Carnegie Mellon University

Computational Models for Defenses against Internet-based Attacks

SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF DOCTOR OF
PHILOSOPHY IN ENGINEERING AND PUBLIC POLICY

BY Li-Chiou Chen

PITTSBURGH, PENNSYLVANIA 15213 USA

August 2003

Carnegie Mellon University

CARNEGIE INSTITUTE OF TECHNOLOGY

THESIS

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

TITLE: Computational Models for Defenses against Internet-based Attacks

PRESENTED BY: Li-Chiou Chen

ACCEPTED BY THE DEPARTMENT OF: Engineering and Public Policy

_____ _____
MAJOR PROFESSOR DATE

_____ _____
DEPARTMENT HEAD DATE

APPROVED BY THE COLLEDGE COUNCIL:

_____ _____
DEAN DATE

ABSTRACT

Internet-based attacks have become an important concern to the government and business since more systems are reliant upon the Internet to exchange information. In particular, distributed denial of service (DDOS) attacks have been used as a prevalent way to compromise the availability of networks or information services. The economic incentives of Internet Service Providers (ISPs) to provide DDOS defenses and the public policy concerns to deploy these defenses have not been formally investigated previously.

Security services, such as Virtual Private Networks, have been provided by ISPs as optional network services to deal with the secrecy of data transportation. In the case of DDOS attacks, ISPs provide DDOS defenses that ensure the availability of the subscribers' online services. This dissertation proposes that ISPs provide DDOS defenses on their network as security services to their subscribers and studies the service models for providing the defenses and the public policies needed to facilitate the provision of the defenses. The focus will be on the DDOS defenses that actively filter out ongoing attack traffic.

This dissertation analyzes how the side effects of defenses influence the provision of the defenses and investigates the economic incentives for the service provision. The contributions of this dissertation are as follows: First, this dissertation categorizes the current defenses that actively respond against DDOS attacks at network routers. The characterization is based on attack detection algorithms and attack responses. Secondly, the service provision model is analyzed based on the performance efficiency of DDOS defenses under various network topologies and various settings in the technology. When providing defenses which are congestion-based and are dynamically enforced, ISPs should design services that focus on adjusting the filtering rate of the attack traffic to meet the needs of different subscribers. When providing defenses which are congestion-based and are dynamically enforced, ISPs should design services that focus on adjusting the filtering rate of the attack traffic to meet the needs of different subscribers. Next, the economic incentives for ISPs to offer defense services are then analyzed based on empirical data. To operate the DDOS defense services cost effectively, ISPs should set the filter location closer to the attack sources and price subscribers based on their willingness to pay. Finally, cooperation among multiple ISPs on providing the defenses is analyzed. In order to improve the quality of the defenses when attacks are distributed, ISPs should cooperate with other highly influential ISPs. Public policies should encourage source filtering and provide incentives for highly influential ISPs to deploy DDOS defenses.

ACKNOWLEDGEMENTS

During the five years of my doctoral study, I have received enormous help from many people. I would like to thank everyone who contributed to my life and my work and who I do not have space to mention in these acknowledgements.

First of all, I would like to gratefully acknowledge the guidance of my thesis advisor, Kathleen Carley. Kathleen was instrumental in developing my theoretical background on computational modeling and serves as a role model for my career. In addition, each of my committee members brought a different perspective to my work. Tom Longstaff has been a great source of inspiration for forming the research problem and for examining the analyses from a practical perspective. Tom also encouraged me during the downtime of writing the dissertation. Benoit Morel's comments on public policy debates in network security were helpful in shaping my work. I enjoyed and learned a lot from my intellectual discussions with him. David Krackhardt's expertise on network analyses brought a different perspective for me to think of the problem outside the box.

Second, I would like to thank the faculty, staff, and students of the Department of Engineering and Public Policy for their support and friendship. Together, the people of EPP create a diverse and pleasant learning environment. In particular, I would like to thank Marvin Sirbu for his advice during my first two years and in my qualifier and for being a great source of information in telecommunications and information technology area. Granger Morgan has been very supportive in helping me to solve many problems in continuing this program.

Moreover, I am indebted to many people in the Networked Systems Survivability Program at the Software Engineering Institute. I am grateful for their patience and help in answering my questions and for providing me expertise on shaping my research.

Finally, I would like to thank my family and my friends. Thanks to my family for their confidence in me, their support and love. Although living more than ten thousand miles away, my parents have constantly reminded me to eat and live well. Their attitude toward work and life has sustained me in pursuing my dream. Thanks to Howard Shih for going though many difficult times with me and for patiently proofreading my work. Howard's positive attitude has transformed my life into sweet pleasure. I also would like to thank many of my friends without mentioning their names for being important ingredients in balancing my life and my work.

Financial support for this work was provided in part by the National Science Foundation ITR 0218466, the National Science Foundation IGERT 9354995 and the Pennsylvania Infrastructure Technology Alliance. Additional support was provided by ICES (the Institute for Complex Engineered Systems) and CASOS (the Center for Computational Analysis of Social and Organizational Systems) at Carnegie Mellon University.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1 INTRODUCTION

Internet-based attacks have become an important concern to the government and business since more systems are reliant upon the Internet to exchange information. Without a secure Internet infrastructure, neither E-commerce such as online purchasing nor E-democracy services such as online voting can be conducted successfully. For business, exploits of attack tools and system unavailability are two major security concerns (InfoSec 2001). For government, preventing Internet-based attacks has been an important issue in national plans to secure critical infrastructure (PCCIP 1997; PCIPB 2003).

Among Internet-based attacks, distributed denial-of-service (DDOS) attacks have emerged as a prevalent way to compromise the availability of online services. These attacks have imposed financial losses for e-commerce businesses. For example, in February 2000, over a period of three days, attackers launched DDOS attacks against several high-profile e-commerce web sites including Yahoo, eBay, and Amazon.com. In some cases, the attackers generated up to 1 gigabit per second of attack traffic, flooding the web sites of these companies (Garber 2000)(Tran 2000). The Yankee Group estimates that the financial losses imposed by the attacks on these companies total more than $1billion (Yankee 2000). The CSI/FBI survey (CSI 2001) shows that 36% of respondents in the last 12-months period have detected DDOS attacks, which imposed more than $4.2 million financial losses.

The scale of DDOS attacks has been increasing in both the number of attack sources and the magnitude of the attack traffic. Since more attack tools now are designed with mechanisms to exploit vulnerabilities automatically, the spread of attack tools is faster and easier. For example, Code-Red worm attacks in August 2001 highlight the potential risk of large-scale DDOS attacks launched from wide spread sources. An empirical study of DDOS attacks estimates that more than 12,000 attacks were launched against more than 5,000 distinct targets in one three-week period (Moore, Voelker et al. 2001).

Many defenses that mitigate the effect of ongoing DDOS attacks have been proposed but none of them have been widely deployed on the Internet infrastructure at this point because of a lack of understanding in the tradeoffs inherent in the complex system consisted of attacks and defenses. Defenses must be compared in a common framework in order to analyze their effectiveness before they are deployed to avoid needless or ineffective spending. Large-scale testing on the Internet is not feasible. Running experiments on small networks is of limited value. It is necessary to develop a framework that can capture the key factors that determine the provision of defenses on large networks.

1.1 PURPOSE

This dissertation proposes a framework to study security services that will provide defenses against Internet-based attacks. In particular, this dissertation focuses on DDOS attacks. This dissertation asks how do Internet Service Providers (ISPs) provide defenses to their subscribers against DDOS attacks? The problem is not just technical but is a management and policy problem as well, involving the setting of policies and meeting the needs of diverse subscribers with different priorities.

The effectiveness of DDOS defenses depends on many factors such that the nature of the network's topology, the specific attack scenario, and the settings of the network routers. Understanding the nature and severity of these tradeoffs will assist attack victims, network providers and public policy makers in making security policy decisions while they are assessing potential defenses against these attacks. This dissertation aims to increase our understanding of these tradeoffs and to derive insights that will enable a more secure infrastructure through the provision of the defenses against Internet-based attacks.

To deploy defenses against DDOS attacks on the Internet infrastructure, ISPs need to configure routers for either tracing, logging or filtering attack traffic before the attack traffic reaches the networks of their subscribers. However, many ISPs hesitate to deploy these defenses due to several practical concerns. First, since each defense has a different mechanism of distinguishing the attack traffic from the legitimate traffic to victims, a defense may mistakenly regard legitimate traffic as attack traffic. It is uncertain that how effective these defenses are in terms of maintaining the network connections available to the legitimate traffic while the defenses mitigate the effect of the attack traffic. Secondly, the overhead imposed by these defenses on routers may be too high. Thirdly, none of the defenses has provided a mechanism for subscribers to inform their ISPs of their preferences in selecting a defense and negotiating parameters in a defense when a tradeoff occurs.

1.2 SCOPE AND CHAPTER OUTLINE

This dissertation promotes that ISPs should provide subscribers with defenses against Internet-based attacks. The provision of defenses involves the technical variables regarding defenses and the economic variables regarding subscribers and providers. This

dissertation provides analytical models to investigate these variables so that the impacts of both these variables can be clarified. Analyses in this dissertation calibrate the models using public available data and provide recommendations for ISPs, potential attack victims and public policy makers.

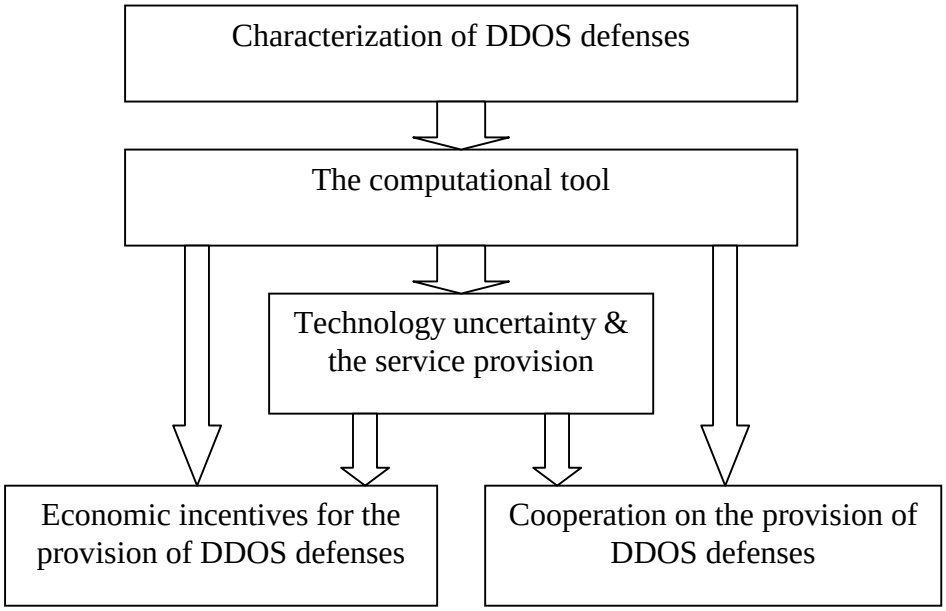

Figure 1.1: The research framework in this dissertation

As illustrated in Figure 1.1, this dissertation studies the provision of DDOS defenses through four sub-problems as follows.

1) What are the technical variables that influence the provision of DDOS defenses? Chapter 3 investigates the status quo of current defenses against DDOS and identifies the technical variables that influence the provision of the defenses. This chapter provides a qualitative study of various DDOS defense mechanisms so that quantitative analyses on the performance and the operational costs of these defenses can be built upon it.

2) How does the performance of the defenses influence the provision of the DDOS defenses? How should providers design the service model for defenses? Based upon the technical variables identified in Chapter 3, Chapter 4 describes the computational tool used to simulate the provision of DDOS defenses on a given network topology, and calculate the performance measures and cost measures of defenses for later analyses in Chapter 5, 6, and 7. Chapter 5 proposes a method for ISPs and their subscribers to define the services for providing the defenses.

3) What are the economic incentives for ISPs to provide defenses at their networks? Chapter 6 analyzes the economic incentives of ISPs for providing the defenses as services.

4) Are there incentives for ISPs to cooperate on providing DDOS defenses when attacks are transported across different administrative domains? Chapter 7 analyzes if there is an economic incentive for ISPs to cooperate on providing DDOS defenses. Policy implications for the problem are also discussed.

This dissertation is devoted to a problem involving the fields of computer security, economics and social network analysis. It is intended to help ISPs and subscribers to consider the benefits of providing DDOS defenses and to realize the tradeoffs in DDOS defenses. In addition, the results from the analyses are expected to aid public policy makers in setting security policy for computer networks to ensure a more secure infrastructure.

Chapter 2 THE PROVISION OF THE DEFENSES AGAINST DDOS ATTACKS

This chapter provides background information for a better understanding of the terminology and the research problem in this dissertation. Section 2.1 provides background information on the Internet infrastructure and IP routing. Section 2.2 describes Internet-based attacks, of which distributed denial-of-service (DDOS) attacks are one type. Since the focus of the dissertation is on the services that provide the defenses against DDOS attacks, Section 2.3 describes DDOS attacks, Section 2.4 describes the defenses and Section 2.5 explains the services.

2.1 THE INTERNET INFRASTRUCTURE

The Internet infrastructure consists of backbone networks, access networks and end user premises. End user premises, such as personal computers, connect to an access network through dial-up lines, cables, DSL or Ethernet. An access network then connects to the point of presence (POP) of backbone networks. Organizations that need dedicated network communications usually build their own access networks and connect them to backbone networks. For example, Internet access providers, such as AOL, need to connect their end users. Internet content providers, such as Yahoo, a university or a large corporation would need dedicated network connections to provide network services.

Backbone networks exchange network traffic through Network Access Points (NAPs) or private peering points. Figure 2.1 illustrates the Internet infrastructure. Throughout this dissertation, network providers (or Internet Service Providers, ISPs) refer to the network operators of backbone networks. Subscribers refer to the network operators of access networks since they subscribe dedicated network connections to backbone networks from network providers.

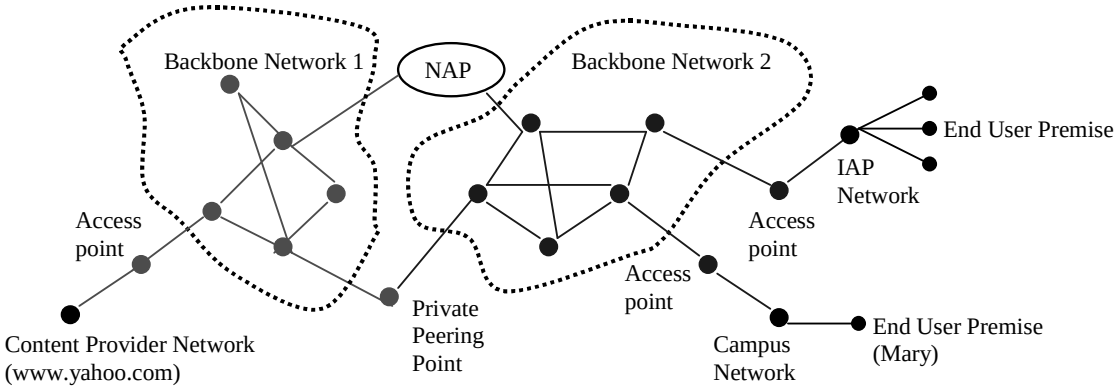

Figure 2.1: The Internet infrastructure

The current Internet utilizes packet switching to transport data. All data transported on the Internet are capsulated as IP packets, which contain the source addresses, the destination addresses, and the contents of the data. Once a computer sends out IP packets, the closest router examines the destination addresses of the IP packets and decides the next router to send the packets based on the routing table that the router maintains. The IP packets are then forwarded by routers one stop by one stop until they reach the destination addresses. The information in a routing table is determined by the routing protocols among routers, which exchange forwarding information periodically.

For example, as in Figure 2.1, if Mary sends out requests from her desktop on campus to access Yahoo's web pages, the requests are sent as IP packets and forwarded by the border router of the campus network to the access point of backbone network 2. The packets are then forwarded by the routers in backbone network 2 to backbone network 1 and finally reach the network where Yahoo's web servers are located.

2.2 INTERNET-BASED ATTACKS

A security incident is a group of attacks that can be distinguished from other attacks because of the distinctiveness of the attackers, attack, objectives, sites, and timing. An attack is an event that occurs on a computer or network as part of a series of steps intended to result in something that is not authorized to happen (Howard 1998). Based on the previous definition, this dissertation further defines an Internet-based attack as an attack that is launched from one or many computers connected to the Internet and that compromises the availability, integrity, or confidentiality of attack victims. An attack victim is defined as the target of attacks, which could be a network, a computer, an information system or an online service.

The number of Internet security incidents is growing as more computers are connected to the Internet. Figure 2.2 shows the trend of the number of hosts connected to the Internet and the number of Internet security incidents handled by the Computer Emergency Response Team/Coordination Center (CERT/CC)[1]. Many recent Internet-based attacks, such as Code-Red worms, Nimda worms and Slammer worms, have utilized

[1] CERT/CC (Computer Emergency Response Team /Coordination Center) is located in Software Engineering Institute at Carnegie Mellon University. Since November 1988, it provides the Internet community an organization that can coordinate responses to the security incidents on the Internet.

automatic mechanisms to propagate attack tools. That is, the attack tools are spread as computer viruses/worms in order to take over as many vulnerable computers as possible in a short period of time. Strategies to propagate countermeasures against the spread of computer viruses have been studied in (Chen and Carley 2003) and will not be discussed further in this dissertation.

The CERT/CC has identified six attack trends (CERT/CC 2002). 1) The level of automation in attack tools continues to increase. Attack tools are easier to use. 2) The sophistication of attack tools is increasing. As a result, it has become increasingly difficult to distinguish attack signatures from legitimate network traffic. 3) The number of newly discovered vulnerabilities reported to the CERT/CC continues to more than double each year. It is difficult for administrators to keep up to date with patches. 4) More technologies are designed to bypass typical firewall configurations. 5) Asymmetric threat is increasing. A single attacker can relatively easily employ a large number of distributed systems to launch devastating attacks against a single victim. 6) The threat of infrastructure attacks is increasing. Infrastructure attacks are attacks that broadly affect key components of the Internet.

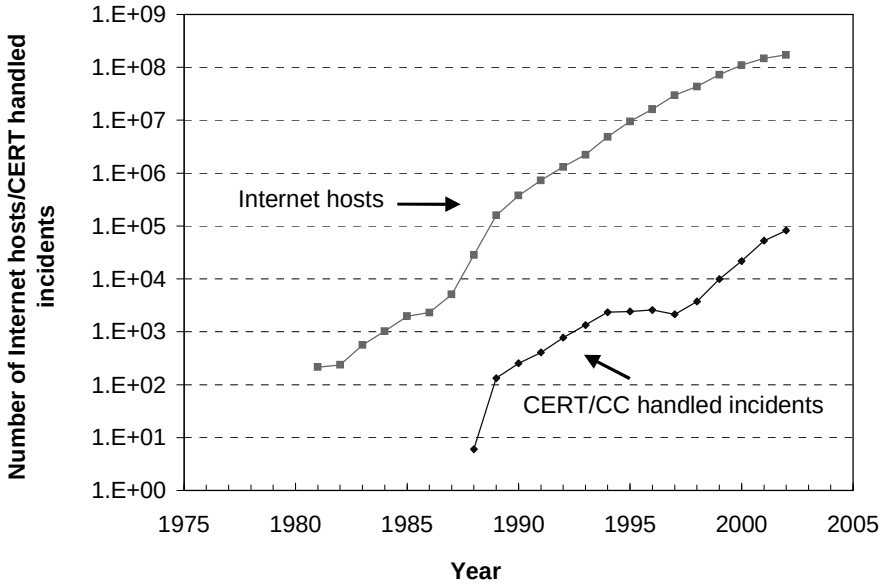

Figure 2.2: The growing trend of Internet hosts and Internet security incidents[2]

2.3 DDOS ATTACKS

DDOS attacks are an Internet-based attack that aims at compromising the availability of computers or network resource. A denial-of-service attack is considered to take place only when access to a computer or network resource is intentionally blocked or degraded as a result of malicious actions taken by another user. These attacks do not necessarily damage data directly, or permanently, but they intentionally compromise the availability of the resource (Howard 1997). In a distributed denial-of-service attack, an attacker could trigger tens of thousands of concurrent attacks on either one or a set of targets by using unprotected Internet nodes around the world to coordinate these attacks (CERT/CC 1999).

[2] The source of the Internet security incidents is the CERT/CC (www.cert.org). The source of the Internet hosts is Internet Software Consortium (www.isc.org).

Figure 2.3 shows the attack flow of a typical DDOS attack system. The "intruder" controls a small number of "masters," which in turn control a large number of "daemons." These daemons are used to launch packet flooding or other attacks against "victims" targeted by the intruder. This dissertation will focus on the attack traffic that is sent from "daemons" to "victims".

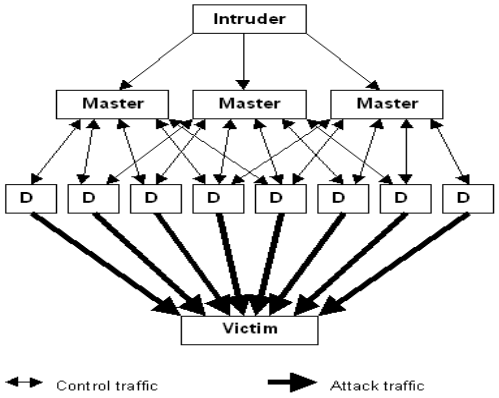

Figure 2.3: A typical DDOS attack system (CERT 1999)

The following terminology is used in describing DDOS attacks:

- Attack source (or sources): An attack source refers to the computer that is utilized by attackers to generate attacks against victims. That is, the computer that is installed with "daemons" as in the Figure 2.3.

- Source network: A source network refers to the access network that the attack sources are located.

- Attack victim (or victims): An attack victim refers to the port/online service, the computer, or the network that is the target of an attack.

- Victim network: A victim network refers to the access network that the attack victims are located.

- Upstream network and downstream network: An upstream network refers to the backbone network that an access network sends out network traffic. A downstream network refers to the service network directly connected to the access network of the destinations.

- Attack traffic: Attack traffic refers to the attack packets sent out by the daemons of a DDOS attack system against attack victims.

- Legitimate traffic: Legitimate traffic refers to regular network traffic which is not generated by a DDOS attack system.

A DDOS attack can unfold in the following way. Referring to Figure 2.1, suppose that DDOS attacks are launched against Yahoo's web servers from computers connected to the DSL line provided by the IAP network in the backbone network 1. Theses computers are attack sources and the IAP network is the source network while Yahoo's web servers are attack victims and Yahoo's network is the victim network. In this example, to maintain the availability of Yahoo's web servers during such an attack, the mitigation strategy is to detect and filter out the attack traffic at some points of the routing path from the IAP network to Yahoo's network.

Several reasons have made tracing and filtering DDOS attacks difficult. First, IP spoofing conceals the true origins of attacks. IP spoofing means attackers use false source IP addresses in attack packets to conceal their origins. The source addresses of IP packets

are not required for IP routing since the routers need only the destination addresses in order to forward the IP packets. Senders of IP packets can forge the source addresses in order to hide their true identities. The forged source addresses make it difficult to trace and to determine the true origins of DDOS attack traffic within the current IP routing environment.

Secondly, tracing and filtering attacks is not only a technical problem but also a policy and economic problem since attack sources can be distributed across multiple administrative domains. Since vulnerability-scanning tools have been automated as mentioned earlier, attackers can exploit the vulnerable computers across the Internet and utilizes them as attack sources. As a result, attack sources can be distributed across multiple administrative domains. In this case, the attack tracing and blocking is more difficult since it involves the cooperation of multiple network providers and subscribers. Under this circumstance, the attack tracing and filtering is a policy and economic problem among various network providers.

Thirdly, filtering attack traffic has a side effect on legitimate traffic because attack tools utilize various vulnerabilities in IP protocols that make it harder to distinguish attack traffic from legitimate traffic. Many tools have been used to launch DDOS attacks, such as Trinoo, TFN, Stacheldraht, and Mstream (Dietrich, Long et al. 2000; Dittrich 2001). Several characteristics in these attack tools making it hard to distinguish attack traffic from legitimate traffic (Houle and Weaver 2001). 1. These tools usually have options to control the packet rate of attacks. 2. The same tool can be used to conduct various flooding attacks such as UDP floods, TCP SYN floods and ICMP echo request floods. 3. TCP and UDP based packet flooding attack tools sometimes alter source and/or destination port numbers

to make reacting with packet filtering by service more difficult. 4. Variants of the attack tools are created based on the same exploit methods used to avoid detection of a specific attack signature. 5. Most of these tools can be used to forge source addresses in attack packets.

Finally, automatic responses against attacks are needed because DDOS attacks can severely damage the availability of the victims in a short period of time before appropriate manual responses can take place. The availability of the attack victims can be compromised in a short period of time once a DDOS attack is launched. It takes time for attack victims and downstream network providers to figure out what has happened and how to react against ongoing attack traffic. In addition, if the attacks are launched across multiple administrative domains, the downstream network provider could not filter attacks effectively without the cooperation of upstream network providers.

2.4 DDOS DEFENSES

In responding to ongoing DDOS attacks, a variety of defenses have been proposed. This section provides an overview of all current solutions to DDOS attacks. Since the focus of this dissertation is on the provision of DDOS defenses from network providers to their subscribers, the defenses evaluated in the later chapters will be focused on network-based defenses that are designed to actively mitigate ongoing attack traffic. Chapter 3 will provide a detail description and characterization for these network-based active defenses.

1) Reaction points: network-based vs. host-based

Reaction points refer to where the responses against attacks take place. Reaction points could be network-based such as those on network routers or host-based such as those on servers that the attack targets. Host-based defenses refer to the defenses that are deployed on the machines that are potential targets of attacks, and defenses are used to increase the tolerance of the targets to the attacks. The methods proposed in (Spatscheck and Peterson 1998; Yan, Early et al. 2000) are in this category. These methods can only mitigate the impact of attacks on the services that the attack targets provide but not block attacks. When attack traffic is large enough to deplete the resources used for mitigating the attacks, additional methods for blocking attacks are needed. Network-based methods are deployed on the points where packets route through the network connections to the targets, such as routers or proxy servers (Ferguson and Senie 1998; Bellovin 2000; Burch and Cheswick 2000; Savage, Wetherall et al. 2000; Stone 2000; Mahajan, Bellovin et al. 2001; Park and Lee 2001b; Ioannidis and Bellovin 2002). These methods are used to either trace or block attack traffic. This dissertation focuses on network-based defenses.

2) Type of response: active vs. passive

A few defenses are designed to actively respond to the attack traffic while the majority are designed to passively trace/log attack traffic. Tracing back to the real sources of attacks has been an established part of DDOS defense studies (Bellovin 2000; Burch and Cheswick 2000; Savage, Wetherall et al. 2000; Park and Lee 2001a; Snoeren, Partridge et al. 2001; Song and Perrig 2001). These methods could facilitate future liability assignments if source IP addresses of attack packets are forged. These methods are for identifying the sources of attacks, not for stopping ongoing attack traffic. In contrast, other defenses are designed to actively reduce the amount of ongoing attack traffic (Ferguson and Senie 1998;

Mahajan, Bellovin et al. 2001; Park and Lee 2001b; Ioannidis and Bellovin 2002; Sung and Xu 2002; Yaar, Perrig et al. 2003). This dissertation focuses on the ones that actively reduce ongoing attack traffic although methods of tracing attack traffic will be discussed to explain how attacks are detected.

3) Attack traffic sampling: probabilistic sampling vs. check-everything

Since examining every packet that goes through a router may impose an enormous storage or computational power requirement, some defenses sample network packets probabilistically to reduce the number of packets to be examined and logged (Huang and Pullen 2001). This dissertation focuses on defenses that check everything once they are triggered.

4) Reaction timing: constant vs. event-triggered

Some defenses needed to be active all the time in order to detect suspicious packets. Egress(SANS 2000) and ingress filtering (Ferguson and Senie 1998) are deployed at local edge routers to examine all incoming and outgoing packets. However, if a defense can be automatically turned on whenever an attack is launched, the overhead could be limited to a certain time period. However, it is difficult to determine the exact timing to trigger a defensive response. A few defenses are triggered based on the congestion level of network links (Huang and Pullen 2001; Mahajan, Bellovin et al. 2001; Xiong, Liu et al. 2001; Ioannidis and Bellovin 2002). This dissertation will model both constant- and event-triggered responses.

5) Detection criteria: attack signatures, congestion pattern, protocols, or source IP addresses

It is hard to distinguish attack packets from legitimate packets especially when both types of packets are sent to the same destination. Many different criteria have been examined. Each criterion has a tradeoff in terms of the number of false positives and false negatives associated with the outcome. Moreover, some criteria are only effective at identifying certain types of attack packets. For example, most intrusion detection systems detect attacks based on anomaly pattern matching or statistical measures of attack signatures (Debar, Dacier et al. 1999). The pushback method treats traffic aggregates as attack flows (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002). In (Schuba, Krsul et al. 1997), a revised TCP state machine is used to identify TCP SYN packet flood. A route-based method detects attack packets with spoofed source IP addresses based on the knowledge of the network's topology on core routers (Park and Lee 2001b). In the next chapter, defenses will be further characterized based on detection criteria.

6) Deployment location: a single point, attack path, or distributed points

Deployment location refers to where a defense is placed and triggered. If a defense is placed at the firewall or the proxy server in a subscriber's network (Schuba, Krsul et al. 1997), it will help the subscriber to discover attacks[3] but will not be effective when the bandwidth of the subscriber's network is saturated. The pushback method triggers filters along the path that traffic aggregates travel (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002) if the routers on this path have deployed such a defense in advance. A

[3] False positive here means the rate of mistakenly regarding normal packets as attack packets.

defense can be gradually deployed at distributed locations across a network (Schnackenberg and Djahandari 2000; Park and Lee 2001b; Ioannidis and Bellovin 2002). To prevent the attack detection from slowing down the backbone network, CenterTrack routs suspicious traffic to an additional overlay network (Stone 2000). In the next chapter, defenses will be further characterized based on the deployment location.

2.5 THE PROVISION OF DDOS DEFENSES

This dissertation proposes that ISPs provide defenses at their network as security services to their subscribers. Security services, such as Virtual Private Networks, have been provided by ISPs as optional network services to deal with the secrecy of data transportation. In this case, the services that provide DDOS defenses ensure the availability of an online service or a network. In this dissertation, the security services proposed are called the network defense services, which actively filter out attack traffic that is detected. In Figure 2.4 a policy framework that describes the context in which DDOS defenses are deployed is shown. An ISP can provide the network defense services to its subscribers along with network connection services. ISPs and subscribers can define how the DDOS defense is provided using a service level agreement (SLA). When attackers launch DDOS attacks on one of the ISP's subscribers, the ISP responds to the attack based on its cost concerns and the requirements of subscribers defined in the SLA.

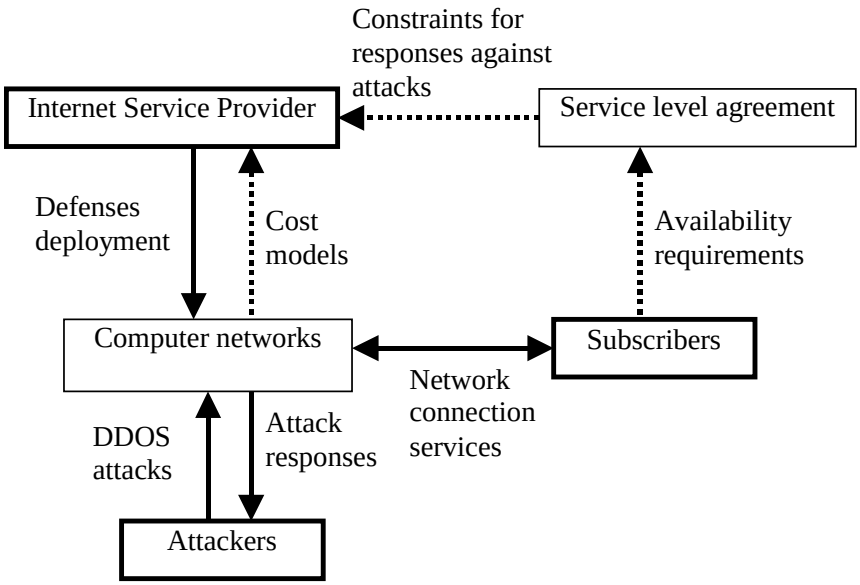

Figure 2.4: Context for the provision of DDOS defenses

The SLA is a legal contract. In principle, the definition of DDOS defenses in the SLA should be simple and flexible enough that an ISP can adjust defenses to minimize the cost imposed on its network while subscribers would remain satisfied with the effectiveness of defenses during attacks. However, the lack of a systematic understanding of DDOS defenses and the inherent cost and performance tradeoffs makes this system ineffective. This dissertation intends to use a computational model of this system to evaluate these tradeoffs and the underlying parameters to enable more effective definitions of DDOS defenses in the SLA. For example, each defense is only effective against a certain attack scenario for any given network topology. Consequently, the cost model of a network for an ISP and the availability requirements from subscribers should be different for each combination of attack scenario and network topology. This dissertation will analyze the cost imposed on a network by different known and expected attack scenarios and defenses. In addition, this dissertation will provide a systematic approach for analyzing future attacks and defenses.

2.6 SUMMARY

This chapter explains the terminology used in this dissertation, provides background information for Internet infrastructure, Internet-based attacks, DDOS attacks and defenses. The next chapter will further categorize the defenses that actively react against DDOS attacks.

Chapter 3 CHARACTERIZATION OF DDOS DEFENSES

Many defenses have been proposed to defend against distributed denial-of-service (DDOS) attacks. In order to provide insights to network operators so that they will know which defenses should be taken under what circumstances, it is necessary to categorize these defenses based on factors that will influence the performance and the deployment of these defenses.

This chapter focuses on a qualitative study of various DDOS defenses. The purpose is to identify the uncertain factors in these defenses so that they can be used in a later qualitative study on the performance and the deployment costs of these defenses. To achieve this purpose, this chapter reviews defenses that have appeared in the literature and characterizes them based on their attack detection algorithms and attack responses. An attack detection algorithm refers to the procedures which a defense uses to identify attacks based on available network information. An attack response refers to the mitigation strategies that a defense triggers once an attack is identified. This chapter categorizes DDOS defenses based on attack detection algorithms and attack responses because the effectiveness and operational costs of a defense are highly dependent on them.

Using this categorization, Internet Service Providers (ISPs) can consider a framework to provide defenses as network services to their subscribers. The next Chapter will utilize this categorization to develop a simulation tool for quantitative analyses on the performance and the deployment costs of various defenses. Chapter 5 and Chapter 6 will analyze the performance efficiency of these defenses and the economic feasibility of providing these defenses as network services respectively.

This chapter is organized as follows. The next section explains the methods of characterization. Section 3.2 categorizes defenses in terms of attack detection algorithms. Section 3.3 categorizes defenses in terms of attack responses. Section 3.4 summarizes the characterization.

3.1 METHODS OF CHARACTERIZATION

Both firewall technology (Cheswick and Bellovin 1994; Zwicky, Cooper et al. 2000) and intrusion detection systems (Mukherjee, Heberlein et al. 1994; Debar, Dacier et al. 1999; Axelsson 2000) have been developed to detect and to respond against various kinds of Internet-based attacks. However, defenses which are specifically designed to respond against DDOS attacks have not drawn much attention until recent years. Since 1999, various automatic DDOS tools have been created (CERT/CC 1999). In particular, large scale DDOS attacks in February 2000 against multiple e-commerce web sites (Tran 2000; Yankee 2000) highlights the potential risk and the severe impacts of DDOS attacks. Current literature on the characterization of DDOS defenses is very limited, and each of the current works serves a different purpose than this chapter.

Most of the available literature, which propose new defenses, review existing defenses. Among these, Savage et al. describes the pros and cons of various defenses (Savage, Wetherall et al. 2001) most extensively, but their purpose is to compare these defenses with a proposed IP traceback method. The most similar work to this chapter is the taxonomy of DDOS defense mechanisms (Mirkovic, Martin et. al. 2002). This taxonomy reviews DDOS defense mechanisms in general, and focuses on finding new features in the DDOS attack problems that have not been solved. The purpose of this chapter is not to provide a complete list of DDOS defenses, but to identify factors that influence the performance and the deployment costs of defenses through a qualitative analysis of various defenses. This purpose has not been addressed in the previous works. The characterization in this chapter is expected to provide ISPs insights on the design of the service provision framework for these defenses, and to provide a foundation for quantitative analyses of the problems associated with the service provision.

To serve this purpose, this chapter will only characterize the defenses that have the following two properties:

1) Reaction points which are network-based: Reaction points to attacks could be network-based such as those on network routers or host-based such as those on servers of the attack victims. Network-based methods are deployed on the points where packets route through network connections, such as routers or proxy servers. Host-based defenses are deployed on the machines that are potential targets of attacks. Host-based methods (Spatscheck and Peterson 1998; Yan, Early et al. 2000) could increase the victims' capability to stay available during attacks but cannot to filter out attack traffic before it reaches victims. These methods increase

the tolerance of victims to attacks and can be used together with the network-based methods. However, this chapter characterizes only defenses in which the reaction points are network-based since the deployment of these defenses requires the cooperation of ISPs and the host-based methods cannot overcome bandwidth saturation attacks.

2) Attack responses which are active: Some defenses are designed to passively trace/log attack traffic. Tracing back to the real sources of attacks has been an established part of DDOS defense studies (Bellovin 2000; Burch and Cheswick 2000; Savage, Wetherall et al. 2000; Park and Lee 2001b; Snoeren, Partridge et al. 2001, 2002; Song and Perrig 2001). These defenses could facilitate future liability assignments but cannot mitigate the impacts of ongoing attack traffic. These defenses have been analyzed previously (Lipson 2002). This chapter only focuses on the defenses which are configured with automatic responses against attacks once they are identified because these defenses can actively mitigate the impact of ongoing attacks on victims.

To characterize the network-based active defenses, the chapter adopts the following two aspects:

1) Attack detection algorithms: Attack detection algorithms are methods to determine whether or not the network traffic monitored should be regarded as attack traffic based on predefined characteristics. Attack detection algorithms can be classified into three categories: congestion based, anomaly based and source validation based. These categories specify different granularity of

attacks, different characteristics to detect attacks, and thus determine different false positive rates under various circumstances. These categories can be used to distinguish the performance tradeoff in defenses caused by the false positive rates.

2) Attack responses: Attack responses are mitigation strategies that a defense triggers in responding against attack traffic. Attack responses can be divided into two categories based on the direction of network traffic that is monitored. These two categories are destination filtering and source filtering. Destination filtering refers to responses that are triggered when attacks are detected in the inbound traffic to the victim networks. Source filtering refers to responses that are triggered when attacks are detected in the outbound traffic from the source networks. These categories can be used to define service provision to different subscribers.

The characterization is based only on information from the current literature that documents attack detection algorithms and attack responses in enough detail. Although many commercial products (Arbor 2002; Asta 2002; Recoourse 2002) satisfy the two properties mentioned above, the technical details in the public available documents are not enough to create a characterization.

3.2 ATTACK DETECTION ALGORITHMS

An attack detection algorithm analyzes network traffic information of the monitored links to determine if the packets transmitted through the links are legitimate. Network traffic information used to identify attack traffic include:

- network packet headers,
- packet rates of network flows/connections, or
- information on dropped packets.

Fields of the network packet headers used to identify attacks include:

- source IP addresses, which indicate the hosts that send the packets,
- destination IP addresses, which indicate the hosts that will receive the packets,
- IP protocol type, such as TCP, UDP, ICMP.
- TCP and UDP source and destination ports, which indicate the port number that the sender and receiver of a specific application use to communicate with each other.

Attack detection algorithms use one or a combination of fields in packets or network traffic information to determine if suspect traffic matches some characteristics of attack traffic, such as the congestion level of links caused by network flows, anomaly TCP connection behavior, or spoofed source IP addresses. Based on these characteristics, this section categorizes defenses as being one of the three categories: congestion-based, anomaly-based, and source validation based. These three categories along with features of various attack detection algorithms are described in the following sub-sections. Since none of these methods can perfectly identify attack traffic without raising false alarms, the effectiveness of the methods should be specified by a false-positive rate and a false-

negative rate[4]. False-positive rate refers to the ratio that legitimate traffic is determined as attack traffic. False-negative rate refers to the ratio that attack traffic is determined as legitimate traffic. Table 3.1 summarizes the characterization based on attack detection algorithms.

3.2.1 CONGESTION-BASED

Defenses in this category determine if there is an attack based on the congestion level of the monitored network links. Once the monitored network links are congested, the attack detection algorithm identifies the type of network flows/connections that contribute to the congestion. These methods identify attack traffic effectively only when attack traffic induces congestion of the monitored links, and the congestion can be observed.

Aggregate-based congestion control (ACC) (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002) has been proposed to reduce DDOS attack traffic and flash crowds based on congestion level. DDOS attack traffic is defined as a high-bandwidth aggregate, which is a collection of packets from one or more flows that have the same destination address prefix. The detection algorithm in ACC determines the destination addresses of the victim machines based on the destination network prefix of packets dropped at the observed router during a very short period. If the number of dropped packets of a certain destination address is larger than average, ACC puts the destination address on a list. The destination addresses in this list are then clustered into 24-bit or longer network prefixes. The arrival rate of each network prefix is estimated from the number of dropped packets. If the arrival rate of a network prefix exceeds a threshold, ACC regards all traffic

[4] False-positive rate refers to the ratio that legitimate traffic is determined as attack traffic. False-negative rate refers to the

to this network prefix as DDOS attack traffic and responds to all incoming traffic sent to this network prefix. The setting of the threshold and the responses will be discussed later in Section 3.3.

Many other studies (Sterne, Schnackenberg et al. 2001; Huang and Pullen 2001; Xiong, Liu et al. 2001) have suggested network congestion level as an indicator of DDOS attacks. These studies focus on attack responses with an implicit assumption that the responses are triggered when link congestion is observed. However, the methods used to determine congestion has not been specified in these studies.

3.2.2 ANOMALY-BASED

Defenses in this category detect anomalous patterns in network traffic. Once pre-defined anomaly patterns are detected in the monitored network links, the attack detection algorithm identifies the type of network flows/connections that contribute to the anomaly.

3.2.2.1 TCP SYN anomaly detection

TCP SYN flood attacks is one type of DDOS attacks that exploit half-open TCP connections to deplete the memory of receiver machines. To initiate a normal TCP connection, a sender first sends a "SYN" packet, and the receiver then sends back a "SYN ACK" packet to acknowledge the sender. The sender replies with an "ACK" packet to complete the initialization. In a TCP SYN flood attack, the senders do not reply "SYN ACK" packets. A TCP connection to which the sender has not responded is called a "half-open" TCP connection. The receiver machines stores the connections in system memory

---

ratio that attack traffic is determined as legitimate traffic.

and wait for replies. Since the replies never come, the “half-open” TCP connections eventually deplete the memory of the receiver machines and the machines can no longer serve further connections.

An active monitoring tool has been developed to monitor and to reduce TCP SYN flood attacks (Schuba, Krsul et al. 1997). The active monitoring method monitors TCP traffic at several points on a local network and utilizes a state machine to determine attack traffic. A new source address that sends TCP SYN is recorded and is assigned to a “new” state. The source addresses that do not reply SYN ACK are assigned to a “bad” state. Any SYN packets from the source addresses in the “bad” state are regarded as attack traffic. However, if attackers forge and randomize the source addresses of attack packets even if they are sent out from the same machine, the memory of the receiver machine can still be depleted by a large amount of TCP SYN packets. Section 3.2.3 will discuss methods that deal with attacks using false source addresses.

3.2.2.2 Asymmetric TCP communications

MULTOPS (Gil and Poletto 2001) has proposed to detect TCP SYN floods at network routers based on TCP packet rates. In a normal TCP connection, receivers acknowledge packets from senders at a constant rate so that the number of the packets received is proportional to the number of packets sent between the two parties of a connection. In TCP SYN flood attacks, attack sources send out a large amount of SYN packets but receivers will probably not to reply to the SYN packets. Based on this pattern, Gil and Poletto assume that the packet rate for the traffic to a network prefix is proportional

to the packet rate from the same network prefix. If the proportional pattern changes, the network prefix is either the source of an attack or the destination of an attack.

3.2.2.3 Normal models of network flows

D-WARD (Mirkovic, Prier et al. 2002) proposes to detect DDOS attack traffic by matching network traffic information with predefined normal flow models. This approach monitors both inbound and outbound traffic of a source network, and is intended to stop attack traffic originating from a network at the border of the source network. Attack flows are identified if they mismatch the normal flow models. Since TCP peer acknowledges every packet it receives, the proposed TCP normal model is defined by a maximum allowed ratio of the number of packets sent and received in the aggregate TCP flow to the peer. The proposed ICMP normal model is defined by a maximum allowed ratio of the number of ICMP request and reply packets, since each normal ICMP message should be paired with a corresponding reply. Since UDP peer is not required to reply to a UDP message, the normal UDP flow model can only be defined by a set of thresholds on UDP packets sent. Although the system is currently underdeveloped, the D-WARD proposal illuminates a new way to detect DDOS attacks at their sources. The false positive rates of this approach will depend on the calibration of the proposed normal flow models.

3.2.2.4 MIB variable correlation

Network management information can be used to detect DDOS attacks (Cabrera, Lewis et al. 2001). SNMP is a network management protocol that stores information about network devices in local databases each of which called a Management Information Base (MIB)(Waldbusser 2000). Local SNMP agents update variables in MIB periodically.

Network administrators can retrieve MIB variables at a central location to monitor the traffic sent to local network devices. The assumption is that some MIB variables may indicate attacks if these variables from receiver machines and from sender machines have some correlation on a sequential time line. For example, in ICMP ping flood, attackers send out ICMP Echo requests in which the IP variable in MIB is "ipOutRequest", and later the receivers will reply with an ICMP Echo in which the same set of variables contains "icmpInEchos." The detection algorithm queries the values of several specific MIB variables from local network devices periodically and correlates the relationship of these values. The purpose of the correlation is to reduce the false positive rate of identifying attack traffic.

3.2.3 SOURCE VALIDATION BASED

Since the current IP protocol permits source hosts to alter source addresses in IP packets, attackers are able to send out IP packets with empty or false source addresses. Although IPSEC[5] (Kent 1998a, 1998b), a transport layer authentication scheme, can be used to authenticate the source addresses of IP packets, but this method is not widely adopted at this point. The false source address problem will still be a big problem in detecting and filtering DDOS attacks in the short term. Because of this reason, the attack victims cannot rely on the source addresses in attack packets to distinguish them from legitimate packets. In this category, various ways are designed to validate the sources of IP packets. Once the source of a packet is determined to be from an attack source, the packet is filtered out.

[5] In particular, authenticated header (AH) is used to provide connectionless integrity and data origin authentication for IP datagrams and to provide protection against replays (Kent 1998b).

3.2.3.1 Ingress filtering

Ingress filtering (Ferguson and Senie 1998) determines false source addresses at edge routers based on the valid IP address range internal to the network. However, a false source address within the valid source addresses within the same network will not be detected by this method. For example, if a packet is sent out from host A with the source address of host B, ingress filtering will not regard it as a false source address if the source address of B is valid in this network. In addition, network traffic from a legitimate mobile IP address has to be tunneled to avoid ingress filtering. This method is enough for tracing the attack traffic to the network but not the computers that originate them.

3.2.3.2 Route-based

Route-based filtering proposes filtering packets of spoofed source IP addresses based on routing information on backbone border routers (Park and Lee, 2001b). A border router maintains a routing table that contains fixed routes to all other domains by exchanging routing information with its neighboring routers in Border Gateway Protocol (BGP)(Rehter and Li, 1998). The proposal suggests using routing information to determine if a packet comes from the correct network device. If the source address of the packet is not consistent with the network device from which it is sent, the packet is regard as an attack packet and should be filtered out. However, current core routers maintain only a forwarding table (a list of destination network prefixes and the corresponding forwarding network interface) but do not maintain an incoming table (a list of source network prefixes and the corresponding incoming interface). Although the forwarding table in a router may indicate the routes that a packet will be forwarded, the routes are not necessary reversible because routing on the Internet is not completely symmetric (Paxson 1996). In addition, there is no

way to determine where the packet comes from when multiple routes are present. SAVE is a protocol being proposed to build up incoming tables in routers (Li, Mirkovic et al. 2001). This protocol proposes that routers propagate their incoming address space to their forwarding destinations.

3.2.3.3 Web connection authentication

A cryptographic method has been proposed to protect a web server from TCP SYN attacks with spoofed source addresses (Xu and Lee 2003). This method drops the first TCP SYN packet from the sender and sends back a HTTP redirection with an encrypted message. The sender who uses real source addresses would receive the encrypted message and include it in the next TCP SYN request. By doing so, all TCP SYN packets with spoofed source addresses will be drop before they reach the web server.

3.2.3.4 IP traceback-based

Methods in this category mitigate DDOS attack traffic by using IP traceback and packet filtering. Packet marking (Savage, Wetherall et al. 2001; Park and Lee 2001a; 2002; Song and Perrig 2001, Sung and Xu 2002; Yaar, Perrig et al. 2003) identifies the paths that attack traffic comes from by inserting marks in packets. Among these methods, currently only IP traceback-based intelligent packet filtering (Sung and Xu 2002) and Pi (Yaar, Perrig et al. 2003) have designed to filter out ongoing attack traffic.

The basic idea of packet marking is that the routers on the path from attack sources to victims insert marks in the IP identification field of ongoing packets, and the victims reconstruct the paths that the attack packets traverse by the marks in the packets. The

problem is that the IP identification field is only 16 bits, which is not enough for storing the entire path since the average path length is roughly 15 (Yaar, Perrig et al. 2003). Certain coding schemes have to apply to shorten the length of marks. Since the marks are not the unique identifier of an attack path after coding, the false positive rate occurs when the legitimate packets have traversed the paths that are coded as the same identifier as the paths that attack packets have traversed. IP traceback-based intelligent packet filtering (Sung and Xu 2002) proposes a preferential filtering to filter out packets with different types of marks with different probabilities. Pi (Yaar, Perrig et al.) proposes to filter packets at edge routers at a certain threshold if the packets have marks that indicate they are from attack sources. Since the mark under this scheme is not unique to every path, the threshold filtering allows the victim to lower the false positive rate at the expense of raising the false negative rate. Both methods allow attack victims to know the true origins of the network traffic but they need to be combined with other methods for identifying the patterns of attack traffic.

3.3 ATTACK RESPONSES

Attack responses in defenses are triggered once attack traffic is detected. To implement attack responses, contemporary routers usually have the functionalities to process network traffic flows based on a set of access rules that defines the characteristics of attack traffic (CISCO 2000). This section first describes the categorization of attack responses. A discussion of several features of attack responses that influence the performance of defenses and the deployment costs of the defense follows. Table 3.2 summarizes the categorization and these features.

3.3.1 CATEGORIZATION OF ATTACK RESPONSES

Attack responses can be applied on either inbound traffic or outbound traffic of a network. Defenses can be categorized into the following two categories based on the direction of network traffic to which attack responses are applied.

1) Destination filtering are attack responses that are triggered when attacks are detected in the inbound traffic of some destination networks. Defenses in this category monitor the network traffic received by some destination networks, and mitigate the impacts of ongoing attack traffic to these destinations. As in Figures 3.1 and 3.2, when subscriber 1 (in ISP 1's network) originates attacks on subscriber 2 (in ISP 2's network), the attack responses are deployed in ISP 2's network. In this case, ISP 2 (the downstream ISP) can only trace back the sources of attacks within the administrative boundary of its network, such as the access router connecting to the subscriber as in Figure 3.1 or the border of its network as in Figure 3.2. Proposed responses that fall in this category include Pushback (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002), Active Responses (Sterne, Schnackenberg et al. 2001, 2002), TCP anomaly detection (Schuba, Krsul et al. 1997), MIB correlation (Cabrera, Lewis et al. 2001), preferential filtering (Sung and Xu 2002) and threshold filtering (Yaar, Perrig et al.).

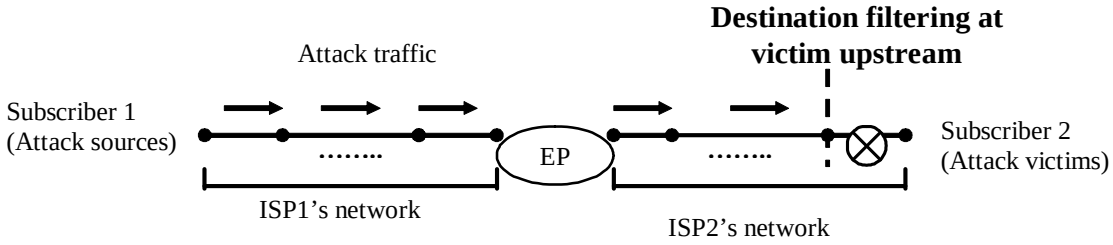

Figure 3.1: An illustration of destination filtering[6] (at victim upstream)

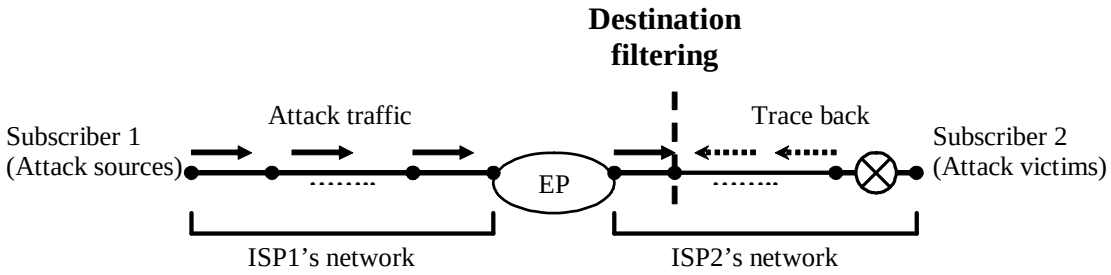

Figure 3.2: An illustration of destination filtering

2) Source filtering occurs when attack responses are triggered when attacks are detected in the outbound traffic of some destination networks. Defenses in this category monitor the network traffic sent from some source networks, and mitigate the impacts of ongoing attack traffic originating from these sources. Since the attacks are filtered out at the sources before they are sent to the downstream subscribers, this method decreases the observable number of attacks at downstream ISPs. Figure 3.3 illustrates an example where ISP 1 places filters at the upstream routers of subscribers 1 so that the attack traffic is filtered out before it is sent to subscriber 2. Defenses in this category are ingress filtering (Ferguson and Senie 1998) and D-WARD (Mirkovic, Prier et al. 2002). Both MULTOPS (Gil and Poletto 2001) and route-based filtering (Park and Lee 2001b) can be either implemented as destination filtering or source filtering.

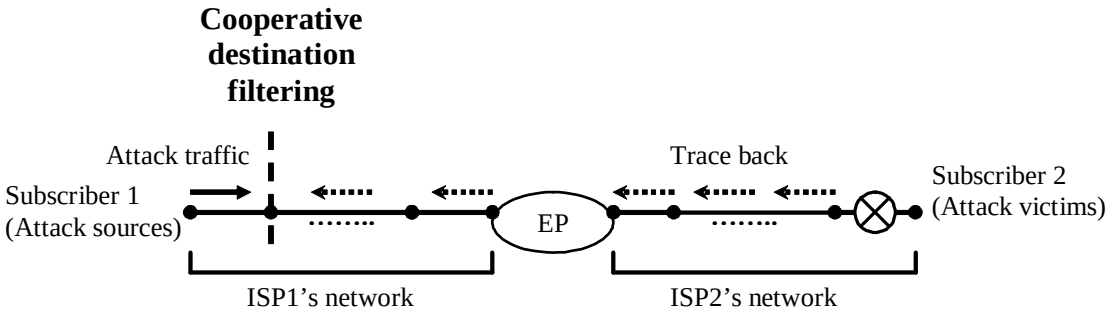

Figure 3.3: An illustration of source filtering

3.3.2 THE TYPE OF ATTACK RESPONSES

Packet filtering and rate limiting are two mechanisms to implement responses in access rules of routers. Either one of these methods is used in the defenses described earlier to implement filtering. Packet filtering either drops or accepts the packet being examined. The granularity of attacks in these two mechanisms is different. Packet filtering detects attacks based on per-packet information while rate limiting limits the transmission rate of the traffic flows to which the packet belongs.

Packet filtering is the action that a device takes to selectively control the flow of data to and from a network. Packet filters allow or block packets, usually while routing them from one network to another. To accomplish packet filtering, network administrators have to establish a set of rules that specify what types of packets are to be allowed and what types are to be blocked. Packet filtering may occur in a router, in a bridge, or in an individual host (Zwicky, Cooper et al. 2000).

Rate limiting is the function that allows a router to control the transmission rate of a specific traffic flow. Rate limiting is a traffic-policing tool used to control network

[6] EP refers to the exchange point that exchanges the network traffic between two backbone networks.

congestion. In the case of protecting against DDOS attacks, an attack detection algorithm identifies the characteristics of the traffic flow that will be policed. Once the characteristics are determined, the rate limiting function will guarantee that the transmission rate of the traffic flow will be lower than a certain rate, which means packets that arrive at a higher rate will be queued or dropped at the router.

Both packet filtering and rate limiting are mechanisms to respond against the DDOS attack traffic; however, they control the attack traffic in different ways. Packet filtering discards all packets that match the characteristics of attack traffic. In contrast, rate limiting allows some network traffic regarded as attack traffic to pass through, but it is limited by a transmission rate. Because of the difference, packet filtering is usually used with an attack detection algorithm that can detect attacks by packet headers, such as anomaly based and source validation based, and rate limiting is used with congestion based attack detection algorithms in which the attack traffic cannot be distinguished from legitimate traffic sent to the same destination.

3.3.3 ATTACK RESPONSE GENERATION

Attack response can be generated either statically or dynamically. Attack response is generated statically in ingress filtering against spoofed source IP addresses (Ferguson and Senie 2000) before the attack detection begins. For example, the network prefix of a local network is 204.69.207.0/24. The attack response can be defined to drop all packets in which source IP addresses is outside 204.69.207.0/24, which is generated statically once the network IP prefix of the local network is determined.

Attack response can be also generated dynamically when attack traffic is detected. For example, the outbound link of the above network is 2Mpbs. An attack is detected when the attack source sends 5Mbps TCP SYN packets to port 80 of the host 204.69.207.9 An attack response to limit the transmission rate of TCP packets to this machine can be generated dynamically to limit the packet rate of the network traffic sent to the host 204.69.207.9 to be much lower than 2Mbps.

3.3.4 DECISION LOCATIONS

Decision locations refer to where attack responses are generated if they are generated dynamically. In order to generate an attack response, an attack detection algorithm needs to collect network traffic information from the decision locations. Theoretically, attack response generation can be deployed at either one of the following locations:

- Attack sources (L1): edge routers of the local network from where the hosts send out packets.
- Source upstream (L2): access routers of an ISP that connect to subscribers' edge routers.
- Backbone routers (L3): core routers that transport network traffic.
- Victims (L4): edge routers of the local network where hosts will receive packets.
- Victim upstream (L5): access routers of an ISP that connect to edge routers of the victims' network.

In practice, attack response generation is rarely deployed at backbone routers (L3) since it is difficult under current technology to monitor high-speed backbone peering links and to analyze the information from these links for attack detection. Studies have been done on monitoring OC-48 peering links (Claffy, Miller et al. 1998; Sager 1998; Fraleigh, Moon et al. 2001). No current published study has monitored links higher than OC48.

Instead of deploying a defense at backbone routers, edge routers are another choice. In order to protect a local network against attack traffic from other networks, network administrators have an incentive to deploy attack detection tools at edge routers to examine inbound network traffic. All anomaly based detection algorithms described in Section 3.2.2 generate attack responses either at the victims (L4) or at the victim upstream (L5) to examine inbound network traffic.

Generating attack responses automatically at the attack sources or the source upstream (L1 or L2) is hard due to three reasons. First, the sources of attack traffic can be spoofed so that victims cannot identify the real sources of attacks. Secondly, even if the genuine sources of attacks can be identified, these sources can be located at many different administrative network domains. In this case, cooperative attack detection and response are necessary. Chapter 7 will investigate the incentives for the cooperative attack filtering and detection. Thirdly, technical difficulties occur for generating attack responses at the sources of attacks. In particular, it is hard to distinguish DDOS attack traffic from legitimate traffic at the sources of attacks since the volume of attack traffic is usually small and only aggregates at certain points close to destinations. Congestion based attack detection algorithms are not effective in this case since attack tools usually do not cause congestion

at the sources. Anomaly based algorithms, such as D-WARD, and source validation based algorithms are able to generate attack responses at attack sources.

3.3.5 ENFORCEMENT LOCATIONS

Enforcement locations refer to where on a network the attack responses will be applied. Once an attack response is enforced on a certain network router, all network packets that pass through the router/links will be examined. If network packets are determined to be attack traffic, the responses will be applied these packets.

Possible enforcement locations are the same as decision locations, which are L1-L5 in Section 3.3.4. The difference is that enforcement locations in practice are not as restrictive as decision locations. Once attack responses are generated, they can be distributed to other locations to be enforced. The overhead imposed by enforcing attack responses only occurs when an attack is detected if the attack responses is enforced and generated dynamically during attacks. In DDOS attacks, appropriate allocation of enforcement locations may enhance the performance of defenses and reduce the overhead imposed. The “appropriate allocation” of enforcement locations for various defenses will be analyzed in Chapter 5.

Both Pushback (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002) and Active Responses (Sterne, Schnackenberg et al. 2001, 2002) can be enforced at all location discussed above (L1-L5). However, the performance of the defense decreases when it is enforced closer to the victim networks.

3.3.6 COMMUNICATING PROTOCOLS

Communicating protocols refer to the protocols used to send control messages between various nodes of a network to coordinate attack detection or attack responses. These control messages are either attack patterns sent from attack detectors to attack response decision locations or attack responses sent from decision locations to enforcement locations. Sending control messages has been done manually which imposes high managerial overhead and has a longer lag time. To reduce the managerial overhead and lag time, communicating protocols have been studied to manage the generation and the distribution of attack responses in distributed locations. Three communicating protocols are explained in detail below.

First, pushback messages (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002) are used to distribute congestion patterns observed at congested links to trigger rate limiting in routers along the path that attack packets have traveled. The “pushback-request” message used to trigger rate limiting includes congestion signature, bandwidth limit, expiration time, depth (how many hops away from congested links), and message type.

Second, the Intruder Detection and Isolation Protocol (IDIP) is an application layer protocol that coordinates attack detection and response at distributed locations. In IDIP, attack detectors send descriptions of suspicious attack events to the Discovery Coordinator, which determines responses and sends out its decisions to nodes that will enforce the decisions (Schnackenberg and Djahandari 2000; Sterne, Schnackenberg et al. 2001, 2002).

Third, the Common Intrusion Detection Framework (CIDF) proposes a language called Common Intrusion Specification Language (CISL) for intrusion detection systems to

communicate attack responses (Staniford-Chen, Tung et al. 1998). CISL provides a common platform for communicating filter policy and attack detection patterns between heterogeneous intrusion detection systems located at distributed locations.

Finally, not all defenses require additional control messages. If each network node can detect attacks autonomously based on the information that a network node collects periodically from its neighboring nodes, attack detection can be implemented without additional communicating protocols.

3.3.7 ADDITIONAL OVERHEAD OF RESPONSES

Defenses mitigate the impact of the attack traffic on the victim network but may impose an additional overhead on the networks that implements them. The additional overhead includes computational overhead imposed by attack detection and attack response enforcement; storage requirement to save logs for attack detection; and communications overhead used to send control messages to distributed locations of a network. The overhead is described below in detail.

First, the computational overhead from attack detection is imposed on a regular basis while the overhead from filter policy enforcement is imposed when a filter policy is enforced. Once filter rules are enforced to examine network packets, a per-packet delay will occur for matching filter rules. Minimizing the per-packet delay is a packet classification problem in router performance optimization. Although most commercial routers are optimized for routing, the per-packet delay of matching filter rules depends on the number of filter rules, the number of characteristics used to identify attacks, and the updating frequencies of the filter rules (Feldmann and Muthukrishnan 2000).

Second, the storage requirement for attack detection depends on the capacity of the network device that the attack detection algorithm monitors and the information needed to determine attack patterns. To monitor high-speed network links, the storage requirement is usually very large. Current technology can scale up to 10Gbps link speed without losing much information on IP packets. To reduce the storage requirement and to catch network packets from high throughput routers, sampling and processing of packet data dynamically will be needed in the future (Iannaccone, Diot et al. 2001).

Third, control messages to coordinate attack detection are an additional overhead to network transmission. If communication occurs between network routers, it is important to know if such communication will result in abnormal behavior of routers. Since most commercial routers are optimized for routing, it is not certain if additional communications among routers will impose additional delay on routers or not. Since this issue is beyond the scope of this thesis, details about the communication overhead caused by communicating protocols will not be discussed further.

3.4 CONCLUSIONS

Categorizing DDOS defenses based on attack detection algorithms help to identify the factors that influence the performance tradeoff of defenses. In the congestion-based defenses, attack detection is based on link congestion and rate limiting is used to respond against attacks. False positives for these defenses occur when both attack traffic and legitimate traffic happen to have the same destination IP prefix. In the anomaly-based defenses, attack detection is based on the anomaly patterns of network traffic, and packet filtering is used to drop attack packets. False positives occur when legitimate traffic shows

anomaly patterns in some rare cases. In the source validation based defenses, attack detection is based on false source IP addresses. False positives occur only when the criteria for determining false IP addresses cannot distinguish it from true source addresses. However, the detection rate for attack traffic depends on how many attack packets contain false source IP addresses since this method cannot prevent attack packets with true source addresses.

Categorizing defenses based on attack responses can help potential subscribers in selecting a defense. ISPs can utilize the distinction between destination filtering and source filtering to design the service provision for either attack sources or attack victims. The locations where attack responses are generated and enforced determine the number of locations needed to deploy defenses, and thus influence the deployment costs.

Both Chapter 5 and Chapter 6 will conduct quantitative analyses on the defenses categorized in this chapter. Chapter 5 will compare the performance tradeoff of various defenses and Chapter 6 will analyze the economic incentives for ISPs to provide these defenses. In order to conduct the analyses in Chapter 5 and 6, a computational tool to simulate an attack-defense complex system is needed. The next chapter will describe this computational tool. This tool utilizes the factors described in this chapter about the defenses and the characteristics described in Chapter 2 about the DDOS attacks.

Category	DDOS defenses	Granularity of attack traffic	Network information needed to monitor	Characteristics of attack traffic	Sources of false positives	Limitations
Congestion based	ACC&pushback (Internet draft expired) (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002)	Flow	Destination IP prefix, transmission rate of network traffic	Network flows that cause link congestion	Legitimate traffic that contributes to the congestion	1.Legitimate traffic is punished the same as attack traffic. 2.False positive increases when the enforcement locations of responses closer to the victims 3.Can only identify attack traffic when congestion occurs
	Automatic responses & IDIP (Sterne, Schnackenberg et al. 2001, 2002)	Not specified	Not specified			
Anomaly based	TCP SYN anomaly (Schuba, Krsul et al. 1997)	Connection	IP protocol type (TCP SYN), source IP address	Expired TCP SYN half-open connections	Connections with longer transmission time will not be served	Can only apply on TCP SYN attacks
	MULTOPS (Gil and Poletto 2001)	Connection	IP protocol type (TCP), TCP packet rate, Source IP address or destination IP address	Asymmetric number of TCP packets to and from one source or destination	IP Routing is not necessary symmetric (inbound and outbound traffic may from different border routers)	Can only apply on TCP SYN attacks
	D-WARD (Mirkovic, Prier et al. 2002)	Flow or connection	IP protocol type, packet rate, source IP address, destination IP address	Packet rates to and from one source (TCP and ICMP) or a maximum sending rate(UDP).		Has to determine the threshold of packet rates for TCP and ICMP, and the maximum sending rate for UDP.
	MIB variables correlation (Cabrera, Lewis, et al. 2001)	Packet	Source IP address, destination IP address, MIB variables	Specific values in MIB variables	Some legitimate traffic has the same correlation	Can only apply within a network that is administrated by SNMP and MIB database
Source validation based	Ingress filtering (RFC 2267) (Ferguson and Senie 1998)	Packet	Source IP address, Valid source IP range	Spoofed source IP address	Traffic from an mobile IP that is not tunneled	1.Can not identify the attack traffic that does not utilize spoofed source IP 2.Need wide deployment
	Route-based filtering (Park and Lee 2001b)	Packet	Source IP address Valid source IP range	Spoofed source IP address	Forwarding tables in core routers do not provide enough information	1.Not apply to attacks that do not utilize spoofed source IP 2.Mechanism for core routers are currently underdeveloped
	Preferential filtering (Sung and Xu 2002), Threshold filtering (Yaar, Perrig et al. 2003)	Packet	IP identification (marks by intermediate routers)	Packets with marks considered as attack paths	Legitimate packets may contain the same marks as attack packets	Intermediate routers have to be reconfigured to insert marks.

Table 3.1: Characterization of DDOS defenses in terms of attack detection algorithms

Category	DDOS defenses	Response generation	Response mechanism	Decision locations	Enforcement locations	Topology dependent	Com. protocol	Overhead
Destination filtering (police inbound traffic of subscribers)	ACC&pushback (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002)	Dynamic	Rate limiting	Edge routers of destinations or upstream access routers (L4, L5)	All locations (L1-L5)	Yes	Yes. Network layer	Controls messages to push responses
	Automatic responses & IDIP (Sterne, Schnackenberg et al. 2001, 2002)	Dynamic	Rate limiting & packet filtering	The Discovery coordinator (single point on a network)	All locations (L1-L5)	Yes	Yes. Application layer	Control messages for coordination
	TCP SYN anomaly (Schuba, Krsul et al. 1997)	Dynamic	Packet filtering	Edge routers of destinations (L4)	Edge routers of destinations (L4)	Not specified	No	States of connections
	MIB variables correlation (Cabrera, Lewis, et al. 2001)	NA	NA	NA	NA	Not specified	SNMP for retrieving MIB variables	SNMP messages
	MULTOPS (Gil and Poletto 2001)	NA	NA	Edge routers of destinations or upstream access routers(L4, L5)	Edge routers of destinations or access routers(L4, L5)	Not specified	No	Hash table to store TCP connection info.
	Route-based filtering (Park and Lee, 2001b)	Static	Packet filtering	Core routers (L3)	Vertex cover set of Core routers (L3)	Yes	No	Route information
	Preferential filtering (Sung and Xu 2002), Threshold filtering (Yaar, Perrig et al. 2003)	Dynamic	Packet filtering	Edge routers of destinations (L4)	Edge routers of destinations or access routers (L4, L5)	Yes	No	Mark insertion in intermediate routers
Source filtering (police outbound traffic of subscribers)	MULTOPS (Gil and Poletto 2001)	NA	NA	Edge routers of destinations or upstream access routers (L1, L2)	Edge routers of sources (L1) or upstream access router of sources (L2)	Not specified	No	Hash table to store TCP connection info.
	Ingress filtering (Ferguson and Senie 1998)	Static	Packet filtering			Not specified	No	Access lists
	D-WARD (Mirkovic, Prier et al. 2002)	Dynamic	Rate limiting			Not specified	No	Hash table to compute flow measures
	Route-based filtering (Park and Lee 2001b)	Static	Packet filtering	Core routers (L3)	Core routers (L3)	Yes	No	Route information

Table 3.2: Characterization of DDOS defenses in terms of attack responses
(Both MULTOPS and Route-based filtering can be applied on either inbound or outbound traffic)

Chapter 4 A COMPUTATIONAL TOOL FOR SIMULATING ATTACKS AND DEFENSES ON THE INTERNET

During a DDOS attack, both attack sources and legitimate clients are sending network traffic to victims. The impact of attacks on victims depends on the capacity of its link, packet rates of both the attack traffic and the legitimate traffic, and the defenses deployed. These factors form a complex system where attacks and defenses interact to determine the impact on victims. In order to provide defenses as services to their subscribers, network providers have to realize the uncertainty in this complex system. By doing this, they can tune variables in defenses to meet the needs of their subscribers, and they can estimate the cost of operating the services.

This chapter describes a computational tool to simulate this complex system. This tool is intended to generate results for quantifying the performance tradeoff made by various defenses, and the economic costs of operating the services. Using the categorization and variables identified in the previous chapter, this tool can be used to generate different attack scenarios on a given network topology where a given defense is deployed. For each attack scenario, the tool can calculate measures to quantify both the performance impacts of the attacks and the defenses. The purpose of this tool is not to catch the dynamics of network traffic transportation from packet layer but to do the first order of magnitude

estimation on the impact of attacks under a certain circumstance. The results generated from this tool will be used in the next two chapters, which analyze the performance impact of defenses, and the economic costs and benefits from the services.

The next section describes the purposes of developing the computational tool. Section 4.2 reviews previous models and tools that simulate the complex attack-defense system. Section 4.3 provides an overview of this tool. Section 4.4 explains the algorithms used in this tool.

4.1 PURPOSES

The computational tool simulates a complex attack-defense system that describes how attack traffic along with other network traffic routes through a given network topology during a DDOS attack when a defense mechanism is deployed. This tool simulates the system at the abstract level and can generate quantitative measures for the performance and the operational costs of the defenses. These results can be used for solving various management and policy problems regarding the deployment of the defenses and the provision of the services. With this computational tool, the interaction of the uncertain variables in defenses and attacks can be estimated in the context of an entire network, rather than for a given point on that network. To orient the design of the tool, this chapter asks the following two research questions:

1) What is the impact of the topology of a network on the performance of defenses?

To mitigate the impact of DDOS attacks on the victim network, the defenses have to be deployed at some points of the routing path between the sources and the destinations

so that attack responses can react against the attack traffic before it is transported to the victim network. Hence, the relative topological locations of attack source networks to the victim networks determine the possible enforcement locations of attack responses, in which the performance of the defenses vary with enforcement locations or types of defenses.

2) What is the relative economic cost of operating services to provide defenses for a given network?

As discussed in the previous chapter, attack responses have additional overhead on routers, which may cause delay of serving other subscribers. The overhead varies with the number of filters triggered to defend the victim network. However, the additional capacity recovered from filtering out attack traffic offsets some of the overhead for attack response filters.

4.2 PREVIOUS MODELS AND TOOLS

Previous attack-defense models have demonstrated how computational tools could be used to reconstruct attack scenarios and victim responses. Cohen's model (Cohen 1999) simulates attack processes and defenses based on a predefined computer network topology. Cohen's model is an attacker-defender game, which could be useful for individual companies to evaluate their reaction time if attacked. Cohen concludes that the timing of acquiring attack or threat information is important for a defender. Moitra's work (Moitra 2000) uses a stochastic model to analyze CERT incident records. From simulation analysis, a correlation is confirmed between the probability of an incident and the damage an incident does. Based on the assumption that defense cost is correlated to the change of the functionality of a system, his work suggests that the survivability, defined as the probability

that a system is functional, increases rapidly at first and then more slowly as the defense cost increases. However, more data is needed to support this conclusion. Chaturvedi et al. (Chaturvedi, Gupta et al. 2000) design a multi-agent based model to study human decisions of taking risk in a simulated online bank operation. The preliminary results show that test subjects have different levels of risk tolerance. Red Teaming is a computer security attack simulation project developed in Sandia National Laboratories Information Design Assurance Red Team (IDART). Red Teaming uses human experts to attack real information systems in order to identify vulnerabilities of these systems and observe the behavior of attackers (Wood and Duggan 1999).

Previous models are designed for different research purposes. They do not include the variables associated with DDOS defenses and Internet topology. They are also not appropriate to investigate security policy issues related to DDOS attacks. However, previous models do show that computational modeling can be a powerful approach in security incident research because this type of research problem involves a complex system, and therefore conducting real world experiments is very difficult.

The complex attack-defense system described in this chapter is implemented in ANSI C. With the same sets of algorithms and parameters, this system can be implemented in other tools as well. Both agent-based simulation tools such as Repast[7] or modeling languages for distributed systems such as Easel[8] can be used to implement this system. The computational tool is grounded by theories in computational modeling (Carley and

[7] The description about Repast is available at repast.sourceforge.net

[8] The description about Easel is available at www.cert.org/easel

Gasser 1999). The next two sections describe the inputs, outputs and the algorithms that are in this computational tool.

4.3 OVERVIEW

4.3.1 PROPERTIES

The computational tool developed in this chapter has the following properties:

1) Abstract level: The computational tool simulates network traffic transportation based on a given network topology. Suppose that an ISP network *G* (Figure 4.1) is the backbone network of an ISP. Each node in the network represents a POP (Point of Presence) of the network, where the networks of subscribers connect. In Figure 4.1, circles represent POPs of the network. Squares represent the networks of subscribers. The networks of subscribers are either legitimate source networks ("x") that generate legitimate traffic, or attack source networks ("a" in Figure 4.1) that originate attacks, to the victim network ("v"). The tool records information about the relative locations of each node in the network, the packet rates of incoming and outgoing traffic for each node, and the capacity of the victim networks.

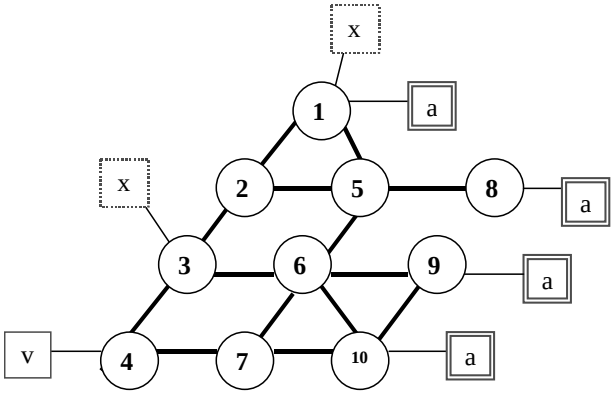

Figure 4.1: An example network

2) Complexity: The complexity of calculations increases with the number of nodes in the network, the number of concurrent attack sources simulated, and the number of concurrent attack victims simulated. In order to reduce the complexity of calculations, the tool uses Monte Carlo sampling to randomly pick attack source networks and to approximate the output measures.

3) Validation: Empirical grounding (Carley 1996) is used to validate this model. This validation approach includes establishing the reasonableness of the simulation model and initializing variables of the model by setting their upper bound, lower bound, and mean value from previous empirical studies. In the next two chapters, attack scenarios will be validated based on data from empirical studies of attack tools, propagation methods (Moore, Voelker et al. 2001), and observable historical data from computer virus propagation (Moore 2001). The types of network topologies will be validated through empirical backbone network topology data (BW 2001).

4.3.2 COMPONENTS

Figure 4.2 is an overview of the components in this computational tool. The computational tool consists of four sets of input parameters, including parameters that quantify the network scenario, the attack scenario, the attack detection, and the attack response. The network scenario parameters model how network traffic is transported on a network. Attack scenario parameters decide the number of victim networks and attack source networks for a scenario. The attack detection parameters and attack response parameters describe a given defense mechanism.

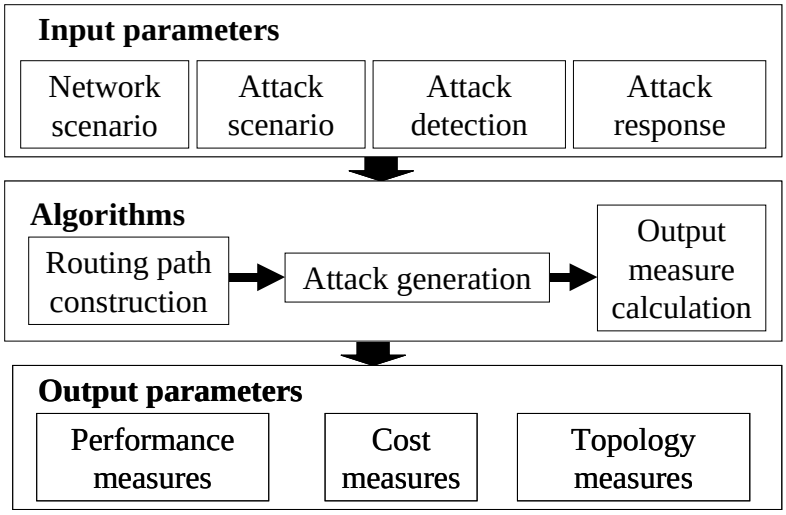

Figure 4.2: The overview of the components in the computational tool

Three sets of output parameters are generated from this tool, which includes performance measures, cost measures, and topology measures. Performance measures are for the analysis on the performance impact of the defenses. Cost measures are for the analysis on the economic cost of operating the service. Topology measures are for the analysis on the correlation between network topology and other output measures.

The tool has three sets of algorithms. During a simulation, the attack generation algorithm sets the packet rate of attack traffic, selects attack sources networks, legitimate

source networks and victim networks. After simulated attacks are determined, the routing path construction algorithm calculates the routing path between attack source networks, legitimate source networks and victim networks. At the end, for each attack scenario and each defense, the output measure calculation algorithm calculates performance measures and cost measures for the further analyses in the next three chapters.

4.4 PARAMETERS AND ALGORITHMS

The computational tool simulates attacks on a parameterized ISP network, which are described by several sets of input parameters. Using the example network described in the last section, this section describes the input parameters, the output parameters, and the algorithms in more details.

4.4.1 INPUT PARAMETERS

4.4.1.1 Network scenario

Network scenario refers to the set of parameters that describe an ISP network from the perspective of potential victim networks when attacks are not present. These parameters include:

- Topology (G): G represents how nodes in a network are connected to one another, as in Figure 4.1. G consists of 10 nodes. The tool records the neighboring nodes for each network node. $S = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$.
- Routing algorithm (R): R refers to the algorithm for determining the routes (the routing path) of transporting the network traffic between any two nodes on a network. For example, using the shortest path algorithm, the routes between

node 1 and node 4 in the example network = {1, 2, 3, 4}. The distance between node 1 and node 4 = 3.

- Capacity of links (C): C refers to the maximum packet rate that is allowed on the link from one node to another, which may vary link by link.
- Utilization of links (U): U refers to the ratio of the actual packet rate of the link from one node to another to the capacity of that link.
- Upstream nodes of legitimate source networks (S_x): S_x refer to the set of POPs where legitimate source networks connect. A legitimate source network refers to the network that originates legitimate traffic. $S_X \subset S$. Suppose that legitimate source networks are connected to every POP. Then, S_X={1, 2, 3, 4, 5, 6, 7, 8, 9, 10} in the example network.
- Number of legitimate source networks (x): x is a counter for S_x.
- Packet rates of legitimate traffic to victim networks (X): X refers to the packet rates of the legitimate traffic that are sent from legitimate source networks to victim networks.

4.4.1.2 Attack scenario

Attack scenario refers to the set of parameters that describe the magnitude and the distribution of attacks. The parameters include:

- Upstream nodes of victim networks (V): V refers to the set of POPs where victim networks connect. A victim network refers to the network that is the target of attacks. $V \subset S$. A victim machine refers to the IP address of the computer that is targeted by the attack traffic. A non-victim machine refers to a computer that is on

the victim network but is not targeted by attack traffic. In the example network, the square with a "v" is the victim network. $V = \{4\}$.

- Number of victim networks (v): v is a counter for V. In Figure 4.1, v=1.
- Upstream nodes of attack source networks (S_A): S_A refers to the set of POPs where attack source networks connect. An attack source network refers to the network that originates DDOS attack traffic. $S_A \subset S$. In Figure 4.1, $S_A = \{1, 8, 9, 10\}$. Squares with an "a" represent the networks from where attackers launch DDOS attacks..
- Number of attack source nodes (a): a is a counter for S_A. In the example, a=4.
- Packet rates of attack traffic (A): A refers to the packet rates of attack traffic originated from each attack source network.
- Attack duration (τ): τ represents how long the attack traffic will be sent.
- Protocol type of attack traffic (P): P refers to the protocol type of attack traffic, which determines the packet size of attack traffic. For example, TCP SYN attack is about 40 bytes per packet.

4.4.1.3 Attack detection

The computational tool does not implement the attack detection algorithm in detail since the purpose of the tool is not to evaluate the effectiveness of the attack detection algorithm. The purpose of the tool is to compare the performance tradeoff of various defenses on a given network topology when the detection rate and the false positive rate of detecting attack traffic can be estimated from other studies.

The tool quantifies the attack detection algorithm using two parameters. They are:

- The filtering rate of attack traffic (f_a): f_a represents the detection rate of the attack detection algorithm at a certain filter node.
- The false positive rate (f_x): f_x represents the false positive rate of the attack detection algorithm at a certain filter node. f_x indicates how much legitimate traffic is filtered as a side effect of filtering attack traffic.

4.4.1.4 Attack response

Attack response parameters describe the deployment of responses in defenses. The parameters include:

- Filter locations (L): L refers to the set of POPs where the attack responses are enforced to react against attack traffic. $L \subset S$. Suppose that a defense to filter out DDOS traffic is enforced at POPs 3 and 7, then $L = \{3, 7\}$.
- Timing of enforcing responses (F): Timing of enforcing responses is either static or dynamic.

4.4.2 OUTPUT PARAMETERS

The computational tool calculates three sets of output measures based on various input parameters. Performance measures quantify the performance tradeoff that more legitimate traffic is dropped due to the side effect of filtering more attack traffic. Cost measures quantify both the transport distance saved by filtering attack traffic preemptively and the number of routers that will be influenced due to the deployment of filters. Topology measures quantify the characteristics of a given network topology.

4.4.2.1 *Performance measures*

The computational tool calculates three measures that can be used to quantify the performance of a defense. They are:

- The number of filter nodes that the attack traffic passes through (α),
- The number of filter nodes that the legitimate traffic passes through (β),
- The proportion of the legitimate traffic bypassing filter nodes (k),
- The link utilization of the connection to victim networks by attack traffic (U_a), and
- The ratio of legitimate traffic received by victim networks to legitimate sent (R_x).

4.4.2.2 *Cost measures*

The computational tool calculates two measures to quantify the economic costs caused by the variation of a network topology. These measures are:

- The total number of filters (H), and
- The total transport distance saved (D).

4.4.2.3 *Topology measures*

In order to study the variation of the characteristics of the topology on the deployment of defense mechanisms, the computational tool calculates several measures to distinguish one topology from others.

- Number of nodes: The number of nodes in a network quantifies the scope of a network.

- Density: Density measures the connectivity of a network, which is defined as the number of edges of a network divided by the largest possible number of edges of this network (Wasserman and Faust 1994).
- Average path length: A path refers to a sequence of nodes that network traffic is routed through based on a given routing algorithm that sends the traffic from source networks to destination networks. Average path length refers to the average number of nodes on the paths for all pairs of nodes in a network.
- Diameter: The maximum of the shortest path length between any two nodes in a network.
- Clustering coefficient: Clustering coefficient measures the cliquish of a network. Node clustering coefficient is defined as the connectivity of the neighbors of a node. Clustering coefficient is the average of node clustering coefficients in a network (Watts and Strogatz 1998).
- Degree centralization: Degree centralization measures the differences of the connectivity among nodes, which takes the average of the difference of individual node connectivity and the average node connectivity (Wasserman and Faust 1994).

4.4.3 ALGORITHMS

This section describes the algorithms used in the computational tool to construct routes, to generate simulated attacks, and to calculate output measures.

4.4.3.1 Routing path construction

Based on a given network routing algorithm, the computational tool calculates routing paths between any given two nodes in the network. Currently, the computational tool has implemented the Dijkstra's shortest path algorithm (Dijkstra 1959) to calculate routing paths, in which multiple paths are allowed. Both OSPF and RIP uses this algorithm to select routes. In addition, the computational tool has the capability to import fixed routing tables for each node if the network does not use the shortest path algorithm, such as BGP.

4.4.3.2 Attack generation

The relative locations of attack source networks and victim networks on a topology influence how many nodes that the attack traffic will go through. The computational tool provides two algorithms to generate these relative locations to evaluate the average performance tradeoff that an ISP has to make when deploying defenses. The two algorithms are source-victim enumeration and source-victim random sampling. For each run of the simulation, a combination of the upstream nodes of victim networks (V) and the upstream nodes of attack source networks (S_A) is picked from the set of the nodes in the network (S). In source-victim enumeration (Figure 4.3), all possible combinations for V and S_A are run and the average values of output measures from all combinations are calculated. In the source-victim random sampling, instead of running all possible combinations, the model randomly selects a sufficient number of combinations to approximate the average values of various output measures.

Input: Given a network topology G with a set of nodes *S*. There are *k* nodes in *S*. S={ s_1, s_2, …s_k}. The attack scenario has *n* victim networks, *m* attack source networks.
Output: *V*, S_A (to output measure calculation)
Procedures:

1. V = { s_1…s_n}, i=0.
2. S_A= { s_1…s_m}, j=0.
3. i=i+1, j=j+1, go to output measure calculation (V, S_A).
4. if j< C_m^{k-n} {

   S_A = enumerate the next combination of *m* nodes from the *k* nodes in *S*, $\vee s_j \in S_A$, $s_j \notin V$.

   go to 3 }
5. if i< C_n^k {

   V = enumerate the next combination of *n* nodes from the *k* nodes in *S*.

   go to 3 }

Figure 4.3: Source-victim enumeration

4.4.3.3 Output measures calculation

Output measure calculation can be run in two modes: computation mode and simulation mode. In the computational mode (Figure 4.4), the output measures α, β, k, D, and H are calculated based on filter locations and the relative locations of attack source networks and victim networks. These measures will be used in the next two chapters to estimate the variation of performance tradeoffs and economic costs of defense mechanisms. The packet rates of legitimate traffic and attack traffic are assumed to be constant in this mode so that these output measures can be calculated without considering the variation of attacks in time. Several simpler formulas are used in the next two chapters to quantify the variation of attacks based on these measures.

In the simulation mode (Figure 4.5), the output measures U_A and R_X are calculated over time. The purpose of this mode is to observe the variation of the impact of attack traffic on legitimate traffic over time. In addition, this mode can observe how network traffic is transported between intermediate nodes. In the case that DDOS attacks induce the saturation of backbone routers, observing the intermediate nodes will be necessary. However, this case is out of the boundary of this thesis and, therefore, will not be discussed in the next two chapters.

Inputs:

Network scenario = {*R, X*}, Attack scenario = {*G, V, S_A, A*}. Attack detection = {f_a, f_x}.

Attack response = {*C, L, F*}

Outputs: α, β, k, D, and H

Procedures:

1 Generate routing path *RP* based on *R*

2 Get S_A and *V* from attack generation.

3 Locate filter nodes based on *C, L* and *F*.

4 H=0; D=0; α=0; β=0; k=0

1. For each node s_i in *S*

   If ($s_i \in$ F) H=H+1

2. For each path *p* from S_A to *V*

   For each node s_i on *p* {

   w=1; w2=0;

   If ($s_i \in$ L) {

   $\alpha=\alpha+1$

   w1= w1(1-f_a) } //calculate the proportional of attack traffic that pass through filters

   w2=w2+w1

   D=D+(distance between S_A and *V* – w2)

3. $\alpha=\alpha$/number of paths between S_A and *V*

4. For each path *p* from S_X to *V*

   For each node s_i on *p*

   if ($s_i \in$ F) $\beta=\beta+1$

   If ($\vee s_i$ on p $\notin$ F) $k= k$ +packet rate of legitimate traffic go through *p/X*

5. $\beta=\beta$/number of paths between S_X and *V*

6. Output α, β, k, D, and H.

Figure 4.4: Output measure calculation (calculating α, β, k, D, and H)

Inputs:

Network scenario = {R, X}, Attack scenario = {G, V, S_A, A, τ, P}. Attack detection = {f_a, f_x}.

Attack response = {L, F}

Outputs: U_a, R_x

Procedures:

7. Generate routing path RP based on R
8. Get S_A and V from attack generation.
9. Locate filter nodes based on V, L and F.
10. For (t=1 to τ) {

For each node s_i in S {

Generate attack traffic based on packet rate A/a if $s_i \in S_A$.

Generate legitimate traffic based on packet rate X/x from if $s_i \in S_X$.

If ($s_i \in L$) {

if the network traffic is attack traffic,

packet rate allowed to the next node

= packet rate received from the previous node *(1-f_a)

if the network traffic is legitimate traffic,

packet rate allowed to the next node

= packet rate received from the previous node *(1- f_x)

} else packet rate allowed to the next node as it is

If packet rate allowed to the next node > the bandwidth of the link {

Drop rate = (current packet rate-bandwidth)/bandwidth

packet rate allowed to the next node= packet rate allowed to the next node*(1-drop rate)

}

transport the network traffic based on its allowed packet rate and routing path RP.

}

output

U_a= packet rate of attack traffic received at the link to the victim network/capacity of the link

output

R_x= packet rate of legitimate traffic received at the link to the victim network/capacity of the link

}

Figure 4.5: Output measure calculation (calculating U_a and R_x)

4.5 CONCLUSIONS

The computational tool described in this chapter is fundamentally interdisciplinary and draws on work from computer network, computer security and social networks analysis. This approach is necessary to adequately understand and evaluate the impact of attacks on critical infrastructure – in this case just one aspect, the Internet.

There are a large number of possible benefits of the tool. First, the tool generates results for further analyses that help network providers and subscribers to consider the benefits of providing DDOS defenses and to recognize the tradeoffs in DDOS defenses. The computational tool provides a systematic framework for thinking through the tradeoffs in defense strategies in this complex system. Thus, this work has direct bearing on security policy decisions at the router level for critical infrastructure. Secondly, this research provides a new technology to help evaluate the costs imposed by various attack scenarios and defenses since it is neither cost effective nor ethical to conduct real world experiments of DDOS attacks on a large network. Finally, the topology measures used in this research could be useful for studies of other large-scale topologies.

Because the goal of the computational tool is to facilitate the analysis on the provision of DDOS defenses, this tool has several limitations on its applications. First, the simulation analysis provides an order of magnitude cost comparison among defenses. However, the real dollar value of the cost will depend on the actual implementation of these defenses. Thus, while the model suggests relative effects in terms of cost, it will not provide real costs. Secondly, the cost measures are based on bandwidth consumption and

router overhead costs for either attack traffic or defenses. Other implementation costs will not be examined since this research focuses on the additional benefit and cost achieved from operating the defenses. Third, there is a limited amount of data available for validating models such as this. To obtain a more precise analysis, network providers can use their own data in the later analyses built upon this tool. Fourth, the tool does not assume intelligent attackers who change their attacks in response to the defenses. However, the tool does provide a foundation for creating a more complex model that will handle adaptive attacks and defenses. Finally, the computational tool developed in this research is limited to analyzing DDOS defenses. This tool would need further revision to analyze defenses for other types of Internet security incidents.

By simulating the complex attack-defense system, the computational tool described in this chapter will facilitate the analyses in the next two chapters for providing performance measures, cost measures, and topology measures. Chapter 5 will use the performance measures and topology measures to quantify the performance tradeoffs in various defenses. Chapter 6 will use the cost measures and topology measures to discuss the economic incentives of ISPs for providing the services.

Chapter 5 THE IMPACT OF TECHNOLOGY UNCERTAINTY ON THE PROVISION OF DDOS DEFENSES

During a DDOS attack[9], the online servers being targeted suffer from the loss of availability. The online servers cannot serve its legitimate clients normally because either the servers cannot handle the excess number of concurrent connections or the network capacity to the servers has been saturated. As discussed in Chapter 3, defenses have been developed to mitigate the impacts of the attacks by enabling attack detection and attack responses. By deploying the defenses at some points on the Internet infrastructure, network providers are able to detect attack traffic and filter out attack traffic preemptively before is sent to their subscribers.

As described in Chapter 2, the tools to launch DDOS attacks vary and are usually automated in order to utilize various vulnerabilities in software and network protocols, and to interfere with attack detection. A single defense deployed at one point of a network can not react against all means of DDOS attacks. To react against different attacks, network providers should have security policies that are flexible enough to tune defenses that are

[9] This chapter will be focused on the DDOS attacks that target hosts in an access network. In some cases, DDOS attacks may target network routers, or cause abnormal behavior or congestion at backbone routers. These cases are beyond the scope of this paper.

effective for their network topology and for the needs of their subscribers. Using the computational tool described in Chapter 4, this chapter analyzes the variables identified in Chapter 3 in both the DDOS defense technology and in the network topology. The goal of this chapter is to provide network providers insights into setting security policies by which to select and tune defenses. In addition, this chapter is also intended to identify critical variables that need to be considered when creating the contract between subscribers and providers for deploying defenses.

This chapter is outlined as follows. Section 5.1 defines the variables in DDOS defense technology. Section 5.2 quantifies the influence of topology on deploying DDOS defenses Section 5.3 quantifies the performance of the defenses. Section 5.4 explains the methodology used to calibrate the uncertain variables using current backbone network topologies. Section 5.5 analyzes the impact of uncertain variables in attack detection and attack responses. Section 5.6 analyzes the impact of uncertain variables in network topology. Section 5.7 provides recommendations for setting security policies to provide DDOS defenses as network services. Section 5.8 concludes this chapter.

5.1 TECHNOLOGY UNCERTAINTY IN DDOS DEFENSES

To shape the security policies for providing DDOS defenses, network providers need to understand what factors influence the performance of the defenses. In this chapter, three quantitative variables are used. They include: 1) the false positive rate in attack detection, which quantifies how well the attack detection algorithm can distinguish attack traffic from legitimate traffic, 2) the filtering rate for attack traffic in attack responses, which quantifies how much attack traffic is dropped proportional to all network traffic received, and 3) the

filter location for enforcing attack responses, which quantifies the locations that the filtering takes place. These variables are described in detail below.

1. The false positive rate in attack detection (f_x)

The false positive rate is the probability that legitimate traffic is determined by an attack detection algorithm to be attack traffic at one node of the network (such as a router). For example, when the attack detection algorithm is congestion-based, such as in Pushback (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002) or in active network response (Schnackenberg and Djahandari 2000; Sterne, Schnackenberg et al. 2001, 2002), rate limiting is applied on the flow of congested network traffic. The false positive rate represents the ratio of the legitimate traffic being regarded as a part of the congested traffic. In the defenses based on anomaly detection, the false positive rate represents the probability that legitimate packets/connections share some characteristics with attack packets, such as asymmetric TCP SNY packets.

2. The filtering rate for attack traffic in attack responses (f_a)

The filtering rate for attack traffic is the ratio of attack traffic being reduced by an attack response at one node of a network. For example, the filtering rate is 0.9 if the packet rate of an attack is reduced from 500 packets per second to 50 packets per second. The filtering rate here is for the convenience of analyzing the uncertainty in attack responses. It is not explicitly defined in DDOS attack responses and it may vary with the packet rate of attack traffic. For example, rate limiting is only triggered when the burst packet rate or average packet rate exceeds an upper bound.

3. The filter location for enforcing attack responses

The filter location refers to the links of a network that enforce attack responses during an attack. Attack responses are enforced either statically before attacks (called "static filters" in this chapter) or dynamically when attacks have been detected (called "dynamic filters" in this chapter). Static filters are enforced on one node to monitor outbound links in TCP anomaly detection (Schuba, Krsul et al. 1997) or in Ingress filtering (Ferguson and Senie 1998). Static filters are also enforced on the vertex cover set of a network in route-based filtering (Park and Lee 2001b). Dynamic filters are enforced in Pushback or in active network responses, in which attack responses are pushed hop by hop toward the attack sources. The filter locations being analyzed in this chapter includes static filters at the upstream POP of the victim's network (denoted as "victim"), static filters at the upstream of attack sources (denoted as "attack sources"), static filters at minimum vertex cover set of a network (denoted as "vc"), and dynamic filters at various number of hops away from the victim's network.

5.2 TOPOLOGY UNCERTAINTY IN DEPLOYING DDOS DEFENSES

In additional to uncertain variables in defenses, the topology of a specific ISP's network poses an uncertain impact on the performance of the defenses as well. Consider a network $G=(S, E)$ has S nodes and E edges. $\forall s \in S$ represents a point of presence (POP) in a backbone network to which access networks connect. The filter nodes $F \subseteq V$ denotes a set of nodes where a filter policy is enforced. Once network traffic is transported through a filter node, it is examined and an attack response will be triggered when attack characteristics in the traffic is detected. The attack sources $S_A \subseteq S$ represent a set of nodes

where the source networks of attacks connect. The legitimate sources $S_X \subseteq S$ represent a set of nodes where the access networks of legitimate clients connect. The victims $V \subseteq S$ represent a set of nodes where the victim networks connect. Routing algorithm R refers to the algorithm for selecting routes of network traffic from sources to destination.

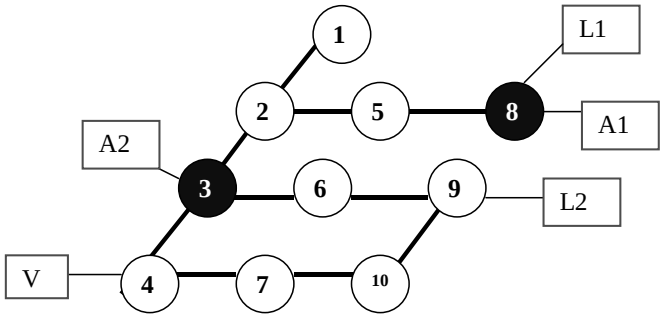

Figure 5.1: An example network

Figure 5.1 shows an example. Circles represent the routers of a network. Squares are networks of subscribers, in which “A1” and ”A2” denote attack sources, “V1” and “V2” denote victims, and “L1” denotes legitimate clients. In this example, S = {1, 2, 3, 4, 5, 6, 7, 8, 9, 10}. S_A = {3,8}. S_X = {8,9}. V = {4}. Filters are triggered at the router upstream of the attack sources, F={3, 8}. If R adopts the shortest path first algorithm, there are two routes from L2 to V: 9 to 6 to 3 to 4 and 9 to 10 to 7 to 4.

Both the attack traffic and the legitimate traffic may pass more than one filter node when they are transported to the victims. If a filter is enforced at a location that cut through the routing path of the attack traffic but not the routing path of the legitimate traffic, the legitimate traffic can bypass the filter so that the attack response does not has a significant impact on the legitimate traffic. Three parameters are considered when determining the number of filter nodes to use:

- α denotes the number of filter nodes enforced on the path from attack sources to victims.

α is an non-negative integer. For dynamic filters, *α is less than or equal to the number of attack sources* because filters have to be triggered at choke points, the locations that can cut off the attack traffic to the victims. For static filters, α is less than or equal to the diameter of the network because static filters are deployed before attacks occur, when attack sources are uncertain.

- β denotes the number of filters enforced on the path from the legitimate clients to victims.

β is an non-negative integer. For dynamic filters, *β is less than or equal to α*. When the legitimate traffic originates from the same node as attack traffic, *β is equal to α*. However, when the legitimate traffic originates from a different node from the attack traffic, β is smaller or equal to α because the choke points of the attack traffic are not necessarily the same as the choke points of the legitimate traffic. For static filters, β is equal to α. Since static filters are deployed before attacks occur, the legitimate traffic would encounter the same number of filters as the attack traffic.

- *k* denotes the proportion of legitimate traffic that is able to bypass the filters.

k is a floating point number. For dynamic filters, *k is non-negative and less than or equal to 1*. For a given legitimate client, the proportion of the legitimate traffic to the victims that bypasses filters can be estimated as the ratio of the total number of routes from the legitimate client to the victims that have been deployed filters and the total number of

routes from the legitimate client to the victims. For static filters, *k is equal to 0* when static filters are deployed at the vertex cover set, because all edges are covered by filters. *k is equal to the proportion of legitimate traffic that do not originate from the same nodes as the attack sources* when static filters are deployed at attack sources.

For example, in Figure 5.1, when the shortest path routing algorithm is used, from A2 to V, the attack traffic will pass through one filter node (node 3), so α is equal to 1. Similarly, α is equal to 2 between A1 and V. However, the deployment of filters at node 3 and node 8 has a side effect on the legitimate traffic from L1 and L2. From L2 to V, the network traffic can route through either node 9, 10, 3, and 4, or node 9, 10, 7, and 4. If network traffic is evenly distributed (based on OSPF) between two routes, a half of the network traffic from L2 will pass through one filter (node 3) and another half of network traffic will not pass through any filter. In this case, β is equal to1 and k is equal to 0.5. Similarly, from L1 and V, β is equal to 2 and k is equal to 0.

This example demonstrates that α depends how the attack traffic is routed through the network and where the filters are enforced. That is, α depends on G, S_A, V, and R. Similarly, both β and k are determined by G, L, V, R and S_X. In Section 5.4, the three parameters (α, β and k) will be estimated using backbone network topologies.

5.3 QUANTIFICATIONS FOR PERFORMANCE MEASURES

Using the parameters quantified in the previous two sub-sections, this section defines two measures for the performance impact on the victims of DDOS attacks given a certain defense and a certain attack scenario. The two measures are:

1) Attack traffic utilization (U_a)

U_a quantifies the proportion of the network capacity utilized by the attack traffic. The network capacity refers to the capacity of the network connection between the victim network and its upstream backbone POP. U_a is a ratio of the attack traffic packet rate received at the victim network to its network capacity C. A denotes the total packet rate of the attack traffic arriving at the upstream POP of the victim network, which is the aggregate of the traffic sent by attack sources distributed on the Internet. U_a quantifies the performance impact imposed by the attack traffic, which is formatted as:

$$U_a = \frac{A(1-f_a)^{\alpha}}{C}, A \geq 0, f_a = [0,1], U_a = [0,1] \quad (1).$$

2) Legitimate traffic arrival rate (R_x)

R_x represents the ratio of the legitimate traffic received by the victim network to the legitimate traffic actually sent to the victims. X denotes the packet rate of the legitimate traffic arriving at the upstream POP of the victim network, which is the aggregate of the traffic sent by legitimate clients distributed on the Internet. R_x quantifies the performance benefit from implementing defenses, which is represented as:

$$R_x = \frac{X[k+(1-k)(1-f_x)^{\beta}]}{X} = [k+(1-k)(1-f_x)^{\beta}] \quad (2).$$
$$f_x = [0,1], R_x = [0,1]$$

When the network capacity is saturated, both the attack traffic and the legitimate traffic will be dropped at the same rate. R_x is at its highest value when the total network traffic reaches its capacity and it decreases as U_a increases when the capacity is saturated

(the proof is in Appendix 5.C). This situation will not be further discussed in this chapter since the defenses have failed to prevent the network connection from being saturated.

Changes in U_A and R_X are determined by parameters f_a, f_x, α, β, and k. The attack scenario is represented by parameters C, A and X. Appendix 5.A summarizes the meaning of the parameters and the marginal changes of both U_A and R_X caused by these parameters. The numerical analyses will be described in Section 5.5, 5.6 and 5.7.

5.4 CALIBRATION OF PARAMETERS

This section describes the method and results from calibrating the three parameters describing network topologies: α, β, and k. Thirty-six backbone network topologies were analyzed. AT&T network was chosen for use in the later analyses. The remaining thirty-five backbone networks[10] are listed to illustrate the variation of other networks from the AT&T network. Table 5.1 lists the topology measures of these networks. Comparing to other networks, the AT&T network is a sparse network that is loosely connected and less centralized. The average path length of this network is close to the average of all other networks. Appendix 5.C is a map of the AT&T network topology. Appendix 5.D shows the correlation matrix of topology measures for the thirty-six networks.

[10] The four backbone network topology maps are from Board Watch magazine, Spring, 2001.

Topology measures	AT &T network topology	All 36 network topologies			
		Mean	Standard Deviation	Maximum	Minimum
Number of nodes	61	29	18	70	7
Density[11]	0.04	0.17	0.14	0.57	0.04
Average shortest path length	3.6	3.4	1.6	7.2	1.4
Diameter	7	7	4	17	2
Clustering coefficient[12]	0.06	0.20	0.20	0.79	0
Degree centralization[13]	3.E-03	0.02	0.02	0.10	5.E-04
Number of nodes in VC set	19	14	9	32	1

Table 5.1: Descriptive statistics for the topology measures of the AT&T network and 36 network topologies

5.4.1 ASSUMPTIONS

The settings of other independent variables and corresponding assumptions are described in this section. These settings and assumptions are used throughout this chapter.

1) Since the analysis is focused on the attacks originated from other POPs, the attack source nodes and the victim nodes are assumed to be distinctive. The maximum, average, and minimum values of α, β, and κ are calculated across all combinations of attack source nodes and victim nodes.

2) Both static filters and dynamic filters are evaluated. For static filters, one attack source node and one victim node are picked from a given network topology. For dynamic filters, two attack scenarios are analyzed.

[11] Density measures the connectivity of a network, which is defined as the number of edges of a network divided by the largest possible number of edges of this network (Wasserman and Faust 1994).

[12] Clustering coefficient measures the cliquishness of a network. Node clustering coefficient is defined as the connectivity of the neighbors of a node. Clustering coefficient is the average of node clustering coefficients in a network (Watts and Strogatz 1998).

[13] Degree centralization measures the differences of the connectivity among nodes, which takes the average of the difference of individual node connectivity and the average node connectivity (Wasserman and Faust 1994).

3) Two attack scenarios are analyzed: single source attacks (attacks originated from a single POP) and distributed source attacks (attacks originated from multiple POPs). For single source attacks, the parameters are calculated based on the combination of any two nodes. For distributed source attacks, 10% of the POPs in a network are selected as source nodes. For all cases, legitimate clients are assumed to be uniformly distributed on the network. That is, S_x is uniformly selected from S.

4) Both the packet rate of attack traffic A and the packet rate of legitimate traffic X are normalized by the capacity of the link to victim's network. A is set to 10 and X is set to 1 if they are not specified. This setting allows the analyses to estimate the saturation point of the capacity.

5) The performance measures are calculated at a given point in time but not at a longitude scale.

6) Dijkstra's shortest path algorithm (Dijkstra 1959) is used to find routing paths between two nodes. If multiple routing paths are found, the network traffic is distributed evenly among the multiple paths. This setting is same as most intra-domain routing protocols, such as OSPF (Huitema 2000).

5.4.2 ALGORITHMS

In static filters at the minimum vertex covering set, k is zero and α *is equal to* β for all cases since the filters have covered all edges. In this case, the legitimate traffic from all POPs has to pass through at least one filter node. Figure 5.2 is the algorithm of calculating α (and β) for static filters in minimum vertex covering set. For dynamic filters, the values

of k, α and β depend on the enforcement locations of filters, which are push one hop away each time after the values of k, α and β are calculated. Figure 5.3 is the algorithm of calculating k, α and β for dynamic filters.

In Section 5.5, the values of k, α and β calibrated from the various network topologies will be used to analyze the provision of the defenses. These analyses are based on the best case, the average case and the worst case from various combinations of the attack sources and victims.

Given a network topology G(V,E).

1 Let the set of filter nodes F = the minimum vertex covering set of G.

2 Select one victim node $i \in V$ as the upstream POPs that the victim networks connect to.

3 Select a set of source nodes $j \in V$ and $i \neq j$ as the upstream POPs that the source networks connect to.

4 Generate the routing paths P between any (i,j) based on the routing algorithm R.

5 Let α (and β) = the number of filter nodes on the routing paths P.

6 Repeat steps 2-4 until all combinations of $(i,j) \in V$ and $i \neq j$ have been analyzed.

7 Output maximum, average and minimum α (and β).

Figure 5.2: The algorithm of calculating α (and β) for static filters on the minimum vertex covering set

Given a network topology G(V,E).

1 Select a set of victim nodes $D \subset V$ as the upstream POPs to which the victim networks connect.

2 Select a set of source nodes $S_A \subset V$ and $D \cap S_A = \phi$ as the upstream POPs to which the source networks connect.

3 Initialize the distance of the filter nodes, h =1.

4 Let the set of edges that the filter policies are enforced, F = (h,h-1).

5 Generate the set of edges AP_j in the routing paths between an attack source node j and a victim node i for all (i,j) where $i \in D$ and $j \in S_A$ based on the routing algorithm R.

6 Generate the set of edges XP_z in the routing paths between a legitimate source node z and a victim node i for all (i, z) where $i \in D$ and $z \in V$ based on the routing algorithm R.

7 Let k = (the number of nodes in V that $XP_z \cap F = \phi$)/(the total number of nodes in V)

8 Let α = the number of distinct edges in $AP_j \cap F$ and let β= the number of distinct edges in $XP_j \cap F$.

9 Increment h by 1 and repeat steps 4-9 until h = the maximum distance between (i, j).

10 Repeat steps 1-10 until all combinations of $(i,j) \in V$ have been analyzed.

11 Output maximum, average and minimum k, α and β.

Figure 5.3: The algorithm of calculating k, α, and β for dynamic filters

5.4.3 ESTIMATION OF PARAMETERS

Table 5.2 summarizes the parameters used to quantify static filters at minimum vertex cover set for the AT&T network topology and for all thirty-six backbone networks.

		Average case	Best case	Worst case
Model parameter α (=β)	AT&T	2.9	6	1
	All 36 networks	2.7	11	1

Table 5.2: The parameters for static filter enforcement

For dynamic filters, the filter location is measured in terms of the number of hops away from the upstream POP of the victim network. The filter location is pushed one more hop away from the upstream POP of the victim network each time after the parameters are calculated. For example, in hop 1, the filter is located at the nodes that are one hop away from the upstream POP of the victim network. Table 5.3 and Table 5.4 list the values of parameters for single source attacks and distributed source attacks, respectively. In the next section, the results for these three parameters will be used to analyze the influence of variables for network topology on the performance of the defenses.

		Filter location (L)							
	Parameters	Victim	Hop 1	Hop 2	Hop 3	Hop 4	Hop 5	Hop 6	Attack source
Average case	k	0	0.02	0.04	0.31	0.54	0.82	0.96	0.98
	α	1	1	1	1	1	1	1	1
	β	1	1	1	0.96	0.96	0.97	1	1
Best case	k	0	0.02	0.05	0.38	0.59	0.84	0.97	0.98
	α	1	1	1	1	1	1	1	1
	β	1	1	1	0.5	0.5	0.6	1	1
Worst case	k	0	0.02	0.03	0.25	0.49	0.8	0.95	0.98
	α	1	1	1	1	1	1	1	1
	β	1	1	1	1	1	1	1	1

Table 5.3: The parameters used in analyses for single source attacks (calculated based on the AT&T network topology)

		Filter location (L)							
	Parameters	Victim	1	2	3	4	5	6	Attack source
Average case	k	0	0.02	0.03	0.18	0.32	0.51	0.59	0.6
	α	1	1	1.1	1.2	1.2	1.2	1	1
	β	1	1	1.1	1.1	1.2	1.3	1.3	1.3
Best case	k	0	0.02	0.03	0.25	0.49	0.8	0.9	0.9
	α	1	1	2	3	4	5	5	5
	β	1	1	1	1	0.5	0.5	0.2	0.14
Worst case	k	0	0.02	0.02	0.02	0.02	0.02	0.02	0.02
	α	1	1	1	1	1	1	1	1
	β	1	1	2	3	4	5	5	5

Table 5.4: The parameters used in analyses for distributed source attacks (calculated based on the AT&T network topology)

5.5 THE IMPACT OF UNCERTAINTY FROM ATTACK DETECTION AND ATTACK RESPONSES

Using the parameters estimated from Section 5.4, this section analyzes the changes of U_a and Rx by varying f_a, f_x and the filter location. The results suggest several principles for the design of the DDOS defenses.

1) A filter should be able to increase the filtering rate of the attack traffic flexibly when the attack traffic increases.

As in Figure 5.4, to maintain the attack traffic utilization lower than 0.1, f_a should be at least 0.9 when the attack traffic is as large as the link capacity (A=1) while f_a should be at least 0.99 when the attack traffic is 10 times of the link capacity (A=10). This result suggests that, to reduce the attack traffic received by victims, f_a should increase when the attack traffic increases no matter where the filter location is. In the other words, how much attack traffic the victims will receive depends on f_a.

2) For dynamic filter, if the filter is closer to the attack source, a high false positive rate is acceptable; if the filter is closer to the victim, a low false positive rate is needed. For static filter at the minimum vertex cover set, a low false positive rate is needed.

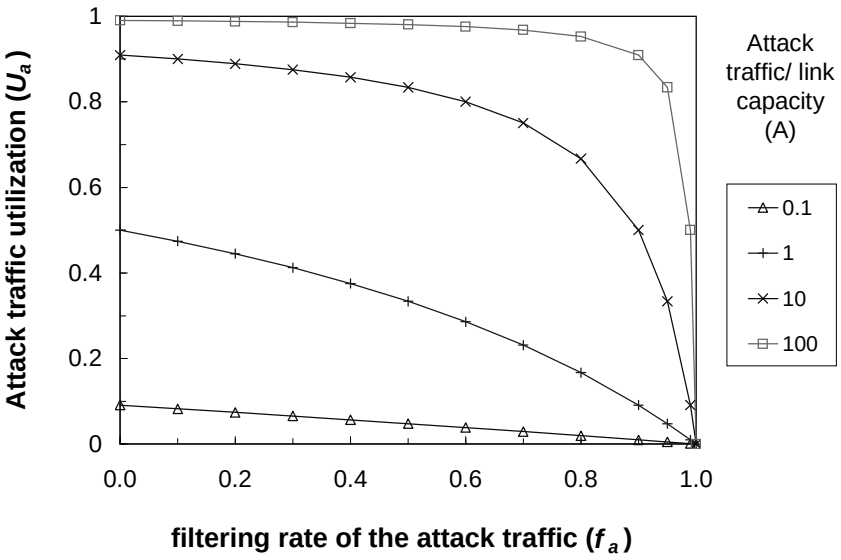

Figure 5.4: Attack traffic utilization (filter location at attack upstream, f_x=0.1)

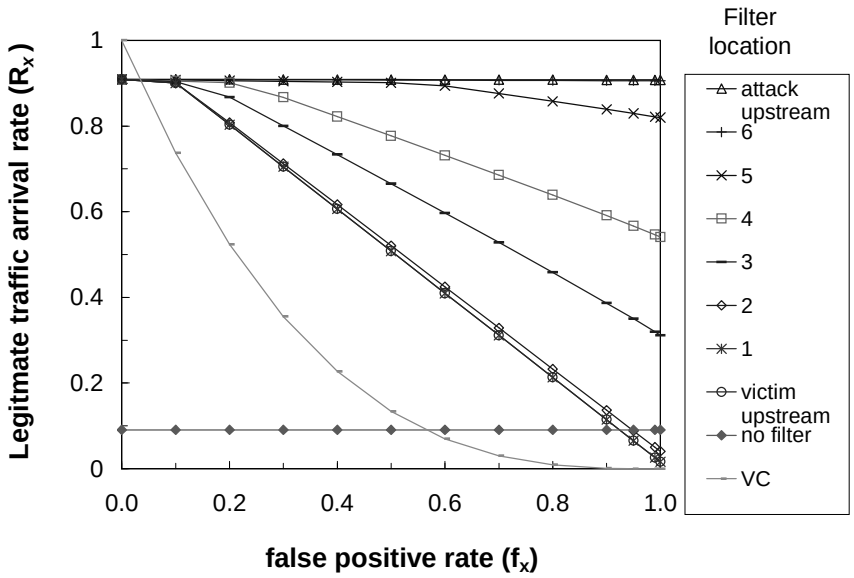

Figure 5.5: Legitimate traffic arrival rate (A=10, f_a=0.99)

As in Figure 5.5, the influence of the filter location increases when the false positive rate of the attack detection increases. For example, when fx=0.9, Rx changes significant when the filter location is closer the attack source. This case explains the sensitive variables in congestion-based attack detection, such as suggested in aggregate-based congestion control (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002) and in many other studies (Sterne, Schnackenberg et al. 2001; Huang and Pullen 2001; Xiong, Liu et al. 2001). For this type of attack responses, an ISP should emphasize both pushing filter locations closer to attack sources and increasing filtering rate but not on reducing false positive rates, when deciding security policies for defense mechanisms. In addition, based on this result, the criteria for any defense that is deployed on attack sources should have high filtering rate of attack traffic. In contrast, when fx=0.1, the Rx does not change significantly when the filter location changes. This case explains that false positive of attack response is the most sensitive variable in anomaly-based attack detection, such as TCP SYN anomaly detection (Schuba, Krsul et al. 1997) and MULTOPS (Gil and Poletto 2001), and in attack detection using MIB variable correlation (Cabrera, Lewis et al. 2001). In particular, when static filters are deployed at the vertex cover set such as suggested in route-based filtering (Park and Lee 2001b), a lower false positive improves Rx more significantly for static filters than for dynamic filters. To decide security policies, ISP should emphasize on reducing false positive of attack detection.

5.6 THE IMPACT OF UNCERTAINTY FROM NETWORK TOPOLOGY

5.6.1 STATIC FILTERS AT MINIMUM VERTEX COVER SET

For static filters, the network topology determines the number of filters that the legitimate traffic would pass through, and therefore determines the performance measures. This impact is explained below:

1) For a given network topology, the relative distance between attack sources and victims determines the performance measures. Figure 5.6 shows the best case, the average case, and the worst case for Rx when f_x varies. The best case occurs when legitimate clients are much closer to victims than attack sources, and the worst case occurs when attack sources are much closer than legitimate clients. An implication of this result is that subscribers should provide online services that are closer to where their clients are located. This strategy shortens the distance between legitimate clients and the online servers (potential victims) when DDOS defenses are implemented.

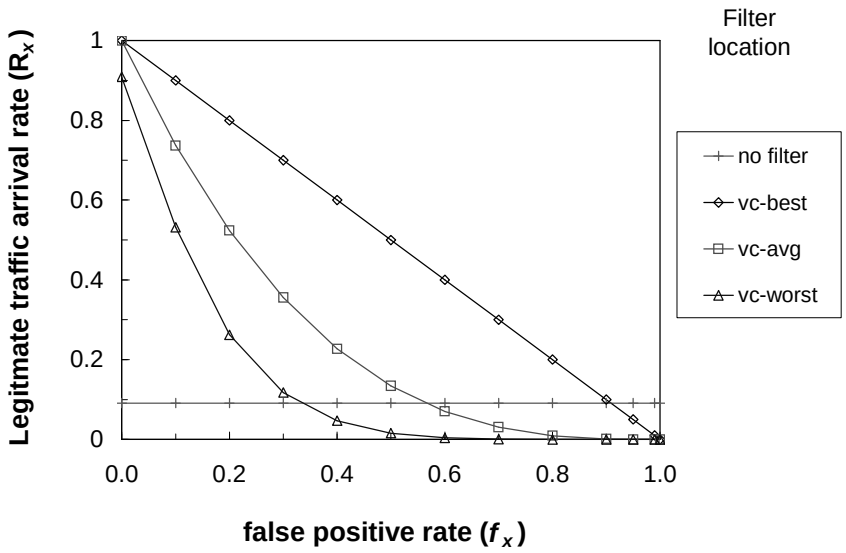

Figure 5.6: Legitimate traffic arrival rate

2) The variation of *Rx* among these three cases is lower when f_x approaches either 1 or 0. When f_x approaches 0, the filters have no impact on legitimate traffic so that the locations of legitimate clients do not matter. When f_x approaches 1, the first filter that legitimate traffic encounters cuts off all legitimate traffic so that the locations of legitimate clients do not matter either.

3) Among various network topologies, when the filtering rate is large, static filters at minimum vertex cover set are better for a dense network with shorter average path length because the first filter that attack traffic encounters cuts off most of attack traffic. When the filtering rate is small, static filters are better for a loosely connected network with longer average path length because attack traffic would encounter more filters.

	Number of filters	Attack responses (f_x=0.01)	
Topology measures	α (=β)	f_a=0.9	f_a=0.99
Number of nodes	0.8	0.4	-0.5
Density	-0.7	-0.5	0.5
Average path length	1.0	0.5	-0.5
Diameter	1.0	0.5	-0.4
Clustering coefficient	-0.4	-0.2	0.4
Degree centralization	-0.7	-0.7	0.3
Number of nodes in VC set	0.9	0.6	-0.3

Table 5.5: Correlation of topology measures of all 36 networks with model parameters and *Rx* for static filters at vertex cover set, average case

Table 5.5 shows the correlation of various network measures with *Rx* at two different filtering rates, in which false positive is fixed to 0.01 (as discussed in Section 5.5). When the filtering rate is relatively large, such as f_a=0.99, *Rx* is negatively correlated to number of nodes in a network, average path length, diameter, and number of nodes in VC set, and positively correlated to density, clustering coefficient and degree centralization. Since *Rx* is lower when the legitimate traffic passes through more filters, network measures and the number of filters are negatively correlated. With a defense set at such settings (f_a=0.99, f_x=0.01), an ISP should deploy filters on one single node, such as the upstream POP of the victim's network, but not on the vertex cover of the network.

When the filtering rate is smaller, such as f_a=0.9, the correlations exhibit the opposite relationships. This result implies that deploying filters at VC has a better performance for a sparse network or a network with a long average path length. The reason is that, in a network with longer paths, more filters on the paths to victims cut off more attack traffic, which compensates for the low filtering rate at a given node. Since the false positive rate is much lower than the filtering rate, the legitimate traffic is not cut off as much as attack traffic.

5.6.2 DYNAMIC FILTERS

For dynamic filters, the impact of network topology varies by the filter location. Figure 5.7 shows *Rx* for the best case, the average case, and the worst at various filter locations. Figure 5.7 is estimated under one attack source and Figure 5.8 is the same estimation under distributed source attacks. When attacks are originated from a single source, the variation of *Rx* for the three cases is negligible when filters are set at the upstream POP of the victim network or filters are close to attack sources. Pushing filters to the upstream POP of attack sources does not degrade performance.

Surprisingly, when attacks are originated from distributed sources, the variation of *Rx* is larger when filters are closer to attack sources. A quick review on the parameters for the estimation can explain this anomaly. As in Table 5.4, the variation of two parameters, κ and β, causes the variation of *Rx* in distributed source attacks. When attacks sources are less distributed, fewer filters are triggered since several attack sources can be cut off from the same node. In this case, the legitimate traffic will pass through fewer filters as well. In contrast, when attack sources are uniformly distributed on the network, more filters are needed to cut them off. In this case, the legitimate traffic will pass more filters and pushing filters to the upstream POP of attack sources is not effective. When attacks are uniformly distributed on multiple sources, detecting attack sources at downstream networks are difficult due to two reasons: 1) the source addresses of the attack traffic can be spoofed to disguise its sources, and 2) significant effort is needed to inform the multiple sources. A better solution is to deploy static filters that police the outbound traffic of a local network to its upstream ISP, such as using source filtering technologies like Ingress Filtering (Ferguson and Senie 1998) and D-WARD (Mirkovic, Prier et al. 2002).

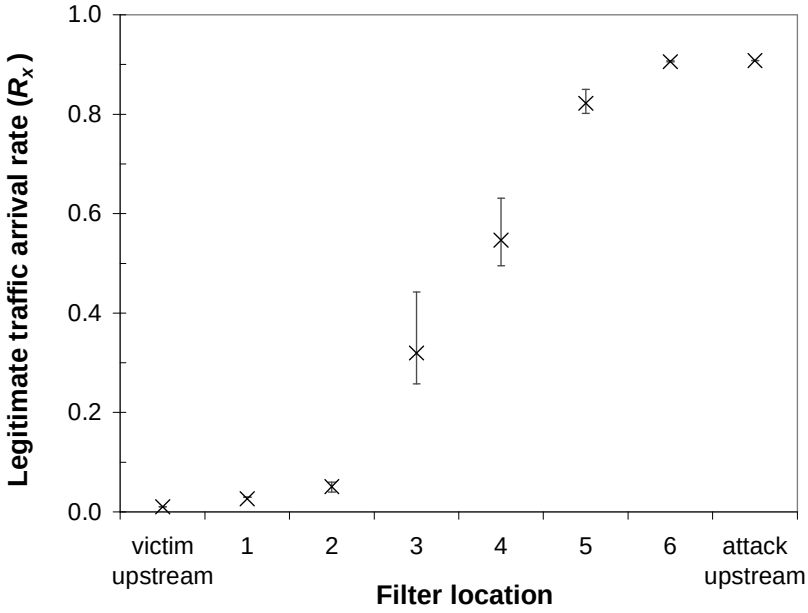

Figure 5.7: R_x for dynamic filters during single source attacks (f_a=0.99, f_x=0.99)

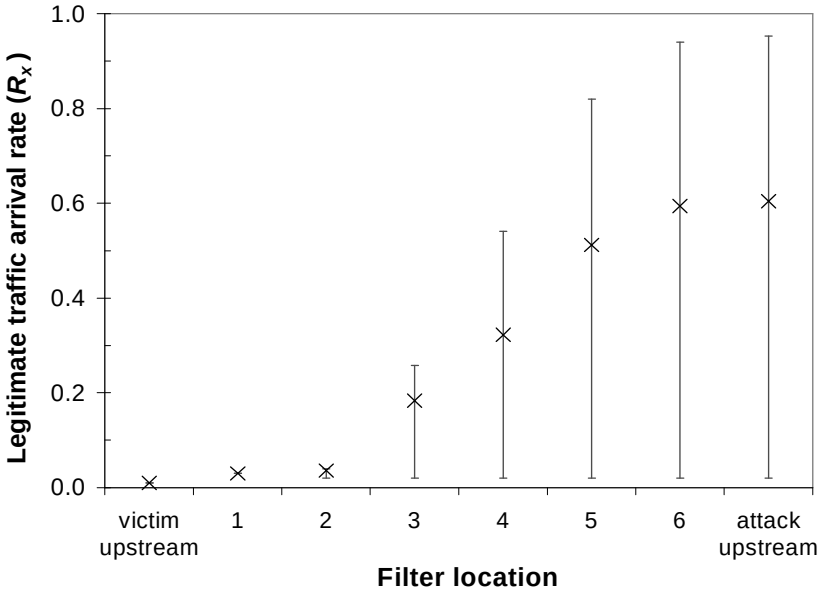

Figure 5.8: R_x for dynamic filters during distributed source attacks (f_a, f_x =0.99)

5.7 SERVICE PROVISION

The packet rates of DDOS attack traffic vary from 250 packets per second to 679,000 packets per second based on an analysis from backscatters (Moore, Voelker et al. 2001). The online servers of different subscribers are usually able to tolerant attacks to a certain extent. To provide DDOS defenses as network services, network providers should design different services that adjust the settings of defenses to meet the needs of their subscribers and to respond against the various packet rates of attack traffic. For adjusting the settings of the defenses this section discusses three possible ways to provide services: "maximum availability", "attack threshold", and "minimum attacks". The three services are discussed as follows.

1. Maximum availability (Maximizing *Rx*).

   This service proposes maximizing the legitimate traffic that a victim network can receive during the attack no matter how much attack traffic arrives at the target servers. In practice, this service adjusts attack responses so that the drop rate of network traffic to the network of a certain subscriber is minimized during attacks. A network subscriber would choose this service when the link capacity is not saturated and the network services provided by the non-target servers are more important than the target servers. For example, while a departmental web server in a campus network becomes an attack target, other services provided by the non-target servers, such as email or other web communications, may be more important to the entire campus community than a single departmental web server.

2. Attack threshold (Keeping U_a under a threshold while maximizing *Rx*).

This service proposes maximizing the legitimate traffic received by a victim network but setting a threshold to limit the attack traffic. In practice, the defense should be flexible enough to adjust attack responses so that the link utilization to the victim's network is under a threshold value defined by subscribers. A network subscriber can define a threshold that the critical online servers can tolerant so that the legitimate traffic still can be transported to both the target servers and the non-target servers during attacks. Internet services are usually designed to handling a certain amount of concurrent requests. For example, high performance web servers (Pai, Druschel et al 1999) and scalable Internet services (Banga, Mogul et al. 1999; Welsh, Culler et al. 2001) have been developed to handle a large amount of concurrent client requests. Tools have been developed to evaluate the capacity of a web server (Banga and Druschel 1997). To maintain the online services available during an attack, it is important to allow as much of the legitimate traffic as possible to pass so the server can process client requests.

3. Minimum attacks (Minimizing U_a).

This service proposes minimizing the attack traffic that a victim network can receive during the attack no matter how much legitimate traffic is sacrificed. In practice, this service means cutting off all suspicious network traffic if it is detected as attack traffic. This service is preferable when the target servers provide critical services to support the victim network internally and the downtime of the target hosts could jeopardize the operation of the victim network.

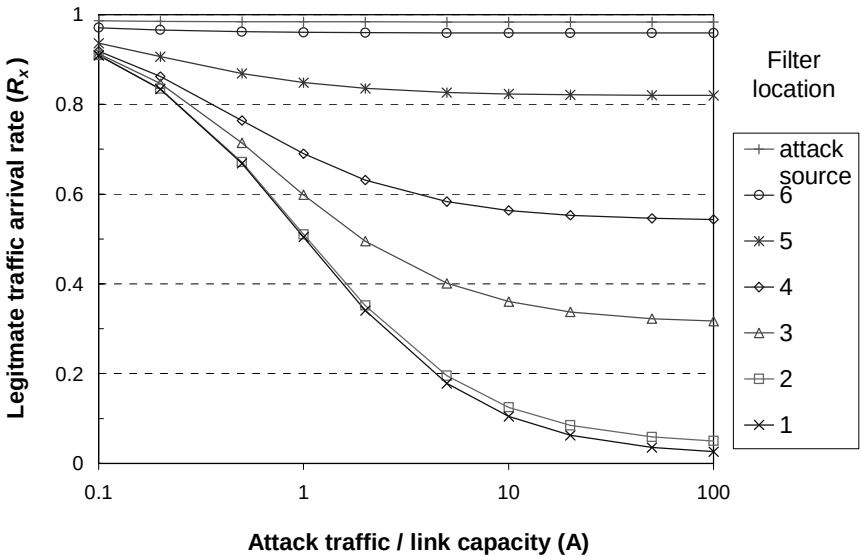

Figure 5.9: Legitimate traffic arrival rate for “maximum availability”

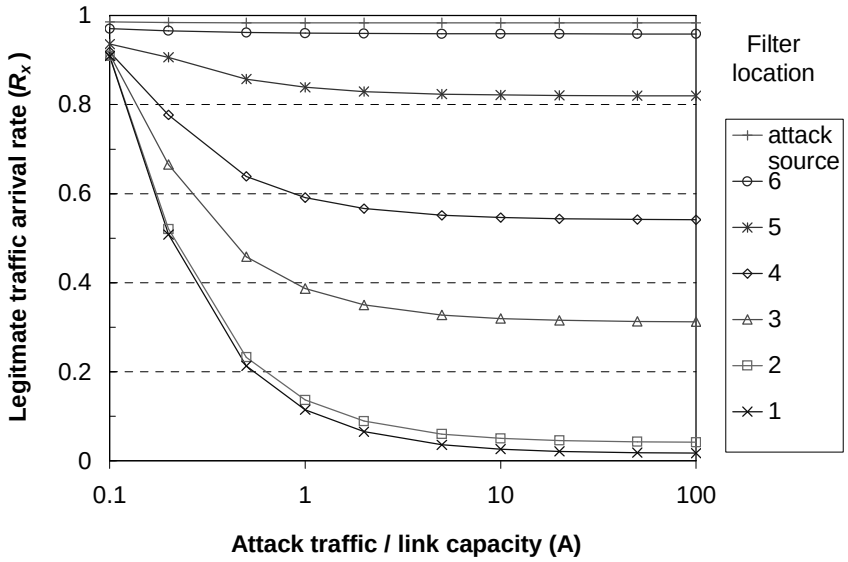

Figure 5.10: Legitimate traffic arrival rate for “attack threshold”

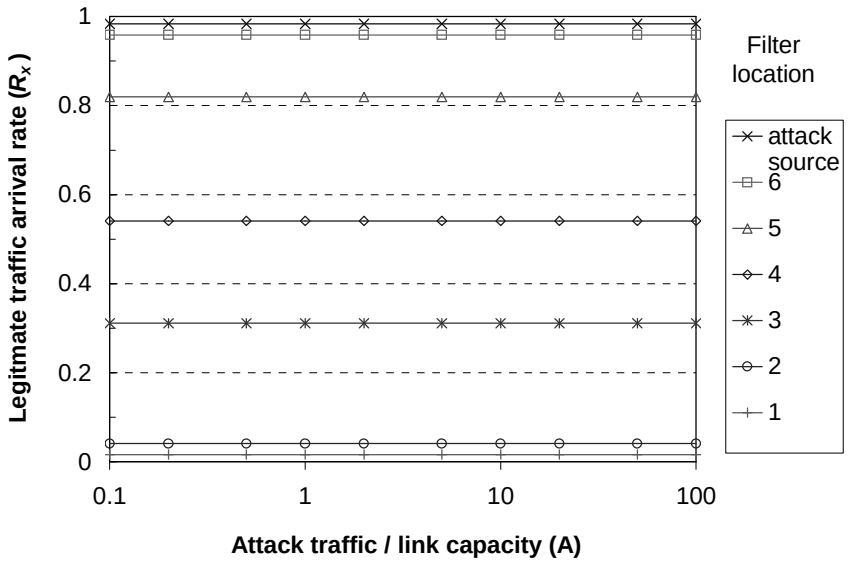

Figure 5.11: Legitimate traffic arrival rate for “minimum attacks”

Using dynamic filters as an example, Figure 9, Figure 10 and Figure 11 show *Rx* when applying the three services, respectively. As shown in all three figures, *Rx* is constant for filter locations that are close to attack sources. This result means that setting attack responses at these locations will resist attacks that generate a large packet rate or with a variable packet rate. However, when network providers are not able to push filters close to attack sources (for example, attacks that originate from another ISP’s network), *Rx* is higher for “maximum availability” than for “attack threshold” and for “minimum attacks”. Under “maximum availability”, the service provision seeks to maximize the amount of legitimate traffic that reaches the victim network. As a consequence, additional attack traffic must be allowed through when the false positive rate is nonzero. Therefore, subscribers have to setup other security guards to ensure the servers can resist the additional payload of attacks. In contrast, “attack threshold” allows less legitimate traffic to

pass with a controllable packet rate of attack traffic and “minimum attacks” allows no attacks to pass. As shown in Figure 11, “minimum attacks” utilizes only the filter location to adjust how much legitimate traffic that victims will receive during attacks and its performance does not change with the packet rate of attack traffic.

5.8 CONCLUSIONS

To ensure the availability of online services during attacks, the provision of DDOS defenses should be designed in a way that clarifies the uncertain variables from the technology and the network topology of providers. This chapter analyzes the impact of these uncertain variables on the performance of defense mechanisms. The results provide recommendations for network providers and subscribers.

For network providers, four recommendations are provided based on the results. 1) The filter location and the filtering rate of attack traffic are the most sensitive variables for defenses in which attack detection is congestion-based and attack responses are dynamically enforced. When providing such defenses, the providers should design services that focus on adjusting the filtering rate of the attack traffic to meet the needs of different subscribers. 2) The false positive rate of attack detection is the most sensitive variable for defenses in which attack detection is anomaly-based and attack responses are statically enforced. To define the service contract or to select defenses, network providers should emphasize the ones that can reduce the false positive rate of attack detection. 3) If the provider has a sparse network or a network with a long average path length, deploying static filters at the minimum vertex cover set is a good strategy since longer paths forces attack traffic through more filters before it arrives at the victim’s network. The only

exception is when the filtering rate of attack traffic is close to 1 and the false positive rate is low. Then, the ISP should deploy filters on one single node, such as the upstream POP of the victim's network, not on the vertex cover of the network. 4) Cooperation among multiple network providers is needed for attacks that originate from distributed sources. When most attacks in the network are originated from a single source, it is a good strategy for network providers to use dynamic filters that push filters to the upstream POP of attack sources. However, the same approach is not effective when most attacks originate from distributed sources. The network provider needs the cooperation of other providers to deploy static filters that police the outbound traffic of a local network to its upstream ISPto more effectively filter attack traffic.

For network subscribers, three recommendations are provided. 1) Since none of the current defenses can filter out attack traffic without posing an impact on the legitimate traffic, subscribers need to determine the attack tolerance of its online servers in order to obtain the availability for its servers during attacks. In particular, when the subscriber has a capacity that is larger than the packet rate of the attack traffic, maintaining a certain tolerance to attacks can avoid any additional dropping of the legitimate traffic. In addition, network providers would be able to tune the defenses based on the availability of the servers to meet the needs of the subscriber's online services. 2) Providing online services that are closer to where their clients are located is a good strategy to maintain the availability of the online service to legitimate clients when DDOS defenses are implemented. 3) Implementing defense mechanisms on the outbound traffic of an access network will ensure the accessibility of the legitimate clients to other online services, which is better than having the victim network filter out legitimate traffic.

The goal of this chapter is to provide a quantitative method to consider the performance impact of uncertain variables for deploying defenses. While the implementation of defense mechanisms should be a community effort, the decisions of upstream network providers would influence the impact of attacks on downstream networks. In the next chapter, the economic incentives of providing defenses will be analyzed. The influence of the compliance of multiple providers will be discussed in Chapter 7.

Appendix 5.A: The marginal change of the performance measures

Parameter (P)	Meaning	Marginal change of the attack traffic utilization, $(\frac{\partial U_A}{\partial P})$	Marginal change of the legitimate traffic arrival ratio, $(\frac{\partial U_X}{\partial P})$
A	Attack traffic packet rate	$\frac{1}{C}(1-f_a)^{\alpha} \geq 0$	0
X	Legitimate traffic packet rate	0	$\frac{1}{C}[k+(1-k)(1-f_x)^{\beta}] \geq 0$
f_a	The filtering rate of attack traffic	$-\frac{A}{C}\alpha(1-f_a)^{\alpha-1} \leq 0$	0
f_x	The false-positive rate of attack detection	0	$-\frac{X}{C}\beta(1-k)(1-f_x)^{\beta-1} \leq 0$
α	The number of filter nodes on the routing path from attack source nodes to victim nodes	$\frac{A}{C}e^{\alpha \ln(1-f_a)} \geq 0$	0
β	The number of filter nodes on the routing path from legitimate source nodes to victim nodes	0	$\frac{X}{C}(1-k)e^{\beta \ln(1-f_x)} \geq 0$
k	The proportion of the legitimate traffic that bypasses the filter nodes	0	$\frac{X}{C}[1-(1-f_x)^{\beta}] \geq 0$

Appendix 5.B: AT&T network topology

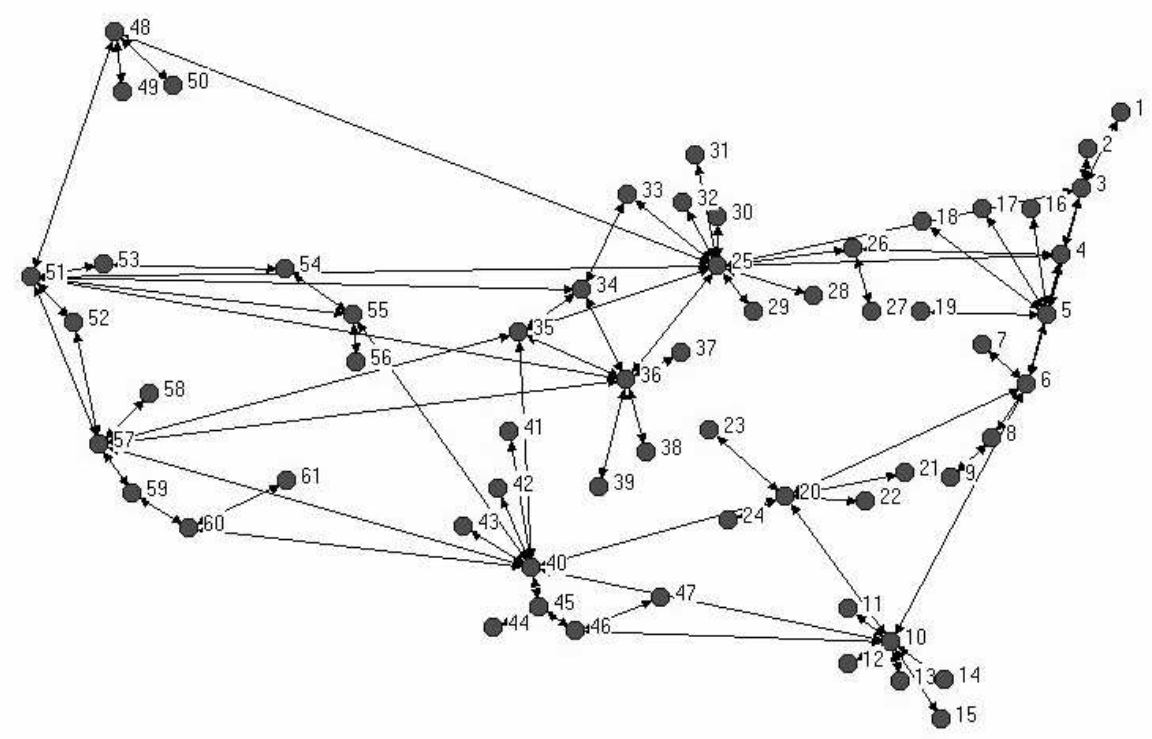

Appendix 5.C:

Prove that R_X is at its maximum when the network connection reaches its capacity.

Proof: Let $T_A = A(1-f_a)^{\alpha}$ and $T_X = X[k+(1-k)(1-f_x)^{\beta}]$.

1) When $T_A+T_X \geq C$, the network connection is saturated. The drop rate of the network connection for all types of traffic is calculated as $d = \frac{T_A+T_X-C}{T_A+T_X}$. The legitimate traffic arrival ratio can be represented as

$$R_X = (\frac{T_X}{X})(1-d) = (\frac{T_X}{X})(\frac{C}{T_A+T_X}) \text{ if } T_A+T_X \geq C .$$

$$\Rightarrow R_X < (\frac{T_X}{X}) \text{ if } T_A+T_X > C .$$

$\Rightarrow R_X$ is maximum when $T_A+T_X = C$.

2) When $T_A+T_X < C$,

$$R_X = \frac{T_X}{X} = \frac{X[k+(1-k)(1-f_x)^{\beta}]}{X} = k+(1-k)(1-f_x)^{\beta} = k+(1-k)(1-\psi(f_a))^{\beta}$$

$$\Rightarrow \frac{\partial R_X}{\partial f_a} = -\beta(1-k)\psi'(f_a)(1-\psi(f_a))^{\beta-1}$$

$\Rightarrow R_X$ is minimum when $\psi(f_a)=1$ and R_X increases when $\psi(f_a)$ decreases.

$\Rightarrow$

Since $T_A+T_X = A(1-f_a)^{\alpha}+X(1-\psi(f_a))^{\beta}$, T_A+T_X increases when $\psi(f_a)$ decreases.

$\Rightarrow R_X$ is maximum when $f_a = f_a^*$ such that $T_A+T_X = C$. QED.

Appendix 5.D: Correlation matrix of topology measures for 36 networks

	Density	Average shortest path length	Clustering coefficient	Degree centralization
Nodes	-0.74	0.77	-0.30	-0.64
Density		-0.68	0.69	0.67
Average shortest path length			-0.46	-0.61
Clustering coefficient				0.21

Chapter 6 THE ECONOMIC INCENTIVES OF PROVIDING DDOS DEFENSES ON THE INTERNET INFRASTRUCTURE

Internet service providers (ISPs) are the front line of the Internet infrastructure protection since they transport network traffic and have direct administration of the infrastructure. What would be the economic incentives of ISPs to provide defense mechanisms against network attacks? This chapter is intended to address this question by analyzing the economic benefits and costs of ISPs to provide defenses on network routers against distributed denial of service (DDOS) attacks. To deploy the defenses against DDOS attacks, ISPs need to configure routers in order to prevent attack traffic from reaching the network connections of their subscribers. Performance efficiency of the services and economic benefits from the services are two important concerns to determine if a defense is a feasible solution or not. The previous chapter has studied the performance efficiency of the defenses. This chapter focuses on evaluating the economic benefits and costs of providing defenses.

This chapter proposes that ISPs should provide DDOS defenses as network services to their subscribers. Security services, such as Virtual Private Networks or firewalls, have been provided by ISPs as optional network services to deal with the secrecy of data transportation. In this case, the services that provide DDOS defenses ensure the availability

of online services. The defense technologies have been discussed in details in Chapter 3. Some of these technologies have been implemented in (Arbor 2002; Asta 2002; Recourse 2002). The purpose of this chapter is to develop an analytical framework that ISPs can use to evaluate their economic benefits and costs when adopting these technologies. This chapter describes mathematical models to quantify the economic benefits and costs of the service provision from the perspectives of both providers and subscribers. By using the models, the chapter examines the benefits and the costs under various filtering methods, filter locations, network topology, and pricing choices.

The next section describes the mathematical models. Section 6.2 uses empirical data of distributed denial of services and information about a backbone network to verify the models empirically. The subsequent sections examine the importance of various factors on the analysis. Section 6.3 investigates the defenses that monitor attacks to victims. Section 6.4 investigates the defenses used to monitor attack sources. Section 6.5 discusses the impact of attacks on network capacity. Section 6.6 considers different attack scenarios. Section 6.7 analyzes network topology. Section 6.8 analyzes the service provision with a different pricing strategy. Section 6.9 analyzes the model in the monopoly market setting, which assumes that the provider is able to maximize its profits by adjusting the service charge. Section 6.10 concludes the chapter.

6.1 MATHEMATICAL MODELS

This section proposes mathematical models to quantify the economic benefits and costs of the DDOS defenses. This section defines the models based on the following

simplifying assumptions. Some of these assumptions will be relaxed in the later sections. These assumptions are:

- DDOS attacks can be traced to their sources within the administrative domain of one network provider. (Chapter 7 will discuss the cooperation among multiple network providers when attacks can be traced across different administrative domains.)

- The attacks saturate the network connections of subscribers to their backbone networks or take down servers inside the network of the subscribers.

- Subscribers would pay based on the utility received from the defense. The utility that a subscriber derives from DDOS defenses is the expected value of losses that would be incurred from DDOS attacks.

- Providers would offer the service to an additional subscriber when the marginal benefit to the provider is larger than the marginal cost to the provider.

- The providers charge all subscribers at a flat rate for a certain time period, such as a month (this assumption will be discussed in Section 6.8).

- The service is offered in a competitive market where the price for the service is determined so that the number of subscribers that are willing to subscribe it is equal to the number of subscribers that the provider would like to offer it (this assumption will be discussed in Section 6.9).

Two categories of DDOS defenses are modeled in this chapter. They are source filtering and destination filtering. Chapter 3 has detail description of these two categories.

Source filtering refers to the defenses that are deployed to monitor the outbound traffic of a subscriber in order to prevent the subscriber from originating attacks. Destination filtering refers to the defenses are deployed to monitor the inbound traffic of a subscriber in order to prevent the subscriber from being attacked.

6.1.1 BENEFITS AND COSTS OF SUBSCRIBERS

What a subscriber is willing to pay for DDOS defenses is assumed to be less than the utility received from the security service. This section uses a linear function to quantify the utility. A similar linear function form has been used to quantify the expected loss associated with the information set being compromised in an attack (Gordon and Loeb 2002) and the utility of subscribers for intermediary services (Bhargava, Choudhary et al. 2000) and digital goods (Bhargava and Choudhary 2001).

The utility that a subscriber derives from DDOS defenses is the expected value of losses that would be incurred from DDOS attacks. The expected loss is quantified by three factors. The attack frequency, $a \in [0,1]$, refers to how often attacks occur. The expected loss per attack, L, refers to how much loss an attack imposes on the subscriber. The quality of the defense, $q \in [0,1]$, quantifies the impact of the performance efficiency on the expected loss, such as the legitimate traffic arrival rate discussed in Chapter 5. Let U denotes the utility function of a subscriber for the service, which is defined as:

$$U = aqL \qquad (1.a).$$

Consider a simplifying situation that only one type of service is offered and the provider charges each subscriber a flat rate p for a certain time period, such as a month.

Based on the assumption that a subscriber is willing to pay less than the utility, the upper bound for the service charge p_d is:

$$P_d \leq aqL \quad \text{(1.b).}$$

Assume that L for all subscribers is proportional to a uniform distribution. q denotes the quality of the service for DDOS defenses, which can be considered as the performance efficiency such as the legitimate traffic arrival rate discussed in Chapter 5. The number of subscribers that will subscribe to the service depends on the distribution of a. $F(a)$ denotes the percentage of the subscribers that have at least a attacks, and assume that L and a are independent variables. Only the subscribers that expect the attack frequency to be larger than $\frac{qL}{P_d}$ would subscribe to the service at P_d. Let M represent the number of subscribers of an ISP. Let N_d denote the number of subscribers that are willing to subscribe to the DDOS defense service. When the price is set at p_d, N_d is calculated as:

$$N_d = F(a)M \quad \text{(1.c).}$$

From (1.c), the lowest attack frequency of all subscribers is a function of N_d, which is:

$$K(N_d) = a = F^{-1}(\frac{N_d}{M}) \quad \text{(1.d).}$$

6.1.2 BENEFITS AND COSTS OF PROVIDERS

This section quantifies the benefits and the costs of providing DDOS defenses. The cost quantification considers only the operational cost of providing DDOS defenses but not

the capital investment for the infrastructure to implement them. Three factors are considered in quantifying the operational cost. They are: 1) fixed cost (C_o), 2) filter overhead (R), and 3) bandwidth saving (W). Both R and W quantify the per-attack operating cost while C_o quantifies the per-subscriber operating cost. Fixed cost (C_o) quantifies the additional cost per subscriber that the provider has to pay in order to set up the service for the subscriber. For example, the cost of additional equipment, such as disk space for logging, or additional administrative overhead. Filter overhead (R) quantifies the per attack overhead of a defense on IP transport due to attack detection and responses. If the provider provides an IP transport service that guarantees a certain quality of service (QoS), the additional overhead imposes an economic cost to the provider. Bandwidth saving (W) quantifies the per attack transport cost saved because attack packets are filtered before they are transported to their destinations.

Filter overhead per attack R is defined to be proportional to the number of filters $H(G)$, the link utilization by legitimate traffic μ_x, and the attack duration τ. Given a network topology G, $H(G)$ is calculated as the number of edges monitored by filters, which are deployed between attack sources and victims. $H(G)$ is influenced by the network topology because filters must be deployed at some cut points between the attack source networks and the victim networks. The model assumes that filters are triggered only when attacks are detected and that the proportional relationship is linear. C_r denotes the unit economic cost of filter overhead and S denotes the number of attack sources, R is defined as:

$$R = \tau \mu_x C_r H(G) \qquad (2.a).$$

Bandwidth saving per attack W is defined to be proportional to transport distance saved $D(G)$, the link utilization by attack traffic μ_a, and the attack duration τ. $D(G)$ is calculated as the transport distance between filters and the victim networks, which is also topology dependent. f_a denotes the attack traffic filtering rate and C_w denotes the unit economic cost of bandwidth. $W(G)$ is defined as:

$$W = \tau \mu_a C_w D(G, f_a) \quad (2.b).$$

The total cost of providing the defense C is the sum of operational cost C_o from all subscribers, and R from all attacks. $\Theta(N_s)$ represents the total number of attacks from all subscribers of the service, which is equal to $\sum_{i=1}^{N} a_i$ where a_i is the attack frequency of i^{th} subscriber. The service is offered to N_s subscribers. The total cost for providing the service to N_s subscribers is calculated as:

$$C = C_o N_s + R\Theta(N_s) \quad (2.c).$$

The total benefit B for providing the service is the sum of the service charge P_s from all subscribers, and W from all attacks:

$$B = P_s N_s + W\Theta(N_s) \quad (2.d).$$

The total profit for providing the services TP is:

$$TP = B - C = P_s N_s + (W - R)\Theta(N_s) - C_o N_s \quad (2.e).$$

By setting $\frac{dTP}{dN_s} = 0$, the lower bound of the service charge (the marginal cost of providing the service to one additional subscriber) is:

$$P_s \geq C_o + [R - W]K(N_s) \quad \text{(2.f).}$$

6.1.3 THE BENEFIT-COST RATIO

Two measures will be used in the later sections to discuss the economic incentives for providing DDOS defense services: 1) the benefit-cost ratio per service (δ_1) and 2) the benefit-cost ratio per attack (δ_2).

The benefit-cost ratio per service (δ_1) measures the ratio of the total benefit to the total cost for the ISP from providing the service under a flat rate pricing scheme based on the equilibrium in the competitive market. When the service is offered in a competitive market, the price $P = P_s = P_d$ is determined at the point that the number of subscribers is equal to $N = N_s = N_d$. The benefit-cost ratio is defined as:

$$\delta_1 = \frac{B}{C} = \frac{PN + W\Theta(N)}{C_o N + R\Theta(N)} \quad \text{(3.a).}$$

The benefit-cost ratio per attack (δ_2) measures the ratio of the bandwidth saving (per-attack operational benefit) to the filter overhead (per-attack operational cost) without considering the fixed cost (per-subscriber operational cost) and the service charge. The benefit-cost ratio per attack is defined as:

$$\delta_2 = \frac{W}{R} = \frac{D}{H}\frac{C_w \mu_a}{C_r \mu_x} \quad \text{(3.b).}$$

The next section will discuss the parameters for a base scenario to provide DDOS defense services. This base scenario calibrates the parameters in the mathematical models using public available data.

6.2 CALIBRATION OF BASE SCENARIO PARAMETERS

6.2.1 EMPIRICAL DATA FOR THE DISTRIBUTION OF THE ATTACK FREQUENCY

The distribution of the attack frequency determines the demand of subscribers for the service of DDOS defenses (as in the equations 1 through 3) when the service is priced at a flat rate. Since no service is currently offered at this point, empirical measures of demand for the service is not available. However, the variation of the demand among individual subscribers can be estimated from empirical data of attack incidents. The variation of the demand can be explained as the difference among the online services that each subscriber operates. For example, the demand for the service from an e-commerce web site such as Yahoo or eBay is higher than a personal web site since the probability of attacks to an e-commerce web site is greater.

This section describes the two empirical data sets: 1) the DDOS data set and 2) the Code-Red data set. They will be used to quantify the demand of individual subscribers in the following sections. The two data sets are used to calibrate *F(a)* in equations 1 through 3. The DDOS data set is used to estimate the distribution of attacks "sent to" one subscriber, and the Code-Red data set is used to estimate the distribution of attacks "originating from" one subscriber.

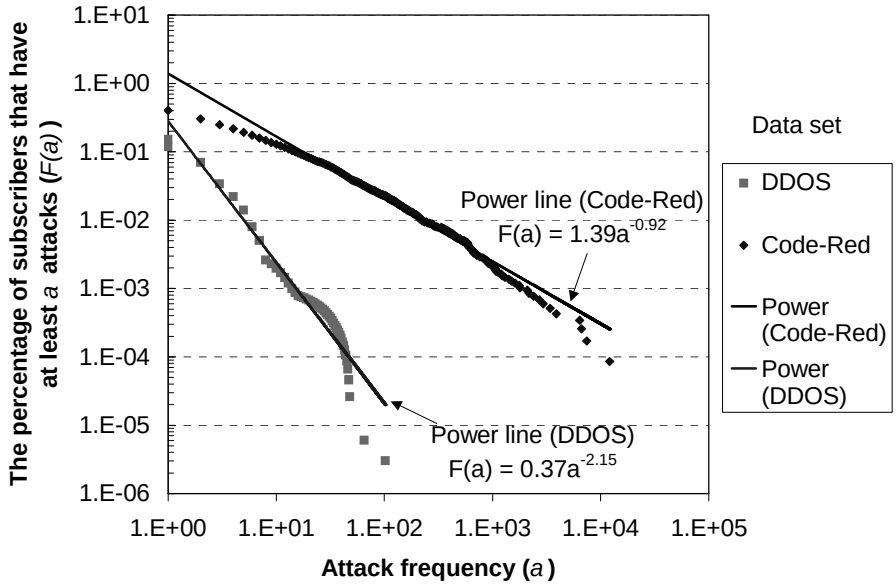

Figure 6.1: The distribution of the attack frequency

Figure 6.1 illustrates the distribution of the attack frequency from these two data sets. The DDOS data set is compiled from the number of distributed denial of services to unique IP addresses estimated in (Moore, Voelker et al. 2001), in which attack events are estimated based on the backscatter data collected from a /8 subnet[14] during a three-week long period. The Code-Red data set is compiled from the number of attack probes caused by the Code-Red worms from unique Autonomous Systems (ASes) estimated in (Moore 2001), in which the attack probes consist of infection attempts to computers in one /8 subnet and two /16 subnets during a 24-hour period when the Code-Red worms started to spread. Computers infected by the Code-Red worm launch DDOS attacks to the White House web site during a certain time period coded in the worm program. Because of this behavior, the Code-Red worm is a propagation program for DDOS attack tools and the Code-Red data set identifies the distribution of DDOS attack sources.

Both studies only report the number of ASes observed versus the attack frequency. To calculate the distribution of the attack frequency, this chapter normalizes the number of ASes involved in attack events in the original data sets by the total number of ASes in August 2001. In addition, the percentage of ASes that are not involved in the attack events is calculated by subtracting the total number of ASes from the number of ASes observed in these two studies. The total number of ASes in August 2001 is 11717. Table 6.1 shows several descriptive statistics of the two data sets.

Data set	Number of samples	a_{max}	$F(a \geq 1)$	$F(a \geq 90\%$ *of total attacks)*
DDOS	10030[15]	102	0.2	0.12
Code-Red	4728[16]	12102	0.4	0.08

Table 6.1: Descriptive statistics of the two data sets

Data set	b_0	b_1	R-Square
DDOS	0.37	-2.15	0.93
Code-Red	1.39	-0.92	0.98

Table 6.2: Parameters for the approximation of the two data sets using a power functional form

[14] In the IPv4 protocol currently used on the Internet, each IP address has 32 bits. A subnet with a *n*-bit network prefix (usually denoted as "/*n*") refers to the first *n* bits of the IP addresses from this subnet are fixed and this subnet can use up to $2^{(32-n)}$ IP addresses.

[15] Number of unique IP addresses

[16] Number of unique ASes

For $a \geq 1$, a power curve functional form $F(a) = b_0 \, a^{b_1}$ is used to approximate the distribution of the attack frequency. These two approximated power curves will be used to calibrate *F(a)* in later analyses. Table 6.2 shows the estimated parameters for the power curve fits.

6.2.2 BANDWIDTH SAVING AND ROUTER OVERHEAD

In (2.a) and (2.b), the number of filters *H(G)* and the transport distance saved *D(G)* are topology dependent factors for estimating the filter overhead *R* and the bandwidth saving *W*, respectively. In order to estimate *R* and *W*, *H(G)* and *D(G)* are calculated using backbone network maps from (BW 2001). Each map describes a core network topology connecting North America cities for a backbone ISP. The detail description of these maps is in Chapter 5.

Both *H(G)* and *D(G)* are calculated based on two attack scenarios: 1) single source attack, and 2) distributed source attack. One node is selected as the attack victim in both scenarios. In a single source attack, one another node is selected as the attack source. In a distributed source attack, 10% of distinct nodes are selected as attack sources. The mean, maximum, minimum and standard deviation of *H* and *D* are calculated by permutations of nodes in each network. The algorithms are detailed in the Chapter 4. Appendix 6.A shows the average values of *D* and *H* for the two attack scenarios using the AT&T network. These values will be used in the following sections.

6.2.3 PARAMETERS FOR THE BASE SCENARIO

The section describes a base scenario using the mathematical models proposed in Section 6.1. Public available data is collected to calibrate the parameters in the models. In the subsequent sections, the parameters for the model analysis are set to the values in the base scenario unless they are otherwise specified. Network providers can utilize the models to estimate their benefits and costs based on their proprietary data for a more precise estimation. This analysis aims at drawing a baseline for the benefits and the costs of providing defenses against DDOS. This base scenario assumes a TCP SYN attack launched at an average packet rate based on data observed from single attack source. Destination filtering is deployed to monitor the inbound traffic to subscribers (victims of attacks). In the base scenario, the unit bandwidth cost is equal to unit filter overhead because this case assumes that the overhead imposed by filtering a packet is equal to the overhead of forwarding a packet.

Table 6.3 is a list of the parameters used in the base scenario. This chapter uses the same AT&T backbone network topology as the one in the Chapter 5 to construct the base scenario.

Category	Notation	Base value	Description
Unit cost	M	2800	Number of subscribers to network connection service. The number of business subscribers for IP transport is estimated from its market share. The estimated market share is 10% and 3.5% for AT&T and Cable & Wireless respectively. Cable & Wireless reported the number of business subscribers is 950. Hence, the estimated number of business subscribers for the AT& T in 2000 is 950*10%/3.5%~2800 (BW2001).
	C_o	$945 /month	Operation cost per subscriber. The operation cost is estimated based on current AT&T security services. AT&T charges a $945 recurring monthly fee for security services in a three-year contract. The recurring monthly fee includes Tunnel Server, 24x7 management and maintenance, help desk support, client software, and 4 hour time to response (BW 2001).
	C_r	$85,025 /month	Unit economic cost of performance overhead. Estimated based on OC3 155Mbps leased line access price from AT&T on Jan. 2001.
	C_w	$85,025 /month	Unit economic benefit of bandwidth saving. Estimated based on OC3 155Mbps leased line access price from AT&T on Jan. 2001
Network topology	$H(G)$	1	Number of edges monitored by filters. H and D are set at the value that dynamic filters are triggered at 7 hops away from the victim network (at the border of the network).
	$D(G)$	7	Distance between filters and the victim networks
Defense	q	1	Performance efficiency (in range [0,1]). The best case for legitimate traffic arrival ratio.
	f_a	0.99	Attack traffic filtering rate (in range [0,1]).
	$L(q)$	$4,080 /attack	Expected loss of an attack. In (CSI 2002), the reported average losses from denial of service for a company annually is $122,389 in 2001. Assume the number of attacks is uniformly distributed among 12 months. The average number of attacks is 2.5 from analysis in Section 6.2.1. The expected loss reduced by filters per attack = $122,389/(12*2.5)~$4,080.
	μ_x	30%	Link utilization of the edge monitored by filters. The link utilization is 20%-35% and 20%-70% in two OC-3 links in a backbone link monitor project described in (Papagiannaki, Moon et al. 2002). 30% is the medium estimation.
Attack	A	60Mb /second	Attack magnitude. It is estimated by 1500 packet per second (pps) and 40 bytes per packet. An attack with 1500 pps is enough to compromise a firewall. In the trace analyzed in (Moore, Voelker et al. 2001), 20% of all attack events had an estimated packet 1500 pps or higher. Minimum TCP packet size which carries TCP acknowledgement but no payload (McCreary, Claffy et al. 2000).
	τ	10 minutes	Duration of an attack. In the trace analyzed in (Moore, Voelker et al. 2001), 20% of attacks ≤ 5 minutes, 50% of attacks ≤ 10 minutes, and 90% of attacks ≤ 1 hour.
	S	1	Number of attack sources.
	$F(a)$		Cumulative distribution of the attack frequency. *"a"* denotes the frequencies of attacks. The DDOS data set is used for the base scenario.

Table 6.3: Parameter setting for the base scenario

6.3 DESTINATION FILTERING

When destination filtering is deployed, the closer the filter can be to the attack source, the more benefit both the provider and the subscriber will have. Figure 6.2 shows that both the provider's benefit and the subscribers' benefit increases when the filter location[17] is closer to the attack source. The provider gains from the increase of the bandwidth saving because attack traffic has been filtered out before it is transported. The subscribers also benefit from an increase of the quality of the service. That is, more legitimate traffic to the victim can bypass the filters.

This result is more significant when the packet rate of an attack is larger than 500pps. A TCP SYN attack with 500 pps is sufficient to overwhelm a server. As observed in (Moore, Voelker et al. 2001), 46% of attacks are larger than 500pps. For a provider, the benefit is more significant when the packet rate is larger than 500pps. In Figure 6.3, when the filter location is further away from the victim network, the benefit-cost ratio per service increases significantly when the packet rate of an attack increases.

[17] Attack upstream means the filter is set at one hop upstream of the network that originates attacks. Victim upstream means the filter is set at the access router to the victim's network. Hop *n* means the filter is set to *n* hops upstream from the access router of the victim's network. For example, hop 1 means that the filter is set one hop upstream from the access router of the victim's network.

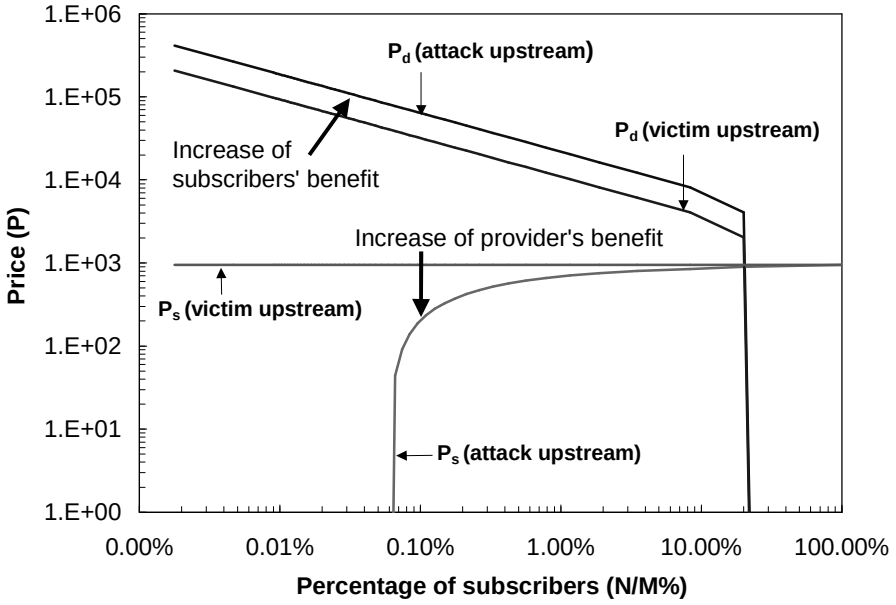

Figure 6.2: Increase on both the provider's benefit and subscribers' benefit by setting filters closer to the attack sources

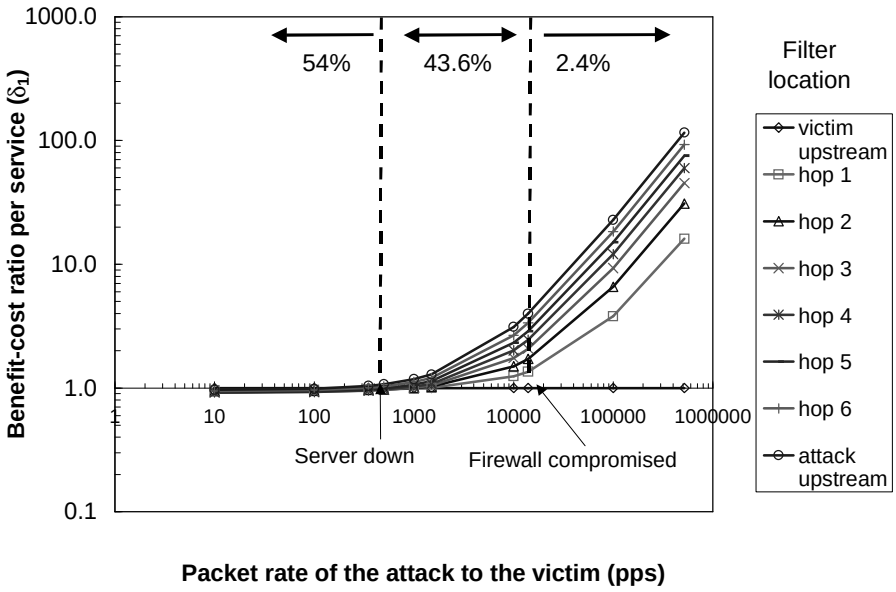

Figure 6.3: The benefit-cost ratio increases when the packet rate of the attack increases if the filter location is further away from victim upstream

6.4 SOURCE FILTERING

This section discusses the benefit-cost ratio per service for source filtering. Source filtering refers to the defenses that are deployed to monitor the outbound traffic of a subscriber in order to prevent the subscriber from originating attacks. The baseline scenario uses destination filtering, which means the defense mechanisms are deployed to monitor the inbound traffic of a subscriber in order to prevent the subscriber from being attacked (Chapter 3 has detail description of these two categories).

Instead of being attack victims, the networks of some subscribers may be exploited by attackers to launch attacks. There are three possible incentives that these subscribers may consider the service of DDOS defenses to prevent their networks being exploited. 1) These subscribers need to maintain the accessibility of their users. The legitimate traffic in their networks is blacklisted as well as the attack traffic if the victims filter out all traffic from their networks that originate attacks. 2) These subscribers may want to avoid liability from being an originator of attacks. Although there is no court case directly pertaining to DDOS at this point, assigning liability to attack sources based on contributory negligence (Kabay 2001) has been promoted as a way to create incentives for source filtering. 3) These subscribers are concerned about their reputation. For example, it could be very embarrassing for high-profile companies such as banks or security consulting firms to be exploited as attack sources.

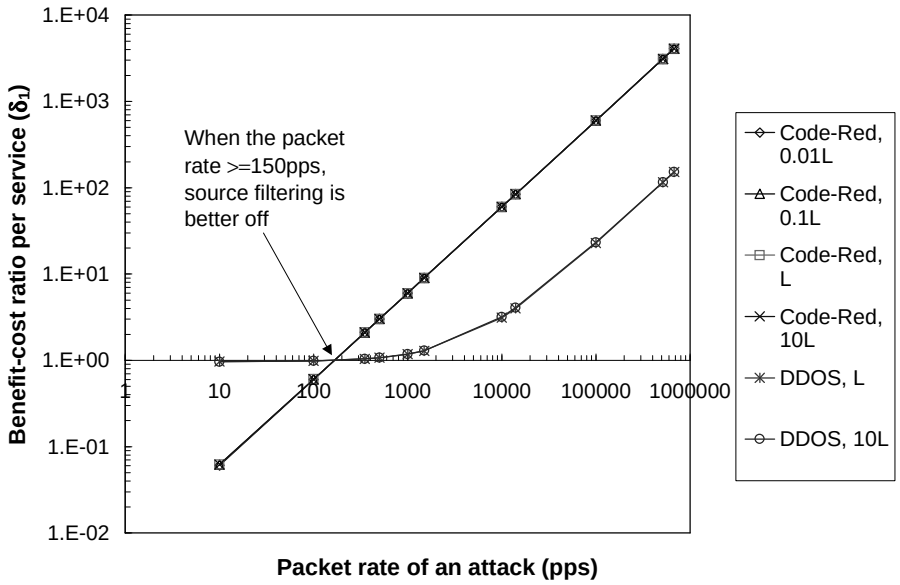

Figure 6.4: Benefit-cost ratio per service for both DDOS and Code-Red data with various levels of expected loss

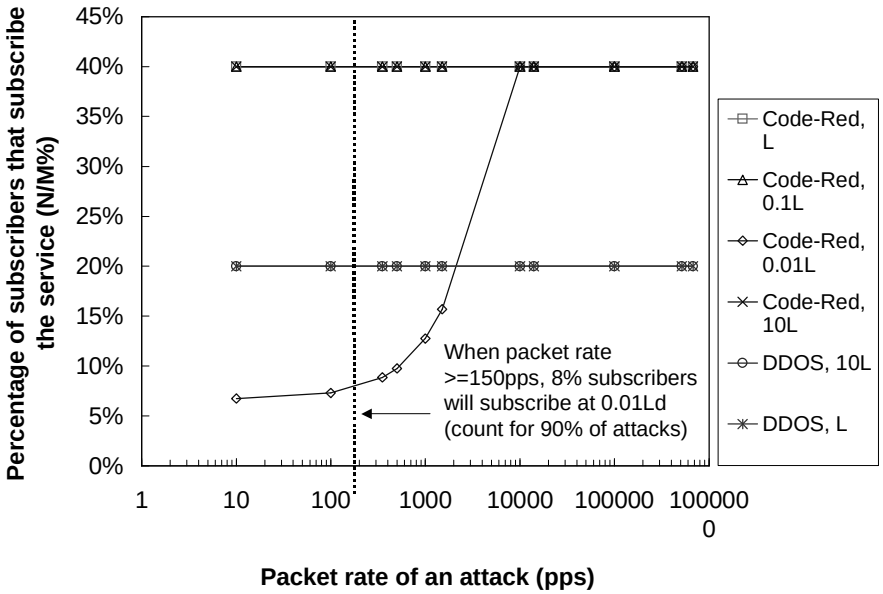

Figure 6.5: Percentage of subscribers for both DDOS and Code-Red data with various levels of expected loss

The Code-Red data set is used to estimate the attack frequency. Figure 6.4 shows the benefit-cost ratio per service and Figure 6.5 shows the percentage of subscribers that subscribe to DDOS defenses. These two figures are estimated for both source filtering (the Code-Red data set) and destination filtering (the DDOS data set) with varying level of expected loss from an attack. This analysis has the following implications:

1) The provider is better off providing source filtering than destination filtering when the packet rate of an attack is higher than a threshold. In Figure 6.4, δ_1 for the Code-Red data set is higher than δ_1 for the DDOS data set when the packet rate exceeds 150pps.

2) Even when the expected loss of attack sources is only 1% of the victims' losses, most attacks can be stopped at sources if source filtering is deployed to monitor the subnets that are more likely to originate attacks. From the Code-Red data set, 90% of attacks originate from 8% of networks. If these 8% of subscribers would subscribe to source filtering, 90% of attacks could be stopped at the sources. In Figure 6.5, this situation occurs when the expected loss of originating attacks is only 1% of the expected loss of the victims. This result implies that a policy is needed to impose a cost on subscribers that originate attacks. Once they suffer from the losses due to originating attacks, they would adopt source filtering and the number of attacks would be reduced.

6.5 NETWORK CAPACITY

When the provider's network is capacity constrained, filtering out attacks that have high packet rates closer their sources would reduce the burst traffic on the congested links. In this case, the bandwidth cost is higher than the filter overhead since the provider has to expand the capacity for the increased traffic. The provision of DDOS defenses is more

beneficial for the provider because they can avoid the capital investment for expanding capacity. Figure 6.6 shows the benefit-cost ratio increases when C_w/C_r increases for both data sets. Since attacks cause only burst traffic during a short period of time, expanding the link capacity may induce excess capacity for the long term. In addition, if attackers intend to cause burst traffic, they can generate attacks with increasingly higher packet rates as capacity is expanded. Deploying filters to prevent attack traffic from consuming capacity is better than expanding capacity for the long-term.

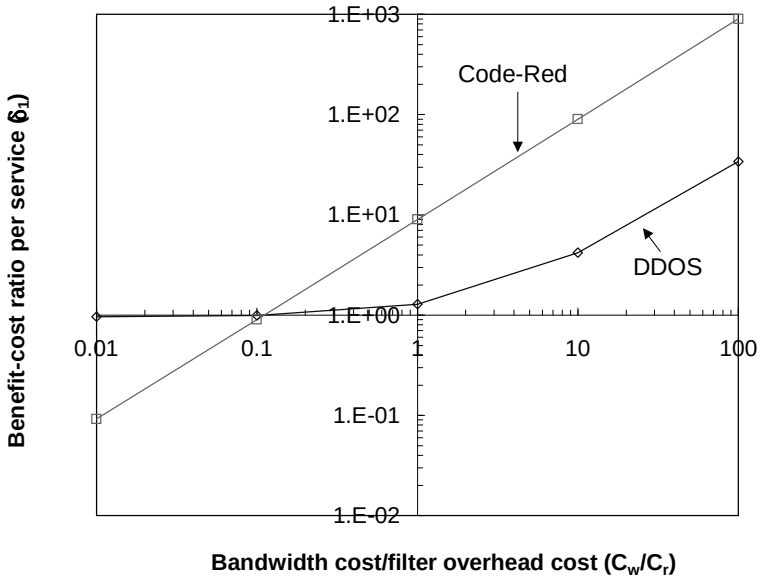

Figure 6.6: The impact of bandwidth cost/filter overhead cost

In addition, when the capacity is more constrained, the source filtering is better than the destination filtering ($C_w/C_r > 0.1$), as shown in Figure 6.6. This result implies that filtering based on destinations of attacks is better than filtering based on sources when the filtering is taking place on core routers where the computational overhead from filtering is a large burden ($C_w/C_r \leq 0.1$).

6.6 DISTRIBUTED SOURCE ATTACKS

In a distributed source attack, a network provider is not better off providing source filtering than providing destination filtering. In Figure 6.7, when the packet rate < 3000pps, δ_1 for the Code-Red data set is smaller than δ_1 for the DDOS data set. When the packet rate > 3000pps, the difference of δ_1 between the two data sets is much smaller than it is during a single source attack.

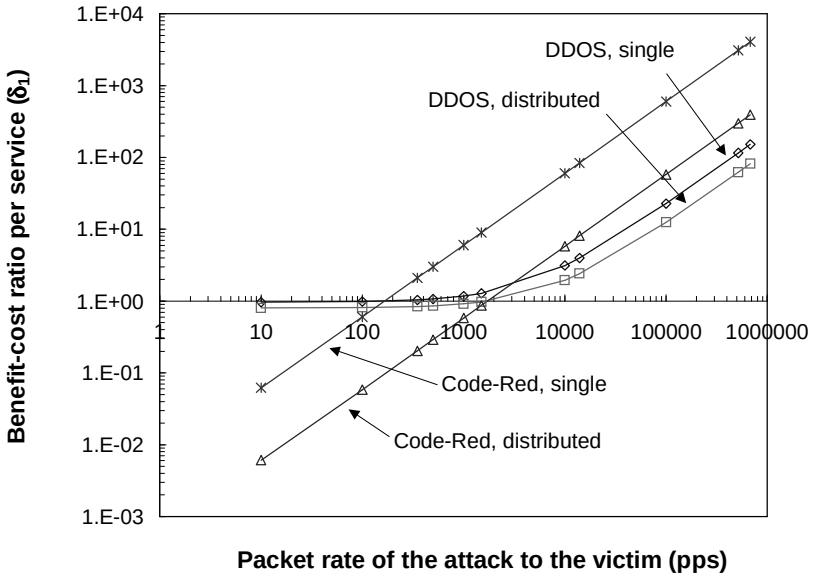

Figure 6.7: Single source attacks vs distributed source attacks for the two data sets

This result implies that the benefit-cost ratio from source filtering decreases more significantly when attack sources are more distributed. This reason for the result is that, for a given packet rate of an attack received by the victim, the packet rate sent from each attack source when the attack is distributed is less than the packet rate sent from a single source when the attack is from one source. The benefit-cost ratio decreases because the bandwidth saving for each attack source decreases during a distributed source attack. At the same

time, the bandwidth saving for the provider that connects to victims decreases less significantly because the attack traffic aggregates when it is closer to the victim.

6.7 NETWORK TOPOLOGY

No two providers have the same network topology. If a network provider has a topology that is different from the AT&T topology used by the base scenario, two variables in the benefit-cost ratio vary. They are: 1) the transport distance for attack traffic (*D*) and 2) the number of filters (*H*) needed to cut off attack traffic. Table 6.4 lists the variation of $\frac{D}{H}$ for 36 backbone network maps. Figure 6.8 plots the variation of the benefit-cost ratio per attack bounded by the variation of $\frac{D}{H}$ listed in Table 6.4.

Although the benefit-cost ratio varies, the results from both Section 6.5 and Section 6.6 hold when network topology changes. In Figure 6.8, δ_1 for the Code-Red data set is still larger than δ_1 for the DDOS data set at a certain packet rate of an attack. There are two constraints by the network topology added to this result. 1) On average, difference between the packet rate for δ_1 from the Code-Red data and the packet rate for δ_1 from the DDOS data (in this case, 500pps) for the group of topologies is larger overall than for the AT&T network topology used in the previous section. 2) The variation of the benefit-cost ratio is larger in source filtering than in destination filtering. In Figure 6.8, the variation of δ_1 for the Code-Red data is larger than δ_1 for the DDOS data.

What are the properties of a network topology that cause the variation? Table 6.5 lists the correlation of the average $\frac{D}{H}$ to several network measures that describe the

topology of the network. Among various measures, average path length has the highest positive correlation to the average $\frac{D}{H}$ in either single source attacks or in the distributed source attacks. The number of nodes in a network has the second highest positive correlation, and density has the highest negative correlation. This result implies that a network that has more nodes and a longer path (lower connectivity) would have a higher benefit-cost ratio. Such a network would benefit more from source filtering.

Attack sources	Maximum	Mean	Minimum	Standard deviation
Single source	15.20	2.54	0.12	1.47
Distributed sources	19.95	2.74	0.30	0.76

Table 6.4: $\frac{D}{H}$ calculated from 36 backbone networks

Attack source	Number of nodes	Density [18]	Average shortest path length	Clustering coefficient[19]	Degree centralization[20]
Single source attacks	0.66	-0.62	0.96	-0.54	-0.48
Distributed source attacks	0.71	-0.67	0.98	-0.55	-0.53

Table 6.5: Correlation between the average $\frac{D}{H}$ (as well as the benefit-cost ratio per attack) and the network measures for all 36 topologies

[18] Density measures the connectivity of a network, which is defined as the number of edges of a network divided by the largest possible number of edges of this network (Wasserman and Faust 1994).

[19] Clustering coefficient measures the cliquishness of a network. Node clustering coefficient is defined as the connectivity of the neighbors of a node. Clustering coefficient is the average of node clustering coefficients in a network (Watts and Strogatz 1998)

[20] Degree centralization measures the differences of the connectivity among nodes, which takes the average of the difference of individual node connectivity and the average node connectivity (Wasserman and Faust 1994).

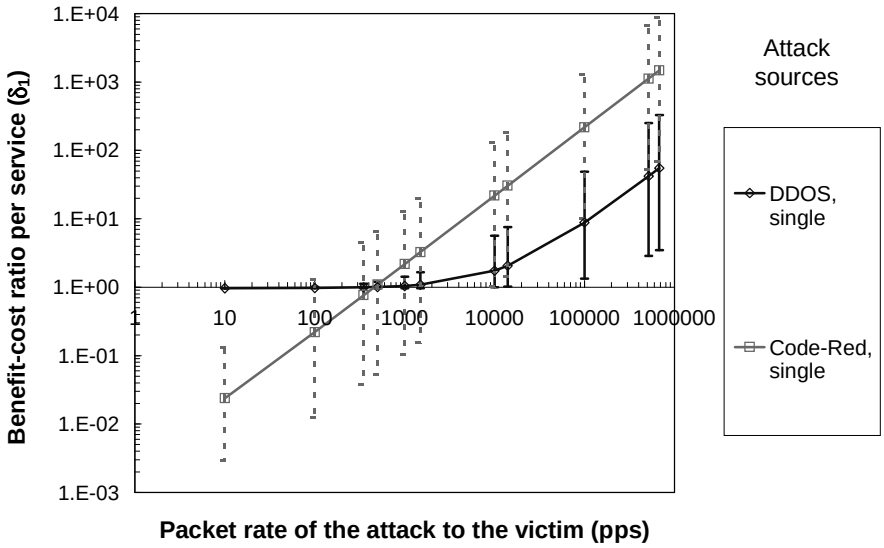

Figure 6.8: The variation of the benefit-cost ratio due to network topology

6.8 PRICING STRATEGY

In the base scenario, a flat rate pricing scheme is assumed for the service. The advantage of the flat rate pricing scheme is that the network provider does not need an accounting mechanism for the number of attacks. However, under such a scheme, the provider still needs to have an idea about the distribution of the attack frequency to/from their subscribers. This section discusses a simpler situation by relaxing the assumption of the flat rate pricing scheme.

This section discusses a different pricing scheme using the benefit-cost ratio per attack (δ_2). δ_2 represents how much benefit over cost that an ISP would obtain purely based on the nature of a defense, the nature of attacks and the topology of the network without considering the payment and the fixed cost from each subscriber. There are two situations when a provider needs to consider this ratio. 1) The provider would like to provide the

service for free in order to attract more subscribers to the IP transport service they offer. As long as the additional benefit from additional subscribers of the IP transport service could cover the fixed cost, the provider should consider providing DDOS defenses. 2) The provider charges the subscribers for only the fixed cost per subscriber so that they do not need to know the distribution of the attack frequency.

Figure 6.9 and Figure 6.10 compare δ_1 and δ_2 for the DDOS and the Code-Red data, respectively. There are two implications from this analysis. 1) For destination filtering (Figure 6.9 based), providing the service for free is more beneficial when the packet rate of an attack exceeds a threshold. This threshold (150pps for attack upstream and 14000pps for hop1) is lower when the filter location is closer to the attack source. As long as the fixed cost per subscriber can be covered from other services, the additional benefit that the provider obtains by offering free DDOS defense is larger than the loss of from the service when the packet rate of an attack is higher than the threshold. 2) For source filtering (Figure 6.10a), the flat rate pricing scheme has the same benefit-cost ratio as the service provided for free if the fixed cost is covered from other services. As in Figure 6.10, δ_1 and δ_2 are approximately equal across all packet rates. The reason for this is that the number of attack frequency is very large in this case so that the benefit per attack is much larger than the benefit from subscription charge. In this case, the impact of the price is negligible.

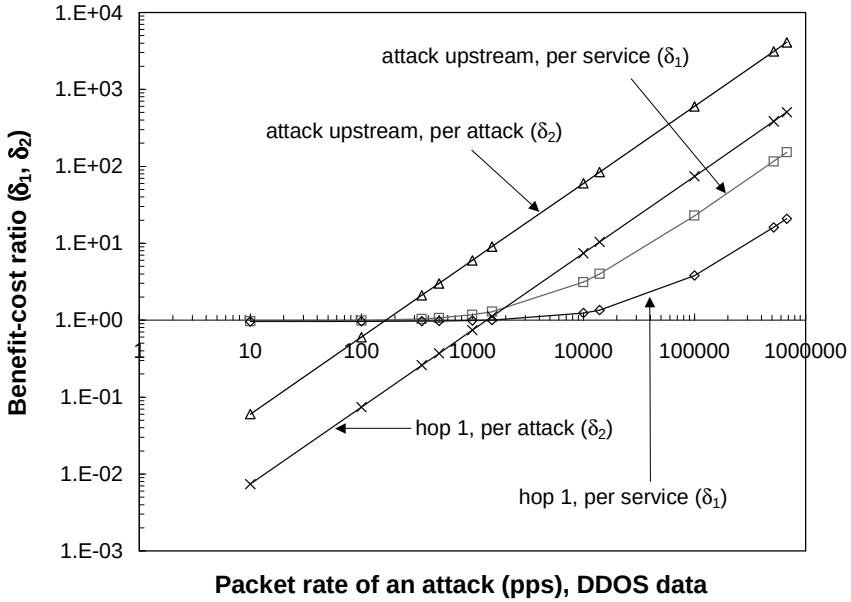

Figure 6.9: Benefit-cost ratio per service vs benefit-cost ratio per attack for DDOS data

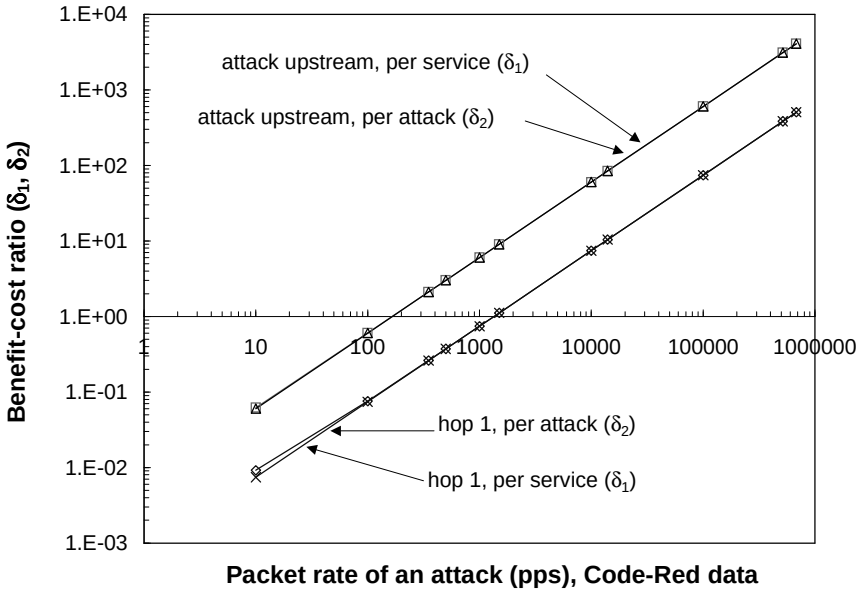

Figure 6.10: Benefit-cost ratio per service vs benefit-cost ratio per attack for Code-Red data

6.9 MONOPOLY MARKET

In the base scenario, the price for providing DDOS defenses is determined based on the competitive market assumption, in which equating subscribers' demand and the marginal cost of providing the service sets the price. However, at the initial stage of deployment for providing automatic defenses on network routers against DDOS attacks, only a few providers are capable of providing the service. In the short term, the providers are able to set the price based on the utility that subscribers receive from the service. Under such conditions, the competitive market assumption may not be true. To describe the service provision in the short term where only a few providers would provide the service, this section discusses the benefit-cost ratio based on the monopoly market assumption.

Only a few subscribers suffering a large number of attacks will be willing to subscribe to the service if the provider sets a flat rate price for all subscribers to maximize its profit (let $p^* = K(n)qL$ and maximize *TP*). The reason is that the flat rate price will be set to attract high profile subscribers because they are willing to pay more for the frequent attacks. Under such a flat rate price, most subscribers are not willing to pay for the service. This situation is not beneficial for the overall security of the infrastructure.

For majority of subscribers, an alternative pricing scheme should be provided under the monopoly market. A possible pricing scheme is to charge subscribers differently based on their individual utility from the service (as equation 1.a). However, the individual utility of the service could be hard to calculate. An alternative is to differentiate the service to several versions for subscribers who have different expected loss. Digital product vertical

differentiation has been studied in (Bhargava and Choudhary 2001). Further study on versioning services for providing DDOS defenses can be drawn.

Figure 6.11 compares two situations. They are 1) the flat rate pricing scheme (assumed in the base scenario) in the competitive market and 2) the differential pricing scheme for individual subscribers (described above) in the monopoly market. The differential pricing considers an extreme case that the provider can price the subscribers based on their individual utility, which is determined by their expected loss and the attack frequency, the benefit-cost ratio increases when the expected loss of a subscriber increases. In particular, when the expected loss increases 10 times of the current level, the benefit cost-ratio is approximately constant across all packet rates. As in Figure 6.12, the Code-Red data shows similar results. This situation is possible in reality. Some subscribers may have higher expected losses than other subscribers because of the online services they provide. For example, an attack on Yahoo or eBay may impose a higher economic cost than an attack on a web site that does not provide E-commerce. Empirical studies (Cavusoglu, Mishra, et al. 2002)(Ettredge and Richardson 2002) have found evidence to support that the stock price of E-commerce firms drop more significantly than conventional firms once a security break is announced.

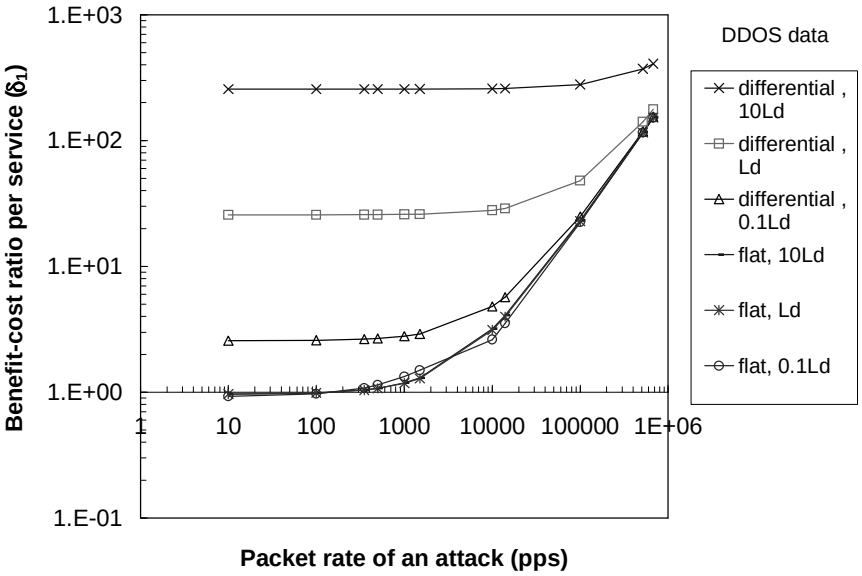

Figure 6.11: Differential pricing in the monopoly market for DDOS data

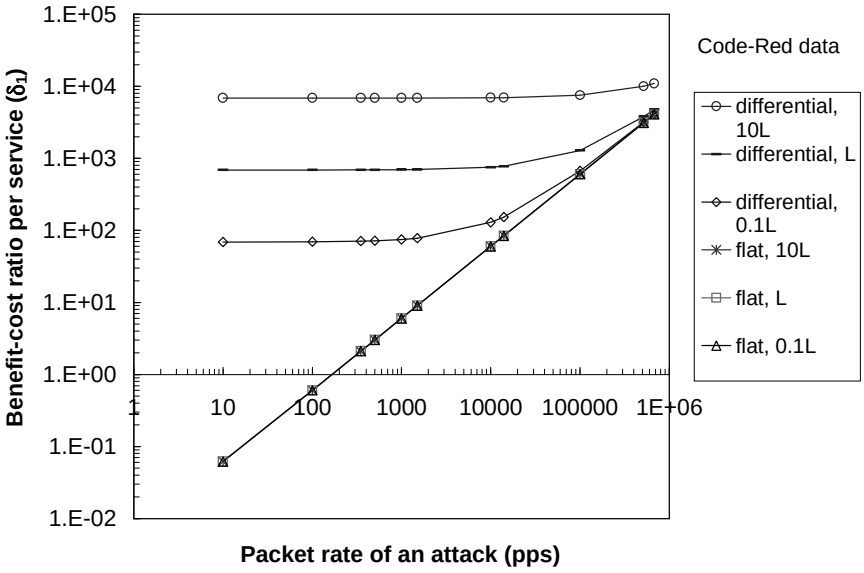

Figure 6.12: Differential pricing in the monopoly market for Code-Red data

6.10 CONCLUSIONS

At this point, none of the automatic defenses against DDOS have been provided as a service because economic incentives for providing such a service are not clear and the technical uncertainty is not well analyzed. The last chapter provided an analysis on the technical uncertainty. This chapter clarifies the economic incentives for providing the service.

To introduce the new service for their subscribers, network providers need to ensure that the operational profit in the long term would justify their capital investment. This chapter has found several reasons to expect that the operational benefit will be higher than the operational cost of the service. First, at the initial stage when few providers are able to deploy the service (monopoly market assumption), the providers should implement a differential pricing scheme. By doing this, the provider can benefit from the different levels of expected loss experienced by subscribers and from the different levels of the attack frequency. Secondly, when more and more providers are able to provide the service (competitive market assumption), no single provider can benefit from the differential pricing since subscribers can have more choices by switching to another provider. In this case, three implications can be drawn from the analysis in this chapter:

1) Setting the filter location closer to the attack source is more beneficial than closer to the victim network for both the subscribers and the providers. This result is more significant when the network of the provider is capacity constrained.

2) Providing source filtering is better for a provider than providing destination filtering when most attacks to its subscribers are launched at high packet rates and when

subscribers that originate attacks suffer losses. Offering source filtering is more beneficial than offering destination filtering since the probability of originating attacks is higher than the probability of being attacked. This result is true even when the loss to originating networks is only 1% of the expected loss of attack victims.

3) Source filtering is more beneficial when the network of the provider is less connected and has a long average path length.

4) The provider is better off providing the destination filtering service for free if the fixed cost per subscribers can be recovered from the additional income from additional subscribers to network transport services in a competitive market.

This chapter shows that the service provision of DDOS defenses can bring economic benefits to providers with an appropriate pricing strategy, some investigation into the expected loss of subscribers, and knowledge on the overall risk level of attacks. This chapter discusses the provision of DDOS defenses that can trace attacks to their sources within the administrative domain of one network provider. Chapter 7 will discuss the cooperation among multiple network providers when attacks can be traced across different administrative domains. As discussed in Chapter 2, the deployment of DDOS defenses needs the cooperation of multiple ISPs when attack traffic is transported across multiple administrative domains. The next chapter will discuss how the cooperation of multiple ISPs would influence the economic incentives of the service provision.

Appendix 6.A

Filter location	Single source			Distributed sources (10% nodes)		
	q	H	D	q	H	D
0	0.50	0.0	0.0	0.50	0	0
1	0.51	1.1	1.0	0.51	2.9	5.7
2	0.52	1.6	1.9	0.51	4.7	10.9
3	0.67	2.2	2.9	0.54	6.3	15.5
4	0.78	2.5	3.8	0.62	7.3	19
5	0.91	2.2	4.8	0.73	7.7	23
6	0.98	1.6	5.7	0.79	7.2	25
7	0.99	1.0	6.7	0.80	6.8	27.4

Table 6.6: Values of *q*, *H* and *D* for destination filtering in the AT&T network

Chapter 7 AN ANALYSIS ON THE COOPERATION OF PROVIDING DDOS DEFENSES

The last chapter investigates the economic incentives of Internet Service Providers (ISPs) for providing DDOS defenses when the attacks can be traced to sources within the administrative domain of one ISP. The cooperation among ISPs is needed in order to improve the performance efficiency of DDOS defenses when the attacks are originated from sources that are located in a different administrative domain from the victims. There are two reasons why cooperation is needed. First, cooperation would help the downstream ISPs (whose networks contain the victims) to distinguish attack traffic from legitimate traffic since the upstream ISPs (whose networks contain the attack sources) are closer to the attack sources when the source addresses of the attack packets are forged (as discussed in Chapter 2). Secondly, cooperation would help to filter out attack traffic closer to their sources. If the upstream ISPs can detect and filter attacks before they are transported, the number of attacks that the victims suffer can be reduced more efficiently.

Although the decisions of the upstream ISPs would influence the performance efficiency of the defenses provided by the downstream ISPs, the upstream ISPs may not be willing to cooperate since they are not the direct targets of the attacks. For this reason, from ISPs' perspective, it is important to investigate under what circumstances they would have an economic incentive to cooperate with others for facilitating DDOS defenses. From a

public policy perspective, it is necessary to know how cooperation among ISPs would influence the network security of the Internet community as a whole. The goal of this chapter is to investigate the economic incentives of ISPs on cooperation to provide DDOS defenses. More importantly, the chapter will discuss what types of policies are needed in order to make cooperation possible.

The next section describes three possible types of cooperation based on current technology. Section 7.2 provides an analytical model to investigate the problem. Section 7.3 describes empirical data for the estimation of the model parameters. Section 7.4 provides the results from the analytical model using the empirical data. Section 7.5 qualitatively discusses the public policy implications of the results. Section 7.6 provides conclusions and recommendations.

7.1 THE TYPES OF THE COOPERATION

The type of cooperation that facilitates DDOS defenses depends on the technology that the network provider chooses. This section identifies three types of cooperation needed for the current technologies.

1. Cooperative attack filtering

In cooperative attack filtering, the upstream ISP assists both the tracing and the filtering of the attacks to the victims. In Figure 7.1, ISP 2 (the downstream ISP) is able to trace back to the attack sources with the cooperation of ISP 1 (the upstream ISP), and ISP 1 places an attack response (filters) at the upstream routers of the attack sources. Automatic mechanisms such as Pushback (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin

2002) are designed to achieve this purpose, which is described in Chapter 3. This method needs the upstream ISP to configure their routers for attack tracing and filtering.

Cooperative attack filtering has two impacts on the service provision. First, cooperation improves the performance efficiency of defense services provided by the downstream ISP to its subscribers. Since filters are set closer to their sources, the attack traffic and the legitimate traffic originated from the upstream ISP can be distinguished more effectively. Secondly, the downstream ISP can be regarded as one of the subscribers of the upstream ISP. The benefit-cost ratio for providing the service can be quantified by the same model as the one in Chapter 6. In this model, the service provision imposes the filter overhead and saves bandwidth for the upstream ISP.

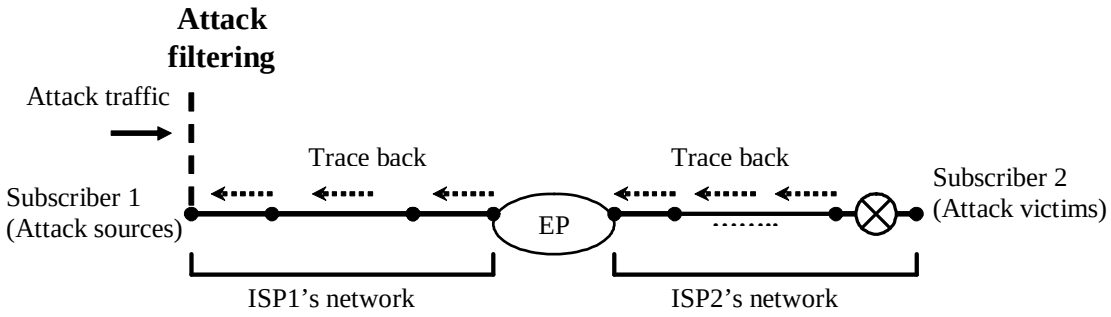

Figure 7.1: An illustration of cooperative attack filtering

2. Cooperative attack detection

In cooperative attack detection, the upstream ISP assists only the attack detection but does not conduct attack filtering. In Figure 7.2, when subscriber 1 (in ISP 1's network) originates attacks to subscriber 2 (in ISP 2's network), the attack responses are deployed in ISP 2's network. In this case, the upstream ISP assists attack detection by inserting additional information in packets or keeps logs of packets. The false positive rate of

distinguishing attacks is reduced because of the assistance of the upstream ISP. Methods such as preferential filtering (Sung and Xu 2002) or threshold filtering (Yaar, Perrig et al. 2003) insert marks in packets for later attack detection. Hash-based IP traceback (Snoeren, Partridge et al. 2001, 2002) leaves traces of packets on routers for later tracing. Chapter 3 has a detail description on these methods.

Cooperative attack detection has two impacts on the service provision. First, similar to cooperative attack filtering, cooperative attack detection improves the performance efficiency of the defenses provided by downstream ISPs. Secondly, the upstream ISPs do not have the benefit from the bandwidth saving because the cooperation does not filter out attack traffic at upstream routers.

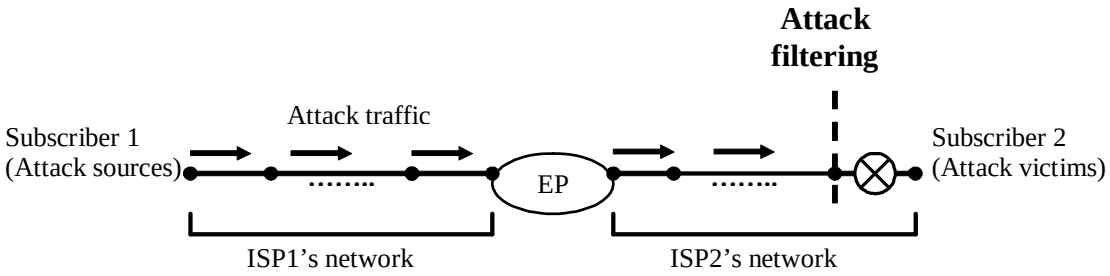

Figure 7.2: An illustration of cooperative attack detection

3. Source filtering

As discussed in Chapter 6, source filtering occurs when the upstream ISP monitors the outbound traffic for attacks sent from its subscribers to others and charges its own subscribers for the service. Since the attacks are filtered out at the sources before it is sent out to the downstream subscribers, this type of cooperation decreases the observable number of attacks at downstream ISPs. Current technologies for source filtering have been

discussed in Chapter 3. Since the provision of source filtering has been discussed in Chapter 6, this chapter will not further analyze source filtering.

7.2 THE ANALYTICAL MODEL

The chapter provides an analytical model to investigate under what circumstances would an ISP decide to cooperate for providing DDOS defenses. The benefit and the cost for the service provision are quantified in the same way as in Chapter 6. The functional form for the decisions of multiple ISPs for cooperation is derived from Critical Mass Theory (Oliver, Marwell, et. al. 1985; Marwell and Oliver 1993), which has been used to quantify the provision of public goods[21] by multiple players.

As discussed in Section 7.1, in either cooperative attack filtering or cooperative attack detection, the effort of cooperation by the upstream ISPs contributes to the performance efficiency of the defenses provided by the downstream ISP. This property meets the concept of public goods (Samuelson 1954). Public goods refer to the goods that are non-excludable and non-rivalry. "Non-excludable" means the provision of the goods by one individual cannot be withheld from the consumption of any other people. "Non-rivalry" means the consumption of the goods by one individual does not reduce the benefit of the goods to any other individual. For example, libraries are public goods. In the DDOS defenses, the upstream ISP cannot exclude other downstream ISPs from having the benefit of the DDOS defenses. In addition, the benefit of the upstream ISP or other downstream ISPs does not change even if other downstream ISPs are benefiting from cooperation. In this problem, the public good is the availability of Internet communication. Theoretical

[21] Public goods were called "collective goods" in the original papers.

analyses using public goods theory have been conducted on understanding the provision of system reliability (Varian 2002).

7.2.1 THE MODEL

Suppose that the self-interest of an ISP is to maximize its total profit. By revising the model in Chapter 6, ISP *i*'s total profit TP_i is the total benefit minus the total cost of operating the service. The decision problem for ISP *i* is:

$$\underset{X_i}{Max}\ TP_i = B_i(P(q)) - C_i(X_i) \tag{1.a}$$

$$P(q) = q^s \tag{1.b}$$

$$q = \sum_{j=1}^{m}[\alpha_j X_j] \tag{1.c}$$

$$B_i, C_i \geq 0;\ q,P \in [0,1];\ s \geq 1;\ \alpha_j \in [0,1];\ X_j = 0,1,\ j = \{1,...,i,...,m\} \tag{1.d}$$

This model is an extension of the model provided in Chapter 6. All the assumptions and variables discussed in Chapter 6 are implicitly included in the model discussed here. The difference is that this model considers the quality of the defense determined by the decisions of all ISPs regarding cooperation. Assuming that the attack is transported across different administrative domains, both B_i and C_i vary based on the decision on whether to cooperate by the ISP and its interconnected ISPs.

The additional uncertainty discussed in this chapter is the decision of an ISP to cooperate with others. Assuming that there are *m* ISPs on the Internet, then $j = \{1,...,i,...,m\}$. Let X_j denotes ISP *j*'s decision. X_j is a binary variable in which "1"

means to cooperate and "0" means not to cooperate. For a particular ISP *i*, the decision problem is to decide if it will cooperate (X_i=1) or it will not cooperate (X_i=0).

The cost C_i represents the cost of ISP *i* to provide the service, which is assumed to be an increasing function over X_i.. The benefit B_i represents the benefit to ISP *i*. *q* denotes the quality of the defense. This parameter indicates the influence of all ISP's decisions on the legitimate traffic that a subscriber can receive during attacks. *q* is assumed to be a linear combination of X_i. The weight parameter a_i denotes how much influence of an ISP's decision has on the quality of the defense. The scale parameter *s* denotes the proportional relationship of *q* to how much the subscribers would like to pay for the service. For example, *s=1* refers to what subscribers would like to pay is linearly proportional to the quality of the defense, as assumed by (1.a) in Chapter 6.

If the attack sources and the victims are in the same administrative domain of an ISP, the total benefit B_i and the total cost C_i is the same as calculated in Chapter 6. This case is a degenerated form of (1.a)-(1.d), in which a_i=1 and a_j=0, $\vee j \neq i$.

7.2.2 THE SOLUTION AND THE BENEFIT-COST RATIO

Assume that an ISP will cooperate only when its total profit from the cooperation is larger than its total profit without the cooperation. The condition for an ISP to cooperate is:

$$X_i = 1, \text{if } B_i(P(\alpha_i + \sum_{j\neq i, j=1}^{m} \alpha_j X_j)) - C_i(X_i = 1) > B_i(P(\sum_{j\neq i, j=1}^{m} \alpha_j X_j)) - C_i(X_i = 0) \quad (2).$$

Assuming that ISPs that decide not to provide services can only choose not to cooperate. Based on this reason, both $B_i(P(\sum_{j\neq i, j=1}^{m} \alpha_j X_j))$ and $C_i(X_i = 0)$ are equal to zero. Thus, the solution for (1.a)-(1.d) is:

$$X_i = \begin{cases} 1 & \text{if } \dfrac{B_i(P(\alpha_i + \sum_{j\neq i, j=1}^{m} \alpha_j X_j))}{C_i(X_i = 1)} > 1 \\ 0 & \text{otherwise} \end{cases} \quad (3).$$

Equation 3 can be interpreted as "an ISP *i* would decide to cooperate when a number of others decide to cooperate such that the benefit is larger than the cost for providing the defense". Let δ denotes the benefit-cost ratio of the service for an ISP to cooperate, which is defined as:

$$\delta = \frac{B_i(P(\alpha_i + \sum_{j\neq i, j=1}^{m} \alpha_j X_j))}{C_i(X_i = 1)} \quad (4).$$

An ISP will cooperate when δ >1. Whether or not δ is larger than 1 depends on the decisions of others (X_j) and the distribution of the weight parameters (α_j) of the ones who decide to cooperate.

7.2.3 CRITICAL MASS FOR THE COOPERATION

From (1.c), *q* is determined by the weights of the ISPs that decide to cooperate. If ISPs who decides to cooperate have diverse weights, *q* varies for each combination of ISPs who cooperate. To simplify the discussion, the chapter discusses a strategy that an ISP can choose with whom to cooperate based on the rank of the weights on their decisions. The

purpose of this strategy is to find the minimal number of ISPs that cooperate such that $\delta > 1$.

Let α_j be an ordered sequence of the weight on ISPs' decisions. "Top *n* ISP" refers to the weights of the *n* ISPs' decisions ranked top *n* in the ordered sequence. In this case, *q* can be interpreted as a cumulative value of the top *n* ISPs' weights, which is formatted as $q = \sum_{j=1}^{n} [\alpha_j X_j]$. At some point of *n*, the benefit is larger than the cost for providing the service. Once the top *n* ISPs cooperate to provide DDOS defenses, other providers will decide to provide the service as well since the benefit for providing the service is larger than the cost. This situation is referred as the critical mass for cooperation.

7.3 DATA ANALYSIS

This section describes two data sets used to estimate the distribution of the weights on ISPs' decision α_j. These two data sets are the Code-Red data set (Moore 2001) and the Route-View data set (Huston 2003), which are used to estimate the distribution of the weights for cooperative attack filtering and cooperative attack detection, respectively. Both data sets are presented as the distribution across multiple Autonomous Systems (ASes). An Autonomous System (AS) is an administrative domain of the Internet. Although sometimes an AS does not directly map to one ISP, it is reasonable to use it here because the purpose is to estimate cooperation of deploying DDOS defenses across administrative domains.

For cooperative attack filtering, the same Code-Red data described in Chapter 6 is used to estimate the distribution of the weights on ISPs' decisions. This data is a distribution of the attack sources from multiple ASes. Figure 7.3 shows the cumulative

ratio of attack sources for the top *n* ASes based on the Code-Red data set. The Code-Red data is used because the decision of an AS that originates more attacks has a higher influence on the performance efficiency than the decision of an AS that originates fewer attacks. The cooperation decision of the former could assist the downstream AS in filtering attacks closer to their sources.

For cooperative attack detection, the Route-View data is used to estimate the distribution of the weights on ISPs' decisions. This data is a distribution of the reachable IP addresses from multiple ASes. . Figure 7.4 shows the cumulative ratio of reachable IP addresses for top *n* ASes. The data set is analyzed by the CIDR report (Huston 2003) based on the data collected by the Route Views project[22] in University of Oregon. The Route-View data is used because the decision of an AS that has more reachable IP addresses has a higher influence on the performance efficiency than the decision of an AS that has less reachable IP addresses. The estimation is based on the assumption that more network traffic is sent out from an AS with a larger reachable IP address range. The more legitimate traffic that an AS originates, the more legitimate packets will be marked or logged within its administrative domain if it decides to cooperate.

[22] The Route View project is funded by Sprint and CISCO systems, which collects Internet routes based on BGP routing tables. The detail description is available at http://www.antc.uoregon.edu/route-views/. The data set used here is collected on April 25, 2003. There are 15269 ASes in the routing system at this point.

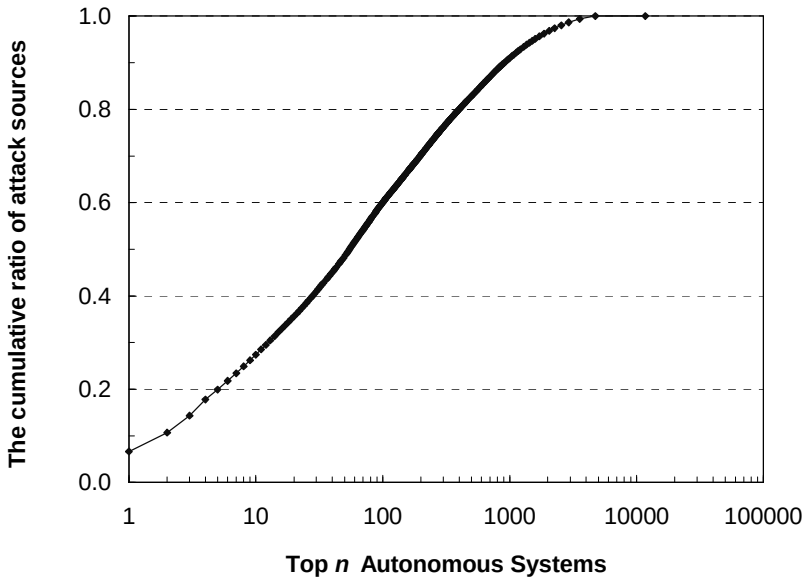

Figure 7.3: The distribution from the Code-Red data

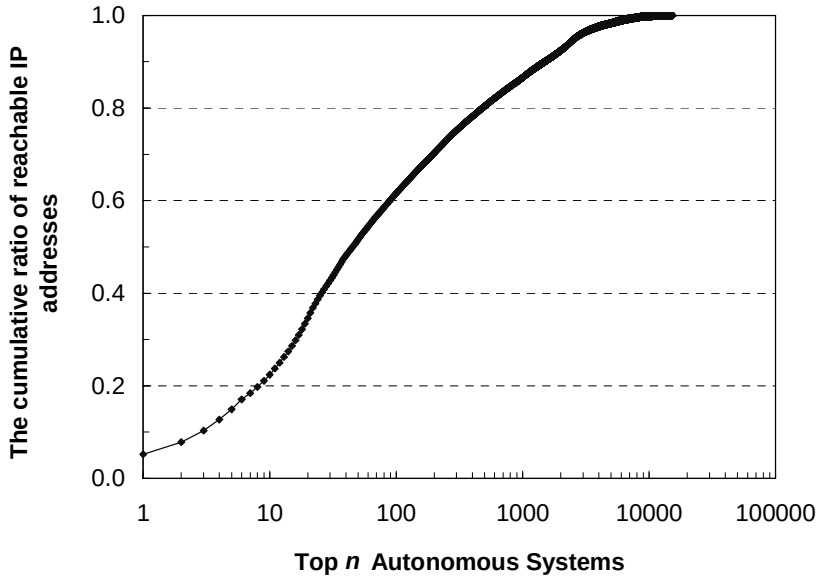

Figure 7.4: The distribution from the Route-View data

7.4 NUMERICAL RESULTS

To illustrate the analytical results in Section 7.3, this section uses the empirical data discussed in the last section to estimate the impact of ISPs' decisions about cooperation on the benefit-cost ratio. Both the flat rate pricing scheme and the differential pricing scheme discussed in Chapter 6 are analyzed here for comparison.

7.4.1 COOPERATIVE ATTACK FILTERING

The number of ISPs who decide to cooperate is a more sensitive variable when the service is provided with differential pricing rather than flat rate pricing. In Figure 7.5, the benefit-cost ratio for differential pricing increases with n while the benefit-cost ratio is a constant in flat rate pricing. Under differential pricing, the prices for the service is charged based on the utility of the subscribers, which is proportional to the quality of the defense q. q is determined by the decisions of ISPs, as described in (1.c).

The implication of this analysis is that the differential pricing scheme is needed at the initial stage of the service provision. When few providers provide DDOS defenses, these providers are able to charge subscribers for the quality of the defense. The additional benefit obtained from the improvement in the performance efficiency of the defense is an incentive for providers to cooperate on attack filtering. After more and more providers join in cooperative attack filtering, the benefit-cost ratio does not vary with the quality of the defense if the flat rate pricing is imposed in the competitive market. Nevertheless, the quality of the defense has already improved because that most providers will provide the defense.

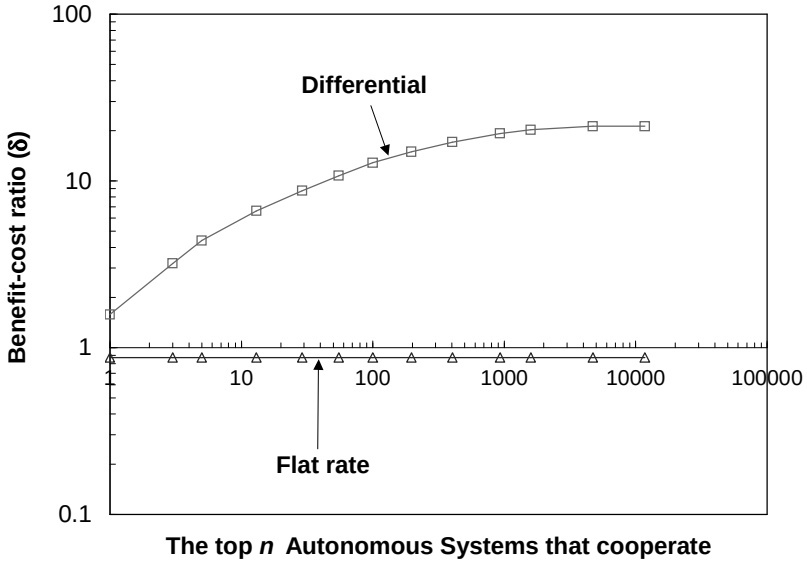

Figure 7.5: The change of benefit-cost ratio for cooperative attack filtering

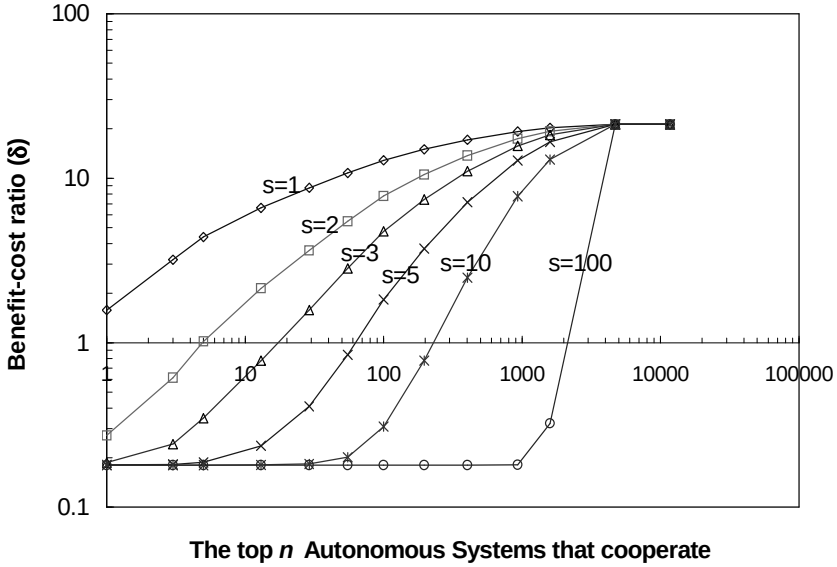

Figure 7.6: The critical mass for cooperative attack filtering

For the critical mass where the benefit-cost ratio is larger than 1, the minimal number of ASes needed increases with the scale parameter *s* under the differential pricing scheme. In Figure 7.6, only the top AS is needed when the subscribers' willingness to pay is linearly proportional to the quality of the defense. However, the proportional relationship could be nonlinear when subscribers value defenses that would provide higher performance efficiency to the legitimate clients. In this case, top 5 ASes are needed for s=2, and top 2000 ASes are needed for an extreme case that s=100.

7.4.2 COOPERATIVE ATTACK DETECTION

For cooperative attack detection, the results have three differences from the results for cooperative attack filtering:

1) The benefit-cost ratio for the flat rate pricing drops significantly when n<10. The result means that the flat rate pricing is not feasible when the majority of the ASes have not deployed the defense. Figure 7.7 shows the result.

2) The number of ASes needed to create the critical mass is more in this case at the same value of the scale parameter. As in Figure 7.8, only the top AS is needed for s=1, top 10 ASes needed for s=2, and top 4000 needed for s=100. This result occurs because the distribution (the Route-View data) used to estimate the decisions of ASes in this case is flatter. When subscribers value more on the performance efficiency to legitimate clients (such as s=2), the influence of a few ASes on cooperative attack detection is not as much as they are on cooperative attack filtering.

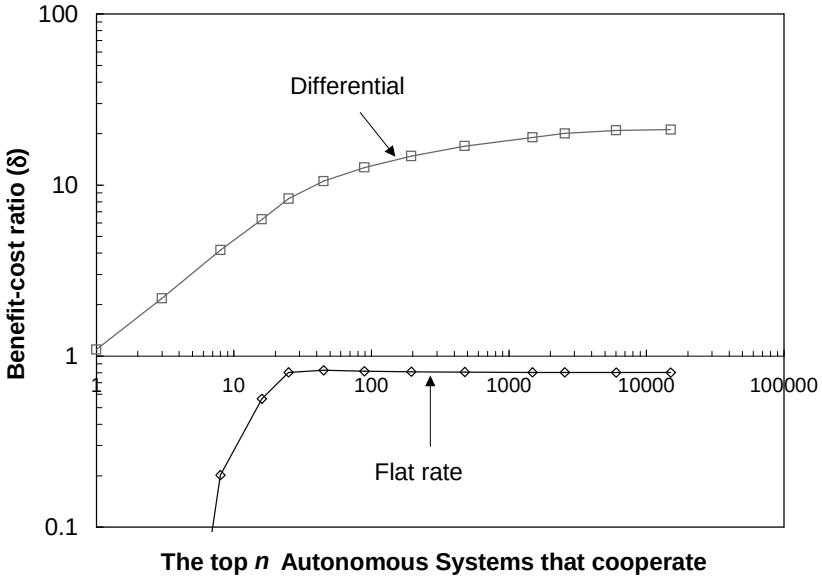

Figure 7.7: The change of benefit-cost ratio for cooperative attack detection

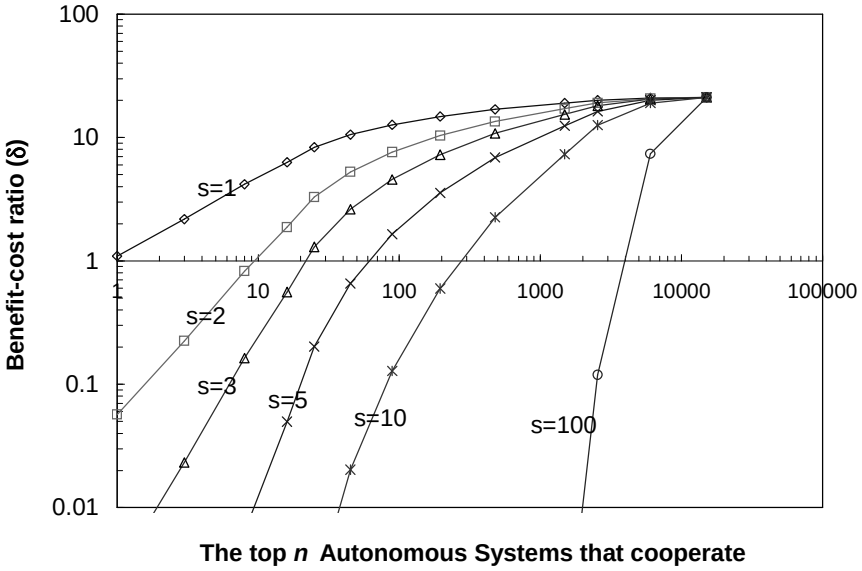

Figure 7.8: The critical mass for cooperative attack detection

7.5 PUBLIC POLICY IMPLICATIONS

Although ISPs have economic incentives to cooperate on providing DDOS defenses once the critical mass for cooperation is reached, several public policy initiatives are necessary to initially facilitate cooperation. Based on the analysis in the previous section, this section discusses the public policy implications below.

1) Widely deployment of the defenses without an appropriate strategy will make reaching the critical mass for cooperation more difficult. To generate the economic incentives for cooperation, a strategy is needed to focus on cooperation among highly influential ISPs.

For cooperative attack filtering, network providers should first cooperate with those who are more likely to originate attacks. In order to generate a large packet rate of attacks, attackers usually look for networks that host many unprotected computers connecting to "big pipes"[23], such as university campus networks. Network providers should cooperate with the providers of those networks so that more attacks can be traced and filtered. As analyzed in Section 7.4.1, to achieve the critical mass for cooperative attack filtering, only the top AS is needed for the linear case and the top 5 ASes needed for the nonlinear case.

For cooperative attack detection, network providers should cooperate with those who have the largest IP address ranges first. As analyzed in Section 7.4.2, the top AS is needed for the linear case and top 10 ASes is needed for the nonlinear case. A wide deployment of this defense is not needed for generating the economic incentives unless the subscribers would demand only the perfect solution (such as the case s=100).

[23] A high bandwidth access to the Internet is usually called a "big pipe."

2) To create an incentive for cooperation, differential pricing needs to be based on the performance efficiency of the defense.

As analyzed in Section 7.4, the incentive for improving the quality of the defense exists only when the providers can price their subscribers based on the performance efficiency. Under such circumstances, ISPs would cooperate on the deployment of DDOS defenses in order to improve the performance efficiency. The policy here should focus on encouraging the development of DDOS defenses to improve the performance efficiency,

3) Liability assignment is needed for creating incentives of source filtering. Both cooperative attack filtering and cooperative attack detection facilitate liability assignment.

As discussed in Chapter 6, liability assignment to the networks that originate attacks provides an incentive for subscribers to subscribe to source filtering. Since sources of attacks are usually forged, either cooperative attack filtering or cooperative attack detection is needed to identify the attack sources. Cooperation is a driver to facilitate liability assignment and the differential pricing scheme discussed in the previous section provides incentives for reaching the critical mass for such cooperation. Before liability assignment can take place, a policy should be developed that focuses on the creation of a critical mass for cooperation so that networks will have the incentive to identify attack sources.

7.6 CONCLUSIONS

The market mechanism is enough to sustain cooperation on deploying DDOS defenses if providers can price their subscribers based on the performance efficiency of the defenses. At the beginning of the service provision, when only a few ISPs have the technology to offer the defenses (the monopoly market), ISPs can determine the price of the service based on how well the defense can provide availability of victims to the legitimate clients. In this case, differential pricing is important to create an incentive for cooperation.

In addition, any policy for creating the critical mass of cooperation in providing DDOS defenses should focus on several highly influential ISPs. For cooperative attack filtering, these highly influential ISPs are the ones that originate the most attacks. For the cooperative attack detection, they are the ones that connect to the most legitimate clients. The cooperation of several highly influential ISPs is enough to create the critical mass for cooperation. When more and more providers are able to cooperate on the provision of DDOS defenses, the quality of the defense will improve.

Chapter 8 CONCLUSIONS

This dissertation investigates under what conditions would ISPs have economic incentives to provide DDOS defenses to their subscribers and studies the service models for providing the defenses and the public policies needed to facilitate the provision of DDOS defenses. To solve this problem, this dissertation proposes that ISPs provide defenses on their network as security services to their subscribers. Security services, such as Virtual Private Networks, have been provided by ISPs as optional network services to deal with the secrecy of data transportation. In this case, the services that provide DDOS defenses ensure the availability of the subscribers' online services. The focus will be on the DDOS defenses that actively filter out ongoing attack traffic.

This dissertation analyzes how the side effects of defenses influence the provision of the defenses and investigates the economic incentives for the service provision. The contributions of this dissertation are as follows: First, this dissertation categorizes the current defenses that actively respond against DDOS attacks at network routers (Chapter 3). The service provision model is analyzed based on the performance efficiency of DDOS defenses under various network topologies and various settings in the technology (Chapter 5). This analysis establishes a technically feasible approach for ISPs to provide the DDOS defenses as services to their subscribers. The economic incentives for ISPs to offer defense services are then analyzed based on empirical data (Chapter 6). Cooperation among

multiple ISPs on providing the defenses is analyzed in terms of two types of cooperation (Chapter7).

This chapter is organized as follows. Section 8.1 reviews the research problem in more detail. Section 8.2 discusses assumptions used throughout the dissertation. Section 8.3 describes the recommendations for subscribers, ISPs, and public policy makers on the provision of DDOS defenses to ensure a more secure infrastructure. Section 8.4 discusses the lessons learned and future research areas based on this dissertation.

8.1 PROBLEM DESCRIPTION

There have been a number of proposals on how to control ongoing DDOS attack traffic automatically at network routers. At this point, most ISPs have not adopted these automatic defenses although few ISPs, such UUNET (Stone 2000), have developed tools to trace and block DDOS traffic hop by hop manually. The lack of defense deployment is due to two reasons: 1) Defenses have side effects on legitimate traffic transportation. Since attack detection algorithms in defenses have a false positive rate of identifying attack traffic, in some occasions, legitimate traffic is regarded as the attack traffic and therefore is filtered out as attack traffic. 2) ISPs currently cannot measure and evaluate the economic incentives to provide defenses. Both developing new defense technology and deploying defenses require initial investment in the infrastructure. In the current market, the revenue that ISPs obtain from IP transport is declining due to excess supply. ISPs need economic incentives to initiate a new investment in network defenses.

The provision of DDOS defenses involves both technological and economical factors. Technically, the effectiveness of DDOS defenses depends on the false positives of

the detection algorithms, the type of network topology, the type of attacks and whether all ISPs are compliant in establishing defenses. Economically, once an ISP decides to deploy the defenses on its network, the provision of the service is influenced by the cost of the provision, the willingness to pay of the subscribers and the cooperation of interconnected ISPs. Since little is known about the interactions among these factors, the service provision model for deploying the defenses is still unclear. This dissertation studies the interactions among these factors to provide recommendations for subscribers, ISPs, and policy makers in the deployment of defenses on Internet infrastructure in the future.

8.2 ASSUMPTIONS

For a better understanding of the limitations of the analyses, several assumptions used throughout the dissertation are listed as follows. The sensitivity analyses are in individual chapters.

- The DDOS attacks saturate the network connections of subscribers to their backbone networks or take down servers inside the network of the subscribers.

- Network subscribers would pay based on the utility received from the defense. The utility that a subscriber derives from DDOS defenses is the expected value of losses that would be incurred from DDOS attacks.

- Providers would like to provide DDOS defenses to their subscribers if the operational benefit is larger than the operational cost.

- Statistical data about DDOS attacks to subscribers of ISPs are hard to obtain due to confidentiality and technical difficulty of data collection. The DDOS

data and the Code-Red data (Moore, Voelker et al. 2001)(Moore 2001) used in this dissertation are the closest approximation to the probability of attacks using publicly available data. However, using their proprietary data, ISPs can adopt the analytical model developed in this dissertation to estimate their benefit and cost of providing defense services. The network topology data is that of ISPs listed in (BW 2001), which is a simplified version of each ISP's actual network topology.

8.3 RECOMMENDATIONS

The dissertation analyzes the benefits and the costs of the stakeholders in the provisioning of DDOS defenses. The stakeholders include the subscribers that originate attacks (attack sources), the ISPs of the attacks sources (upstream ISPs), the subscribers that are victims of attacks (victims), and the ISPs of the victims (downstream ISPs). This section provides recommendations for these stakeholders as well as public policy makers based on evidence found in the dissertation.

8.3.1 RECOMMENDATIONS TO SUBSCRIBERS

Several recommendations are provided for network subscribers when considering the DDOS defenses.

1) Subscribers need to recognize the attack tolerance of their online servers in order to estimate the availability of their servers during attacks. Since none of the current defenses can filter out attack traffic without posing an impact on legitimate traffic, network providers would be able to tune the defenses based on the availability of

the servers to meet the needs of the subscribers. In particular, when the subscriber has a capacity that is larger than the packet rate of the attack traffic, maintaining a certain tolerance to attacks can minimize any additional dropping of the legitimate traffic.

2) Subscribers should provide online services that are closer to where their clients are located when DDOS defenses are implemented in order to maintain the availability of the online service to legitimate clients. For example, distributed content storage systems can provide online content closer to legitimate clients.

3) Subscribers should implement defenses on the outbound traffic of an access network. The defenses will ensure the accessibility of legitimate clients to other online services, which is better than having the victim network filter out legitimate traffic.

8.3.2 RECOMMENDATIONS TO PROVIDERS

To provide the defenses, ISPs need to consider three main issues: 1) service models for dealing with the technological uncertainty in defenses, 2) economic incentives for providing the services, and 3) incentives for cooperation with other ISPs. These issues are explained as follows.

8.3.2.1 Technological uncertainty

To provide DDOS defenses, ISPs should consider the following recommendations regarding technological uncertainty:

1) Network providers should design services that focus on adjusting the filtering rate of the attack traffic to meet the needs of different subscribers when providing defenses which are congestion-based and are dynamically enforced. The filter location and the filtering rate of attack traffic are the most sensitive variables for such defenses.

2) Network providers should design services that focus on the false positive rate of attack detection when providing defenses that are anomaly-based and are statically enforced. The false positive rate of attack detection is the most sensitive variable for such defenses.

3) In order to improve the quality of the defenses when attacks are distributed, network providers should cooperate with highly influential network providers. For attack detection, they should cooperate with administrative domains that have largest reachable source IP addresses. For attack filtering, they should cooperate with the ones that originate the most attacks. Possible incentives for cooperation include the increase in the quality of the defense service, the increase in reputation because conducting the best practice, and economic incentives for providing the services.

8.3.2.2 Economic incentives

To introduce the new service for their subscribers, network providers need to ensure that the operational profit in the long term would justify their capital investment. This dissertation has found several reasons to expect that the operational benefits will be higher

than the operational costs of the service. Here is a sequence of actions for a provider to implement the services of DDOS defenses.

First, at the initial stage when few providers are able to deploy the service (monopoly market assumption), the provider should implement a differential pricing scheme. By doing this, the provider can benefit from the different levels of expected loss experienced by subscribers, from the different levels of the attack frequency, and the different quality of defenses demanded.

Secondly, when more and more providers are able to provide the service (competitive market assumption), no single provider can benefit from differential pricing since subscribers have more choices and can switch to another provider. In this case, the providers should consider the following:

1) Providers should set the filter location closer to the attack source since it is more beneficial for both the subscribers and the providers. This result is more significant when the network of the provider is capacity constrained.

2) Providers should provide the destination filtering service for free if the fixed cost per subscribers can be recovered from the additional income from additional subscribers to network transport services in a competitive market.

3) Providers should provide source filtering when attacks are launched at high packet rates and when subscribers that originate attacks suffer losses, such as losses due to liability assignment. Offering source filtering is more beneficial than offering destination filtering since the probability of originating attacks is higher than the

probability of being attacked. This result is true even when the loss to originating networks is only 1% of the expected loss of attack victims. Source filtering is also more beneficial when the network of the provider is less connected and has a long average path length.

8.3.3 RECOMMENDATIONS TO POLICY MAKERS

The market mechanism is enough to sustain the provision of DDOS defenses. To facilitate cooperation among ISPs to reach a critical mass for providing the DDOS defense service, several recommendations are made for policy makers:

1) Policy makers should set up a program helping the industry to acquire the technologies that can detect and react against attack traffic at sources. The technologies for conducting source filtering at subscribers' network are still underdeveloped. Even though ISPs would like to provide the services to their subscribers, the technologies are not ready at this moment. For example, Ingress filtering may not be feasible in several situations (Ferguson and Senie 1998; CISCO 2003).

2) Policy makers should provide capital incentives for highly influential ISPs to deploy the defenses once new DDOS defenses are available. Capital incentives are necessary to initiate the service provision for DDOS defenses although ISPs have an economic incentive to continue to operate the services. The initiation of the services becomes important for an overall service deployment. It is in the ISPs' interest to cooperate on the provision of the services once a critical mass is created for deploying the defenses.

3) Policy makers should consider laws that assign liability to the attack sources because liability assignment creates an incentive for subscribers to reduce the attacks originating from their networks. In this case, subscribers who subscribe to source filtering should be exempted from liability, since they have conducted the best practice[24]. To whom the liability of Internet-based attacks should be assigned is an on-going debate in both academia and public policy making. In the future, if the liability is assigned to the software companies for buggy programs and if the liability assignment manages to improve the quality of software, the benefit of deploying DDOS defenses would be reduced because the risk of Internet-based attacks would be lower. However, assigning liability to software companies may not necessarily improve the quality of software. Before the debate is resolved, the dissertation proposes to assign the liability to the sources of attacks since the liability assignment is an incentive for cooperation in providing DDOS defenses.

8.4 LESSONS LEARNED AND FUTURE RESEARCH

The dissertation is fundamentally interdisciplinary and draws on work in computer security, microeconomics, and social network analysis. This approach is necessary in order to adequately understand and evaluate the impact of attacks on the critical infrastructure, in this case, the Internet.

There are a large number of possible benefits of the tool. First, the proposed service provision framework for DDOS defenses will help ISPs and subscribers to

[24] Several technical issues about conducting the best practice to prevent DDOS have been documented in IETF RFC2013 (Killalea 2000) and in (Greene, Morrow, et al. 2002).

consider the benefits of providing DDOS defenses and to recognize the tradeoffs in DDOS defenses. Secondly, the computational model developed in this dissertation provides a systematic framework for thinking through the tradeoffs in defense strategies in the complex attack-defense system. Thus, this work has direct bearing on security policy decisions at the router level for a critical infrastructure. Thirdly, the research framework provides a new method to evaluate the costs imposed by various attack scenarios and defenses since it is neither cost effective nor ethical to conduct real world experiments of DDOS attacks on a large network. Finally, this dissertation provides a theoretical basis for evaluating the provision of security service, DDOS defenses in this case.

Because the dissertation focuses on the provision of DDOS defenses, it has several limitations. First, the quantitative analysis in this dissertation provides an order of magnitude benefit and cost comparison among defenses. However, the real dollar value of the cost will depend on the implementation of these defenses. Secondly, the cost model is based on the router overhead and the bandwidth consumption costs by either attack traffic or defenses. Other implementation costs are not examined since this dissertation focuses on examining the operational benefit and the operational cost caused by defenses. Thirdly, there is a limited amount of data available for validating parameters such as the frequency of attacks in the analyses. To obtain a more precise analysis, network providers can use their own data in the models provided here. Finally, the computational model developed in this research is intended to provide decision support for tradeoffs in DDOS defenses only. This model would need further revision to analyze defenses for other types of Internet-based attacks.

In the future, changes in both technology and legislation would inevitably alter the assumptions upon which the conclusions are drawn in this dissertation. Several possible future changes are discussed as follows:

1) High implementation cost would invalidate the model. If new defense mechanisms require substantial upgrade of the network components of providers, the implementation cost would become an important variable for the analyses. A new study has to be conducted in order to evaluate the economic benefit of providing the defenses. Return on investment should be an important factor to consider the further investment in building the infrastructure for new defenses.

2) Development of distributed content delivery has a positive effect on deploying DDOS defenses. Future trends in distributed replications or caching of online services allow the delivery of web content closer to their legitimate clients. In this case, less legitimate traffic will be reduced by network filtering, as estimated in this dissertation.

3) The change of routing protocols would vary the results from the model but will not invalidate the model. If BGP, which is currently used for inter-domain routing, becomes more widely adopted as an intra-domain routing protocol, the routing paths between any two edge routers are most likely to be longer than the ones that are calculated in this dissertation using the Shortest Path First algorithm. In this case, adopting source filtering is even more beneficial for providers as shown in the dissertation.

4) Long response time of attack detection would invalidate the analyses. The model in this dissertation does not consider the response time between when an attack is launched and when the attack is detected. However, if the attack can cause severe damage to the victim before it is detected, the benefit cost analysis in this dissertation would be invalidated. In this case, more variables need to be included in estimating the economic benefit and cost of providing the service.

5) Adaptive attackers would result in more dynamic scenarios of attacks. The model in this dissertation does not consider the situation where attackers change attack sources dynamically during an attack in order to avoid filtering. The model in this dissertation would have to be revised to capture the dynamic strategy of defending attacks that avoid filtering or prevent routers from detecting and filtering attacks.

Several future research areas can be conducted based on this dissertation. First, attacks to network routers or attacks that cause the instability of global routing (Cowie, Ogielski et al. 2001) are another threat to network providers. In this case, the providers are attack victims themselves. The deployment of defenses will bring more obvious performance benefits to network providers in addition to the economic benefits mentioned in this dissertation. Secondly, liability assignment on the attack sources should be considered as a future research issue for cyber laws. Third, calibrating the probability of attacks using security incident records is important for pricing security services. Probability theory, such as extreme value theory, can be considered for calibrating the probability function of attacks instead the power curves used in this dissertation. Finally, the assessment of the utility function of subscribers is important for determining the price of DDOS defenses. Using option theory might be a good direction for future research.

REFERENCES

Arbor (2002). PeakFlow. Waltham, MA, Arbor Networks, Inc. Available at http://www.arbornetworks.com.

Asta (2002). Vantage System. Seattle, WA, Asta Networks, Inc. Available at http://www.astanetworks.com.

Axelsson, S. (2000). Intrusion detection systems: a survey and taxonomy. Goteborg, Sweden, Department of Computer Engineering, Chalmers University.

Banga, G. and P. Druschel (1997). Measuring the capacity of a Web server. USENIX Symposium on Internet Technologies and Systems, Monterrey, CA.

Banga, G., J. Mogul, et al. (1999). A scalable and explicit event delivery mechanism for UNIX. The USENIX 1999 Technical Conference, Monterey, CA.

Bellovin, S. M. (2000). ICMP traceback message, Internet Draft: draft-bellovin-itrace-00.txt.

Bhargava, H. K. and V. Choudhary (2001). "Information goods and vertical differentiation." Journal of Management 18(2): 89-106.

Bhargava, H. K., V. Choudhary, et al. (2000). "Pricing and product design: intermediary strategies in an electronic market." International Journal of Electronic Commerce 5(1): 37-56.

Burch, H. and B. Cheswick (2000). Tracing anonymous packets to their approximate source. LINUX System Administration Conference, New Orleans, LA.

BW (2001). Directory of Internet Service Providers. The Board Watch Magazine.

Cabrera, J. B. D., L. Lewis, et al. (2001). Proactive detection of distributed denial of service attacks using MIB traffic variables - a feasibility study. IEEE/IFIP International Symposium on Integrated Network Management.

Carley, K. L. and L. Gasser (1999). Computational Organization Theory. Distributed Artificial Intelligence. G. Weiss. Cambridge, MA, MIT Press.

Carley, K. M. (1996). Validating computational models. Pittsburgh, PA, Carnegie Mellon University.

Cavusoglu, H., B. Mishra, et al. (2002). The effect of Internet security breach announcements on market value of breached firms and Internet security developers. Workshop on Information Systems and Economics, Barcelona, Spain.

CERT/CC (1999). Distributed denial of service tools. Pittsburgh, PA, CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University.

CERT/CC (2000). Denial of service attacks and the federal response. Pittsburgh, PA, CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University.

CERT/CC (2002). Overview of attack trends. Pittsburgh, CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University.

Chaturvedi, A. R., M. Gupta, et al. (2000). Agent-based simulation approach to information warfare in the SEAS environment. 33rd Hawaii International Conference on System Sciences, Hawaii.

Chen, L.-C. and K. M. Carley (2003). "The impact of countermeasure propagation on the prevalence of computer viruses." IEEE Transactions on Systems, Man, and Cybernetics - Part B, to appear.

Cheswick, W. R. and S. M. Bellovin (1994). Firewalls and Internet Security: Repelling the Wily Hacker, Addison-Wesley Pub Co.

CISCO (2000). Strategies to protect against distributed denial of service, CISCO: 2000.

CISCO (2003). The IP source tracker, CISCO systems. 2003.

Claffy, K. C., G. Miller, et al. (1998). The nature of the beast: recent traffic measurements from an Internet backbone. INET, Geneva, Switzerland.

Cowie, J., A. Ogielski, et al. (2001). Global routing instabilities triggered by Code Red II and Nimda worm attacks, Renesys Corporation.

CSI (2001). CSI/FBI computer crime and security survey. Computer Security Issues & Trend. VI.

Debar, H., M. Dacier, et al. (1999). "Towards a taxonomy of intrusion detection systems." Computer Networks 31(8).

Dietrich, S., N. Long, et al. (2000). Analyzing distributed denial of service tools: the shaft case. USENIX Systems Administration Conference, New Orleans, LA.

Dijkstra, E. W. (1959). "A note on two problems in connection with graphs." Numerische Mathematik(1): 269-271.

Dittrich, D. (2001). Distributed denial of service tools. Available at http://staff.washington.edu/dittrich/misc/ddos/.

Ettredge, M. and V. J. Richardson (2002). Assessing the risk in e-commerce. Proceedings of the 35th Hawaii International Conference on System Sciences, Hawaii.

Feldmann, A. and S. Muthukrishnan (2000). Tradeoffs for packet classification. IEEE INFOCOM.

Ferguson, P. and D. Senie (1998). Network ingress filtering: defeating denial of service attacks which employ IP source address spoofing, IETF RFC2267.

Fraleigh, C., S. Moon, et al. (2001). Packet-level traffic measurement from a Tier-1 IP backbone. Burlingame, CA, Sprint ATL.

Fundenberg, D. and J. Tirole (1991). Game Theory, MIT Press.

Garber, L. (2000). "Denial-of-service attacks rip the Internet." IEEE Computer 33(4): 12-17.

Gil, T. M. and M. Poletto (2001). MULTOPS: a data-structure for bandwidth attack detection. USENIX Security Symposium, Washington, D.C.

Gordon, L. A. and M. P. Loeb (2002). "The economics of information security investment." ACM Transaction on Information and System Security 5(4): 438-457.

Greene, B. R., C. L. Morrow, et al. (2002). ISP security - real world techniques, The North American Network Operators' Group. Available at www.nanog.org.

Houle, K. J. and G. M. Weaver (2001). Trends in denial of service attack technology. Pittsburgh, CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University.

Howard, J. D. (1997). An analysis of security incidents on the Internet. Department of Engineering and Public Policy, Pittsburgh, PA, Carnegie Mellon University.

Howard, J. D. and T. A. Longstaff (1998). A common language for computer security incidents, Sandia National Laboratories.

Huang, Y. and J. M. Pullen (2001). Countering denial-of-service attacks using congestion triggered packet sampling and filtering. 10th International Conference on Computer Communications and Networks.

Huitema, C. (2000). Routing in the Internet. Upper Saddle River, NJ, Prentice-Hall.

Huston, G. (2003). The CIDR report. 2003. Available at http://www.cidr-report.org.

Iannaccone, G., C. Diot, et al. (2001). Monitoring very high speed links. ACM Internet Measurement Workshop, San Francisco.

InfoSec (2001). Information security industry survey, Information Security Magazine.

Ioannidis, J. and S. M. Bellovin (2002). Implementing pushback: router defense against DDoS attacks. Network and Distributed System Security Symposium.

Kabay, M. E. (2001). "Distributed denial-of-service attacks, contributory negligence and downstream liability." ACM Ubiquity.

Kent, S. and R. Atkinson (1998a). Security architecture for the Internet protocol, The IP Security Protocol Working Group , Internet Engineering Task Force.

Kent, S. and R. Atkinson (1998b). IP authentication header, The IP Security Protocol Working Group , Internet Engineering Task Force.

Killalea, T. (2000). Recommended Internet service provider security services and procedures, Internet Engineering Task Force.

Li, J., J. Mirkovic, et al. (2001). SAVE: Source Address Validity Enforcement protocol. IEEE INFOCOM.

Lipson, H. (2002). Tracking and tracing cyber-attacks: technical challenges and global policy issues. Pittsburgh, CERT Coordination Center, Software Engineering Institute. Available at http://www.cert.org/archive/pdf/02sr009.pdf.

Mahajan, R., S. M. Bellovin, et al. (2001). "Controlling high bandwidth aggregate in the network." Computer Communications Review.

Marwell, G. and P. E. Oliver (1993). The critical mass in collective action: a micro-social theory. New York, Cambridge University Press.

McCreary, S. and K. C. Claffy (2000). Trends in wide area IP traffic patterns: a view from Ames Internet Exchange. ITC Specialist Seminar, Monterey, CA.

Medina, A., I. Matta, et al. (2000). "On the origin of power laws in Internet topologies." ACM SIGCOMM Computer Communication Review.

Mirkovic, J., J. Martin, et al. (2002). A taxonomy of DDoS attacks and DDoS defense mechanisms. Los Angeles, Computer Science Department, University of California.

Mirkovic, J., G. Prier, et al. (2002). Attacking DDoS at the source. Proceedings of ICNP, Paris, France.

Moitra, S. D. and S. L. Konda (2000). A simulation model for managing survivability of networked information systems. Pittsburgh, Software Engineering Institute.

Moore, D. (2001). The spread of the Code-Red Worm (CRv2). 2001. Available at http://www.caida.org/analysis/security/code-red/.

Moore, D., G. M. Voelker, et al. (2001). Inferring Internet denial-of-service activity.

USENIX Security Symposium, Washington DC.

Mukherjee, B., L. T. Heberlein, et al. (1994). "Network intrusion detection." IEEE Network 8(3): 26-41.

Nickolson, W. (1995). Microeconomic Theory: Basic Principles and Extensions.

Oliver, P., G. Marwell, et al. (1985). "A theory of the critical mass. I. Interdependence, group heterogeneity, and the production of collective action." American Journal of Sociology 91(3): 522-556.

Pai, V., P. Druschel, et al. (1999). Flash: An efficient and portable Web server. The USENIX 1999 Annual Technical Conference, Monterey, CA.

Papadimitriou, C. H. and K. Steiglitz (1982). Combinatorial Optimization. Englewood Cliffs, NJ, Prentice-Hall.

Papagiannaki, K., S. Moon, et al. (2002). Analysis of measured single-Hop delay from an operational backbone network. IEEE INFOCOM, New York.

Park, K. and H. Lee (2001a). On the effectiveness of probabilistic packet marking for IP traceback under denial of service attack. Proceedings of IEEE INFOCOM.

Park, K. and H. Lee (2001b). On the effectiveness of route-Based packet filtering for distributed DoS attack prevention in power-Law Internet. ACM SIGCOMM'01, San Diego, CA, Department of Computer Science, Purdue University.

Paxson, V. (1996). End-to-end routing behavior in the Internet. ACM SIGCOMM '96.

PCIPB (2003). The national strategy to secure cyberspace, The President's Critical Infrastructure Protection Board.

Peterson, L. L. and B. S. Davie (1996). Computer Networks: A System Approach. San Francisco, CA, Morgan Kaufmann Publishers.

Recourse (2002). ManHunt. Redwood City, CA, Recourse Technologies, Inc. Available at http://www.recourse.com.

Rekhter, Y. and T. Li (1998). A Border Gateway Protocol 4 (BGP-4), Internet Engineering Task Force.

Sager, G. (1998). Security Fun with OCxmon and cflowd. Internet2 Working Group meeting.

Samuelson, P. A. (1954). "The pure theory of public expenditure." Review of Economics and Statistics 36: 387-389.

SANS (2000). Egress filtering v 0.2, SANS Institute. Available at http://www.sans.org/y2k/egress.htm.

Savage, S., D. Wetherall, et al. (2001). "Practical network support for IP traceback." ACM/IEEE Transactions on Networking 9(3): 226-237.

Schnackenberg, D. and K. Djahandari (2000). Infrastructure for intrusion detection and response. DARPA Information Survivability Conference and Exposition (DISCEX).

Schuba, C. L., I. V. Krsul, et al. (1997). Analysis of a denial of service attack on TCP. IEEE Symposium on Security and Privacy.

Snoeren, A. C., C. Partridge, et al. (2001). Hash-based IP traceback. ACM SIGCOMM.

Snoeren, A. C., C. Patridge, et al. (2002). "Single-packet IP traceback." IEEE Transaction on Networking 10(6): 721-734.

Song, D. X. and A. Perrig (2001). Advanced and authenticated marking schemes for IP traceback. IEEE INFOCOM.

Spatscheck, O. and L. L. Peterson (1998). "Defending against denial of service in Scout." Operating Systems Review (Winter).

Staniford-Chen, S., B. Tung, et al. (1998). The Common Intrusion Detection Framework (CIDF). DARPA Information Survivability Workshop, Orlando FL.

Sterne, D., D. Schnackenberg, et al. (2001). Autonomic response to distributed denial of service attacks. Recent Advances in Intrusion Detection conference.

Sterne, D., D. Schnackenberg, et al. (2002). Active network based DDoS defense. DARPA Active Networks Conference and Exposition.

Stone, R. (2000). CenterTrack: An IP Overlay Network for Tracking DoS. USENIX Security Symposium, Denver, CO.

Sung, M. and J. Xu (2002). IP traceback-based intelligent packet filtering: a novel technique for detecting against Internet DDos attacks. Proceedings of the IEEE International Conference on Network Protocols.

Tran, K. T. L. (2000). Yahoo! portal is shutdown by web attack. Wall Street Journal: 6.

Tran, K. T. L. (2000). Hackers attack major Internet sites, temporarily shutting Buy.com, Ebay. Wall Street Journal: 3.

Varian, H. R. (2002). System reliability and free riding. Workshop on Economics and Information Security, Berkeley, CA.

Waldbusser, S. (2000). Remote network monitoring management information base, IETF RFC 2819.

Wasserman, S. and K. Faust (1994). Social Network Analysis: Methods and Applications. Cambridge, Cambridge University Press.

Watts, D. J. and S. H. Strogatz (1998). "Collective dynamics of 'small-world' networks." Nature 393(4).

Welsh, M., D. Culler, et al. (2001). SEDA: An Architecture for Well-Conditioned, Scalable Internet Services,. The Eighteenth Symposium on Operating Systems Principles (SOSP-18), Banff, Canada.

Wood, B. J. and R. A. Duggan (1999). Red teaming of advanced information assurance concepts. DARPA Information Survivability Conference & Exposition.

Xiong, Y., S. Liu, et al. (2001). "On the defense of the distributed denial of service attacks: an on-off feedback control approach." IEEE Transaction on Systems, Man, and Cybernetics - Part A: Systems and Humans 31(4): 282-293.

Xu, J. and W. Lee (2003). "Sustaining availability of web servers under server denial of service attacks." IEEE Transaction on Computers, special issue on Reliable Distributed Systems 25(2): 195-207.

Yaar, A., A. Perrig, et al. (2003). Pi: A path identification mechanism to defend against DDos attack. IEEE conference on security and privacy.

Yan, J., S. Early, et al. (2000). The XenoService - a distributed defeat for distributed denial of service. Information Survivability Workshop.

Yankee (2000). $1.2 billion impact seen as a result of recent attacks launched by Internet hackers, The Yankee Group.

Zwicky, E. D., S. Cooper, et al. (2000). Building Internet Firewalls, O'Reilly & Associates.