

# Florida State University Libraries

---

Electronic Theses, Treatises and Dissertations

The Graduate School

---

2008

## Detecting Spam Zombies by Monitoring Outgoing Messages

Peng Chen



FLORIDA STATE UNIVERSITY  
COLLEGE OF ARTS AND SCIENCES

DETECTING SPAM ZOMBIES BY MONITORING OUTGOING  
MESSAGES

By

PENG CHEN

A Thesis submitted to the  
Department of Computer Science  
in partial fulfillment of the  
requirements for the degree of  
Master of Science

Degree Awarded:  
Fall Semester, 2008

The members of the Committee approve the Thesis of Peng Chen defended on October 17, 2008.

---

Zhenhai Duan  
Professor Directing Thesis

---

Xin Yuan  
Committee Member

---

Zhenghao Zhang  
Committee Member

Approved:

---

David Whalley, Chair  
Department of Computer Science

---

Joseph Travis, Dean, College of Arts and Sciences

The Office of Graduate Studies has verified and approved the above named committee members.

To my family.

## ACKNOWLEDGEMENTS

I would like to express my gratitude to my adviser, Dr. Zhenhai Duan, for his constant guidance and support, which have been invaluable to conduct the research and writing of this thesis. I am very grateful to Dr. Xin Yuan and Dr. Zhenghao Zhang, for their serving as part of the committee of the thesis and their valuable input and feedback.

I also thank my friends who have been supporting and encouraging me for a long time. Especially, I am deeply thankful for my wife who always takes care of my life carefully and tenderly, such that I am able to finish my work. At last, This work is dedicated to my parents in China who give me my life to enjoy all what I have.

— Peng

# TABLE OF CONTENTS

List of Tables . . . . .	vi
List of Figures . . . . .	vii
Abstract . . . . .	viii
<b>1. INTRODUCTION . . . . .</b>	<b>1</b>
<b>2. RELATED WORK . . . . .</b>	<b>4</b>
<b>3. PROBLEM FORMULATION . . . . .</b>	<b>6</b>
<b>4. BACKGROUND ON SEQUENTIAL PROBABILITY RATIO TEST . . . . .</b>	<b>8</b>
<b>5. DETECTING SPAM ZOMBIES . . . . .</b>	<b>12</b>
5.1 SPOT Detection Algorithm . . . . .	12
5.2 Alternative Designs . . . . .	15
5.3 Impact of Dynamic IP Addresses . . . . .	16
<b>6. PERFORMANCE EVALUATION . . . . .</b>	<b>19</b>
6.1 Overview of the Email Trace and Methodology . . . . .	19
6.2 Performance Evaluation of SPOT . . . . .	22
6.3 Performance Evaluation of Alternative Designs . . . . .	25
6.4 Dynamic IP Addresses . . . . .	26
<b>7. DISCUSSION . . . . .</b>	<b>29</b>
7.1 Practical Deployment . . . . .	29
7.2 Possible Evasion Techniques . . . . .	30
<b>8. CONCLUSION . . . . .</b>	<b>31</b>
REFERENCES . . . . .	32
BIOGRAPHICAL SKETCH . . . . .	34

# LIST OF TABLES

6.1	Summary of the email trace. . . . .	20
6.2	Summary of sending IP addresses. . . . .	20
6.3	Summary of virus sending IP addresses. . . . .	21
6.4	Performance of SPOT. . . . .	22
6.5	Performances of CT and PT. . . . .	25

## LIST OF FIGURES

3.1	Network model. . . . .	7
5.1	Average number of required observations when $H_1$ is true ( $\beta = 0.01$ ) . . . . .	18
6.1	Illustration of message clustering. . . . .	22
6.2	Number of actual observations. . . . .	24
6.3	Distribution of spam messages in each cluster. . . . .	27
6.4	Distribution of total messages in each cluster. . . . .	28
6.5	Distribution of the cluster duration. . . . .	28



# ABSTRACT

Compromised machines are one of the key security threats on the Internet; they are often used to launch various security attacks such as DDoS, spamming, and identity theft. In this thesis we address this issue by investigating effective solutions to automatically identify compromised machines in a network. Given that spamming provides a key economic incentive for attackers to recruit the large number of compromised machines, we focus on the subset of compromised machines that are involved in the spamming activities, commonly known as spam zombies. We develop an effective spam zombie detection system named SPOT by monitoring outgoing messages of a network. SPOT is designed based on a powerful statistical tool called Sequential Probability Ratio Test, which has bounded false positive and false negative error rates. Our evaluation studies based on a two-month email trace collected in a large U.S. campus network show that SPOT is an effective and efficient system in automatically detecting compromised machines in a network. For example, among the 440 internal IP addresses observed in the email trace, SPOT identifies 132 of them as being associated with compromised machines. Out of the 132 IP addresses identified by SPOT, 126 can be either independently confirmed (110) or highly likely (16) to be compromised. Moreover, only 7 internal IP addresses associated with compromised machines in the trace are missed by SPOT.

# CHAPTER 1

## INTRODUCTION

A major security challenge on the Internet is the existence of the large number of compromised machines. Such machines have been increasingly used to launch various security attacks including DDoS, spamming, and identity theft [1]. Two natures of the compromised machines on the Internet—sheer volume and wide spread—render many existing security countermeasures less effective and defending attacks involving compromised machines extremely hard. On the other hand, identifying and cleaning compromised machines in a network remain a significant challenge for system administrators of networks of all sizes.

In this thesis we focus on the subset of compromised machines that are used for sending spam messages, which are commonly referred to as spam zombies. Given that spamming provides a critical economic incentive for the controllers of the compromised machines to recruit these machines, it has been widely observed that many compromised machines are involved in spamming [2]. A number of recent research efforts have studied the aggregate global characteristics of spamming botnets (networks of compromised machines involved in spamming) such as the size of botnets and the spamming patterns of botnets, based on the sampled spam messages received at a large email service provider [2, 3].

Rather than the aggregate global characteristics of spamming botnets, we aim to develop a tool for system administrators to automatically detect the compromised machines in their networks in an online manner. We consider ourselves situated in a network and ask the following question: How can we automatically identify the compromised machines in the network as outgoing messages pass the monitoring point sequentially? The approaches developed in the previous work [2, 3] cannot be applied here. The locally generated outgoing messages in a network normally cannot provide the aggregate large-scale spam view required by these approaches. Moreover, these approaches cannot support the online detection

requirement in the environment we consider.

The nature of sequentially observing outgoing messages gives rise to the sequential detection problem. In this thesis we will develop a spam zombie detection system, named SPOT, by monitoring outgoing messages. SPOT is designed based on a statistical method called Sequential Probability Ratio Test (SPRT), developed by Wald in his seminal work [4]. SPRT is a powerful statistical method that can be used to test between two hypotheses (in our case, a machine is compromised vs. the machine is not compromised), as the events (in our case, outgoing messages) occur sequentially. As a simple and powerful statistical method, SPRT has a number of desirable features. It minimizes the expected number of observations required to reach a decision among all the sequential and non-sequential statistical tests with no greater error rates. This means that the SPOT detection system can identify a compromised machine quickly. Moreover, both the false positive and false negative probabilities of SPRT can be bounded by user-defined thresholds. Consequently, users of the SPOT system can select the desired thresholds to control the false positive and false negative rates of the system.

In this thesis we develop the SPOT detection system to assist system administrators in automatically identifying the compromised machines in their networks. We also evaluate the performance of the SPOT system based on a two-month email trace collected in a large U.S. campus network. Our evaluation studies show that SPOT is an effective and efficient system in automatically detecting compromised machines in a network. For example, among the 440 internal IP addresses observed in the email trace, SPOT identifies 132 of them as being associated with compromised machines. Out of the 132 IP addresses identified by SPOT, 126 can be either independently confirmed (110) or are highly likely (16) to be compromised. Moreover, only 7 internal IP addresses associated with compromised machines in the trace are missed by SPOT. In addition, SPOT only needs a small number of observations to detect a compromised machine. The majority of spam zombies are detected with as little as 3 spam messages.

The remainder of the thesis is organized as follows. In Chapter 2 we discuss related work in the area of botnet detection. We formulate the spam zombie detection problem in Chapter 3. Chapter 4 provides the necessary background on SPRT for developing the SPOT spam zombie detection system. In Chapter 5 we provide the detailed design of SPOT. Chapter 6 evaluates the SPOT detection system based on the two-month email trace. We

briefly discuss the practical deployment issues and potential evasion techniques in Chapter 7, and conclude the thesis in Chapter 8.

# CHAPTER 2

## RELATED WORK

In this chapter we discuss related work, focusing on the studies that utilize spamming activities to detect bots.

Based on email messages received at a large email service provider, two recent studies [2, 3] investigated the aggregate global characteristics of spamming botnets including the size of botnets and the spamming patterns of botnets. These studies provided important insights into the aggregate global characteristics of spamming botnets by clustering spam messages received at the provider into spam campaigns using embedded URLs and near-duplicate content clustering, respectively. However, their approaches are better suited for large email service providers to understand the aggregate global characteristics of spamming botnets instead of being deployed by individual networks to detect internal compromised machines. Moreover, their approaches cannot support the online detection requirement in the network environment considered in this thesis. We aim to develop a tool to assist system administrators in automatically detecting compromised machines in their networks in an online manner.

Xie, *et al.* developed an effective tool DBSpam to detect proxy-based spamming activities in a network relying on the packet symmetry property of such activities [5]. We intend to identify all types of compromised machines involved in spamming, not only the spam proxies that translate and forward upstream non-SMTP packets (for example, HTTP) into SMTP commands to downstream mail servers as in [5].

BotHunter [6], developed by Gu *et al.*, detects compromised machines by correlating the IDS dialog trace in a network. It was developed based on the observation that a complete malware infection process has a number of well-defined stages including inbound scanning, exploit usage, egg downloading, outbound bot coordination dialog, and outbound

attack propagation. By correlating inbound intrusion alarms with outbound communications patterns, BotHunter can detect the potential infected machines in a network. Unlike BotHunter which relies on the specifics of the malware infection process, SPOT focuses on the economic incentive behind many compromised machines and their involvement in spamming. Compared to BotHunter, SPOT is a light-weight spam zombie detection system; it does not need the support from the network intrusion detection system as required by BotHunter.

As a simple and powerful statistical method, Sequential Probability Ratio Test (SPRT) has been successfully applied in many areas [7]. In the area of networking security, SPRT has been used to detect portscan activities [8], proxy-based spamming activities [5], and MAC protocol misbehavior in wireless networks [9].

# CHAPTER 3

## PROBLEM FORMULATION

In this chapter we formulate the spam zombie detection problem in a network. In particular, we discuss the network model and assumptions we make in the detection problem.

Figure 3.1 illustrates the logical view of the network model. We assume that messages originated from machines inside the network will pass the deployed spam zombie detection system. This assumption can be achieved in a few different scenarios. First, in order to alleviate the ever-increasing spam volume on the Internet, many ISPs and networks have adopted the policy that all the outgoing messages originated from the network must be relayed by a few designated mail servers in the network. Outgoing email traffic (with destination port number of 25) from all other machines in the network is blocked by edge routers of the network [10]. In this situation, the detection system can be co-located with the designated mail servers in order to examine the outgoing messages. Second, in a network where the aforementioned blocking policy is not adopted, the outgoing email traffic can be replicated and redirected to the spam zombie detection system. We note that the detection system does not need to be on the regular email traffic forwarding path; the system only needs a replicated stream of the outgoing email traffic. Moreover, as we will show in Section 6, the proposed SPOT system works well even if it cannot observe all outgoing messages. SPOT only requires a reasonably sufficient view of the outgoing messages originated from the network in which it is deployed.

A machine in the network is assumed to be either compromised or normal (that is, not compromised). In this thesis we only focus on the compromised machines that are involved in spamming. Therefore, we use the term a *compromised machine* to denote a *spam zombie*, and use the two terms interchangeably. Let  $X_i$  for  $i = 1, 2, \dots$  denote the successive observations of a random variable  $X$  corresponding to the sequence of messages originated from machine

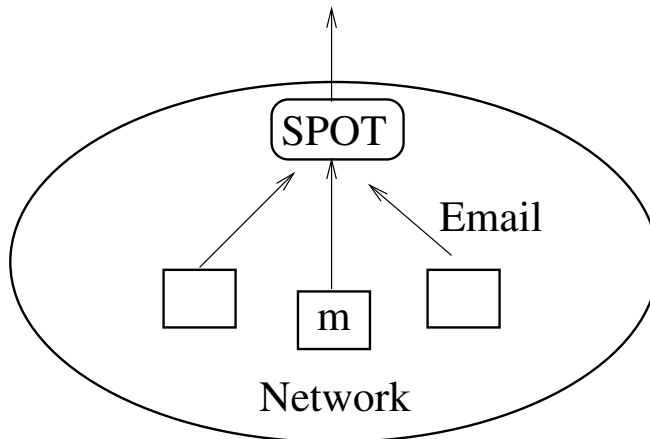


Figure 3.1: Network model.

$m$  inside the network. We let  $X_i = 1$  if message  $i$  from the machine is a spam, and  $X_i = 0$  otherwise. The detection system assumes that the behavior of a compromised machine is different from that of a normal machine in terms of the messages they send. Specifically, a compromised machine will with a higher probability generate a spam message than a normal machine. Formally,

$$Pr(X_i = 1|H_1) > Pr(X_i = 1|H_0), \quad (3.1)$$

where  $H_1$  denote that machine  $m$  is compromised and  $H_0$  that the machine is normal.

We assume that a sending machine  $m$  as observed by the spam zombie detection system is an end-user client rather than a mail relay server. This assumption is just for the convenience of our exposition. The proposed SPOT system can handle the case where an outgoing message is forwarded by a few internal mail relay servers before leaving the network. We discuss practical deployment issues in Chapter 7. We further assume that a (content-based) spam filter is deployed at the detection system so that an outgoing message can be classified as either a spam or nonspam. The spam filter does not need to be perfect in terms of the false positive rate and the false negative rate. From our communications with network operators, an increasing number of networks have started filtering outgoing messages in recent years. Based on the above assumptions, the spam zombie detection problem can be formally stated as follows. As  $X_i$  arrives sequentially at the detection system, the system determines with a high probability if  $m$  has been compromised. Once a decision is reached, the detection system reports the result, and further actions can be taken, e.g., to clean the machine.



# CHAPTER 4

## BACKGROUND ON SEQUENTIAL PROBABILITY RATIO TEST

In this chapter we provide the necessary background on the Sequential Probability Ratio Test (SPRT) for understanding the proposed spam zombie detection system. Interested readers are directed to [4] for a detailed discussion on the topic of SPRT.

In its simplest form, SPRT is a statistical method for testing a simple null hypothesis against a single alternative hypothesis. Intuitively, SPRT can be considered as an one-dimensional random walk with two user-specified boundaries corresponding to the two hypotheses. As the samples of the concerned random variable arrive sequentially, the walk moves either upward or downward one step, depending on the value of the observed sample. When the walk hits or crosses either of the boundaries for the first time, the walk terminates and the corresponding hypothesis is selected. In essence, SPRT is a variant of the traditional probability ratio tests for testing under what distribution (or with what distribution parameters), it is more likely to have the observed samples. However, unlike traditional probability ratio tests that require a pre-defined number of observations, SPRT works in an online manner and updates as samples arrive sequentially. Once sufficient evidence for drawing a conclusion is obtained, SPRT terminates.

As a simple and powerful statistical tool, SPRT has a number of compelling and desirable features that lead to the wide-spread applications of the technique in many areas. First, both the actual false positive and false negative probabilities of SPRT can be bounded by the user-specified error rates. This means that users of SPRT can pre-specify the desired error rates. A smaller error rate tends to require a larger number of observations before SPRT terminates. Thus users can balance the performance and cost of an SPRT test. Second, it has been proved that SPRT minimizes the average number of the required observations for reaching a

decision for a given error rate, among all sequential and non-sequential statistical tests. This means that SPRT can quickly reach a conclusion to reduce the cost of the corresponding experiment, without incurring a higher error rate. In the following we present the formal definition and a number of important properties of SPRT. The detailed derivations of the properties can be founded in [4].

Let  $X$  denote a Bernoulli random variable under consideration with an unknown parameter  $\theta$ , and  $X_1, X_2, \dots$  the successive observations on  $X$ . As discussed above, SPRT is used for testing a simple hypothesis  $H_0$  that  $\theta = \theta_0$  against a single alternative  $H_1$  that  $\theta = \theta_1$ . That is,

$$\begin{aligned} Pr(X_i = 1|H_0) &= 1 - Pr(X_i = 0|H_0) = \theta_0 \\ Pr(X_i = 1|H_1) &= 1 - Pr(X_i = 0|H_1) = \theta_1. \end{aligned}$$

To ease exposition and practical computation, we compute the logarithm of the probability ratio instead of the probability ratio in the description of SPRT. For any positive integer  $n = 1, 2, \dots$ , define

$$\Lambda_n = \ln \frac{Pr(X_1, X_2, \dots, X_n|H_1)}{Pr(X_1, X_2, \dots, X_n|H_0)}. \quad (4.1)$$

Assume that  $X_i$ 's are independent (and identically distributed), we have

$$\Lambda_n = \ln \frac{\prod_1^n Pr(X_i|H_1)}{\prod_1^n Pr(X_i|H_0)} = \sum_{i=1}^n \ln \frac{Pr(X_i|H_1)}{Pr(X_i|H_0)} = \sum_{i=1}^n Z_i \quad (4.2)$$

where  $Z_i = \ln \frac{Pr(X_i|H_1)}{Pr(X_i|H_0)}$ , which can be considered as the step in the random walk represented by  $\Lambda$ . When the observation is one ( $X_i = 1$ ), the constant  $\ln \frac{\theta_1}{\theta_0}$  is added to the preceding value of  $\Lambda$ . When the observation is zero ( $X_i = 0$ ), the constant  $\ln \frac{1-\theta_1}{1-\theta_0}$  is added.

The Sequential Probability Ratio Test (SPRT) for testing  $H_0$  against  $H_1$  is then defined as follows. Given two user-specified constants  $A$  and  $B$  where  $A < B$ , at each stage  $n$  of the Bernoulli experiment, the value of  $\Lambda_n$  is computed as in Eq. (4.2), then

$$\begin{aligned} \Lambda_n \leq A &\implies \text{accept } H_0 \text{ and terminate test,} \\ \Lambda_n \geq B &\implies \text{accept } H_1 \text{ and terminate test,} \\ A < \Lambda_n < B &\implies \text{take an additional observation.} \end{aligned} \quad (4.3)$$

In the following we describe a number of important properties of SPRT. If we consider  $H_1$  as a detection and  $H_0$  as a normality, an SPRT process may result in two types of errors:

false positive where  $H_0$  is true but SPRT accepts  $H_1$  and false negative where  $H_1$  is true but SPRT accepts  $H_0$ . We let  $\alpha$  and  $\beta$  denote the user-desired false positive and false negative probabilities, respectively. There exist some fundamental relations among  $\alpha$ ,  $\beta$ ,  $A$ , and  $B$  [4],

$$A \geq \ln \frac{\beta}{1-\alpha}, \quad B \leq \ln \frac{1-\beta}{\alpha},$$

for most practical purposes, we can take the equality, that is,

$$A = \ln \frac{\beta}{1-\alpha}, \quad B = \ln \frac{1-\beta}{\alpha}. \quad (4.4)$$

This will only slightly affect the actual error rates. Formally, let  $\alpha'$  and  $\beta'$  represent the actual false positive rate and the actual false negative rate, respectively, and let  $A$  and  $B$  be computed using Eq. (4.4), then the following relations hold,

$$\alpha' \leq \frac{\alpha}{1-\beta}, \quad \beta' \leq \frac{\beta}{1-\alpha}, \quad (4.5)$$

and

$$\alpha' + \beta' \leq \alpha + \beta. \quad (4.6)$$

Eqs. (4.5) and (4.6) provide important bounds for  $\alpha'$  and  $\beta'$ . In all practical applications, the desired false positive and false negative rates will be small, for example, in the range from 0.01 to 0.05. In these cases,  $\frac{\alpha}{1-\beta}$  and  $\frac{\beta}{1-\alpha}$  very closely equal the desired  $\alpha$  and  $\beta$ , respectively. In addition, Eq. (4.6) specifies that the actual false positive rate and the false negative rate cannot be both larger than the corresponding desired error rate in a given experiment. Therefore, in all practical applications, we can compute the boundaries  $A$  and  $B$  using Eq. (4.4). This will provide at least the same protection against errors as if we use the precise values of  $A$  and  $B$  for a given pair of desired error rates. The precise values of  $A$  and  $B$  are hard to obtain.

Another important property of SPRT is the number of observations,  $N$ , required before SPRT reaches a decision. The following two equations approximate the average number of observations required when  $H_1$  and  $H_0$  are true, respectively.

$$E[N|H_1] = \frac{\beta \ln \frac{\beta}{1-\alpha} + (1-\beta) \ln \frac{1-\beta}{\alpha}}{\theta_1 \ln \frac{\theta_1}{\theta_0} + (1-\theta_1) \ln \frac{1-\theta_1}{1-\theta_0}} \quad (4.7)$$

$$E[N|H_0] = \frac{(1-\alpha) \ln \frac{\beta}{1-\alpha} + \alpha \ln \frac{1-\beta}{\alpha}}{\theta_1 \ln \frac{\theta_1}{\theta_0} + (1-\theta_1) \ln \frac{1-\theta_1}{1-\theta_0}} \quad (4.8)$$

From the above equations we can see that the average number of required observations when  $H_1$  or  $H_0$  is true depends on four parameters: the desired false positive and negative rates ( $\alpha$  and  $\beta$ ), and the distribution parameters  $\theta_1$  and  $\theta_0$  for hypotheses  $H_1$  and  $H_0$ , respectively. We note that SPRT does not require the precise knowledge of the distribution parameters  $\theta_1$  and  $\theta_0$ . As long as the true distribution of the underlying random variable is sufficiently close to one of hypotheses compared to another (that is,  $\theta$  is closer to either  $\theta_1$  or  $\theta_0$ ), SPRT will terminate with the bounded error rates. An imprecise knowledge of  $\theta_1$  and  $\theta_0$  will only affect the number of required observations for SPRT to reach a decision.

# CHAPTER 5

## DETECTING SPAM ZOMBIES

In this chapter we develop the spam zombie detection system SPOT, which utilizes the Sequential Probability Ratio Test (SPRT) presented in the last chapter. As a comparison, we also give two alternative designs CT and PT. We discuss the impacts of SPRT parameters on SPOT in the context of spam zombie detection. To ease exposition of the algorithm, we ignore the potential impact of dynamic IP addresses [11] and assume that an IP address corresponds to a unique machine. We will informally discuss the impact of dynamic IP addresses at the end of this chapter. We will formally evaluate the performance of SPOT, CT, and PT and the potential impact of dynamic IP addresses in the next chapter, based on a two-month email trace collected on a large U.S. campus network.

### 5.1 SPOT Detection Algorithm

In the context of detecting spam zombies in SPOT, we consider  $H_1$  as a detection and  $H_0$  as a normality. That is,  $H_1$  is true if the concerned machine is compromised, and  $H_0$  is true if it is not compromised. In addition, we let  $X_i = 1$  if the  $i$ th message from the concerned machine in the network is a spam, and  $X_i = 0$  otherwise. Recall that SPRT requires four configurable parameters from users, namely, the desired false positive probability  $\alpha$ , the desired false negative probability  $\beta$ , the probability that a message is a spam when  $H_1$  is true ( $\theta_1$ ), and the probability that a message is a spam when  $H_0$  is true ( $\theta_0$ ). We discuss how users configure the values of the four parameters in the next section. Based on the user-specified values of  $\alpha$  and  $\beta$ , the values of the two boundaries  $A$  and  $B$  of SPRT are computed using Eq. (4.4).

In the following we describe the SPOT detection algorithm. Algorithm 1 outlines the steps of the algorithm. When an outgoing message arrives at the SPOT detection system,

---

**Algorithm 1** SPOT spam zombie detection system

---

```
1: An outgoing message arrives at SPOT
2: Get IP address of sending machine  $m$ 
3: // all following parameters specific to machine  $m$ 
4: Let  $n$  be the message index
5: Let  $X_n = 1$  if message is spam,  $X_n = 0$  otherwise
6: if ( $X_n == 1$ ) then
7:   // spam, Eq. 4.2
8:    $\Lambda_{n+} = \ln \frac{\theta_1}{\theta_0}$ 
9: else
10:  // nonspam
11:   $\Lambda_{n+} = \ln \frac{1-\theta_1}{1-\theta_0}$ 
12: end if
13: if ( $\Lambda_n \geq B$ ) then
14:  Machine  $m$  is compromised. Test terminates for  $m$ .
15: else if ( $\Lambda_n \leq A$ ) then
16:  Machine  $m$  is normal. Test is reset for  $m$ .
17:   $\Lambda_n = 0$ 
18:  Test continues with new observations
19: else
20:  Test continues with an additional observation
21: end if
```

---

the sending machine’s IP address is recorded, and the message is classified as either spam or nonspam by the (content-based) spam filter. For each observed IP address, SPOT maintains the logarithm value of the corresponding probability ratio  $\Lambda_n$ , whose value is updated according to Eq. (4.2) as message  $n$  arrives from the IP address (lines 6 to 12 in Algorithm 1). Based on the relation between  $\Lambda_n$  and  $A$  and  $B$ , the algorithm determines if the corresponding machine is compromised, normal, or a decision cannot be reached.

We note that in the context of spam zombie detection, from the viewpoint of network monitoring, it is more important to identify the machines that have been compromised than the machines that are normal. After a machine is identified as being compromised (lines 13 and 14), it is added into the list of potentially compromised machines that system administrators can go after to clean. The message-sending behavior of the machine is also recorded should further analysis be required. Before the machine is cleaned and removed from the list, the SPOT detection system does not need to further monitor the message sending behavior of the machine.

On the other hand, a machine that is currently normal may get compromised at a later

time. Therefore, we need to continuously monitor machines that are determined to be normal by SPOT. Once such a machine is identified by SPOT, the records of the machine in SPOT are re-set, in particular, the value of  $\Lambda_n$  is set to zero, so that a new monitoring phase starts for the machine (lines 15 to 18).

SPOT requires four user-defined parameters:  $\alpha$ ,  $\beta$ ,  $\theta_1$ , and  $\theta_0$ . In this section we discuss how a user of SPOT configures these parameters, and how these parameters may affect the performance of SPOT. As discussed in the previous chapter  $\alpha$  and  $\beta$  are normally small values in the range from 0.01 to 0.05, which users can easily specify independent of the behaviors of the compromised and normal machines in the network.

Ideally,  $\theta_1$ , and  $\theta_0$  should indicate the true probability of a message being spam from a compromised machine and a normal machine, respectively. However, as we have discussed in the last chapter,  $\theta_1$  and  $\theta_0$  do not need to accurately model the behaviors of the two types of machines. Instead, as long as the true distribution is closer to one of them than another, SPRT can reach a conclusion with the desired error rates. Inaccurate values assigned to these parameters will only affect the number of observations required by the algorithm to terminate. Moreover, SPOT relies on a (content-based) spam filter to classify an outgoing message into either spam or nonspam. In practice,  $\theta_1$  and  $\theta_0$  should model the detection rate and the false positive rate of the employed spam filter, respectively. We note that all the widely-used spam filters have a high detection rate and low false positive rate.

To get some intuitive understanding of the average number of required observations for SPRT to reach a decision, Figures 5.1 (a) and (b) show the value of  $E[N|H_1]$  as a function of  $\theta_0$  and  $\theta_1$ , respectively, for different desired false positive rates. In the figures we set the false negative rate  $\beta = 0.01$ . In Figure 5.1 (a) we assume the probability of a message being spam when  $H_1$  is true to be 0.9 ( $\theta_1 = 0.9$ ). That is, the corresponding spam filter is assumed to have a 90% detection rate. From the figure we can see that it only takes a small number of observations for SPRT to reach a decision. For example, when  $\theta_0 = 0.2$  (the spam filter has 20% false positive rate), SPRT requires about 3 observations to detect that the machine is compromised if the desired false positive rate is 0.01. As the behavior of a normal machine gets closer to that of compromised machine (or rather, the false positive rate of the spam filter increases), i.e.,  $\theta_0$  increases, a slightly higher number of observations are required for SPRT to reach a detection.

In Figure 5.1 (b) we assume the probability of a message being spam from a normal

machine to be 0.2 ( $\theta_0 = 0.2$ ). That is, the corresponding spam filter has a false positive rate of 20%. From the figure we can see that it also only takes a small number of observations for SPRT to reach a decision. As the behavior of a compromised machine gets closer to that of a normal machine (or rather, the detection rate of the spam filter decreases), i.e.,  $\theta_1$  decreases, a higher number of observations are required for SPRT to reach a detection.

From the figures we can also see that, as the desired false positive rate decreases, SPRT needs a higher number of observations to reach a conclusion. The same observation applies to the desired false negative rate. These observations illustrate the trade-offs between the desired performance of SPRT and the cost of the algorithm. In the above discussion, we only show the average number of required observations when  $H_1$  is true because we are more interested in the speed of SPOT in detecting compromised machines. The study on  $E[N|H_0]$  shows a similar trend (not shown).

## 5.2 Alternative Designs

When we first undertook the project, we have also considered two alternative designs in detecting spam zombies, one based on the number of spam messages and another the percentage of spam messages sent from a machine, respectively. For simplicity, we refer to them as the count-threshold (CT) detection algorithm and the percentage-threshold (PT) detection algorithm, respectively.

In CT, the time is partitioned into windows of fixed length  $T$ . A user-defined threshold parameter  $C$  specifies the maximum number of spam message that may be originated from a normal machine in any time window. The system monitors the number of spam messages  $n$  originated from a machine in each window. If  $n > C$ , then the algorithm declares that the machine has been compromised.

PT works in a similar fashion, except that it works on the spam percentage. Formally, let  $N$  and  $n$  denote the total messages and spam messages originated from a machine  $m$  within a window  $T$ , then PT declares machine  $m$  as being compromised if  $\frac{n}{N} > P$ , where  $P$  is the user-defined maximum spam percentage of a normal machine.

In the following we briefly compare them with the SPOT system. The three algorithms have the similar time and space complexities. They all need to maintain a record for each observed machine and update the record as messages arrive from the machine. However,



unlike SPOT, which can provide a bounded false positive rate and false negative rate, and a confidence how well SPOT works, the error rates of CT and PT cannot be a priori specified.

SPOT requires four user-defined parameters  $\alpha$ ,  $\beta$ ,  $\theta_1$ , and  $\theta_0$ . As we have discussed in the previous sections, selecting values for the four parameters are relatively straightforward. In contrast, selecting the “right” values for the parameters of CT and PT are much challenging and tricky. They require a thorough understanding of the different behaviors of the compromised and normal machines in the concerned network and a training based on the history of the two different behaviors in order for them to work reasonably well in the network. Our preliminary studies of the two alternative designs confirm that, unlike SPOT, the performance of the two alternative algorithms is sensitive to the parameters used in the algorithm. They may have either higher false positive or false negative rates.

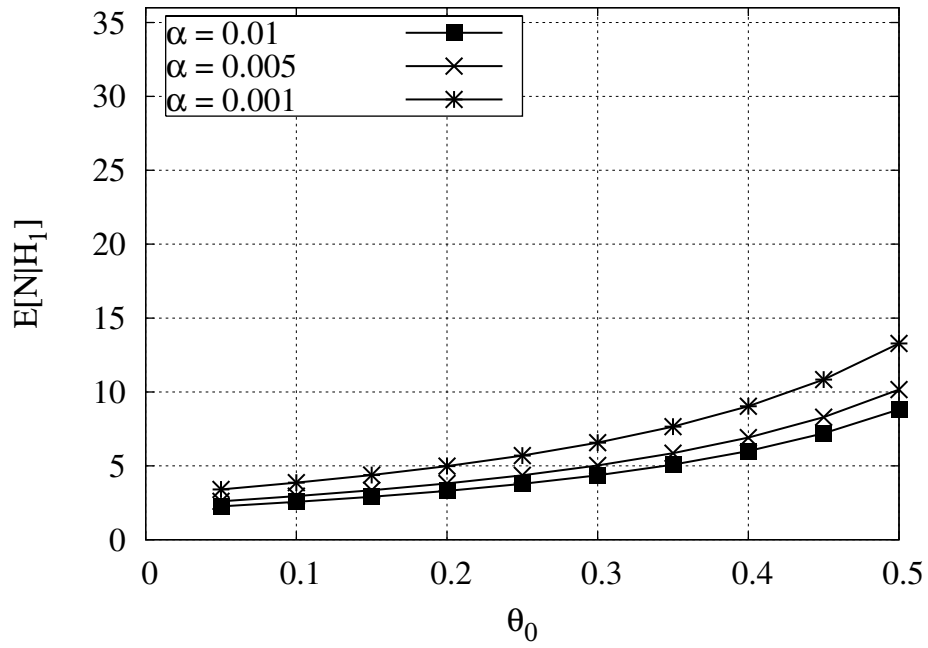
### 5.3 Impact of Dynamic IP Addresses

In the above discussion of the SPOT algorithm we have for simplicity ignored the potential impact of dynamic IP addresses and assumed that an observed IP corresponds to a unique machine. This needs not to be the case for the algorithm to work correctly. SPOT can work extremely well in the environment of dynamic IP addresses. To understand the reason we note that SPOT can reach a decision with a small number of observations as illustrated in Figure 5.1, which shows the average number of observations required for SPRT to terminate. In practice, we have noted that 3 or 4 observations are sufficient for SPRT to reach a decision for the vast majority of cases. If a machine is compromised, it is likely that more than 3 or 4 spam messages will be sent before the (unwitty) user shutdowns the machine. Therefore, dynamic IP addresses will not have any significant impact on SPOT.

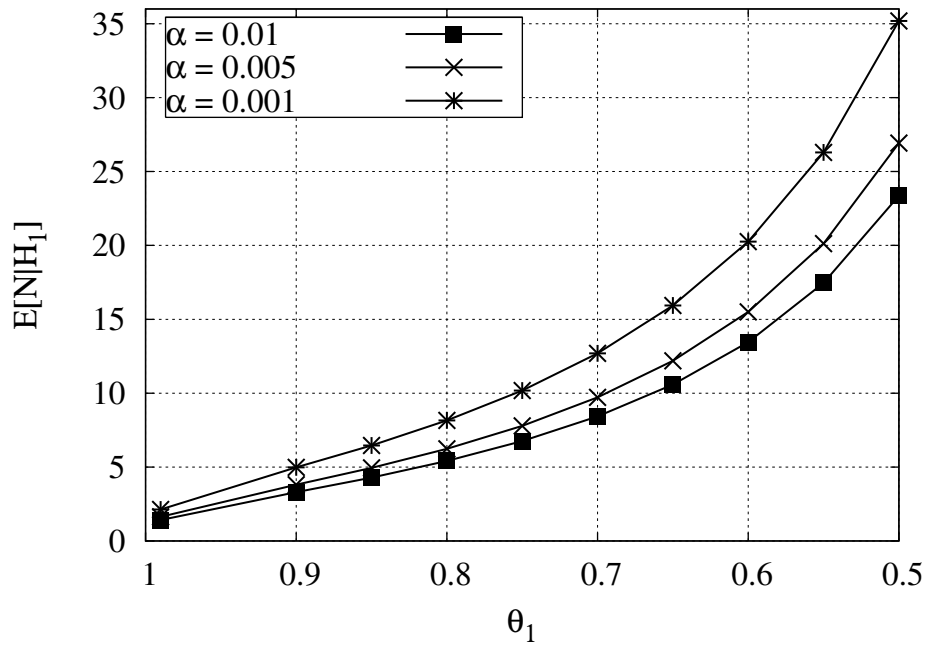
By contrast, CT and PT need to deal with dynamic IP addresses very carefully. We have introduced that both CT and PT need two parameters. For CT, we define a time window and a maximum number of spam messages; for PT, we define a time window and a maximum percentage of spam messages. Let us discuss the time window first. The ideal condition is the length of the time window is equal to the duration of one machine’s life time, but it is impossible for a fix length time window to fit into any different life time of different machines at the same time. If the length of the time window is shorter than the duration of one machine’s life time, this machine’s life time will be splitted to multiple windows. There could be two cases. In the first case, the last window is only occupied by the final part of this

machine's life time. This will give CT or PT a chance to count correctly. In the second case, the last window might be shared by this machine and other machine or even other machines. This must lead to a wrong result. If the length of the time window is longer than the duration of one machine's life time, the situation is similiar to the second case of the above discussion about shorter time window. So, it is possible to get a wrong result. For CT, we also need to set the maximum number of spam messages  $C$ , which is the threshold of counting. If CT counts more than  $C$  spam messages in a time window, it declares a zombie. But, if more than one machine share one time window, CT might count spam messages from different machines together by mistake. The same mistake might happen when PT count messages. Another reason to affect the performances of CT and PT is when they group messages to fix length time window, they might not get enough spam messages in each interval even the total number of spam messages is clearly big enough.

We formally evaluate the impact of dynamic IP addresses on SPOT, CT and PT in the next chapter.



(a)  $\theta_1 = 0.9$



(b)  $\theta_0 = 0.2$

Figure 5.1: Average number of required observations when  $H_1$  is true ( $\beta = 0.01$ )

# CHAPTER 6

## PERFORMANCE EVALUATION

In this chapter we evaluate the performance of the SPOT detection system and two alternative designs CT and PT based on a 2-month email trace collected on a large U.S. campus network. We also study the potential impact of dynamic IP addresses on SPOT.

### 6.1 Overview of the Email Trace and Methodology

The email trace was collected at a mail relay server deployed in the Florida State University (FSU) campus network between 8/25/2005 and 10/24/2005, excluding 9/11/2005 (we do not have trace on this date). During the course of the email trace collection, the mail server relayed messages destined for 53 subdomains in the FSU campus network. The mail relay server ran SpamAssassin [12] to detect spam messages. The email trace contains the following information for each incoming message: the local arrival time, the IP address of the sending machine, and whether or not the message is spam. In addition, if a message has a known virus/worm attachment, it was so indicated in the trace by an anti-virus software. The anti-virus software and SpamAssassin were two independent components deployed on the mail relay server. Due to privacy issues, we do not have access to the content of the messages in the trace.

Ideally we should have collected all the outgoing messages in order to evaluate the performance of SPOT. However, due to logistical constraints, we were not able to collect all such messages. Instead, we identified the messages in the email trace that have been forwarded or originated by the FSU internal machines, that is, the messages forwarded or originated by an FSU internal machine and destined to an FSU account. We refer to this set of messages as the *FSU emails* and perform our evaluation of SPOT based on the FSU emails. We note the set of FSU emails does not contain all the outgoing messages originated

Table 6.1: Summary of the email trace.

Measure	Non-spam	Spam	Aggregate
Period	8/25/2005 – 10/24/2005 (excl. 9/11/2005)		
# of emails	6,712,392	18,537,364	25,249,756
# of FSU emails	5,612,245	6,959,737	12,571,982
# of infected emails	60,004	163,222	223,226
# of infected FSU emails	34,345	43,687	78,032

Table 6.2: Summary of sending IP addresses.

	Non-spam only	Spam only	Mixed
# of IP (%)	121,103 (4.9)	2,224,754 (90.4)	115,257 (4.7)
# of FSU IP (%)	175 (39.7)	74 (16.8)	191 (43.5)

from inside FSU, and the compromised machines identified by SPOT based on the FSU emails will likely be a lower bound on the true number of compromised machines inside FSU campus network.

An email message in the trace is classified as either *spam* or *non-spam* by SpamAssassin [12] deployed in the FSU mail relay server. For ease of exposition, we refer to the set of all messages as the *aggregate* emails including both spam and non-spam. If a message has a known virus/worm attachment, we refer to such a message as an *infected message*. We refer to an IP address of a sending machine as a *spam-only* IP address if only spam messages are received from the IP. Similarly, we refer to an IP address as *non-spam only* and *mixed* if we only receive non-spam messages, or we receive both spam and non-spam messages, respectively. Table 6.1 shows a summary of the email trace. As shown in the table, the trace contains more than 25 M emails, of which more than 18 M, or about 73%, are spam. During the course of the trace collection, we observed more than 2 M IP addresses (2,461,114) of sending machines, of which more than 95% sent at least one spam message. During the same course, we observed 440 FSU internal IP addresses. Table 6.2 shows the classifications of the observed IP addresses. More detailed analysis of the email trace can be found in [13] and [14], including the daily message arrival patterns, and the behaviors of spammers at both the mail-server level and the network level.

In Table 6.3, we show the relation between virus sending IP addresses and spam sending

Table 6.3: Summary of virus sending IP addresses.

	Non-spam only	Spam only	Mixed	Aggregate
# of IP (%)	1,032 (9.9)	6,705 (64.6)	2,648 (25.5)	10,385(100)
# of FSU IP (%)	19 (9.3)	42 (20.6)	143 (70.1)	204(100)

IP addresses. First, only 9.9% of total IP addresses sending virus are non-spam only IP addresses that never send spam messages. The similar percentage(9.3%) is observed for FSU IP addresses. Second, the highest percentage part of IP addresses in total IP addresses is spam only IP addresses. This indicates the correlation between virus sending and spam sending IP addresses. Third, this trend is not verified when we analyze FSU IP addresses. The highest percentage part of IP addresses in FSU IP addresses is mixed IP addresses. The reason is there are many mail relay servers in the FSU network. They merged messages from clients, and hide the real senders. We verified the most part of mixed IP addresses in FSU are these servers. In addition, we can observe that only a small part of IP addresses that send spam messages take part in sending virus. In total amount of IP addresses, there is only 4.2%(10385 out of 2461114) of them sending virus. By contrast, for IP addresses in FSU, the percentage of sending virus, which is 46.4%(204 out of 440), is much bigger than that in the total IP addresses. It is likely because hackers have a local preference when they choose targets [15].

In order to study the potential impacts of dynamic IP addresses on the SPOT system, we obtain the subset of FSU IP addresses in the trace whose domain names contain “wireless”, which normally have dynamically allocated IP addresses. For each of the IP addresses, we group the messages sent from the IP address into clusters, where the messages in each cluster are likely to be from the same machine (before the IP address is re-assigned to a different machine). We group messages according to the inter-arrival times between consecutive messages, as discussed below. Let  $m_i$  for  $i = 1, 2, \dots$  denote the messages sent from an IP address, and  $t_i$  denote the time when message  $i$  is received. Then messages  $m_i$  for  $i = 1, 2, \dots, k$  belong to the same cluster if  $|t_i - t_{i-1}| \leq T$  for  $i = 2, 3, \dots, k$ , and  $|t_{k+1} - t_k| > T$ , where  $T$  is an user-defined time interval. We repeat the same process to group other messages. Let  $m_i$  for  $i = j, j + 1, \dots, k$  be the sequence of messages in a cluster, arriving in that order. Then  $|t_k - t_j|$  is referred to as the *duration* of the cluster. Figure 6.1

illustrates the message clustering process. The intuition is that, if two messages come closely in time from an IP address (within a time interval  $T$ ), it is unlikely that the IP address has been assigned to two different machines within the short time interval.

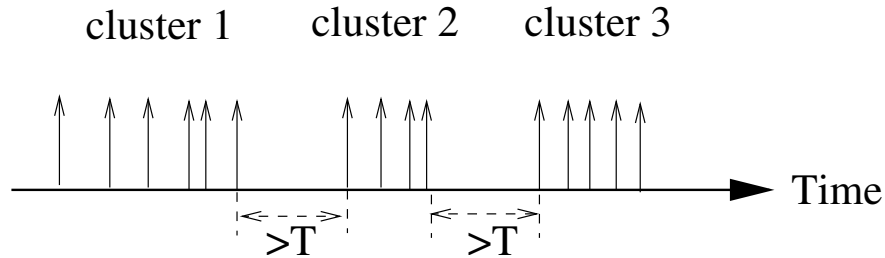


Figure 6.1: Illustration of message clustering.

In the evaluation studies, we whitelist the known mail servers deployed on the FSU campus network, given that they are unlikely to be compromised. If a deployed mail server forwards a large number of spam messages, it is more likely that machines behind the mail server are compromised. However, just based on the information available in the email trace we cannot decide which machines are responsible for the large number of spam messages, and consequently, determine the compromised machines. Chapter 7 discusses how SPOT can handle this case in practical deployment.

## 6.2 Performance Evaluation of SPOT

In this section, we evaluate the performance of SPOT based on the collected FSU emails. In all the studies, we set  $\alpha = 0.01$ ,  $\beta = 0.01$ ,  $\theta_1 = 0.9$ , and  $\theta_0 = 0.2$ . That is, we assume the deployed spam filter has a 90% detection rate and 20% false positive rate. Many widely-deployed spam filters have much better performance than what we assume here.

Table 6.4: Performance of SPOT.

Algorithm	Total # FSU IP	Detected	Confirmed (%)	Missed (%)
SPOT	440	132	126 (94.7)	7 (5.3)

Table 6.4 shows the performance of the SPOT spam zombie detection system. As discussed above, there are 440 FSU internal IP addresses observed in the email trace. SPOT

identifies 132 of them to be associated with compromised machines. In order to understand the performance of SPOT in terms of the false positive and false negative rates, we rely on a number of ways to verify if a machine is indeed compromised. First, we check if any message sent from an IP address carries a known virus/worm attachment. If this is the case, we say we have a confirmation. Out of the 132 IP addresses identified by SPOT, we can confirm 110 of them to be compromised in this way. For the remaining 22 IP addresses, we manually examine the spam sending patterns from the IP addresses and the domain names of the corresponding machines. If the fraction of the spam messages from an IP address is high (greater than 98%), we also claim that the corresponding machine has been confirmed to be compromised. We can confirm 16 of them to be compromised in this way. We note that the majority (62.5%) of the IP addresses confirmed by the spam percentage are dynamic IP addresses, which further indicates the likelihood of the machines to be compromised.

For the remaining 6 IP addresses that we cannot confirm by either of the above means, we have also manually examined their sending patterns. We note that, they have a relatively overall low percentage of spam messages over the two month of the collection period. However, they sent substantially more spam messages towards the end of the collection period. This indicates that they may get compromised towards the end of our collection period. However, we cannot independently confirm if this is the case.

Evaluating the false negative rate of SPOT is a bit tricky by noting that SPOT focuses on the machines that are potentially compromised, but not the machines that are normal (see Chapter 5). In order to have some intuitive understanding of the false negative rate of the SPOT system, we consider the machines that SPOT does not identify as being compromised at the end of the email collection period, but for which SPOT has re-set the records (lines 15 to 18 in Algorithm 1). That is, such machines have been claimed as being normal by SPOT (but have continuously been monitored). We also obtain the list of IP addresses that have sent at least a message with a virus/worm attachment. 7 of such IP addresses have been claimed as being normal, i.e., missed, by SPOT.

We emphasize that the infected messages are only used to confirm if a machine is compromised in order to study the performance of SPOT. Infected messages are not used by SPOT itself. SPOT relies on the spam messages instead of infected messages to detect if a machine has been compromised to produce the results in Table 6.4. We make this decision by noting that, it is against the interest of a professional spammer to send spam messages



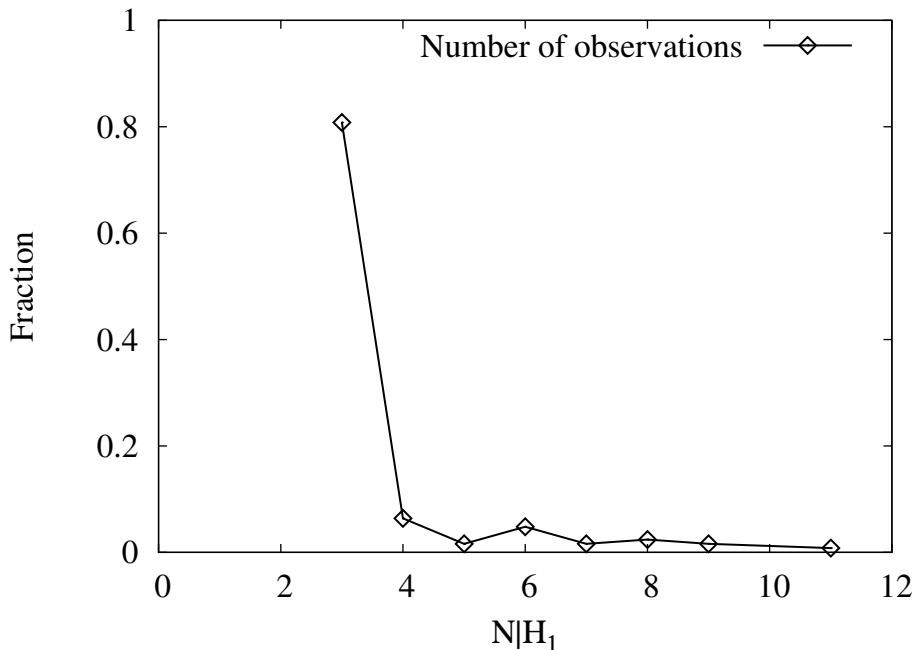


Figure 6.2: Number of actual observations.

with a virus/worm attachment. Such messages are more likely to be detected by anti-virus softwares, and hence deleted before reaching the intended recipients. This is confirmed by the low percentage of infected messages in the overall email trace shown in Table 6.1. Infected messages are more likely to be observed during the spam zombie recruitment phase instead of spamming phase. Infected messages can be easily incorporated into the SPOT system to improve its performance.

We note that both the actual false positive rate and the false negative rate are higher than the specified false positive rate and false negative rate, respectively. One possible reason is that the evaluation was based on the FSU emails, which can only provide a partial view of the outgoing messages originated from inside FSU.

Figure 6.2 shows the distributions of the number of actual observations that SPOT takes to detect the compromised machines. As we can see from the figure, the vast majority of compromised machines can be detected with a small number of observations. For example, more than 80% of the compromised machines are detected by SPOT with only 3 observations. All the compromised machines are detected with no more than 11 observations. This

indicates that, SPOT can quickly detect the compromised machines. We note that SPOT does not need compromised machines to send spam messages at a high rate in order to detect them. Here, “quick” detection does not mean a short duration, but rather a small number of observations. A compromised machine can send spam messages at a low rate (which, though, works against the interest of spammers), but it can still be detected once enough observations are obtained by SPOT.

### 6.3 Performance Evaluation of Alternative Designs

In Chapter 5, we have already mentioned two alternative designs in detecting spam zombies, one based on the number of spam messages(CT) and another the percentage of spam messages sent from a machine(PT). In this section, we evaluate the performance of CT and PT based on the user-defined parameters.

Both CT and PT need a fixed time window  $T$  and an appropriate user-defined threshold. In this evaluation, we set  $T$  to 1 hour. For CT, we set the threshold to 30 messages, which means a machine will be thought to be a zombie if it sends more than 30 spam messages in any one hour window; for PT, we set the threshold to 50%, which means a machine will be thought to be a zombie if it sends more than 50% spam messages in any one hour window. Since SPOT needs at least 3 spam message(when  $\alpha = 0.01$ ,  $\beta = 0.01$ ,  $\theta_0 = 0.2$ , and  $\theta_1 = 0.9$ ) to detect a zombie, to compare PT to SPOT, we require at least 6 messages has been sent from specific IP addresses. So, we ignore all of IP addresses which send less than 6 messages.

Table 6.5: Performances of CT and PT.

Algorithm	Total # FSU IP	Detected	Confirmed (%)	Missed (%)
CT	440	81	79 (59.8)	53 (40.2)
PT	440	84	83 (61.9)	51 (38.1)

The result shows in Table 6.5. CT detects 81 zombies out of 440 total FSU IP addresses, which is 61.36% of what SPOT has detected and PT detects 84 zombies out of total 440 FSU IP addresses, which is 63.63% of what SPOT has detected. Although CT and PT have a high accuracy (for CT, 79 out of 81 machines has been confirmed to be zombie; for PT, 83 out of 84 machines has been confirmed to be zombie), if we consider the number of zombies they missed(for CT, it is 53; for PT, it is 51. Here, we use machines that send virus as

a verification, too.), CT only detects 59.8% of verified zombies; PT only detects 61.9% of verified zombies. Also, We observed all of zombies detected by CT or PT have been detected by SPOT, and all of confirmed zombies that are detected by CT(79) and PT(83) fall into the set of confirmed zombies that are detected by SPOT(126). This proves that SPOT has more detection power over CT and PT.

## 6.4 Dynamic IP Addresses

In order to understand the potential impacts of dynamic IP addresses on the performance of SPOT, CT and PT, we group messages from a dynamic IP address (with domain names containing “wireless”) into clusters with a time interval threshold of 30 minutes. Messages with a consecutive inter-arrival time no greater than 30 minutes are grouped into the same cluster. Given the short inter-arrival duration of messages within a cluster, we consider all the messages from the same IP address within each cluster as being sent from the same machine. That is, the corresponding IP address has not been re-assigned to a different machine within the concerned cluster. (It is possible that messages from multiple adjacent clusters are actually sent from the same machine.)

Figure 6.3 shows the cumulative distribution function (CDF) of the number of spam messages in each cluster. From the figure we can see that more than 90% of the clusters have no less than 10 spam messages, and more than 96% no less than 3 spam messages. Given the large number of spam messages sent within each cluster, it is unlikely for SPOT to mistake one compromised machine as another when it tries to detect spam zombies. Indeed, we have manually checked that, spam messages tend to be sent back to back in a batch fashion when a dynamic IP address is observed in the trace. Figure 6.4 shows the CDF of the number of all messages (including both spam and non-spam) in each cluster. Similar observations can be made to that in Figure 6.3.

Figure 6.5 shows the CDF of the durations of the clusters. As we can see from the figure, more than 75% and 58% of the clusters last no less than 30 minutes and one hour (corresponding to the two vertical lines in the figure), respectively. The longest duration of a cluster we observe in the trace is about 3.5 hours.

Given the above observations, in particular, the large number of spam messages in each cluster, we conclude that dynamic IP addresses will not have any important impact on the performance of SPOT. SPOT can reach a decision within the vast majority (96%) of the

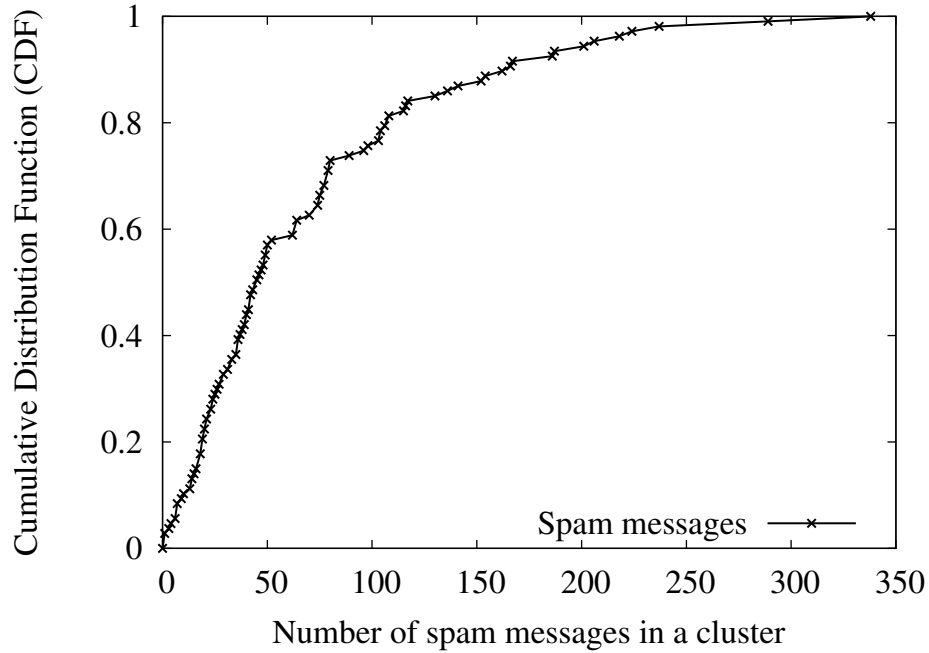


Figure 6.3: Distribution of spam messages in each cluster.

clusters in the setting we used in the current performance study. It is unlikely for SPOT to mistake a compromised machine as another.

After We manually check the result, we are surprised that what we worry about for CT and PT in Chapter 5 does not really happen. We can't find any time window sharing event between clusters. Neither is splitted any big group of spam messages to multiple time window, such that any zombie is missed. That is because that we use 30 minutes to separate clusters, which is probably longer than normal reboot and IP address assignment process. Another reason is related to DHCP scheme and abundance of IP addresses resource. Bigger IP addresses pool you have, less probably a new machine are assigned to an old IP address.

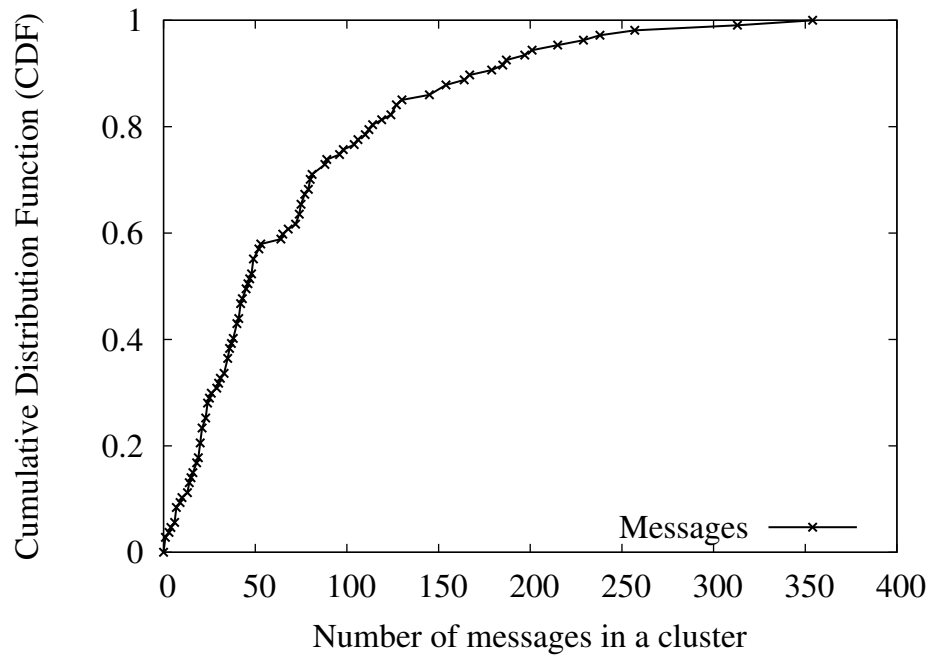


Figure 6.4: Distribution of total messages in each cluster.

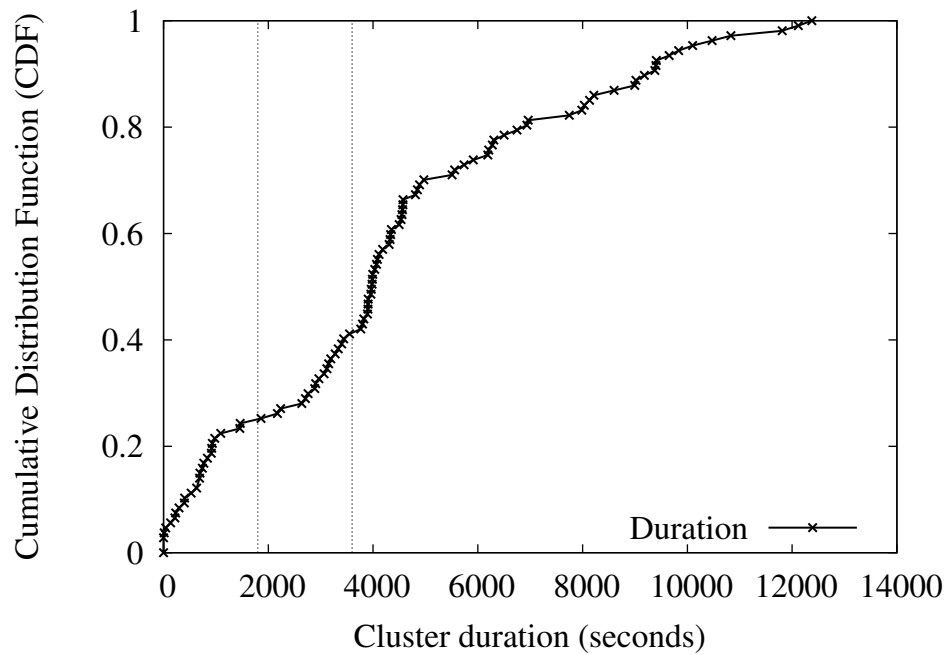


Figure 6.5: Distribution of the cluster duration.

## CHAPTER 7

### DISCUSSION

In this chapter we discuss the practical deployment issues and possible techniques that spammers may employ to evade the SPOT system and two alternative designs CT and PT.

#### 7.1 Practical Deployment

To ease exposition we have assumed that a sending machine  $m$  (Figure 3.1) is an end-user client machine. It cannot be a mail relay server deployed by the network. In practice, a network may have multiple subdomains and each has its own mail servers. A message may be forwarded by a number of mail relay servers before leaving the network. SPOT can work well in this kind of network environments. In the following we outline two possible approaches. First, SPOT can be deployed at the mail servers in each subdomain to monitor the outgoing messages so as to detect the compromised machines in that subdomain. Second, and possibly more practically, SPOT is only deployed at the designated mail servers, which forward all outgoing messages (or SPOT gets a replicated stream of all outgoing messages), as discussed in Chapter 3. SPOT relies on the `Received` header fields to identify the originating machine of a message in the network [16, 17]. Given that the `Received` header fields can be spoofed by spammers, SPOT should only use the `Received` header fields inserted by the known mail servers in the network. SPOT can determine the reliable `Received` header fields by backtracking from the last known mail server in the network that forwards the message. It terminates and identifies the originating machine when an IP address in the `Received` header field is not associated with a known mail server in the network.

## 7.2 Possible Evasion Techniques

Given that SPOT relies on (content-based) spam filters to classify messages into spam and non-spam, spammers may try to evade the developed SPOT system by evading the deployed spam filters. They may send completely meaningless non-spam messages (as classified by spam filters). However, this will reduce the real spamming rate, and hence, the financial gains, of the spammers. More importantly, as shown in Figure 5.1 (b), even if a spammer reduces the spam percentage to 50%, SPOT can still detect the spam zombie with a relatively small number of observations (25 when  $\alpha = 0.01$ ,  $\beta = 0.01$ , and  $\theta_0 = 0.2$ ). So, trying to send non-spam messages will not help spammers to evade the SPOT system. Moreover, in certain environment where user feedback is reliable, for example, feedback from users of the same network in which SPOT is deployed, SPOT can rely on classifications from end users (in addition to the spam filter). Although completely meaningless messages may evade the deployed spam filter, it is impossible for them to remain undetected by end users who receive such messages. User feedbacks may be incorporated into SPOT to improve the spam detection rate of the spam filter. As we have discussed, trying to send spam at a low rate will also not evade the SPOT system. SPOT relies on the number of (spam) messages, not the sending rate, to detect spam zombies.

By contrast, it is sending rate that affects the detection very much for CT and PT. If spammers reduce the number or percentage of spam messages in a time window, the user-defined threshold will be compromised. For example, we set 30 messages as the threshold in CT. Once this information has been figured out by spammers, they are able to break this by sending less than 30 spam messages in any time window. Moreover, if they know the size of a time window, they could send even more spam messages by crossing two windows' boundary, ie. send 29 at the end of the first window, then send another 29 at the beginning of the next window. The similiar scheme can be used for attacking PT.

# CHAPTER 8

## CONCLUSION

Compromised machines are a major security threat on the Internet. Given that spamming provides the critical economic incentive for attackers to recruit the large number of compromised machines, in this thesis we developed SPOT, an effective spam zombie detection system by monitoring outgoing messages in a network. SPOT was designed based on a simple and powerful statistical tool named Sequential Probability Ratio Test to detect the subset of compromised machines that are involved in the spamming activities. SPOT has bounded false positive and false negative error rates. It also minimizes the number of required observations to detect a spam zombie. Our evaluation studies based on a 2-month email trace collected on the FSU campus network showed that SPOT is an effective and efficient system in automatically detecting compromised machines in a network. We have also evaluated two alternative designs based on spam count and spam fraction. The results shows that SPOT is over them in both detection number and detection accuracy. In summary, they are not as effective as SPOT.



## REFERENCES

- [1] J. Markoff. Russian gang hijacking pcs in vast scheme. *The New York Times*, <http://www.nytimes.com/2008/08/06/technology/06hack.html>, August 2008. 1
- [2] Y. Xie, F. Xu, K. Achan, R. Panigrahy, G. Hulten, and I. Osipkov. Spamming botnets: Signatures and characteristics. In *Proc. ACM SIGCOMM*, Seattle, WA, August 2008. 1, 2
- [3] L. Zhuang, J. Dunagan, D. R. Simon, H. J. Wang, I. Osipkov, G. Hulten, and J. D. Tygar. Characterizing botnets from email spam records. In *Proc. of 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats*, San Francisco, CA, April 2008. 1, 2
- [4] A. Wald. *Sequential Analysis*. John Wiley & Sons, Inc, 1947. 1, 4, 4
- [5] M. Xie, H. Yin, and H. Wang. An effective defense against email spam laundering. In *ACM Conference on Computer and Communications Security*, Alexandria, VA, October 2006. 2
- [6] G. Gu, P. Porras, V. Yegneswaran, M. Fong, and W. Lee. Bothunter: Detecting malware infection through ids-driven dialog correlation. In *Proc. 16th USENIX Security Symposium*, Boston, MA, August 2007. 2
- [7] G. B. Wetherill and K. D. Glazebrook. *Sequential Methods in Statistics*. Chapman and Hall, 1986. 2
- [8] J. Jung, V. Paxson, A. Berger, and H. Balakrishnan. Fast portscan detection using sequential hypothesis testing. In *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, May 2004. 2
- [9] S. Radosavac, J. S. Baras, and I. Koutsopoulos. A framework for mac protocol misbehavior detection in wireless networks. In *Proceedings of the 4th ACM workshop on Wireless security*, Cologne, Germany, September 2005. 2
- [10] S. Linfoord. Increasing spam threat from proxy hijacking. <http://www.spamhaus.org/news.lasso?article=156>. 3
- [11] Y. Xie, F. Xu, K. Achan, E. Gillum, M. Goldszmidt, and T. Wobber. How dynamic are ip addresses? In *Proc. ACM SIGCOMM*, Kyoto, Japan, August 2007. 5

- [12] SpamAssassin. The apache spamassassin project. <http://spamassassin.apache.org/>. 6.1, 6.1
- [13] Z. Duan, K. Gopalan, and X. Yuan. Behavioral characteristics of spammers and their network reachability properties. In *IEEE International Conference on Communications (ICC)*, Glasgow, Scotland, June 2007. 6.1
- [14] Z. Duan, K. Gopalan, and X. Yuan. Behavioral characteristics of spammers and their network reachability properties. Technical report, Department of Computer Science, Florida State University, June 2006. Tech. Rep. TR-060602. 6.1
- [15] Z. Chen, L. Gao, and K. Kwiat. Modeling the spread of active worms. In *Proc. INFOCOMM*, San Francisco, CA, April 2003. 6.1
- [16] J. Klensin. Simple mail transfer protocol. RFC 2821, April 2001. 7.1
- [17] P. Resnick. Internet message format. RFC 2822, April 2001. 7.1

# BIOGRAPHICAL SKETCH

## **Peng Chen**

Peng Chen received his Bachelor of Science in Management Information Systems from Tian Jin University in 1999. After working for four years as a Software Engineer, he was admitted to the graduate school by Bei Hang University. In 2006, he got his Master of Engineering degree. Then, he went to Florida and enrolled into the Master's Program in Computer Science at Florida State University in fall 2006. His research areas of interest are on Communications and Routing protocols, Network Security, and Distributed Networks.