THE ECONOMIC INCENTIVES OF PROVIDING NETWORK SECURITY SERVICES ON THE INTERNET INFRASTRUCTURE

LI-CHIOU CHEN
SCHOOL OF COMPUTER SCIENCE AND INFORMATION SYSTEMS
PACE UNIVERSITY
lchen@pace.edu

THOMAS A. LONGSTAFF
SOFTWARE ENGINEERING INSTITUTE
CARNEGIE MELLON UNIVERSITY
tal@cert.org

KATHLEEN M. CARLEY
INSTITUTE FOR SOFTWARE RESEARCH INTERNATIONAL
CARNEGIE MELLON UNIVERSITY
kathleen.carley@cmu.edu

ABSTRACT

Distributed denial-of-service (DDOS) attacks have emerged as a prevalent way to compromise the availability of networks and servers, which imposed financial losses for e-commerce businesses. Many defenses that mitigate the effect of ongoing DDOS attacks have been proposed. However, none of the defenses have been widely deployed on the Internet infrastructure at this point because of a lack of understanding in the economic incentives inherent in providing the defenses as well as uncertainty in current defenses. We propose that ISPs should provide DDOS defenses as network services to ensure the availability of a network or a server when the technology is ready. This paper provides an analytical framework for the proposed service to align the economic incentives. Using empirical data from security incidents, this paper show that the proposed service 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.

Keywords: network security, distributed denial of service, network services, cost-benefit analysis, economic incentive.

INTRODUCTION

Network distributed denial-of-service (DDOS) attacks [12] compromise the availability of victims' networks or servers. Past incidents have caused financial losses of victims [10, 24-25, 27]. Many defenses that mitigate the effect of ongoing DDOS attacks have been proposed and the uncertainty inherent in the technology has been previously studied [6, 14, 17]. Currently, some ISPs have developed methods to trace the sources of attack traffic on

their backbone networks [21, 22] and some ISPs[1] have started to offer services that mitigate the impact of DDOS attacks. Automatic mechanisms on responding against ongoing attack traffic are still underdeveloped in practice. More research effort is still needed to develop the automatic responses. Our purpose here is to assess if any economic incentive would push ISPs towards the development of the automatic mechanisms so that ISPs will further provide them to their subscribers. This 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 [16, 26].

What would be the economic incentives of ISPs to provide defenses against network attacks such as DDOS? This paper is intended to address this question by analyzing the economic benefits and costs of ISPs to provide the defenses at some choke points of the Internet infrastructure, such as network routers/proxy servers. We propose that ISPs should provide network defenses as network security services to their subscribers. Network 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 a network or a server during attacks.

We developed an analytical model to quantify the benefits and costs of the service provision. The model considers both the demand of subscribers (potential attack victims/sources) and the supply of the providers (ISPs) to deploy the network defenses. We analyzed the model analytically and calibrated some parameters using empirical data on network attacks. Based on these results, we provide recommendations on aligning ISPs' economic incentives.
The next section introduces the proposed service and describes the analytical model, followed by a section on the analytical results from the model, another on the empirical calibration and finally the model results are discussed. Conclusions and future works follow.

---

[1] AT&T offers DDOS detection and response services (http://www.att.com/news/2004/06/01-13096) starting from June 2004 but the service does not specify performance in a Service Level Agreement (SLA). Starting from March, 2004, MCI offers DDOS detection service with a SLA that guarantees some link utilization during DDOS attacks. However, this service does not trigger automatic responses against attacks and it provides only attack detection when customers report suspicious attacks
(http://global.mci.com/terms/us/products/internet/sla/).

THE ANALYTICAL MODEL FOR PROVISION OF NETWORK SECURITY SERVICES

We propose that ISPs provide network security services to their subscribers. The services deploy DDOS defenses on some choke points of the Internet infrastructure and react actively to filter DDOS attack traffic during attacks. We consider two types of DDOS defenses: source filtering and destination filtering. Source filtering refers to the defenses that monitor the outbound traffic from a subscriber in order to prevent the subscriber from originating attacks (attack source). Destination filtering refers to the defenses that monitor the inbound traffic to a subscriber in order to prevent the subscriber from being attacked (attack victim). A detail description of the current technologies is in [6]. We define our analytical model based on the following assumptions:

- Attacks: DDOS attacks saturate the network connections of subscribers to their backbone networks or take down servers inside the network of the subscribers. The attacks can be traced to their sources within the administrative boundary of one network provider. Even if the attacks are originated from subscribers of another network provider, the provider of the victims can still trace to the network provider that carries the attack traffic.

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

- Providers: 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.

- Pricing: Providers charge all subscribers at a flat rate for a certain time period for the security service, such as a month. Many ISPs such as AOL currently offer virus scanning and firewall at a flat rate in additional to the network connection service that they provide. We will vary this assumption and analyze other pricing schemes when we discuss the model results.

- Market: 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. We will also discuss the service provision in a monopoly market when we discuss the model results.

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. We use 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 [11] and the utility of subscribers for intermediary services [1] and digital goods [2].

The utility that a subscriber derives from DDOS defenses is the expected loss that would be incurred from DDOS attacks. Economic losses from Internet security breaches have been studied previously [4, 9]. The expected loss is quantified by three factors: the attack frequency, $a \in [0,1]$, referring to how often attacks occur, the expected loss per attack, L, referring to how much loss an attack imposes on the subscriber and the quality of the defense, $q \in [0,1]$, quantifying the impact of the performance efficiency on the expected loss. Let U denote the utility function of a subscriber for the service, which is defined as:

$$U = aqL \tag{1.a}$$

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:

$$P_d \leq aqL \tag{1.b}$$

Assume that L for all subscribers is proportional to a uniform distribution. Let q denote the quality of the service for DDOS defenses, which can be considered as a network performance measure, such as the arrival rate of legitimate traffic. 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. As a result, only the subscribers that expect the attack frequency to be larger than $\dfrac{qL}{P_d}$ would subscribe to the service at P_d. Let M represent the total number of subscribers of an ISP. Let N_d denote the number of subscribers that are willing to subscribe to the network security service. When the price is set at P_d, N_d is calculated as:

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

From (1.c), the lowest attack frequency expected by the subscribers of the network security service is a function of N_d, which is:

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

Benefits and Costs of Providers

The cost quantification considers only the operational cost of providing DDOS defenses but not the capital investment on the infrastructure. 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. On the contrary, bandwidth saving (W) reduces the cost, which quantifies the per-attack transport benefit. This benefit comes from filtering attack packets 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 choke points between the attack sources and the victims. 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:

$$R = \tau \mu_x C_r H(G) \tag{2.a}$$

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:

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

The total cost of providing the defense C is the sum of operational cost C_o from all subscribers, and R from all attacks. Let $\Theta(N_s)$ represent 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. When the service is offered to N_s subscribers, the total cost for providing the service is calculated as:

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

The total benefit for providing the service is calculated as:

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

The total profit for providing the services TP is:

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

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

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

ANALYTICAL RESULTS

From (1.a)-(1.d) and (2.a)-(2.f), the price range of the security service obtained is the following:

$$C_o + [R - W]K(n) \leq p \leq K(n)qL$$
(3.a)

How a provider sets the price within this range depends on the market (its competitors) and its pricing strategy. In the short term, if all providers have the same marginal cost, the equilibrium price and the equilibrium number of subscribers in a competitive market can be calculated by equaling (3.b) and (3.c). The equilibrium number of subscriber n^* will satisfy

$$C_o + [R - W - qL]K(n^*) = 0$$
(3.b).

The equilibrium price is

$$p^* = K(n^*)qL = C_o + [R - W]K(n^*)$$
(3.c).

The total provider's benefit is equal to its profit, which is

$$TP = p^* n^* - [R - W]\Theta(n) - C_o n^*$$
(3.d).

The total subscribers' benefit is

$$CS = qL\Theta(n) - p^* n^*$$
(3.e).

The total social benefit is

$$SB = TP + CS = [qL - R + W]\Theta(n) - C_o n^*$$
(3.f).

Table 1 lists the impact of each variable on TP, CS and SB. We summarized two major findings as follows:

1) *When the capacity of the network is constrained, providers have more benefits over costs of providing defense mechanisms using flat rate pricing.* When the capacity of the ISP's network is constrained, the bandwidth saving is larger than the filter overhead ($R<W$). During a DDOS attack, an ISP's network capacity can be constrained because attackers intend to cause burst traffic. Even if the ISP expands its network capacity, attackers can still generate attacks with increasingly higher packet rates. In this case, all TP, CS and *SB* increase with bandwidth saving and decrease with filter overhead so that the provider's interest is aligned with the subscribers' interests.

2) *When the capacity of the network is not constrained, providers have more costs over benefits of providing defense mechanisms using flat rate pricing in a competitive market. In this case, other pricing strategies should be considered.* When the capacity of the ISP's network is not constrained, the bandwidth saving is smaller than filter over head ($R>W$). In this case, providers have losses from providing the defense mechanisms because the flat rate price cannot fully recover the cost. Subscribers that have low probability of being attacked will not pay for the service because they simply expect less loss from the attacks than the service fee. Under this circumstance, the providers should consider other pricing strategies.

Table 1: The impacts[2] of variables on provider' benefit, subscribers' benefit and social benefit

| Variables | | R=W (TP=0) | | | R<W (TP>0) | | | R>W (TP<0) | | |
Name	Increase in	TP	CS	SB	TP	CS	SB	TP	CS	SB
Operational cost	C_o	0	↓	↓	↓	↓	↓	↑	↓	↓
Reduced expected loss	L, q	0	↑	↑	↑	↑	↑	↓	↑	↑
Router overhead	$R(=_x, C_r, H)$	0	↓	↓	↓	↓	↓	↑	↓	↓
Bandwidth saving	$W(=_a, C_w, D)$	0	↑	↑	↑	↑	↑	↓	↑	↑
Attack duration	τ	0	0	0	↑	↑	↑	↑	↓	↓

EMPIRICAL EVIDENCE FOR PARAMETER CALIBRATIONS

We estimated the variation of the demand among individual subscribers using empirical data of network attacks. The variation can be explained as the variation in the attack risk of subscribers' online services. 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.

We used two data sets to calibrate the probability of attacks $F(a)$ since $F(a)$ determines the shape of the demand function. These two empirical data sets are: 1) the DDOS data set [18] and 2) the Code-Red data set [19]. The DDOS data set is used to estimate the distribution of attacks "sent to" subscribers (for destination filtering), and the Code-Red data set is used to estimate the distribution of attacks "originating from" subscribers (for source filtering). Figure 1 shows that both data sets can be modeled by a power curve functional form (R-square = 0.93 and 0.98,

respectively). We will use the two estimated functional form to calibrate $F(a)$ in the next section.

We calculated R and W using an AT&T backbone network map from [3]. This map describes a core network topology connecting North America cities for AT&T network. In addition, we collected public available data to calibrate parameters of a base scenario (Table 2). In the next section, the parameters for the model analysis are set to the values in this base scenario unless they are otherwise specified. 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). 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. A detail description of the data sets and the topology calculation is in [5].

---

[2] "0" denotes no influence, "↓" denotes an increase on the parameter will decrease TP, CS or SB, and "↑" denotes an increase on the parameter will increase TP, CS or SB.

Table 2: Parameter setting for 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 [3].
	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 [3].
	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 [8], the reported average annually losses from denial of service for a company 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 prior analysis. 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 [20]. 30% is the medium estimation.
Attack	μ_a	60Mb /second	Attack magnitude. It is estimated by 1500 packet per second (pps) and 40 bytes per packet [7]. An attack with 1500 pps is enough to compromise a firewall. In the trace analyzed in [18], 20% of all attack events had an estimated packet 1500 pps or higher. Minimum TCP packet size which carries TCP acknowledgement but no payload [15].
	τ	10 minutes	Duration of an attack. In the trace analyzed in [18], 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.

Figure 1: The empirical data of network attacks

COST AND BENEFIT ANALYSIS BASED ON EMPIRICAL EVIDENCE

The empirical calibration is to clarify three issues that can not be determined by the analytical results alone. 1) When the capacity of the network is constrained, how do we choose from different defense technologies? 2) What are the factors that influence the capacity constraint during an attack? 3) If the flat rate pricing cannot support the security services, what are the alternatives? Each of the following sub-sections will address each of the three questions, respectively. To avoid presenting absolute monetary values of the benefits and costs, we will use a benefit-cost ratio ($\frac{B}{C}$) to present the empirical results.

Filtering Technology

What defense technologies that a network provider should adopt when bandwidth cost is a concern of the operation? Here we discuss two types of technologies: 1) destination filtering: filtering inbound traffic of subscribers to prevent the subscribers from being attacked, and 2) source filtering: filtering outbound traffic of subscribers to prevent the subscribers from sending out attack traffic. We used the DDOS data to calibrate the demand for destination filtering and the Code-Red data to calibrate the demand for source filtering.

When destination filtering is deployed, the closer the filters can be to the attack sources, the more benefit both the provider and the subscriber will have. Figure 2 shows that both the provider's benefit and the subscribers'

benefit increases when the filter location[3] 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 in the quality of the service. That is, more legitimate traffic to the subscribers can bypass the filters.

Some subscribers may be exploited by attackers to launch attacks. When subscribers suffer losses from originating attacks, the network provider will be better off to adopt source filtering than destination filtering. This result occurs when the packet rate of an attack is larger than a threshold, 150pps for our scenario (Figure 3). This point is where the network capacity is constrained (W>R) as we discussed in the analytical results from the model. This result implies that a policy is needed to impose a cost on subscribers that originate attacks. Possible ways of imposing such a cost include blacklisting the subscribers that originate attacks, assigning liability to attack sources [13], or revealing the origins of the attack sources.

---

[3] 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.

Figure 2: Increase on both the provider's benefit and subscribers' benefit
by setting filters closer to the attack sources

Figure 3: Benefit-cost ratio per service for source filtering and destination filtering

Capacity Constraints

What is the impact of other factors on the network capacity constraints? Here we discuss two factors in our model: the ratio of bandwidth cost and filter overhead and the distribution of attacks sources.

First, the network capacity becomes constrained when the unit bandwidth cost is 10 times of the unit filter overhead.

In this case, source filtering is more beneficial for the provider. Figure 4 shows that the benefit cost ratio in source filtering exceeds its value in destination filtering when $C_w/C_r > 0.1$.

Second, the packet rate for the capacity constraint increases when the number of attack sources increases and when the attack sources are distributed. As in Figure 5, when the packet rate < 3000pps, the benefit-cost ratio for the source filtering data set is smaller than it is for the

destination filtering. When the packet rate > 3000pps, the difference of the benefit-cost ratio between the two approaches is much smaller than it is during a single source attack. This reason for the result is that, for a given packet rate of an attack received by the victim, the packet rate from one attack source when the attack is distributed is less than the packet rate from one attack source when the attack is from one source.

Figure 4: The impact of bandwidth cost/filter overhead cost

Figure 5: Single source attacks vs distributed source attacks

Pricing Strategies

The advantage of the flat rate pricing scheme is its simplicity. However, under such a scheme, the provider will not have incentives to provide the service if the network is not capacity constrained. We will relax this flat rate assumption in this section. For comparison, we analyzed two other strategies: 1) free bundling and 2) differential pricing.

We will discuss the free bundling pricing scheme using the benefit-cost ratio per attack, which represents how much benefit over cost that an ISP would obtain without considering the payment and the fixed cost from each subscriber. This situation happens when providers would like to attract more subscribers to the IP transport service or when providers charge the subscribers for only the fixed cost per subscriber. Using source filtering (Figure 6) as an example, the flat rate pricing scheme has the approximately same benefit-cost ratio as the free bundling scheme if the fixed cost is recovered from other services. The reason for this is that the number of attack frequency is very large in our Code-Red data set so that the benefit per attack is much larger than the benefit from service charge. In this case, the impact of the service charge is negligible. In addition, if the benefit from network connection services

is larger than the fixed cost, the free bundling scheme is even more beneficial for the provider than the flat rate scheme since the provider obtain both the bandwidth saving and the additional gains from other services.

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 estimate in practice. An alternative is to differentiate the service to several versions for subscribers who have different expected loss. Similar schemes have been used in digital product vertical differentiation [2]. Figure 7 compares the flat rate pricing scheme and the differential pricing scheme for individual subscribers. 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. Across all packet rates, the differential pricing scheme is more beneficial for the provider than the flat rate scheme. The analysis on the differential pricing here is preliminary. Further mechanisms are needed for aligning subscribers with different prices since it is hard in practice to evaluate the expected loss of subscribers.
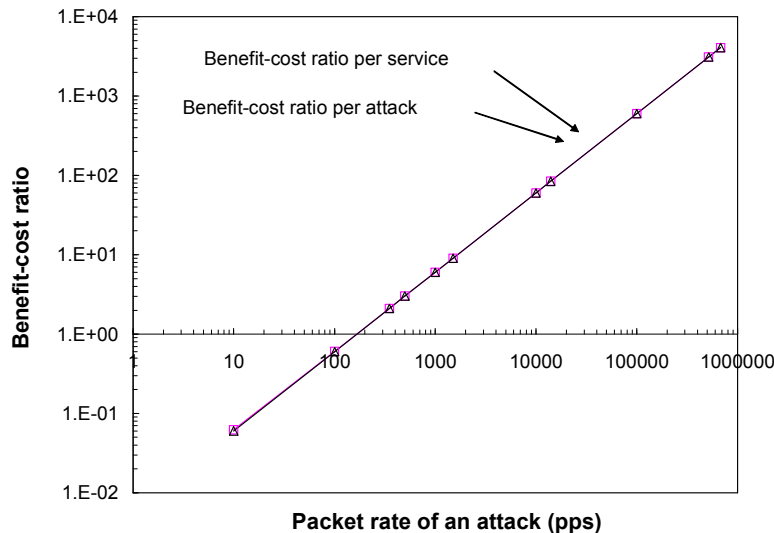

Figure 6: Benefit-cost ratio per service vs benefit-cost ratio per attack
for source filtering at the upstream router of attack sources

Figure 7: Differential pricing vs. flat rate pricing in the monopoly market for source filtering

CONCLUSIONS

We proposed a quantitative method to investigate the economic incentives for providing services to respond against ongoing DDOS attack traffic. 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. We found several factors that will influence the operational profit.

At the initial stage, when few providers are able to deploy the service (monopoly market), 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. When more and more providers are able to provide the service (competitive market), no single provider can benefit from the differential pricing scheme since subscribers can have more choices by switching to another provider. In this case, three implications can be drawn:

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.

3) The provider is better off providing the destination filtering service for free if the fixed cost per subscribers can be recovered from the additional revenue brought by new subscribers to network transport services.

We provided an analysis on the economic incentives of providing DDOS defenses. With an appropriate pricing strategy and some investigation into the expected loss from attacks, network providers can benefit from providing the security services and align their interests with subscribers. This work is just our first step to investigate this problem. Future work on estimating subscribers' expected loss and collecting data on attack incidents are needed to facilitate our proposal.

REFERENCES

[1] Bhargava, H.K., V. Choudhary, and R. Krishnan, Pricing and product design: intermediary strategies in an electronic market. *International Journal of Electronic Commerce*, 2000. Vol. 5,No. 1: pp. 37-56.

[2] Bhargava, H.K. and V. Choudhary, Information goods and vertical differentiation. *Journal of Management*, 2001. Vol. 18,No. 2: pp. 89-106.

[3] BW, Directory of Internet Service Providers, The Board Watch Magazine. 2001.

[4] Cavusoglu, H., B. Mishra, and S. Raghunathan. The effect of Internet security breach announcements on market value of breached firms and Internet security developers. *Workshop on Information Systems and Economics*. 2002. Barcelona, Spain.

[5] Chen, L.-C., Computational Models for Defenses against Internet-based Attacks, Department of Engineering and Public Policy. 2003, Carnegie Mellon University: Pittsburgh.

[6] Chen, L.-C., T.A. Longstaff, and K.M. Carley, Characterization of defense mechanisms for distributed of denial of service attacks. *Computers & Security*, 2004. Vol. 23,No. 8: pp. 665-678.

[7] Claffy, K.C., G. Miller, and K. Thompson. The nature of the beast: recent traffic measurements from an Internet backbone. *INET*. 1998. Geneva, Switzerland.

[8] CSI, CSI/FBI computer crime and security survey, Computer Security Issues & Trend. 2001.

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

[10] Garber, L., Denial-of-service attacks rip the Internet. *IEEE Computer*, 2000. Vol. 33,No. 4: pp. 12-17.

[11] Gordon, L.A. and M.P. Loeb, The economics of information security investment. *ACM Transaction on Information and System Security*, 2002. Vol. 5,No. 4: pp. 438-457.

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

[13] Kabay, M.E., Distributed denial-of-service attacks, contributory negligence and downstream liability. *ACM Ubiquity*, 2001. Vol. No.

[14] Lipson, H., Tracking and tracing cyber-attacks: technical challenges and global policy issues. 2002, CERT Coordination Center, Software Engineering Institute: Pittsburgh.

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

[16] McCurdy, D., The DHS Infrastructure Protection Division: Public-Private Partnerships to Secure Critical Infrastructures. 2004, ISAlliance.

[17] Mirkovic, J. and P. Reiher, A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Computer Communication Review*, 2004. Vol. 34,No. 2: pp. 39 - 53.

[18] Moore, D., G.M. Voelker, and S. Savage. Inferring Internet denial-of-service activity. *USENIX Security Symposium*. 2001. Washington DC.

[19] Moore, D., C. Shannon, and J. Brown. Code-Red: a case study on the spread and victims of an Internet worm. *ACM SIGCOMM/USENIX Internet Measurement Workshop*. 2002. Marseille, France.

[20] Papagiannaki, K., et al. Analysis of measured single-Hop delay from an operational backbone network. *IEEE INFOCOMM*. 2002. New York.

[21] Snoeren, A.C., et al. Hash-based IP traceback. *ACM SIGCOMM*. 2001.

[22] Stone, R. CenterTrack: An IP overlay network for tracking DoS. *USENIX Security Symposium*. 2000. Denver, CO.

[23] Symantec, Symantec Internet security threat report. 2004, Symantec.

[24] Tran, K.T.L., Hackers attack major Internet sites, temporarily shutting Buy.com, Ebay, *Wall Street Journal*. 2000. pp. 3.

[25] Verton, D., Teen hacker 'Mafiaboy' pleads guilty to 55 charges, *ComputerWorld*. 2001.

[26] WH, The national strategy to secure cyberspace. 2003, The White House.

[27] Yankee, $1.2 Billion Impact Seen as a Result of Recent Attacks Launched by Internet Hackers. 2000, The Yankee Group.

ACKNOWLEDGMENTS

contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation, the Commonwealth of Pennsylvania or the U.S. government.

AUTHORS' BIOGRAPHIES

Dr. Li-Chiou Chen received her Ph.D. from Carnegie Mellon University in Engineering and Public Policy. She is an assistant professor at the Department of Information Systems in the School of Computer Science and Information Systems, Pace University. Her research interests are focused on combining artificial intelligence and agent-based modeling to conduct technological and policy analysis in the area of information security. Specific areas includes countermeasures against the propagation of computer viruses, computational modeling for defenses against distributed denial of service attacks and agent-based simulations on policies to counter the spread of epidemics.

Dr. Thomas A. Longstaff received his PhD in 1991 at the University of California, Davis in software environments. He is a senior member of the technical staff in the Network Situational Awareness Program at the Software Engineering Institute (SEI), Carnegie Mellon University. He is currently managing research and development in network infrastructure security for the program. His publication areas include information survivability, insider threat, intruder modeling, and intrusion detection.

Dr. Kathleen M. Carley received her Ph.D. from Harvard. She is a professor at the Institute for Software Research International, Carnegie Mellon University. Her research combines cognitive science, social networks and computer science. Specific research areas are dynamic network analysis, computational social and organization theory, adaptation and evolution, computational text analysis, and the impact of telecommunication technologies and policy on behavior and disease contagion within and among groups. Her models meld multi-agent technology with network dynamics and empirical data. Illustrative large-scale multi-agent network models she and the CASOS team have developed are: BioWar -- city, scale model of weaponized biological attacks; OrgAhead -- a strategic and natural organizational adaptation model; and DyNet -- a change in covert networks model.