MODELING DISTRIBUTED DENIAL OF SERVICE ATTACKS AND DEFENSES

Li-Chiou Chen
Kathleen M. Carley

PROJECT SUMMARY

February 2002

Distributed denial of service (DDOS) attacks have emerged as a prevalent way to take down web sites and have imposed financial losses to companies. The CSI/FBI survey (CSI 2001) shows that 36% of respondents in the last 12-months period have detected denial of service, which imposed more than $4.2 million financial losses. The effectiveness of DDOS defenses depends on many factors such that the nature of the network's topology, the specific attack scenario, and various characteristics of the network routers. However, little research has focused on the tradeoffs inherent in this complex system. We propose to develop a computational testbed to study security policies and the associated technologies that provide defenses against DDOS attacks. We will then use this framework to evaluate various policies and technologies. In this work we draw on research in the areas of computer science, information science, organizational theory and social networks.

There have been a number of proposals on how to control the on-going DDOS attack traffic. None have been widely deployed. The effectiveness of DDOS defenses depends on many factors, such as the type of network topology, the type of attacks and whether all ISPs are compliant in establishing defenses. However, little is known about the interactions among these factors. Knowing what tradeoffs will occur as these factors vary will assist stakeholders in making security policy decisions and adjusting for the chance that others may not make the same decisions. The research proposed in this project will illuminate these tradeoffs and lead to a computational model for examining various DDOS defenses and attack scenarios at the router level.

To orient this research we focus on two basic research questions. First, how do ISPs provide DDOS defenses at the lowest cost while their subscribers remain satisfied with the availability of network connections during attacks? A cost-performance analysis of the effectiveness of DDOS defenses will be conducted using results from the computational model. This cost-performance analysis will aid ISPs and local network administrators in their evaluation of DDOS defenses. Second, we ask where are the critical points in a network to deploy defenses? We examine the impact of network topology on the deployment location of defenses. Graph level indices and models from social network studies will be used to categorize network topologies and to select deployment locations for defenses. This analysis will provide guidance to decision makers.

Benefits of the proposed research include the following. First, the policy framework proposed in this research will help ISPs and subscribers to consider the benefits of providing DDOS defenses and to realize the tradeoffs in DDOS defenses. Results from this study will aid decision makers in setting security policy for computer networks. Thirdly, since it is costly and unethical to conduct real world experiments of DDOS attacks on a large network, this research provides a testbed to evaluate the costs imposed by various attack scenarios and defenses. Moreover, topological measures developed in this research could be useful for studies of other large-scale topologies. This will extend social network measures typically used on small person-to-person networks to large scale computer networks. Finally, this research will provide a theoretical basis for evaluating DDOS defenses building on interdisciplinary studies from the fields of computer science, information science, organizational theory and social network analysis.

I. Introduction

Computer-based attacks on critical infrastructure have become a great concern to the US government as the nation has become increasingly reliant upon the Internet to exchange information among various systems (PCCIP 1997). To mitigate the risk of computer security incidents, evaluating the effectiveness of defenses becomes an important issue. The purpose of this research is to propose and utilize a framework to study security policies that will provide defenses against distributed denial of service (DDOS) attacks. We take it as axiomatic that the solution is not just technical but involves the setting of policies and meeting the needs of diverse users with different priorities.

Distributed denial of service (DDOS) attacks have emerged as a prevalent way to take down web sites and have imposed financial losses to companies. In the DDOS attacks to Yahoo, eBay and Amazon.com in February 2000, the attacker generated 1GB of requests per second flooding the web sites of these companies (Tran 2000). The Yankee Group estimates that the financial losses imposed by these attacks on these companies are more than $1billion (Yankee 2000). The CSI/FBI survey (CSI 2001) shows that 36% of respondents in the last 12-months period have detected denial of service, which imposed more than $4.2 million financial losses. In addition, the scale of DDOS attacks is increasing. Code-Red worm attacks in August 2001 also highlight the potential risk of large-scale distributed denial-of-service attacks from wide spread sources (Moore 2001). An empirical study of distributed denial of service 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[1] have been proposed to control the on-going DDOS attack traffic (Ferguson and Senie 1998; Stone 2000; Mahajan, Bellovin et al. 2001; Park and Lee 2001). The effectiveness of DDOS defenses depends on many factors such that the nature of the network's topology, the specific attack scenario, and various aspects of the network routers. However, little research has focused on the tradeoffs inherent in this complex system. Understanding the nature and severity of these tradeoffs will assist stakeholders[2] in making security policy decisions even while they are considering defenses against these attacks. The proposed research will increase our understanding of these tradeoffs and derive insights that will enable a more secure infrastructure.

We propose to develop and utilize a computational model for evaluating tradeoffs between various factors in DDOS defenses at the router level under diverse attack scenarios and network topologies. This research will focus on two research questions. First, how do Internet Service Providers (ISPs) provide DDOS defenses at the lowest cost while their subscribers remain satisfied with the availability of network connections during attacks? A cost-performance analysis of the effectiveness of DDOS defenses will be conducted using results from the computational model. This cost-performance analysis will aid ISPs and local network administrators in their evaluation of DDOS defenses. Second, we ask where are the critical points in a network to deploy defenses? We examine the impact of network topology on defense location. Graph level indices and models from social network studies will be used to categorize network topologies and to select deployment locations for defenses. This analysis will provide guidance to decision makers.

This project description is divided into five sections. Next the background information and research questions are described. In section 3 a policy framework for providing DDOS defenses is described. In section 4 the computational model, virtual experiments and validation approach is laid out. In section 5 the benefits and limitations of the proposed research are discussed.

II. Background

This section will review previous models that describe the complex systems of Internet security incidents, and explains research questions that we will investigate.

2.1 Previous models

Computational modeling approach has been used to model Internet security incidents, in particular, to model attack scenarios and victim responses. Cohen's model (Cohen 1999) simulates attack processes

[1] Defenses refer to methods that can mitigate the effect of attacks.

[2] The stakeholders involved include backbone ISPs, regional ISPs, attack victims, administrators of other Internet hosts and government decision makers.

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 estimate their reaction time to a known attack scenario. 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, the 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, the probability if 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. Gupta et al. (Gupta, Chaturvedi, 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 to cyber attacks. Red Teaming is a computer security attack simulation project developed in Sandia National Laboratories Information Design Assurance Red Team (IDART). To verify its models, Red Teaming collects attacker behavior information and identifies vulnerabilities of information systems by using human experts to attack real information systems (Wood and Duggan 1999).

Previous models show that computational modeling could be a powerful approach in research regarding Internet security incident as the problem involves tradeoffs in a complex system, where conducting real world experiments is difficult, costly and in some cases, unethical. However, previous models enable only a limited investigation of security policy problems associated with the DDOS attacks as they do not focus on the router level of computer networks. We build on these earlier approaches, particularly the work on risk tolerance and the cost of defense and extend them by enabling the analysis relative to specific topologies.

2.2 Problem description

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; Howard and Longstaff 1998). A distributed denial of service (DDOS) attack is usually launched from multiple hosts on the Internet to saturate the bandwidth of victims' network connections. In a DDOS 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).

A number of defenses against DDOS attacks have been proposed. These are reviewed in Section III. Typically, each defense is designed to solve a different part of the DDOS attack problem. This research will focus on evaluating the defenses that can be automatically triggered at network routers to reduce the on-going DDOS attack traffic.

To deploy defenses against DDOS attacks at network routers, ISPs need to configure routers in order to prevent attack traffic from reaching the network connections 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 normal traffic to victims, a defense may mistakenly regard the normal traffic as the attack traffic. It is uncertain that how efficient these defenses are in terms of maintaining the network connections available to the normal traffic of victims and non-victims while these defenses control the attack traffic. Secondly, the overhead imposed by these defenses on routers is uncertain. Thirdly, none of the defenses has provided a mechanism for ISPs to know the preferences of their subscribers in selecting a defense and negotiating parameters in a defense if in fact a tradeoff occurs. Nevertheless, in this complex system such tradeoffs are inevitable.

A key problem then is "How good of a defense really exists at the global level if only some of the ISPs are compliant and install defenses?" As will be discussed the answer is likely to depend on the deployment strategy, the technology available, the attack scenario, the underlying network topology, and the compliance of other ISPs. Whether ISPs will be compliant in art depends on the cost-performance tradeoff. However, this tradeoff is also affected by the topology, attack and defense strategies. What is

needed to address this problem is a framework for systematically evaluating the interactions in this complex socio-technical-policy system.

The proposed research will provide such a framework for ISPs to determine what information they will need from their subscribers to deploy DDOS defenses, meet the demands of their subscribers, and choose the appropriate tradeoff in building a defense. As part of this research we will develop a computational model for systematically investigating and characterizing these tradeoffs given different defenses, attack scenarios and network topologies.

The computational model will then be used to investigate two questions. First, how do ISPs provide DDOS defenses at the lowest cost while their subscribers remain satisfied with the availability of network connections during attacks? What is the cost performance tradeoff? What are the security policies for implementing defenses at the network router[3] level? Security policies here refer to the decisions or rules regarding initialization, configuration and termination of defenses. A set of defenses at the router level will be examined to identify the set of parameters that can be negotiated between ISPs and subscribers. Given these parameters, a cost-performance analysis on various types of DDOS defenses will be conducted using the computational model. Virtual experiments that vary all key parameters will be conducted to provide an overview of the impact of different decisions.

Second, we ask where are the critical points in a network to deploy defenses? Essentially, we will be examining the impact of different computer network topologies, and for each of these identifying the characteristics of critical points for deploying defenses. Related questions include, how does a router propagate security policies to other routers (what network paths are followed) so that it can defend against a new attack? What topological factors will influence the effectiveness of defenses or the cost imposed by attacks and defenses? The research will analyze the impact of network topology on defense deployment. Graph level indices and measures from social network studies will be used to identify deployment location in this analysis. This analysis will help decision makers identify those critical points in a network where defenses are needed.

III. The Policy Framework of DDOS Defense Provision

The basic policy framework (Figure 1) guides the decision maker in determining the cost effectiveness of various defense strategies. In this section we explain how the cost model of attacks and defenses can be useful for ISPs and their subscribers. We have developed a preliminary version of the cost model that describes the cost imposed by attacks and defenses on a network. The proposed research will further extend and refine this model taking specific aspects of the underlying technologies and actual networks into account. In the proposed research the cost model will take into account the nature of the attack and the specific deployed defenses. This provides a more accurate assessment of costs. We now discuss the characteristics of different attack scenarios, defenses, and network topology present in the literature.

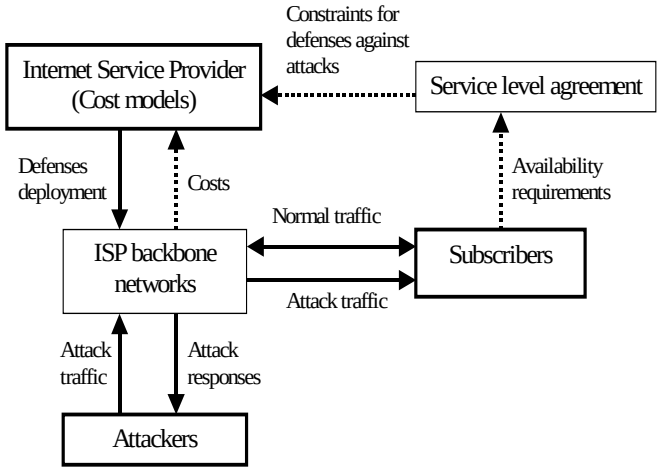

Figure 1: Context for DDOS defense provision

[3] Network routers could be Internet autonomous system (AS) inter-domain border routers or intra-domain routers.

3.1 The policy framework of providing defenses against DDOS

In figure 1 a policy framework that describes the context in which DDOS defenses are deployed is shown. An ISP can provide DDOS defense service 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 to 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.

However, 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 make this system ineffective. We intend to use a computational model of this system to evaluate these tradeoffs and the underlying parameters and so enable more effective definitions of DDOS defenses in the SLA. For example, each defense is only effective against a certain attack scenario and given a particular network topology. Consequently, the cost model of a network for an ISP and the availability requirements from subscribers should be different. This research will analyze the cost imposed on a network by different known and expected attack scenarios and defenses. In addition, this research will provide a systematic approach for analyzing future attacks given future defenses (both in terms of technology and network topology).

3.2 Characteristics of DDOS attacks

Many tools have been used to launch DDOS attacks, such as Trinoo, TFN, Stacheldraht, and Mstream (Dietrich, Long et al. 2000; Dittrich 2001). These tools have several common characteristics. 1. These tools usually have options to control the rate of generating attack packets. 2. The same tool can be used to conduct different types of attacks, including TCP SYN, UDP, and ICMP flood. 3. Variants of the attack tools are created based on the same exploit methods used to avoid detection of a specific attack signature. 4. Most of these tools could forge source IP addresses in attack packets. In all cases, the nature of the attack can be characterized in terms of the sources, severity (number of victims), and impact (compromised protocols and packet rates).

This research will generate simulated attack scenarios based on four key variables: the sources of attacks, the number of victims, compromised protocols and packet rates. In all cases, we will use known and commonly expected attacks to inform our choice of values for these variables. The sources of attacks will be randomly distributed among network domains, from one single domain and from several major domains. The sources of attacks are determined by the propagation methods of the attack tools. The number of victims will depend on the intention of the attackers – we will examine at least three levels of severity low, medium and high. Compromised protocols are network protocol types of packets that a compromised computer has sent out. Packet rates are the amount of network traffic sent out by attack tools during a certain period of time. Packet rates are determined by both the bandwidth of the network links that connect to attack sources and the performance of computers compromised by attack tools.

The sources of attacks are influenced by the propagation methods of attack tools. Attack tools could be installed manually by attackers or automatically by propagation mechanisms that are coded in attack tools. For example, since Code-Red worm is a self-propagated computer virus that is coded to conduct DDOS attacks (Moore, Voelker et al. 2001), it can generate DDOS traffic from more widely spread sources than tools installed manually by attackers can. We will generate different distributions of attack sources based on the propagation methods of attack tools. Automatic propagation methods of computer viruses have been studied (Chen and Sirbu 2000; Chen and Carley 2001), thus providing preliminary work for this research.

3.3 Characteristics of defenses against DDOS

In response to on-going attacks, a variety of defenses have been proposed. Since the purpose of this research is to provide the policy framework for ISPs to deploy DDOS defenses on their network routers, the defenses evaluated in the later simulations will be focused on network-based defenses that are

designed to actively reduce the amount of on-going attack traffic. To provide an overview of current solutions to DDOS attacks, we summarize the characteristics of all DDOS defenses as the following.

Reaction points: network-based vs. host-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 targets. Network-based methods are deployed on the points where packets route through network connections, 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 2001; Ioannidis and Bellovin 2002). Host-based defenses are deployed on the machines that are potential targets of attacks. These methods (Spatscheck and Peterson 1998; Yan, Early et al. 2000) could increase the victims ability to tolerate attacks but not stop them. We will model both types of reaction points.

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 2001; Snoeren, Partridge et al. 2001; Song and Perrig 2001). These methods could facilitate future liability assignments since most source IP addresses of attack packets can be forged. These methods are for identifying the sources of attacks; not for stopping on-going attack traffic. In contrast, other defenses are designed to actively reduce the amount of on-going attack traffic (Ferguson and Senie 1998; Mahajan, Bellovin et al. 2001; Park and Lee 2001; Ioannidis and Bellovin 2002). We will focus on the later one: reduction in on going attack traffic.

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 to the number of packets to be examined / logged (Savage, Wetherall et al. 2000; Huang and Pullen 2001). We will model two types of sampling: probabilistic and check everything. Further, the virtual experiments will vary the rate of sampling to determine its relative efficacy compared to other factors such as the type of defense.

Reaction timing: constant vs. event-triggered

Some defenses needed to be turn on all the time in order to detect the 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 an attack response. A few defenses trigger attack response 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). We will model both constant and event triggered (high congestion) response. Further, in the virtual experiments the level of congestion needed to trigger a response will be varied.

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

It is hard to distinguish attack packets from normal 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[4] associated with the outcome. Moreover, some criteria are only effective at identifying a certain type 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). TCP SYN packet flood can be identified by a state machine (Schuba, Krsul et al. 1997). Attack packets with spoofed source IP addresses can be identified given knowledge of the networks topology (Park and Lee 2001). In our simulations we will characterize the detection criteria in terms of the likelihood of false positives and false negatives rather than trying to model the wide range of exact characteristic of specific detection routines.

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

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 but will not be effective when the bandwidth of the subscriber's network is saturated. The pushback method triggers filters along the path that 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 defense can be gradually deployed at distributed locations across a network (Schnackenberg and Djahandari 2000; Park and Lee 2001; 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 this case, we will model three location strategies for deployment: single point, attack path and distributed points. In the case of distributed points we will run a sensitivity analysis across the number of points as a function of the size and density of the underlying network.

3.4 Characteristics of network topology

Previous studies have confirmed that topology is a dominant factor influencing network routing (Paxson 1996; Govindan and Reddy 1997; Medina, Matta et al. 2000; Palmer, Siganos et al. 2001) and attack tolerance (Faloutsos, Faloutsos et al. 1999; Albert, Jeong et al. 2000). However, it remains uncertain to what extent network topology could influence DDOS attacks and the deployment locations of defenses. Topological factors that have been shown in other contexts to be key include: the size of the network (Anderson, Butts and Carley, 1999), the density of the network (number of connections / number of possible connections) (Anderson, Butts and Carley, 1999), degree of clustering, degree of hierarchy, and average path length (Carley, Lee and Krackhardt, forthcoming).

Using our computational model we will evaluate deployment locations of DDOS defenses by varying the type of network topology. Graph level indices and measures from social network analysis (Scott 1991; Wasserman and Faust 1994) will be used to characterize the type of network topology. A variety of measures from social network analysis are particularly useful for locating critical points (Anderson, Butts et al. 1999), such as "degree centrality" (Freeman 1979) and "betweenness centrality" (Bonacich 1987; Freeman, Borgatti et al. 1991). Another set of measures are critical for measuring rates of flows through nodes in a network (Freeman, Borgatti et al. 1991). Still other measures are used for hierarchy and clustering.

This research plans to uses the following measures to categorize the network topologies:

1. Size: The number of nodes. We examine size to determine whether response strategies scale.
2. Density: The number of links relative to the number of possible links. We examine size to determine whether response strategies scale.
3. Node centrality: In order to determine the minimum number of nodes needed to deploy defenses on Internet topology, we will identify central nodes that control the flow of information through the network topology (Freeman, Borgatti et al. 1991; Wasserman and Faust 1994). As noted, defenses can be deployed to these critical points instead of to every single point in the network topology. We will examine whether such selected deployment is more cost effective, and whether degree, betweenness or information centrality is the key parameter for determining node criticality. We will examine node centrality both from a global (entire network) and a local (sub-graph) perspective.
4. Clustering: An important factor influencing the speed with which information flows through the network and its robustness in the face of errors is the extent to which is composed of sub-clusters (Carley, Lee and Krackhardt, forthcoming). We will first cluster Internet AS topologies into sub-graphs based on node connectivity. Graph level indices can be calculated later for each sub-graph. Given a global network we will use both structural equivalence (Concor, Wasserman and Faust 1994) and clique formation to cluster the nodes. We use these techniques as the focus is on whether nodes actually link to each other rather than do they happen to access the same information. For networks we generate with the computational model we will vary the degree of clustering (number of links within clique compared to the number between cliques) and the number of clusters.

5. Hierarchy: The extent to which there are cycles in the overall topology. Drawing on the work in organization theory we note that the more hierarchical the network the easier it may be to defend against attack but the slower it might respond. To determine the relevance of this tradeoff for Internet AS topologies we will vary the degree of hierarchy in these networks. Various measures of hierarchy will be used including Krackhardt's graph hierarchy and Freeman's hierarchy.
6. Future changes in Internet topology: One of the key features of the Internet AS topology is that it is constantly changing. Changes occur almost daily in both the size and degree of connectivity. We will use historical empirical data on Internet AS topology and use this to extrapolate a pattern and rate of change. Using this information we will allow the networks in the computational model to evolve over time in a comparable fashion. This will enable us to examine the extent to which the various defense strategies will degrade or respond gracefully to changes in the underlying topology. We note that a defense strategy that cannot withstand topology changes is very brittle and of less long-term value.

3.5 The cost model of DDOS attacks and defenses

Both the attack traffic and the defenses impose costs on a network. We have developed a cost model and will use it to quantify these additional costs during an attack for a network. The total cost is the summation of the transport cost that is used to transmit DDOS attack packets, referred as the link cost, and the computational cost of routers that is impose by defenses to examine packets, referred as the node cost. Measuring cost will make it possible to examine tradeoffs for defense strategies in terms of cost, availability requirements of subscribers, and the overall brittleness of the system. In some cases, ISPs need to make decisions on the tradeoffs between costs, requirements of subscribers, and long-term benefits. Subscribers, e.g., in choosing how much they are willing to spend on defense need to make decisions about the tradeoffs between availability of network connections to normal traffic and the tolerance to attack traffic. For example, if the location of a defense is close to the attack sources, the transport cost of attack packets could be lower. However, attack sources of an on-going attack are usually hard to determine. Instead, if the location of a defense is close to victims, the transport cost is higher since attack packets have been traveled along the network. In addition, the number of deployment points needed to trigger a defense could be lower if the defense is placed at several central points of the network. However, these central points could be backbone routers that process a high volume of traffic, the computational cost of routers to examine packets will be higher. Both defenses may be equally brittle. We now describe this model in greater detail.

Notations

Consider a network $G = (S, E)$ with vertices S and edges E.

Attacks:

Attack nodes are edge routers of an ISP. Attack nodes connect to the networks of subscribers that launch DDOS attack traffic. $A \subset S$ denotes to the set of attack nodes.

Defenses:

Filter nodes are routers that are deployed a defense that will filter out attack packets. $F \subset S$ denotes to the set of filter nodes. $P \subset S$ denotes the set of nodes on which the traffic to victim nodes has to go through any one of the filter nodes. $Q \subset S$ denotes the set of nodes on which the traffic to victim nodes has not to go through any one of the filter nodes. $P \cap Q = \phi$.

Victims:

Victim nodes are edge routers of an ISP. Victim nodes connect to the networks of subscribers that are attacked by the DDOS attack traffic. A victim network refers to the connections of the network that are saturated by the attack traffic. A victim machine refers to the IP address of the computer is the destination IP address targeted by the attack traffic. A non-victim machine refers to the computer that is in the victim network but is not targeted by attack traffic. $V \subset S$ denotes the set of victim nodes.

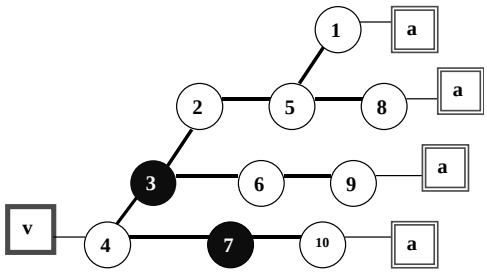

Figure 2. An Example Network

Figure 2 shows an example of a network. Circles are routers of a network. S = {1, 2, 3, 4, 5, 6, 7, 8, 9, 10}. Squares with an "a" are local networks from where attackers launch DDOS attacks. The set of attacker nodes, A, = {1, 8, 9, 10}. The square with a "v" is the victim's network. The set of victim nodes, V, = {4}. Suppose that a defense to filter out DDOS traffic is placed at routers 3 and 7. The set of filter nodes, F, = {3, 7}. P = {1, 2, 3, 5, 6, 7, 8, 9, 10} and Q = {4}.

Definitions

Equation (1) quantifies the link cost from attack nodes to filter nodes, which is the unit transport cost of a link multiplied by the total attack traffic that is routed through the links between attack nodes and filter nodes. C_l refers to the unit transport cost of a link. L(i,j) is the distance of the shortest path between any two distinct nodes i and j, which is measured by the number of hops between two nodes. $X_a(i,j)$ is the rate of the attack traffic from an attack node i to a victim node j.

$$LC_{af} = \sum_{i\in A, i\in P} \sum_{f\in F} [L(i,f) * C_l(i,f) * \sum_{j\in V} X_a(i,j)] \quad (1)$$

Equation (2) quantifies the link cost from filter nodes to victim nodes, which is the unit transport cost of a link multiplied by the total attack traffic that is allowed to pass through the filter nodes. $\delta_a(i)$ denotes the filtering rate of the attack traffic at node i.

$$LC_{fv} = \sum_{f\in F} \sum_{j\in V} [L(f,j) * C_l(f,j) * \sum_{i\in A, i\in P} X_a(i,j) * (1-\delta_a(i))] \quad (2)$$

Equation (3) quantifies the link cost from attack nodes to victim nodes for the attack traffic that does not go through any one of the filter nodes, which is the unit transport cost of a link multiplied by the total attack traffic that does not go through filter nodes.

$$LC_{av} = \sum_{i\in A, i\in Q} \sum_{j\in V} [L(i,j) * C_l(i,j) * X_a(i,j)] \quad (3)$$

Equation (4) is the sum of all link costs, including (1), (2) and (3).

$$LC = LC_{af} + LC_{fv} + LC_{av} \quad (4)$$

Equation (5) quantifies the node cost that is imposed by packet filters to examine packets at filter nodes, which is the unit computational cost multiplied by the total traffic that goes through filter nodes to victims. $C_n(f)$ is the unit computational cost of the router f which examines packets. $X_v(i,j)$ is the rate of normal traffic from node i to victim machines in a victim node j. $X_{nv}(i,j)$ is the rate of normal traffic from node i to non-victim machines in a victim node j. Y(f) is the rate of normal traffic at the filter node f that goes through the same links as normal traffic to victim nodes but its destination is not victim nodes.

$$NC = \sum_{f\in F} \{C_n(f) * [\sum_{i\in A, i\in p} \sum_{j\in V} X_a(i,j) + \sum_{i\in P} \sum_{j\in v} (X_v(i,j) + X_{nv}(i,j)) + Y(f)]\} \quad (5)$$

Equation (6) represents the total cost of the network G with a defense triggered at F, which is the sum of the link cost, (4), and the node cost, (5).

$$TC = LC + NC \quad (6)$$

A network operator, such as an ISP, would like to minimize TC. However, subscribers will require ISP to reduce attack traffic and to maintain a certain level of availability of their network connections during an attack. Following equations will describe these requirements as constraints of minimizing TC.

Consider a victim node j. Equation (7) calculates the total attack traffic that is sent to victim j. It is the sum of the attack traffic that goes through filter nodes but does not be filtered out, and the attack traffic that does not go through filter nodes. $\delta_a(f)$ is the filtering rate of attack traffic. 1- $\delta_a(f)$ is the false-negative rate that filter nodes mistakenly regard attack traffic as normal traffic.

$$T_a(j) = \sum_{i\in A, i\in P} \sum_{f\in F} X_a(i,j) * (1-\delta_a(f)) + \sum_{i\in A, i\in Q} X_a(i,j) \quad (7)$$

Equation (8) calculates the total normal traffic that is sent to victim machines in node j. It is the sum of the normal traffic to victim machines that goes through filter nodes and does not be filtered out, and the normal traffic to victim machines that does not go through filter nodes. $\delta_v(f)$ is the false-positive rate that filter nodes mistakenly regard that normal traffic as attack traffic.

$$T_v(j) = \sum_{i \in S, i \in P} \sum_{f \in F} X_v(i, j) * (1 - \delta_v(f)) + \sum_{i \in S, i \in Q} X_v(i, j) \quad (8)$$

Equation (9) calculates the total normal traffic that is sent to non-victim machines in node j. As the same reasons in (8), it is the sum of the normal traffic to non-victim machines that goes through filter nodes and does not be filtered out, and the normal traffic to non-victim machines that does not go through filter nodes. $\delta_{nv}(f)$ is the false-positive rate that filter nodes mistakenly regard that normal traffic as attack traffic.

$$T_{nv}(j) = \sum_{i \in S, i \in P} \sum_{f \in F} X_{nv}(i, j) * (1 - \delta_{nv}(f)) + \sum_{i \in S, i \in Q} X_{nv}(i, j) \quad (9)$$

Assume the inter-arrival time of these three types of traffic is uniformly distributed. Routers in this network maintain a drop-tail queue. The edge routers to the victim's link will drop the three types of packets equally when the link is congested. Equation (10) represents the drop rate of the victim's link due to congestion, d(j).

$$d_j = (T_a(j) + T_v(j) + T_{nv}(j) - 1)/(T_a(j) + T_v(j) + T_{nv}(j)) \quad (10)$$

Assume that j can tolerant attack traffic to a certain level. The utilization of the victim's link by attack traffic, $U_a(j)$, should have a upper bound, $U_a^{up}(j)$. C is the capacity of the victim's link.

$$U_a(j) = (1/C) * T_a(j) * (1 - d_j) \le U_a^{up}(j) \quad (11)$$

In addition, the victim will require that the throughput of normal traffic is higher than a threshold. The utilization of the victim's link by normal traffic to victim machines, $U_v(j)$, should have a lower bound, $U_v^{low}(j)$.

$$U_v(j) = (1/C) * T_v(j) * (1 - d_j) \ge U_v^{low}(j) \quad (12)$$

In the meantime, the network provider should maintain the network available to non-victim machines in the victim network. The utilization of the victim's link by non-victim machines, $U_{nv}(j)$, should have a lower bound, $U_{nv}^{low}(j)$.

$$U_{nv}(j) = (1/C) * T_{nv}(j) * (1 - d_j) \ge U_{nv}^{low}(j) \quad (13)$$

The objective is to minimize (6) that is constrained by (7)-(13). Table 1 is a listing of decision variables in this cost model.

Notation	Meaning	Factors that will determine the value of the variable
F	Filter deployment location	ISP's cost consideration, ISP backbone network topology, and technology feasibility of a defense
$\delta_a(f)$	The filtering rate at filter f for the attack traffic, $f \in F$	Technology uncertainty of a defense
$\delta_n(f)$	The filtering rate at filter f for the normal traffic to victim machines, $f \in F$	
$\delta_{nv}(f)$	The filtering rate at filter f for the normal traffic to non-victim machines, $f \in F$.	
$U_a^{up}(j)$	Utilization of the victim's link by the attack traffic	Requirements imposed by subscribers in SLA.
$U_v^{low}(j)$	Utilization of the victim's link by the normal traffic to victim machines	
$U_{nv}^{low}(j)$	Utilization of the victim's link by the normal traffic to non-victim machines	

Table 1: Decision variables in the cost model

IV. Simulation Model Development

This research will develop a computational model to investigate the two research questions that were previously described. Using this model, a number of virtual experiments will be run. We describe the two key virtual experiments. In addition to these, a number of sensitivity analyses will be run to try and estimate the characteristics of the model's response surface. Virtual experiment 1 will evaluate the cost-performance tradeoffs for the various defenses for ISPs, victims and non-victims. Using these results we will create a cost-performance mapping of the various defenses at the router level. Virtual experiment 2 will be used to locate critical points for defense deployment. Results from this experiment will inform strategy selection for defenses.

4.1 The computational model

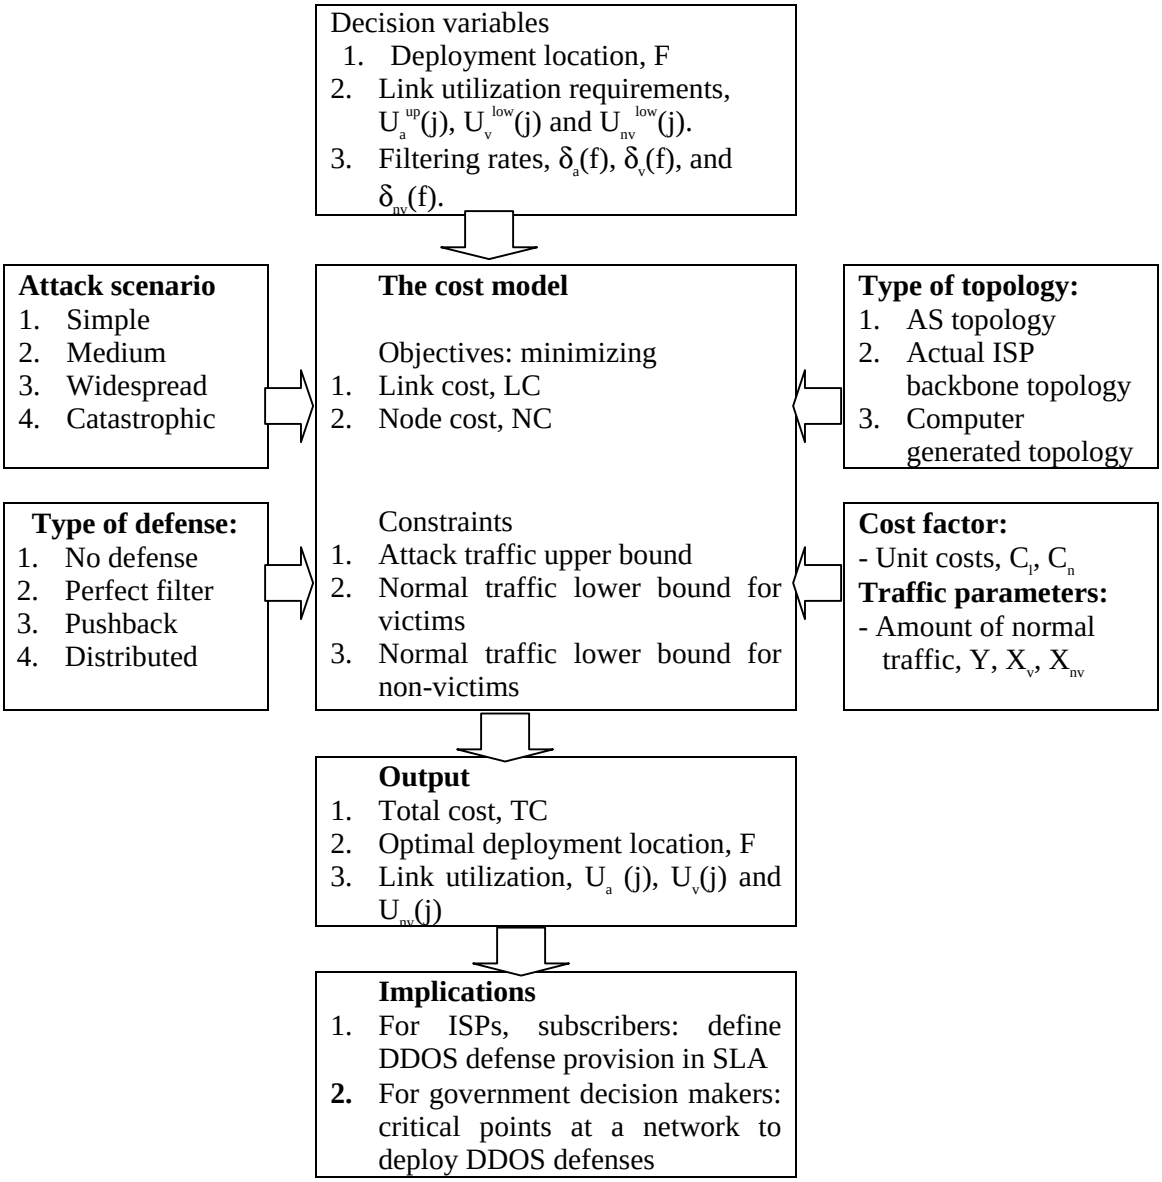

Figure 3: The computational model

Within the computational model the total cost for a network is calculated given a particular attack scenario and defense. Figure 3 contains a top-level view of this computational model. Decision variables include deployment location, link utilization requirements, and filtering rates for different types of traffic. The input variables are attack scenario, type of defense, type of topology, cost factor and traffic parameter. The output variables are the total cost, optimal deployment location and link utilization for different types of traffic.

Using this computational model we will begin with two virtual experiments described in the following sections. It is important to note that these two experiments do not span the space of all possible experiments. Nor does running these two experiments enable a complete characterization of the response surface. We begin with these two experiments because the results have direct bearing on key research and policy questions. We intend to do a number of other experiments given this system and to develop a more complete understanding of the response surface.

4.2 Preliminary results

To show how the simulation model works, we run the model based on the example network described in section 3.5. Figure 2. In this example, we run five different sets of filter location and compare the results with the no filter case. The "pushback" type of defense is placed. The capacity of the victim's link is B_v. Each attack node sends $1B_v$ traffic to the victim node. Each node sends $0.05B_v$ normal traffic to the victim machines, and $0.05B_v$ normal traffic to the non-victim machines. Filtering rates $\delta_a(f)= \delta_v(f)=0.99$, and $\delta_{nv}(f) = 0$ for $\forall f \in F$.

Figure 4 and 5 shows the results of running this example. Figure4 shows that F={1,8,9,10} is the best filter deployment location because filter on this location reduce the most of the attack traffic and remain the highest availability of the link to the normal traffic. The link utilization by the attack traffic is reduced from 80% to only 4%. The link utilization by the normal traffic to victim machines is improved from 10% to 35%. The link utilization by the normal traffic to the non-victim machines is improved from 10% to 50%.

However, the total cost of the ISP is the highest when F={1,8,9,10}. Figure 5 shows the total cost of the ISP. If a subscriber requires that the link utilization by the attack must be lower than 5% and the link utilization by the total normal traffic must be larger than 60%. F={5,6,7} is a better choice for the ISP because its cost is lowest among all feasible solutions.

In this project, we will design virtual experiments to analyze the changes of the link utilization and the total cost of an ISP by varying decision variables, attack scenarios, defenses and network topology. Based on these results, we expect to provide ISPs some insights into offering DDOS attack defenses.

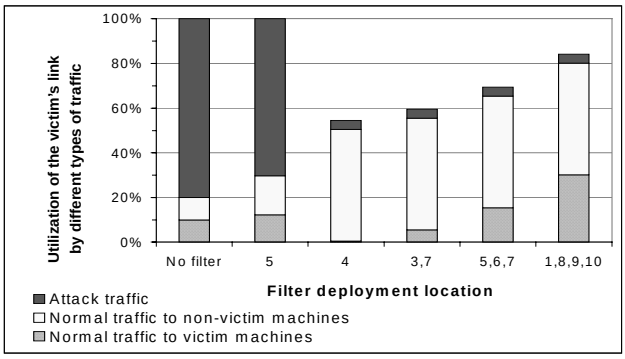

Figure 4: The link utilization by different types of traffic for different filter deployment location

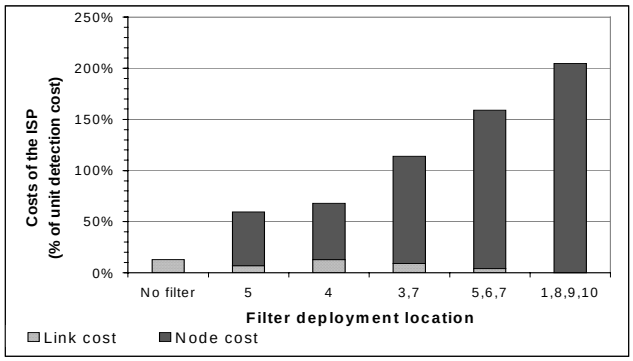

Figure 5: The total cost of an ISP for different filter deployment location

4.3 Virtual experiment 1: the tradeoffs in different defenses

What are the tradeoffs in different defenses? Virtual experiment 1 is designed to investigate this question. Two independent variables, attack scenario and type of defense, will be manipulated in this experiment. Dependent variables are the link cost and the node cost. A set of randomly generated network topologies will be generated using Monte Carlo estimation techniques. In generating these networks the factors – size, density, clusters, hierarchy and rate of change will be varied as previously described.

Attack scenarios

Variables to describe an attack include number of victim nodes, number of attack nodes, total normal traffic and packet rates. Packet rate is the amount of the attack traffic sent by one attack node as the percentage of a victim's link capacity. For example, if a victim's link capacity is 2Mbps, the packet rate is 10% if an attack node sends out 200kbps attack traffic. We assume constant bit rate for all attack traffic, which complies with empirical observation (Moore, Voelker et al. 2001). Number of packets for a certain packet rate will depend on compromised protocols, such as TCP SYN or UDP flood.

Based on equation (10) in Section 3.5, we calculate how many attack nodes are needed to saturate a victim's network link. Figure 6 shows the number of attack nodes needed to achieve 80% drop rate of the total normal traffic to a victim. A simple attack is originated from few attack nodes, which is similar to the DDOS attack against CERT/CC web site (Schwartz 2001). A widespread attack is originated from more 50 attack nodes, which can be achieved by attack propagation tools like Code-Red Worms (Moore 2001). We add a medium attack in between these two scenarios as a benchmark. In addition, we develop a future scenario that attacks may be launched from more attack nodes and target at more than one victim at the same time.

The total normal traffic that is sent by other nodes to a victim node is also a variable to determine how many attack nodes are needed. This variable depends on the usage behavior of a victim's network. We represent the total normal traffic in terms of the percentage of a victim's link capacity. We vary the total normal traffic based on peak rate (100%), medium usage (50%) and low usage (10%).

Attack scenario	Number of victim nodes, |V|	Number of attack nodes, |A|	Total normal traffic	Packet rate, $X_a(i)$, $i \in V$
Simple	1	5	100%	100%
Medium	1	15	50%	30%
Widespread	1	50	10%	10%
Catastrophic	50% of the ISP's network nodes	|V| * number of attack nodes need to saturate a victim	50%	100%

Table 2: Attack scenarios

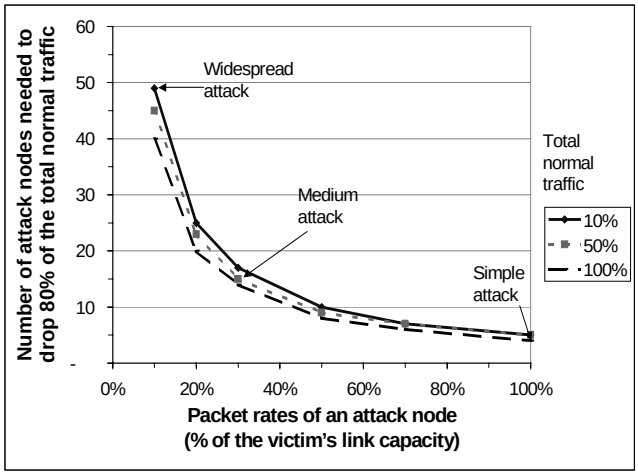

Figure 6: The number of attack nodes needed to saturate a victim's network link

Type of defense

Type of defense refers to the defenses that will be investigated. Four different types of defense will be simulated. These defenses are:

1. No defense: No defense is deployed, which is a benchmark for all other defenses.
2. Perfect filter: special situations in which filters can 100% distinguish the attack traffic from the normal traffic.
3. Pushback: the attack filters are triggered along the attack path, such as mechanisms in (Mahajan, Bellovin et al. 2001; Ioannidis and Bellovin 2002).
4. Distributed: the attack filters are distributed at some points of the network, such as mechanisms in (Ioannidis, Keromytis et al. 2000; Schnackenberg and Djahandari 2000; Park and Lee 2001).

In summary, virtual experiment 1 will include 4 attack scenarios and 3 types of defense. Table 3 summarizes scenarios in virtual experiment 1. Each scenario will be run several hundreds of times on randomly generated topologies to generate a mean estimate. The exact number of runs per cell will be determined by running a sample test on estimate the expected size of the standard errors around this mean then the final number of runs will chosen so that the expected standard errors will be sufficiently small that a reasonable point estimate of the mean is produced. The means and standard errors of each scenario will be compared graphically and statistically (using both analysis of variance techniques and other non-parametric tests). Sensitivity analyses will be conducted to determine the extent to which the control variables (factors controlling the topology) impact the results.

Independent variables	Number	Values
Attack scenario	4	simple, medium, widespread, catastrophic
Type of defense	4	no defense, perfect filter, pushback, distributed
Total number of scenarios = 4*4 = 16		

Table 3: Scenarios in virtual experiment 1

4.4 Virtual experiment 2: the impact of network topology

Where are the critical points in a network to deploy DDOS defenses? Virtual experiment 2 will investigate this question. There are three independent variables in this experiment: 1. attack scenario, 2. type of network topology and 3. deployment location. Dependent variables are the link cost and the node cost. In this experiment the defense type will be "distributed" as that is the most likely defense to occur in the real world. We will Monte Carlo across the number of defense sites. The set of attack scenarios will be the same as those used in virtual experiment 1.

Type of network topology

As previously noted there are a large number of possible networks. In this experiment, we focus our analysis by looking at three specific networks. As this focuses our analysis in a small portion of the space of topologies we will do a second experiment looking at other random networks where we systematically vary the size, density, clustering, hierarchy and rate of change. However, in this experiment two we will focus on just three topologies:

1. em-as: an empirical topology of Internet AS from BGP routing tables (CAIDA 1997),
2. gen-r: a computer-generated router level topology within an AS (size comparable to the empirical topology), and
3. gen-random: a computer-generated random network topology with the links distributed in a uniform random fashion (size and density comparable to the empirical topology).

Deployment location

Deployment location represents the routers that will deploy a defense, which is the set of the filter nodes. Deployment location that will be simulated is the following:

1. Centrality: at the ten most central nodes (Freeman 1979; Bonacich 1987; Freeman, Borgatti et al. 1991). This will be done three times using degree, betweenness and information.
2. Upstream: at nodes that are one step (hop) away from victim nodes.
3. Vertex covering: at nodes that are listed in the vertex covering set of nodes.

In summary, virtual experiment 2 will include 4 attack scenarios, 3 types of network topology, and 3 deployment locations. Table 4 summarizes scenarios in virtual experiment 2. As in virtual experiment 1, each scenario will be run as many times as needed to generate a mean estimate. All other variables will be Monte Carlo'd over and treated as controls.

4.5 Validations

To validate this model we will use empirical grounding (Carley 1996). 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. Attack scenarios will be validated based on data from empirical studies on attack tools, propagation methods (Moore, Voelker et al. 2001), and observable historical data from computer virus propagation (Moore 2001). Type of network topology will be validated through empirical AS topology data (CAIDA 1997) and empirical studies on router topology (Paxson 1996; Palmer, Siganos et al. 2001).

Independent variables	Number	Values
Attack scenario	4	simple, medium, widespread, catastrophic
Type of network topology	3	em-as, gen-r, gen-random
Deployment location	5	Centrality (3 types), upstream, vertex covering
Total number of scenarios = 4*3*5 = 60		

Table 4: Scenarios in virtual experiment 2

V. Benefits and Limitations

The proposed research is fundamentally interdisciplinary and draws on work in computer science, information science, social networks, and organization theory (particularly computational organization theory). This approach is necessary if we are to adequately understand and evaluate the impact of attacks on the critical infrastructure – in this case just one aspect of it – the Internet.

There are a large number of possible benefits of the proposed research. First, the policy framework proposed in this research will help ISPs and subscribers to consider the benefits of providing DDOS defenses and to realize the tradeoffs in DDOS defenses. The computational model provides a systematic framework for thinking through the tradeoffs in defense strategies in this complex system. Results from the specific experiments outlined provide guidance as to how to defend when many potential victims refuse to take action and for how to grow the network topology to make it less subject to catastrophic attacks. Thus, this work has direct bearing on security policy decisions at router level for a critical infrastructure. Since it is neither cost effective nor ethical to conduct real world experiments of DDOS attacks on a large network, this research provides a new technology to help evaluate the costs imposed by various attack scenarios and defenses. Moreover, the topological measures developed in this research could be useful for studies of other large-scale topologies. Finally, this research will provide a theoretical basis for evaluating DDOS defenses.

Limitations of the proposed research include the following. First, simulation analysis provides an order of magnitude cost comparison among defenses. However, the real dollar value of the cost will depend on the implementation of these defenses. Thus, while the model suggests relative effects in terms of cost, it will not provide real costs. Secondly, the cost model is based on bandwidth consumption costs

by either attack traffic or defenses. Other implementation costs will not be examined since this research focuses on the additional benefit and cost caused by defenses. Third, despite working with CERT, there is a limited amount of data available for validating models such as this. Fourth, the analysis we will engage in first is, in a sense, static; i.e., we do not assume intelligent attackers who change their attack in response to the defense strategy. This is clearly important and would be an important next step. Such a step, we note, requires the fundamental research described herein as a foundation. Finally, the computational model developed in this research is intended to provide decision support for tradeoffs in DDOS defenses. This model would need further revision to analyze defenses for other types of Internet security incidents.

REFERENCES

Albert, R., H. Jeong, et al. (2000). "Error and attack tolerance of complex networks." Nature **406**(27 July).

Anderson, B. S., C. Butts and K.M. Carley (1999). "The interaction of size and density with graph level measures" to social networks." Social Networks(21): 239-267.

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

Bonacich, P. (1987). "Power & centrality: A family of measures." American Journal of Sociology(92): 1170-1182.

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

Butts, C. and K.M. Carley (2001). Multivariate Methods for Inter-structural Analysis. CASOS Working Paper, CMU, www.casos.ece.cmu.edu.

CAIDA (1997). Global ISP interconnectivity by AS number, San Diego Supercomputer Center, University of California, San Diego. **2001**.

Carley, K. M. (1998). Organizational Adaptation. Annals of Operations Research. 75, 25-47.

Carley, K. M. (1997) Organizations and Constraint Based Adaptation. In Raymond A. Eve, Sara Horsfall & Mary E. Lee (Eds.) Chaos, Complexity and Sociology: Myths, Models and Theories, (Pp. 229-242) Thousand Oaks, CA: Sage.

Carley, K. M. (1996). Communicating New Ideas: The Potential Impact of Information & Telecommunication Technology, Technology in Society, 18(2), 219-230.

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

Carley, K. M. (1995) Communication Technologies & their Effect on cultural homogeneity, consensus, & the diffusion of new ideas. Sociological Perspectives , 38(4), 547-571.

Carley, K. (1991a). Growing Up: The Development and Acquisition of Social Knowledge. In Howard J. & Callero P. (Eds.) The Self-Society Dynamic: Cognition, Emotion, and Action (Pp. 72-105) Cambridge, England: Cambridge University Press.

Carley, K. (1991b). A Theory of Group Stability. American Sociological Review , 56(3), 331-354.

Carley, K. M. (1990). Group Stability: A Socio-Cognitive Approach. In E. Lawler, B. Markovsky, C. Ridgeway & H. Walker (Eds.) Advances in Group Processes: Theory & Research . (Pp. 1-44), Vol. VII. Greenwhich, CN: JAI Press.

Carley, K. M. & D. Krackhardt (1996). Cognitive inconsistencies and non-symmetric friendship. Social Networks, 18, 1-27.

Carley, K. M., J. Lee and D. Krackhardt (forthcoming). Destabilizing Networks. Connections.

Carley, K.M. & J. Lee (1998) “Dynamic Organizations: Organizational Adaptation in a Changing Environment.” Ch. 15 (pp. 269-297) in Joel Baum (Ed.) Advances in Strategic Management, Vol. 15, Disciplinary Roots of Strategic Management Research. JAI Press. Pp. 269-297.

Carley,K. M. &. D. M Svoboda (1996) Modeling Organizational Adaptation as a Simulated Annealing Process. Sociological Methods and Research, 25(1), 138-168.

CERT/CC (1999). Results of the distributed-systems intruder tools workshop. Pittsburgh, Pennsylvania, USA, CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University.

Chen, L.-C. and K. M. Carley (2001). A computational model of computer virus propagation. Conferences in Computational Social and Organizational Science.

Chen, L.-C. and M. Sirbu (2000). The diffusion of Internet security incidents and its policy implications. Pittsburgh, Department of Engineering and Public Policy, Carnegie Mellon University.

Cohen, F. 1999. 'Simulating Cyber Attacks, Defenses, and Consequences' : Fred Cohen & Associates.

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.

Dittrich, D. 2001. 'Distributed Denial of Service Tools', University of Washington, Seattle, WA. Available at http://staff.washington.edu/dittrich/misc/ddos/.

Faloutsos, M., P. Faloutsos, et al. (1999). On power-law relationships of the Internet topology. ACM SIGCOMM '99 conference on Applications, technologies, architectures, and protocols for computer communications, Cambridge, MA.

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

Freeman, L. C. (1979). "Centrality in social networks: conceptual clarification." Social Networks(1): 215-239.

Freeman, L. C., S. P. Borgatti, et al. (1991). "Centrality in valued graphs: a measure of betweenness based on network flow." Social Networks(13): 141-154.

Govindan, R. and A. Reddy (1997). An analysis of Internet inter-domain topology and route stability. IEEE INFOCOM, Kobe, Japan.

Gupta, M., Chaturvedi, A.R. and Mehta, S. 2000. 'The Experimental Analysis of Information Security Management Issues for Online Financial Services' International Conference for Information Systems.

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 Longstaff, T.A. (1998). 'A Common Language for Computer Security Incidents'. SAND98-8667, Sandia National Laboratories, Albuquerque, NM and Livermore, CA.

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.

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

Ioannidis, S., A. Keromytis, et al. (2000). Implementing a distributed firewall. ACM Conference on Computer and Communications Security, Athens, Greece.

Kaufer, D. S. & Carley, K. M. (1993) Communication at a Distance: The Effect of Print on Socio-Cultural Organization & Change, Hillsdale, NJ: Lawrence Erlbaum Associates.

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

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

Moitra, S.D. and Konda, S.L. 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. San Diego, CA, CAIDA, Supercomputer Center, University of California at San Diego.

Moore, D., G. M. Voelker, et al. (2001). Inferring Internet Denial-of-Service Activity. USENIX Security Symposium, Washington DC.

Palmer, C., G. Siganos, et al. (2001). The connectivity and fault-tolerance of the Internet topology. Infocom '01.

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

Park, K. and H. Lee (2001). 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.

PCCIP 1997. 'Critical Foundations: Protecting America's Infrastructure' : President's Commission on Critical Infrastructure Protection. Executive Office of the President of the United States.

Ren, Y. (2001) A Computer Simulation Model of Transactive Memory Systems. CASOS Working Paper, CMU, www.casos.ece.cmu.edu.

SANS (2000). Egress filtering v 0.2, SANS Institute.

Savage, S., D. Wetherall, et al. (2000). Practical network support for IP traceback. The 2000 ACM SIGCOMM Conference, Stockholm, Sweden.

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.

Schwartz, J. (2001). Computer vandals clog antivandalism web site. New York Times**:** 5.

Scott, J. (1991). Social Network Analysis, SAGE Publications.

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

Song, D. X. and A. Perrig (2001). Advanced and Authenticated Marking Schemes for IP Traceback. IEEE Inforcom.

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

Stone, R. (2000). CenterTrack: an IP overlay network for tracking DoS. USENIX Security Symposium, Denvor, CO.

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

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

Wood, B.J. and Duggan, R.A. 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.

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.