normal for the observed network and then deciding which activities should be flagged as abnormal. The main advantage of anomaly detection is that it does not require prior knowledge of intrusion and can thus detect new intrusions. The main disadvantage of ADSs is their inability to provide details about the attack [8].

In this report, the performance of the two most common anomaly based intrusion detection systems, Maximum Entropy Anomaly Detector (MaxEnt) and Network Traffic Anomaly Detector (NETAD), is evaluated. Results are presented which show that without effective discrimination between P2P and attack traffic, the attack detection capabilities of these ADSs are not effective.

1.2 Problem Statement

Denial of Service (DOS) attacks are popular due to their ability to significantly affect networks. DOS, as the name implies, exhausts the resources of the target network by sending invalid traffic. A DOS attack when carried out by multiple devices is known as a Distributed Denial of Service (DDOS) attack. DDOS is considered to be the biggest threat for networks. Thus, training ADSs to detect DDOS attacks is a critical requirement for current and future networks. The presence of a huge volume of P2P traffic in current internet traffic is considered to be a significant barrier for ADSs to provide accurate results [5], [6]. To overcome this barrier, ADSs must discriminate between P2P and attack traffic for effective protection against DDOS attacks.

Relevant work includes the development of several ADSs over the last decade [7]-[10]. Some traffic features were identified in [2] that do not overlap between P2P and attack traffic. Later, these features were used to propose a P2P traffic preprocessor. This preprocessor improved the overall detection accuracy of ADSs. However, the effectiveness of the P2P traffic preprocessor was tested only against Portscan, an attack to scan for open ports on the targeted system, and DOS attacks. Traffic analysis and anomaly detection was performed by classifying traffic according to their distinguishing features in [13]. Evaluation of different anomaly detection algorithms in a Software Defined Network (SDN) was done in [14]. The MaxEnt, NETAD, Threshold Random Walk with Credit Based rate limiting (TRW-CB), and Rate Limiting (RL) ADSs were evaluated in [14]. Implementation of ADSs and their performance evaluation for Transmission Control Protocol (TCP) flood attacks, User Datagram Protocol (UDP) flood attacks, and Portscan attacks were performed in [1]-[3]. Policies were introduced in [16] to protect a P2P network from DDOS attacks.

The goal of this project is to evaluate the performance of MaxEnt and NETAD ADSs. These ADSs use different detection frameworks to identify anomalies. They are tested under normal operating conditions (normal) and in combination with a P2P preprocessor (adaptive) to detect DDOS attacks in the presence of P2P traffic. The performance of these ADSs is also evaluated against TCP and UDP flood DOS attacks.

1.3 Report Organization

Chapter 1 provided an introduction to DOS and DDOS attacks, anomaly-based intrusion detection systems, and the presence of P2P traffic in current networks. The impact of P2P traffic on the performance of anomaly detection systems was also described.

Chapter 2 introduces Distributed Denial of Service (DDOS) attacks, its different types, impact, and evaluation.

Chapter 3 describes ADSs and the P2P traffic preprocessor used in this project for the detection of attacks.

Chapter 4 presents a description of the datasets used in this project, and the operations carried out to make the datasets suitable for testing and evaluation. It also provides the system configuration on which the tests were conducted. Receiver Operating Characteristics (ROC) curves are presented to evaluate the performance of MaxEnt and NETAD ADSs against DDOS, and TCP flood and UDP flood DOS attacks.

Chapter 5 concludes the report and provides suggestions for future work.

Chapter 2

Distributed Denial of Service Attacks

Distributed Denial of Service (DDOS) attacks overwhelm the capacity of a target network to handle traffic and provide services to legitimate users. The idea is to send enough invalid network traffic to the target to exhaust its resources. It differs from traditional DOS attacks as the attack is launched from multiple devices, called botnets, that are distributed over the internet as illustrated in Figure 2.1. This form of attack is tough to deflect, mostly due to the sheer number of devices. DDOS attacks also enable remote administration of the attack.

In 2014, DDOS attacks were the most popular amongst attackers [4]. Improvements to increase the potential of DDOS attacks have made it possible for an attacker to reach an attack traffic volume of 100 GBps. The current trend is to attack connected resources instead of attacking the target directly. The most obvious approach for this is to attack the Domain Name System (DNS) server responsible for resolving IP addresses. If these servers do not respond, then users will not be able to access the resources associated with the servers as they do not know the corresponding IP addresses.

A prime example of a DNS server attack is the recent attack against Dyn, a company that provides core internet services for Twitter, SoundCloud, Spotify, and Reddit. This is the most disruptive DDOS attack in history. Attackers used the Mirai botnet to compromise and leverage Internet of Things (IoT) CCTV cameras and then used them to attack the DNS servers of Dyn.

DDOS attacks are launched using botnets to flood the target. Recently, attackers

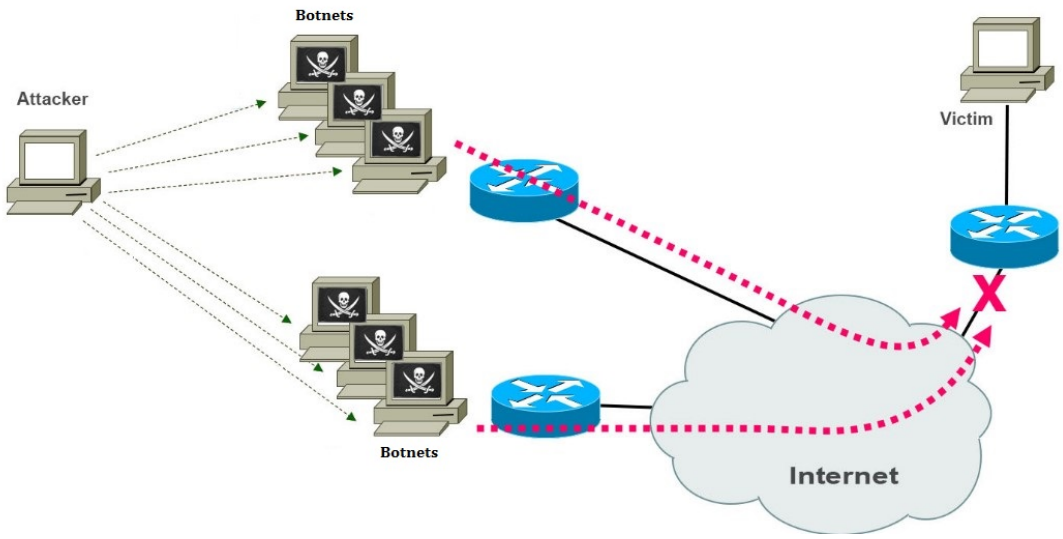

Figure 2.1: An example of a Distributed Denial of Service (DDOS) attack.

used botnets which can generate attacks with high bandwidth to increase the impact of the attack. Furthermore, most attackers now use P2P networks to make their attack infrastructure highly resilient to detection. To make it tough for signature-based detection to filter the traffic, new attack scripts have been developed which randomize every part of the attack traffic.

Attackers are interested in attacking the target using unprotected or less secure resources available over the internet such as IoT security cameras. The attack on Dyn is an example of this. Attackers have also gained an increased interest in attacking mobile phone devices such as flooding them with inbound telephone calls.

2.1 Common DDOS and DOS Attacks

This section presents descriptions of some common DDOS and DOS attacks that have occurred over the past few years.

2.1.1 ICMP Flood Attack

An Internet Control Message Protocol (ICMP) flood attack is one of the primary methods used by attackers in 2015 as reported by Symantec's Global Intelligence Network [17], as shown in Table 2.1. ICMP flood attacks come in different forms, namely, generic ping broadcast, generic ICMP flood, and generic ICMP unreachable.

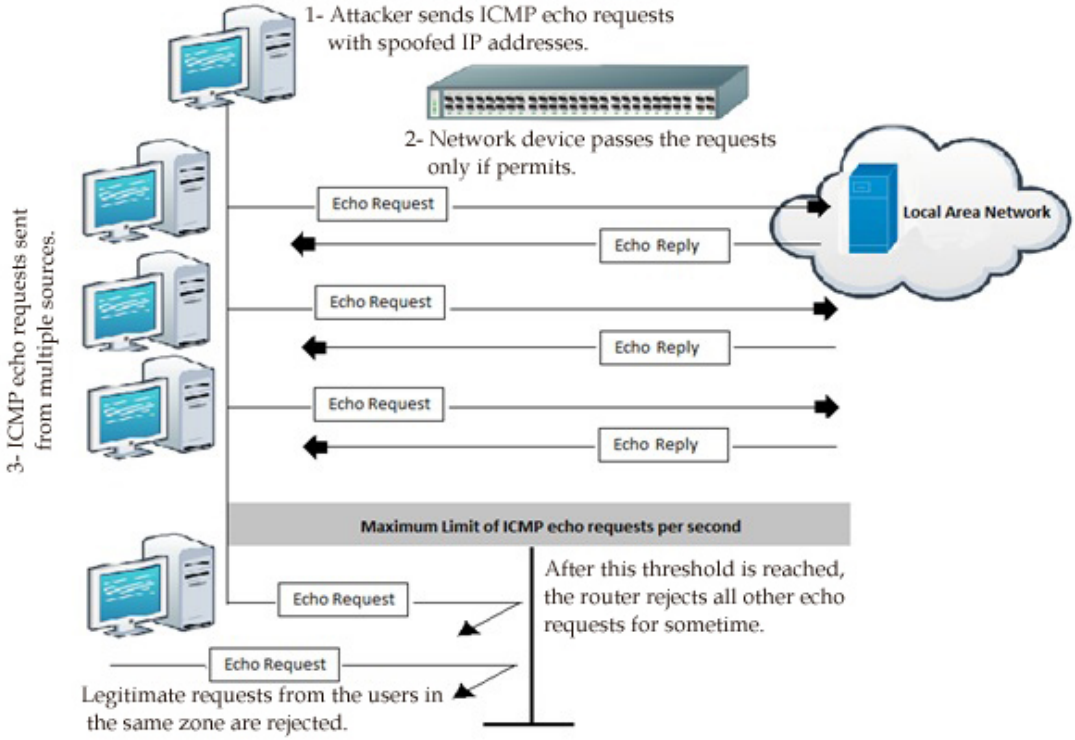

Figure 2.2: An ICMP flood attack from multiple devices to the target.

An ICMP flood attack sends large numbers of IP packets containing echo requests using fake source IP addresses as illustrated in Figure 2.2. These echo requests quickly consume the bandwidth of the network as the target expends its resources to respond which makes it unable to process legitimate network traffic.

2015 Attacks	**Attack Rate**	**2014 Attacks**	**Attack Rate**
Generic ICMP Flood	85.7%	DNS Amplification	29.4%
Generic TCP SYN Flood	6.4%	Generic ICMP Flood	17.2%
Generic Ping Broadcast	2.1%	Generic Ping Broadcast	16.8%
Generic Teardrop	2.0%	Generic ICMP Unreachable	5.7%

Table 2.1: DDOS attack trends from Symantec's Global Intelligence Network [17].

2.1.2 DNS Amplification Attack

A DNS amplification attack is carried out by sending DNS name lookup requests to a DNS server. These lookup requests are sent with a fake source IP address which is the IP address of the target. The DNS server sends the responses to these requests to the target which floods the target. It is extremely difficult to prevent these attacks because the response traffic is legitimate traffic.

2.1.3 TCP Flood Attack

A Transmission Control Protocol (TCP) flood attack is a type of DOS attack. It exploits the normal TCP three-way handshake mechanism to consume the resources of the target and render it unresponsive to legitimate users. The TCP three-way handshake process for establishing a connection between a client and server is as follows:

- client generates a connection request by sending a SYN message to the server;
- server acknowledges the request by sending a SYN-ACK message to the client;
- client responds with an ACK message to the SYN-ACK sent by the server.

In a TCP flood attack, the attacker sends TCP connection requests faster than the targeted server can process them. Flooding these requests causes the targeted server to become unresponsive to legitimate users.

2.1.4 UDP Flood Attack

User Datagram Protocol (UDP) flood is a type of DOS attack in which the attacker overwhelms random ports on the targeted server with IP packets containing UDP datagrams. The target checks for applications associated with these datagrams and if it does not find any association, a destination unreachable packet is sent to the sender. As the number of UDP datagrams received and answered increases, the targeted server becomes overwhelmed and unresponsive to legitimate users.

Chapter 3

Anomaly Detection Systems and the P2P Preprocessor

This chapter provides a description of the anomaly detection systems being evaluated and the P2P preprocessor.

3.1 Network Traffic Anomaly Detector (NETAD)

NETAD [10] operates on rule-based filtered traffic. These rules are modeled on a subset of common protocols. It computes a packet score depending on the time and frequency of each byte of the packet. A threshold is applied to the packet score to detect anomalous packets [2],[3].

NETAD uses two stages, a filtering stage to select the start of inbound client sessions, and a modeling phase. The filtering of traffic is based on the premise that the first few packets of a connection request are sufficient for anomaly detection. Thus, the filtering stage removes most of the traffic, significantly reducing the load on the modeling stage, and passing only the traffic most likely to contain evidence of attacks. The filtering stage removes the following packets.

- All non-IP packets, e.g. ARP.
- All outgoing packets.
- All TCP streams that begin with SYN-ACK.
- UDP packets to a port number higher than 1023.

- TCP packets with sequence numbers more than 100 past the initial sequence number.
- Packets addressed to any address/port/protocol combination if more than 16 have been received within the previous 1 minute interval (to limit bursts of UDP or ICMP traffic).

The second stage of NETAD models nine subsets of filtered traffic corresponding to the 9 most common packet types, for a total of $9 \times 48 = 432$ rules. The nine models represent commonly used and commonly exploited protocols while 48 denotes the first 48 bytes of the packet. Each of these bytes is treated as a nominal attribute with 256 possible values. The anomaly score for a packet is the sum of the anomaly scores reported for each of the 432 rules [10]. The NETAD anomaly score for a packet is then

$$\Sigma t n_a (1 - r/256)/r + t_i/(f_i + r/256).$$

where
t is the time since the attribute was last anomalous in training or testing;
n_a is the number of training packets from the last anomaly to the end of training;
r is the number of allowed values (up to 256);
i is a value from 0 to 255;
t_i is the time since the value i (0-255) was last observed;
f_i is the number of times i was observed in the training packets.

3.2 Maximum Entropy Anomaly Detector (MaxEnt)

The Maximum Entropy anomaly detector [9] estimates the legitimate traffic distribution using maximum entropy estimation. Training traffic is divided into 2,348 packet classes, and maximum entropy estimation is then used to develop a baseline distribution for each packet class. Packet classes are derived from two dimensions. The first dimension contains four classes:

1. TCP;
2. TCP SYN;
3. TCP RST;
4. UDP.

In the second dimension, each of these four classes is split into 587 subclasses based on destination port numbers.

MaxEnt uses a sliding window detection approach. The packet class distributions observed in real-time windows of duration t seconds are compared with the baseline distribution using the Kullback-Leibler (KL) divergence measure. KL divergence measures the difference between two distributions and is used as a distance measure. An alarm is raised if, for a particular packet class, the KL divergence exceeds a threshold ηk more than h times in the last window of w time slots.

3.3 P2P Traffic Preprocessor

A P2P traffic preprocessor was proposed in [2] to increase the efficiency of general purpose ADSs. This preprocessor was designed after identifying traffic features which successfully discriminate P2P traffic from attack traffic. These distinguishing traffic features are discussed below.

3.3.1 Probability of Success Times the Average Volume Per Connection

The probability of establishing a successful connection (P_s) is very high for legitimate traffic, moderate for P2P traffic, and relatively low for attack traffic [2]. P_s helps in detecting attack generating hosts by classifying them as malicious on the basis of lower values of P_s. However, an attack can be hidden easily by generating a significant number of successful connections from the same host in parallel. Furthermore, P2P traffic has about 40% failed connection attempts [2] which makes it hard to discriminate between P2P and attack traffic using P_s. Similarly, there is very low volume of data per connection (V) for attack traffic as compared to the high volume in P2P traffic. Thus, the product of P_s and V can be used to discriminate between P2P and attack traffic. The probability of success times the average volume per connection is

$$P_s \times V.$$

A host or time slot is marked as malicious when the value of this product is less than or equal to ℓ_1 for that host or timeslot

$$P_s \times V \leqslant \ell_1.$$

where ℓ_1 is set to 10.0.

3.3.2 Probability of Connection Failure Times the Probability of Failure on a Particular Destination Port/Destination IP

Attack traffic has a very high probability of connection failure (P_f), while P_f is low for legitimate traffic and moderate for P2P traffic. Thus, it is easy to identify attack generating hosts with the value of P_f. However, P2P traffic also generates a significant number of failed connections which makes it harder for ADSs to identify it as non-malicious. Classification of hosts as malicious using low values of P_f results in a high false positive rate. The probability of connection failure on a particular destination port/destination IP address (P_{f_o}) can be used with P_f to help discriminate between P2P and attack traffic as P2P connection failures are not targeted to specific port/IP addresses. Denial of Service (DOS) attacks usually targets multiple ports and multiple IP addresses simultaneously to minimize P_{f_o}. However, this will result in a high value of P_f, lower P_s, and lower V. Therefore, P_{f_o} can be used with P_f to more accurately discriminate between attack and P2P traffic. The probability of failure times probability of failure on a specific destination port/IP address is

$$P_f \times P_{f_o}.$$

The P2P preprocessor marks a host as malicious when the value of $P_f \times P_{f_o}$ is greater than or equal to ℓ_2

$$P_f \times P_{f_o} \geq \ell_2.$$

Where ℓ_2 is 0.2. As the value of ℓ_2 approaches zero, the detection rate approaches 100%, while there is no detection when $\ell_2 = 1$.

3.3.3 Attack Sustainability Factor

All failed connections in a window of w timeslots are counted and sorted as a list. The median value of the sorted list is defined as the Sustainability Factor (S_F). A time interval or host is considered malicious when the value of S_F exceeds a threshold value ℓ_3

$$S_F \geq \ell_3.$$

where ℓ_3 is set to 0.

3.3.4 Preprocessor Model

The P2P preprocessor communicates with the ADS about the maliciousness of a host or time interval. A host is identified as malicious or legitimate by using the network traffic features described above. The P2P preprocessor maintain a record entry ($HostEntry$) for each source host. The $HostEntry$ record contains the values of P_s, V, P_f, P_{f_o}, and S_F for the current window [2].

A circular queue (Q) is also maintained to keep the status (0 for benign traffic and 1 for malicious traffic), for each timeslot in the window. The queue entries have the following form:

HostEntry
unsigned short P_s, V, P_f, P_{f_o};
unsigned long S_F;
unsigned short $[w]$ Q;

The P2P preprocessor communicates with the ADS about the maliciousness of a host when one of the following events occurs.

Connection Success or Failure The network traffic features identified above are updated for each connection. If the connection fails, P_f, P_{f_o}, and S_F are updated. Otherwise, only the value of P_s is updated. The value of V is updated for all connections.

Window Timeout The P2P preprocessor marks the host as malicious or legitimate for each window. This is based on the network traffic features identified above. This information is updated in the circular queue of the host. After updating Q, the values of the network features are reset for the next window.

A window of w timeslots is marked as malicious when

$$P_s \times V \leqslant \ell_1,\ P_f \times P_{f_o} \geq \ell_2 \text{ and } S_F \geq \ell_3.$$

Chapter 4

Performance Evaluation and Results

In this chapter, the detection accuracy of NETAD and MaxEnt are evaluated for both normal and adaptive operation. The construction of the dataset is also described. Testing of the ADSs was carried out with DDOS attacks, TCP flood attacks, and UDP flood attacks. A dataset containing only legitimate traffic was used for training the ADSs. All tasks were performed using a laptop with the hardware and software configuration shown in Table 4.2. Performance results in the form of ROC curves and tables are presented later in the chapter.

4.1 Data Set Description

Three types of data are used, legitimate traffic, P2P traffic, and attack traffic as described below.

4.1.1 Legitimate Traffic

The legitimate traffic data was captured at the National University of Science and Technology (NUST) Pakistan [18] using different applications including file transfer, web browsing, instant messaging, and video streaming. Legitimate traffic does not contain any P2P or attack traces. The mean packet rate for the legitimate traffic is 3168 pkts/sec.

Client Name	**Traffic Volume**	**Throughput (MBps)**
Flashget	60.7 MB	1.2
UTorrent	1.08 GB	2.6
BitTorrent	1.59 GB	2.81
Halite	400 MB	1.76

Table 4.1: P2P file sharing application statistics.

4.1.2 P2P Traffic

The P2P traffic data used in this project is from BitTorrent and Kadmelia protocols. These two protocols were chosen based on their huge traffic volume generation capabilities. Different torrent clients were used as they might differ in their behavior and abilities to handle P2P file sharing. Multiple torrent files are used for the transfer of data. Statistics for the collected P2P traces from the file sharing applications are presented in Table 4.1.

4.1.3 Attack Traffic

The dataset containing DDOS traces was obtained from the Center for Applied Internet Data Analysis (CAIDA) [19]. This dataset contains one hour of anonymized traces from a DDOS attack on August 4, 2007. The trace is split into 5 minute packet capture (pcap) files. The total size of the dataset is approximately 5 GB. Legitimate and any other type of traffic other than from the DDOS attack was removed. Thus, only traffic from the attacker and response of the victim is included in the dataset. The attack dataset contains a large portion of ICMP flood and DNS amplification attacks. The datasets containing TCP flood and UDP flood attacks were obtained from the DARPA Intrusion Detection Evaluation Group at MIT Lincoln Laboratory [20].

Manufacturer	Dell Corporation
System Type	64-bit
Operating System	Windows 10 Professional
Processor	Intel Core I7-6700HQ
Processor Speed	3.5 GHZ (Turbo)
Memory	16 GB
Number of Cores	4
Java Platform IDE	Eclipse Mars 4.5.2
Java Development Kit	JDK 1.8
Network Protocol Analyzer	Wireshark 2.2.1

Table 4.2: Hardware and software configuration of the system.

4.2 Data Set Operations

Datasets which were obtained different features and time stamps. DDOS dataset otained from CAIDA was in the form of unknown hardware addresses. Furthermore, the datasets are from different sources so they must be merged to make a single dataset with consistent time stamps. The following operations were performed to make the dataset suitable for testing purposes.

- Attack traffic was separated from the DARPA dataset to have explicit TCP and UDP flood traces.
- IP addresses and MAC addresses were added to the dataset.
- After training of the ADSs on the legitimate traffic dataset, datasets of benign, P2P, and attack traffic was merged to form one single dataset.
- The time stamps were changed in dataset to be consistent over a particular time period.

4.3 Evaluation Metrics

For any ADS, one of the main requirement is a high detection accuracy. Inaccurate detection results, i.e. such as P2P traffic counted as malicious, could lead to misleading conclusions when planning, designing, and implementing network security solutions. The metrics used for evaluating ADS performance are as follows.

Detection Rate The percentage of malicious packets in the dataset detected successfully as malicious.

False Positive Rate The percentage of legitimate packets marked as malicious.

4.3.1 Maximum Entropy Parameters

As mentioned in Chapter 3, MaxEnt uses a sliding window detection approach and it raises an alarm if the KL divergence exceeds a value of threshold ηk, more than h times in a window of w time slots. The parameters used here are $h = 30$, $w = 60$, and ηk is varied from 1 to 6 in increments of 0.2. Results are presented in the form of ROC curves and tables for normal and adaptive MaxEnt.

4.3.2 NETAD Parameters

NETAD calculates an anomaly score for each packet. A threshold d is applied to the packet score for marking it as malicious or legitimate. The value of the threshold d is varied from 2^1 to 2^{13}. Results are presented in the form of ROC curves and tables for normal and adaptive NETAD.

4.4 Results

This section presents results of the evaluations carried out for different attacks. The data is presented in the tables used to generate ROC curves for the performance comparison of MaxEnt and NETAD anomaly detection systems under normal and adaptive conditions. Results are presented separately for DDOS, TCP flood, and UDP flood attacks.

4.4.1 Distributed Denial of Service Attack

Figure 4.1 shows the performance of normal and adaptive MaxEnt to detect DDOS attacks. Figure 4.1 is based on the data in Table 4.3. It can be seen that MaxEnt normal was able to detect most of the DDOS attack, but at the cost of a 57% false positive rate at a maximum detection rate of 97%. Using MaxEnt adaptive did not significantly improve the performance.

NETAD detected most of the DDOS attack, as shown in Figure 4.2, with a lower false positive rate. Figure 4.2 is based on Table 4.4. NETAD adaptive significantly improved the detection rate while reducing false positive rate. The reason for the effectiveness of NETAD in contrast to the MaxEnt for DDOS detection is its ability to limit bursts of ICMP traffic by filtering out packets addressed to any port/protocol combination if more than 16 have been received in the last 60 seconds.

4.4.2 TCP Flood Attack

Figure 4.3 illustrates the accuracy of the MaxEnt ADS to detect a TCP flood attack. The data in Table 4.5 was used to generate the ROC curves in Figure 4.3. MaxEnt achieves a maximum detection rate of 94%, but at the cost of a 79% false positive rate. MaxEnt adaptive improved the detection rate but did not significantly reduce the false positive rate.

NETAD performed considerably better than MaxEnt with TCP flood attacks. However, the improvement of NETAD with the P2P preprocessor is not significant as shown in Figure 4.4 which is based on Table 4.6. This is because it already has a very high TCP flood attack detection rate.

4.4.3 UDP Flood Attack

Figure 4.5 shows the ROC curves for the performance of MaxEnt in detecting UDP flood attacks. Figure 4.5 is based on the data in Table 4.7. MaxEnt adaptive was successfully able to detect 94% of the attack traffic with a 73% false positive rate. The improvement in the false positive rate with the P2P preprocessor is not very significant, but the detection rate is slightly better.

NETAD normal was not able to detect most of the UDP flood attacks as shown in Figure 4.6 which is based on the data in Table 4.8. NETAD was not able to detect almost 50% of the UDP flood attack traffic. Efforts to improve the detection accuracy by using P2P preprocessor did not have a huge impact on the performance.

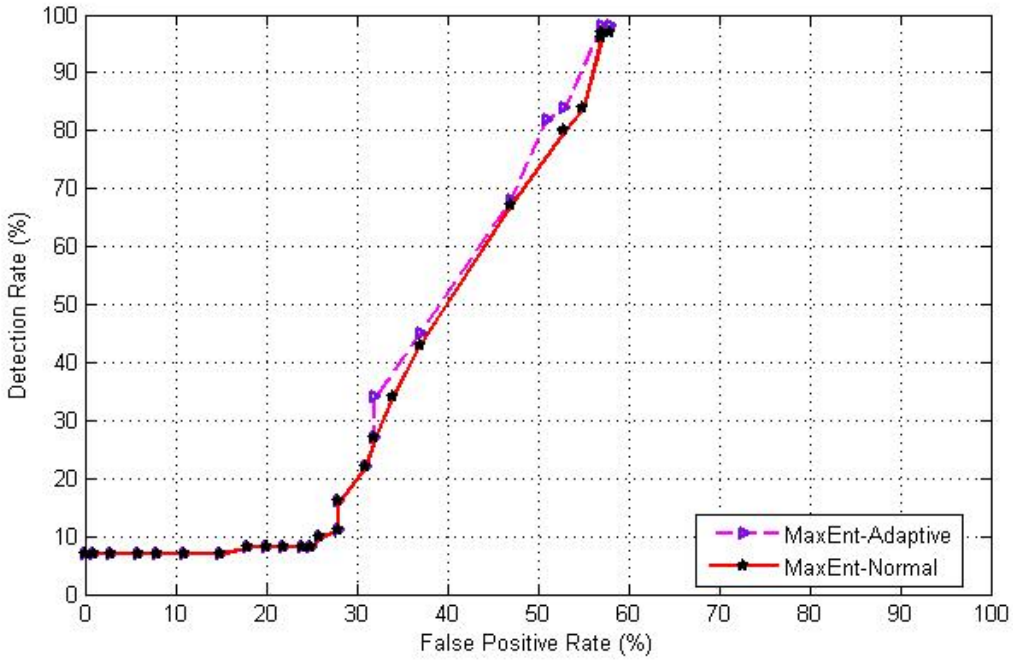

Figure 4.1: ROC curves for normal and adaptive MaxEnt with a DDOS attack.

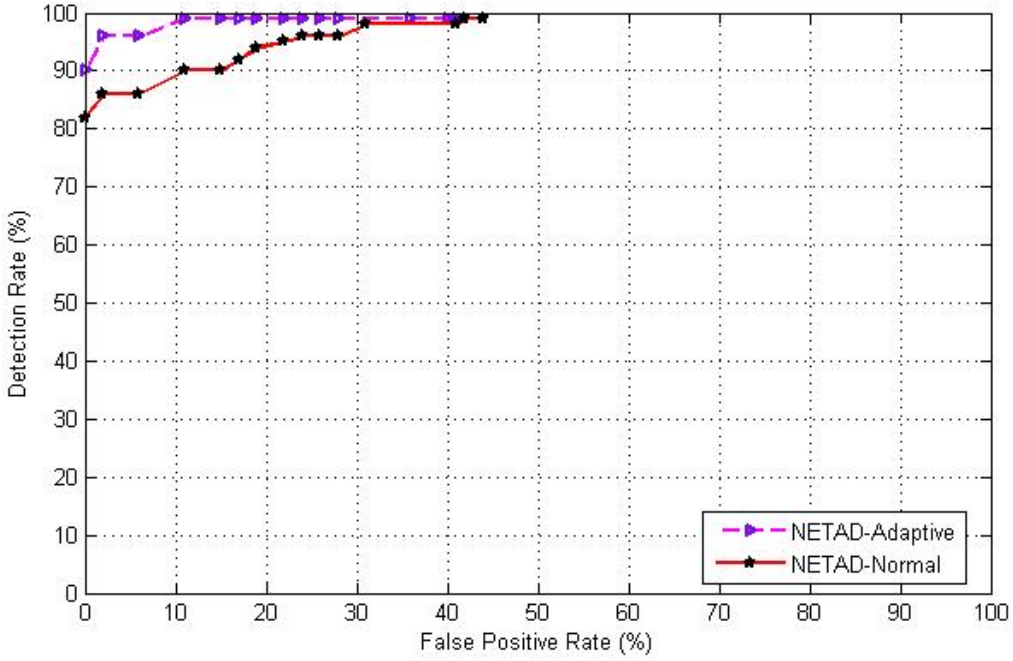

Figure 4.2: ROC curves for normal and adaptive NETAD with a DDOS attack.

Threshold	**Normal Detection Rate (%)**	**Normal False Positive Rate (%)**	**Adaptive Detection Rate (%)**	**Adaptive False Positive Rate (%)**
1.0	97	58	98	58
1.2	97	57	98	57
1.4	96	57	98	57
1.6	96	57	98	57
1.8	84	55	84	53
2.0	80	53	82	51
2.2	67	47	68	47
2.4	43	37	45	37
2.6	34	34	34	32
2.8	27	32	27	32
3.0	22	31	22	31
3.2	16	28	16	28
3.4	11	28	11	28
3.6	10	26	10	26
3.8	8	25	8	25
4.0	8	24	8	24
4.2	8	22	8	22
4.4	8	20	8	20
4.6	8	18	8	18
4.8	7	15	7	15
5.0	7	11	7	11
5.2	7	8	7	8
5.4	7	6	7	6
5.6	7	3	7	3
5.8	7	1	7	1
6.0	7	0	7	0

Table 4.3: Detection and false positive rates of normal and adaptive MaxEnt with a DDOS attack.

Threshold	**Normal Detection Rate (%)**	**Normal False Positive Rate (%)**	**Adaptive Detection Rate (%)**	**Adaptive False Positive Rate (%)**
2	99	44	99	41
4	99	44	99	41
8	99	44	99	40
16	99	42	99	40
32	98	41	99	36
64	98	31	99	28
128	96	28	99	28
256	96	26	99	26
512	96	24	99	24
1024	95	22	99	22
2048	94	19	99	19
4096	92	17	99	17
8192	90	15	99	15

Table 4.4: Detection and false positive rates of normal and adaptive NETAD with a DDOS attack.

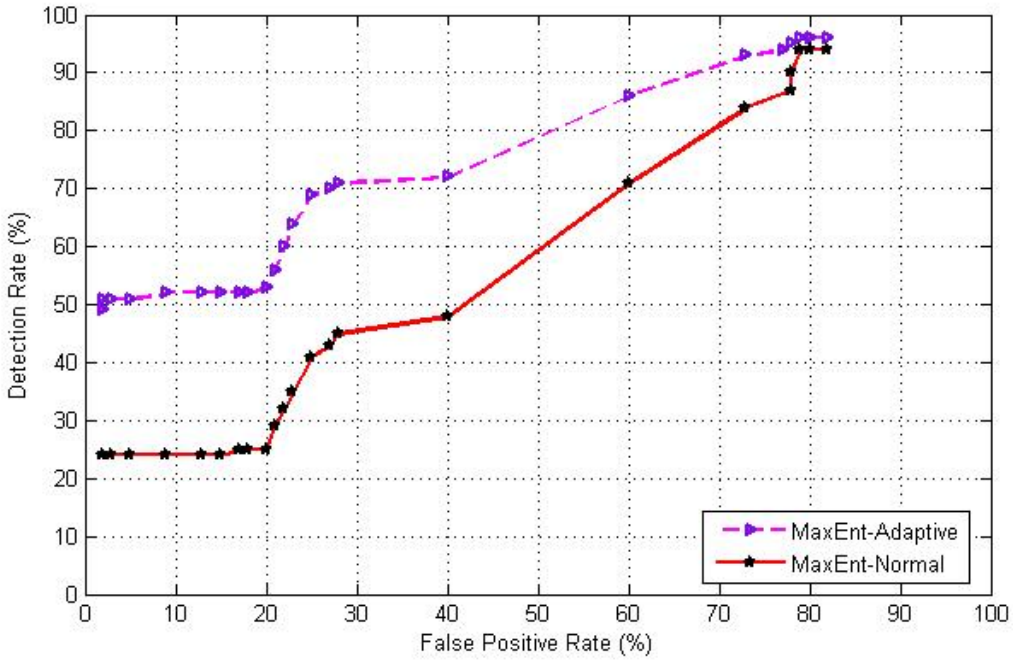

Figure 4.3: ROC curves for normal and adaptive MaxEnt with a TCP flood attack.

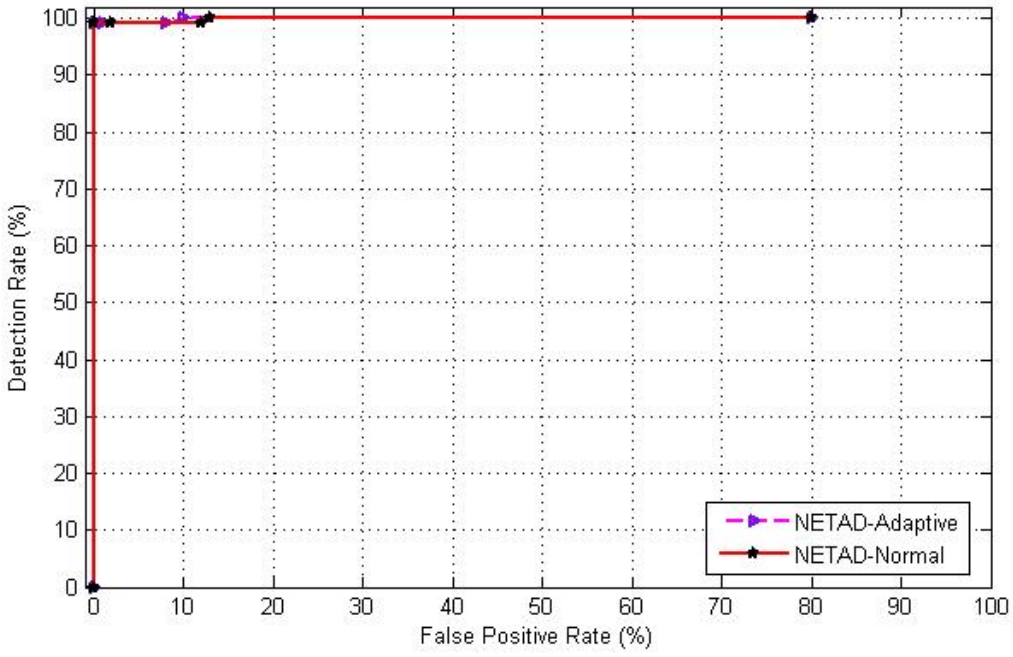

Figure 4.4: ROC curves for normal and adaptive NETAD with a TCP flood attack.

Threshold	**Normal Detection Rate (%)**	**Normal False Positive Rate (%)**	**Adaptive Detection Rate (%)**	**Adaptive False Positive Rate (%)**
1.0	94	82	96	82
1.2	94	80	96	80
1.4	94	79	96	79
1.6	90	78	95	78
1.8	87	78	94	77
2.0	84	73	93	73
2.2	71	60	86	60
2.4	48	40	72	40
2.6	45	28	71	28
2.8	43	27	70	27
3.0	41	25	69	25
3.2	35	23	64	23
3.4	32	22	60	22
3.6	29	21	56	21
3.8	25	20	53	20
4.0	25	18	52	18
4.2	25	17	52	17
4.4	24	15	52	15
4.6	24	13	52	13
4.8	24	9	52	9
5.0	24	5	51	5
5.2	24	3	51	3
5.4	24	3	51	3
5.6	24	2	51	2
5.8	24	2	51	2
6.0	24	2	49	2

Table 4.5: Detection and false positive rates of normal and adaptive MaxEnt with a TCP flood attack.

Threshold	**Normal Detection Rate (%)**	**Normal False Positive Rate (%)**	**Adaptive Detection Rate (%)**	**Adaptive False Positive Rate (%)**
2	100	80	100	80
4	100	13	100	10
8	99	12	99	8
16	99	12	99	8
32	99	2	99	1
64	99	0	99	0
128	99	0	99	0
256	99	0	99	0
512	99	0	99	0
1024	99	0	99	0
2048	99	0	99	0
4096	99	0	99	0
8192	0	0	0	0

Table 4.6: Detection and false positive rates of normal and adaptive NETAD with a TCP flood attack.

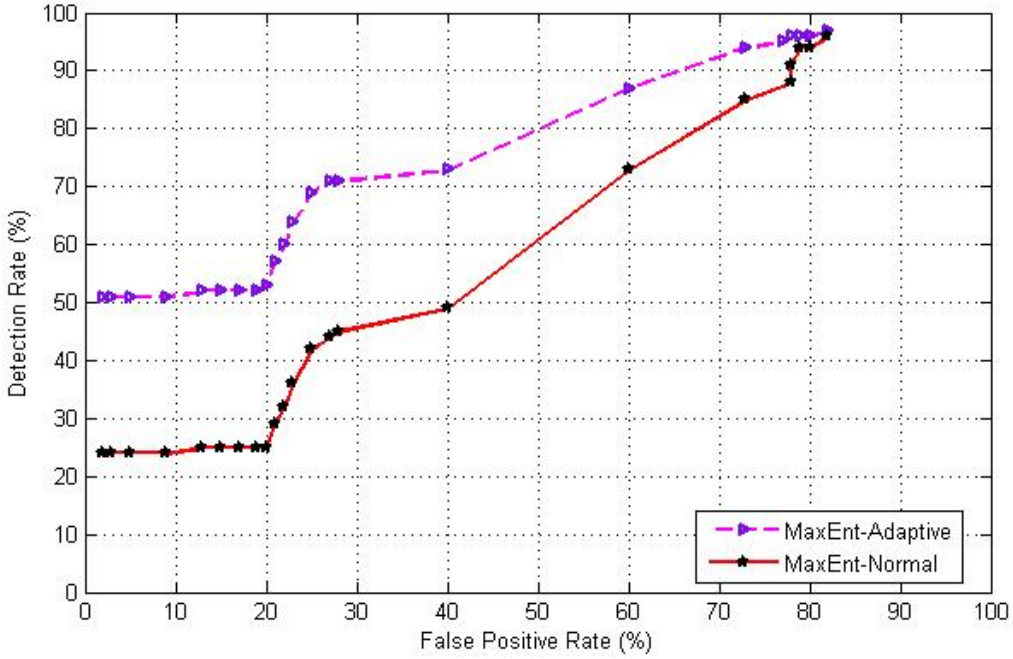

Figure 4.5: ROC curves for normal and adaptive MaxEnt with a UDP flood attack.

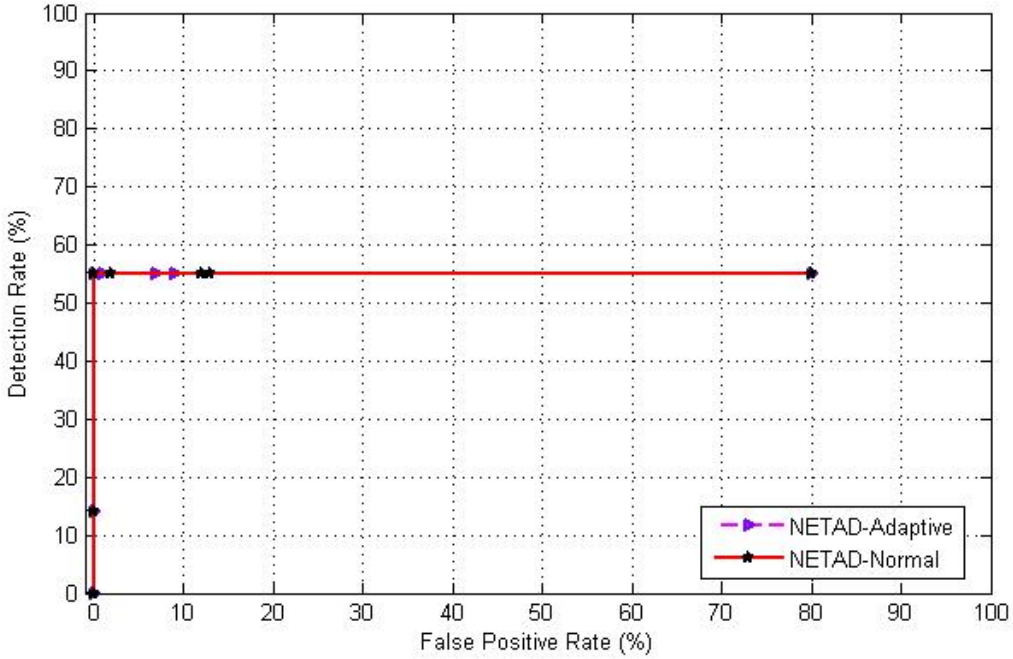

Figure 4.6: ROC curves for normal and adaptive NETAD with a UDP flood attack.

Threshold	**Normal Detection Rate (%)**	**Normal False Positive Rate (%)**	**Adaptive Detection Rate (%)**	**Adaptive False Positive Rate (%)**
1.0	96	82	97	82
1.2	94	80	96	80
1.4	94	79	96	79
1.6	91	78	96	78
1.8	88	78	95	77
2.0	85	73	94	73
2.2	73	60	87	60
2.4	49	40	73	40
2.6	45	28	71	28
2.8	44	27	71	27
3.0	42	25	69	25
3.2	36	23	64	23
3.4	32	22	60	22
3.6	29	21	57	21
3.8	25	20	53	20
4.0	25	19	52	19
4.2	25	17	52	17
4.4	25	15	52	15
4.6	25	13	52	13
4.8	24	9	51	9
5.0	24	5	51	5
5.2	24	3	51	3
5.4	24	3	51	3
5.6	24	2	51	2
5.8	24	2	51	2
6.0	24	2	51	2

Table 4.7: Detection and false positive rates of normal and adaptive MaxEnt with a UDP flood attack.

Threshold	**Normal Detection Rate (%)**	**Normal False Positive Rate (%)**	**Adaptive Detection Rate (%)**	**Adaptive False Positive Rate (%)**
2	55	80	55	80
4	55	13	55	9
8	55	12	55	9
16	55	12	55	7
32	55	2	55	1
64	55	0	55	0
128	55	0	55	0
256	55	0	55	0
512	55	0	55	0
1024	55	0	55	0
2048	55	0	55	0
4096	14	0	14	0
8192	0	0	0	0

Table 4.8: Detection and false positive rates of normal and adaptive NETAD with a UDP flood attack.

Chapter 5

Conclusion and Future Work

Distributed Denial of Service (DDOS) attacks are destructive for communication networks. With the rapid increase in P2P traffic and the diversity of applications, the problem of distinguishing P2P traffic from attack traffic has become important. Thus, there is active interest in distinguishing P2P traffic to improve the performance of Anomaly Detection Systems (ADSs). Furthermore, there is a need for ADSs to detect DDOS attacks with a low false positive rate.

In this project, two commonly used anomaly detectors, namely Network Traffic Anomaly Detector (NETAD) and Maximum Entropy Anomaly Detector (MaxEnt), were evaluated with and without a P2P traffic preprocessor for the detection of DDOS attacks. The performance of MaxEnt and NETAD was also evaluated for TCP flood and UDP flood attacks.

The performance parameters examined were detection rate and false positive rate. The detection accuracy of normal ADSs was compared with their adaptive forms. Experimental results show that the P2P traffic preprocessor improves the overall accuracy of ADSs to detect attacks. In contrast, NETAD has better detection rate with lower false positive rate than MaxEnt with DDOS attacks.

In future, the P2P traffic preprocessor can be optimized to increase its robustness. The preprocessor can be tuned to keep an additional record entry for hosts with a high number of failed connection attempts and mark them as malicious. The preprocessor can then communicate with ADS to share the information from its record entry to make decision whether a host is malicious or legitimate. This will allow ADSs to detect DDOS attacks more efficiently. Moreover, an in-depth investigation is required to determine the effectiveness of adaptive ADSs for application layer attacks.

Bibliography

[1] The Security Division of NETSCOUT. *Global DDOS attack data for 1H 2016.* Arbor Networks Inc., 2016.

[2] S. Ali, K. Wu, and H. Khan. *Traffic anomaly detection in the presence of P2P traffic.* Proceedings of the IEEE Conference on Local Computer Networks, pp. 482–485, 2014.

[3] M. Meeker. *Internet trends report 2016.* Kleiner Perkins Caufield Byers, 2016.

[4] Cisco Press. *Defeating DDOS Attacks.* White Paper: Cisco Press. 2014.

[5] Cisco Press. *Forecast and Methodology.* White Paper: Cisco VNI Forecast and Methodology. 2014.

[6] Cisco Press. *The Zettabyte EraTrends and Analysis.* White Paper: Cisco VNI Forecast and Methodology. 2016.

[7] K. Wang and S. J. Stolfo. *Anomalous payload-based network intrusion detection.* Proceedings of the International Workshop on Recent Advances in Intrusion Detection, pp. 203–222, 2004.

[8] G. Maselli, D. Luca, and S. Suin. *Design and implementation of an anomaly detection system: An empirical approach.* Proceedings of the Terena Networking Conference, 2003.

[9] Y. Gu, A. McCallum, and D. Towsley. *Detecting anomalies in network traffic using maximum entropy estimation.* Proceedings of the ACM SIGCOMM Conference on Internet Measurement, 2005.

[10] M. V. Mahoney. *Network traffic anomaly detection based on packet bytes.* Proceedings of the ACM Symposium on Applied Computing, pp. 346–350, 2003.

[11] M. M. Williamson. *Throttling viruses: Restricting propagation to defeat malicious mobile code*. Proceedings of the Computer Security Applications Conference, pp. 61–68, 2002.

[12] S. E. Schechter, J. Jung, and A. W. Berger. *Fast detection of scanning worm infections*. Proceedings of the International Workshop on Recent Advances in Intrusion Detection, pp. 59–81, 2004.

[13] A. Lakhina, M. Crovella, and C. Diot. *Mining anomalies using traffic feature distributions*. Proceedings of the ACM SIGCOMM Computer Communication Review, 2005.

[14] A. B. Ashfaq, M. J. Robert, A. Mumtaz, M. Q. Ali, A. Sajjad, and S. A. Khayam. *A comparative evaluation of anomaly detectors under portscan attacks*. Proceedings of the International Workshop on Recent Advances in Intrusion Detection, pp. 351–371, 2008.

[15] S. A. Mehdi, J. Khalid, and S. A. Khayam. *Revisiting traffic anomaly detection using software defined networking*. Proceedings of the International Workshop on Recent Advances in Intrusion Detection, pp. 161–180, 2011.

[16] M. Qi. *P2P network-targeted DDOS attacks*. Proceedings of the International Conference on the Applications of Digital Information and Web Technologies, 2009.

[17] C. Wueest. *The continued rise of DDOS attacks*. White Paper: Security Response, Symantec Corporation, 2014.

[18] NUST School of Electrical Engineering and Computer Science. *WisNet Traffic Datasets*. National University of Science and Technology, Pakistan.

[19] Center for Applied Internet Data Analysis. *The CAIDA DDOS Attack 2007 Dataset*. The CAIDA UCSD "DDOS Attack 2007" Dataset, 2007.

[20] DARPA Intrusion Detection Evaluation Group of MIT Lincoln Laboratory. *DARPA Intrusion Detection Data Sets*. Cyber Systems and Technology, 2000.