

Minimizing Collateral Damage by Proactive Surge Protection

Jerry Chou[†], Bill Lin[†], Subhabrata Sen[‡], Oliver Spatscheck[‡]

[†]University of California San Diego, [‡]AT&T Labs-Research

ABSTRACT

Existing mechanisms for defending against distributed denial-of-service (DDoS) attacks are generally reactive in nature. However, the onset of large-scale bandwidth-based attacks can occur suddenly, potentially knocking out substantial parts of a network before reactive defenses can respond. Even for traffic flows that are not under direct attack, significant *collateral damage* will result if these flows pass through links that are common to attack routes. This paper presents a *proactive-surge-protection* (PSP) mechanism that aims to provide a broad first line of defense against DDoS attacks. Our solution aims to minimize collateral damage by providing *bandwidth isolation* between traffic flows. This isolation is achieved through a combination of traffic forecasting, proportional allocation of network capacity, metering and tagging of packets at the network perimeter, and preferential dropping of packets inside the network. Our solution is readily deployable using existing router mechanisms. Simulations across three large backbone networks show that up to 95.5% of the network could suffer collateral damage without protection, but our solution was able to reduce the amount of collateral damage by 60.5-97.8%, even with a coarse-grained forecasting scheme.

Categories and Subject Descriptors

C.2.6 [Computer-Communication Networks]: Internetworking

General Terms

Design, Security

Keywords

Denial of service, malicious attacks, admission control

1. INTRODUCTION

The Internet has become increasingly the main communication channel for the developed world, as it continues its phenomenal growth in traffic volumes and reach. In addition, the types of applications carried over the Internet are also changing from basic

applications such as web browsing to application with real time constraints such as IP telephony. Unfortunately, the increased reliance on the Internet has also increased the risk that a single attack by a rogue nation, terrorist organization, or disgruntled individual could seriously disrupt communications for a huge population. In particular, an attacker can potentially disable a network by flooding it with traffic. Such attacks are also known as bandwidth-based distributed denial-of-service (DDoS) attacks and are the focus of our work.

Depending on the network operator, the provider network may be a small-to-medium regional network or a large core network. For small-to-medium size regional networks, this type of bandwidth-based attacks has certainly disrupted service in the past. For core networks with huge capacities, one might argue that such an attack risk is remote. However, as reported in the media, [6] botnets with more than 100,000 bots exist in the Internet today. These large botnets combined with the prevalence of high speed Internet access, give attackers multiple tens of Gb/s of attack capacity. Moreover, core networks are oversubscribed – e.g., in the Abilene network, [2] some of the core routers have an incoming capacity of larger than 30 Gb/s from access networks, but only 20 Gb/s of outgoing capacity to the core. Although commercial ISPs do not publish their oversubscription levels, they are generally substantially higher than the ones found in the Abilene network due to commercial pressures of maximizing the return on core investments.

Considering these insights, one might wonder why we have not seen multiple successful bandwidth-based attacks to large core networks in the past. The answer to this question is difficult to assess. Partially, attacks might not be occurring because the organizations which control the botnets are interested in making money by distributing SPAM, committing click frauds, or extorting money from mid-sized websites. Therefore, they would have no commercial interest in disrupting the Internet as a whole. Another reason might be that network operators are closely monitoring their traffic and actively trying to intervene. Nonetheless, recent history has shown that if such an attack possibility exists, it will eventually be exploited. For example, [4] described SYN flooding attacks years before the attacks were used to disrupt servers in the Internet.

The goal of our work is to mitigate the impact of bandwidth-based attacks on provider networks. Most existing proposed mechanisms for defending against distributed denial-of-service (DDoS) attacks are reactive in nature – i.e., they are only triggered after an attack has been detected in an effort to limit the damage. However, the onset of large-scale bandwidth-based attacks can occur almost instantaneously, causing potentially a huge *surge* in traffic that can effectively knock out substantial parts of a network before reactive defense mechanisms have a chance to respond. This paper presents a *proactive-surge-protection* (PSP) mechanism that aims to provide

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

LSAD'07, August 27, 2007, Kyoto, Japan.

Copyright 2007 ACM 978-1-59593-785-8/07/0008 ...\$5.00.

a broad first line of defense against DDoS attacks when they happen. In particular, under a flooding attack, traffic loads along attack routes will exceed link capacities, causing packets to be dropped indiscriminately. Without proactive protection, even for traffic flows that are not under direct attack, substantial packet loss will occur if these flows pass through links that are common to attack routes, resulting in significant *collateral damage*.

Our solution is based on providing *bandwidth isolation* between traffic flows so that the collateral damage to traffic flows not under direct attack is substantially reduced. This bandwidth isolation is achieved through a combination of forecasting of expected traffic, proportional allocation of network capacity based on traffic forecast, metering and tagging of packets at the network perimeter into two differentiated priority classes based on capacity allocation, and preferential dropping of packets in the network when link capacities are exceeded. It is important to note that our approach has no impact on the regular operation of the network if no link is overloaded. It therefore introduces no penalty in the common case.

Our solution is deployable using existing router mechanisms that are already available in modern routers, which makes our approach scalable, feasible, and cost effective. Simulation results show that our approach is robust even when used with a coarse-grained forecasting method. In particular, simulation results across three backbone networks show that up to 95.5% of the network could suffer collateral damage without protection. Our solution was able to reduce the amount of collateral damage by 60.5-97.8%.

In addition, our solution is resilient to IP spoofing as well as changes in the underlying traffic characteristics such as the number of TCP connections. This is due to the fact that we focus on protecting traffic between different ingress-egress interface pairs in a provider network and both the ingress and egress interface of an IP datagram can be directly determined by the network operator. Therefore, the network operator does not have to rely on unauthenticated information such as a source or destination IP address to tag a packet.

The rest of this paper is organized as follows. Related work is discussed in Section 2. PSP scheme is described in Section 3, with greater details on the proportional scaling algorithm are presented in Section 4. Our experimental setup and results are presented in Sections 5 and 6, respectively. Finally, conclusions are drawn in Section 7.

2. RELATED WORK

DDoS protection has received substantial attention in the literature. The oldest approach, still heavily in use today, is typically based on coarse-grain traffic anomalies detection [17, 3]. Trace-back techniques [25, 21, 22] are then used to identify the true attack source, which could be disguised by IP spoofing. After detecting the true source of the DDoS traffic the network operator can block the DDoS traffic on its ingress interfaces by configuring access control lists or by using DDoS scrubbing devices such as [5]. Although these approaches are practical, they do not allow for an instantaneous protection of the network. As implemented today, these approaches require multiple minutes to detect and mitigate DDoS attacks, which does not match the time sensitivity of today's applications.

More recently, the research community has focused on enhancing the current Internet protocol and routing implementations. For example, multiple proposals have suggested to limit the best effort connectivity of the network using techniques such as capabilities models [19, 26], proof-of-work schemes [15], filtering schemes [16] or default-off communication models [7]. The main focus of these papers is the protection of customers connecting to the core net-

work rather than protecting the core itself, which is the focus of our work. To illustrate the difference, consider a scenario in which an attacker controls a large number of zombies. These zombies could communicate with each other, granting each other capabilities or similar rights to communicate. If planned properly, this traffic is still sufficient to attack a core network. The root of the problem is that the core cannot trust either the sender or the receiver of the traffic to protect itself.

One proactive solution was presented in [24]. Similar to the proposals limiting connectivity cited above, it focuses on protecting individual customers. This leads again to a trust issue in that a service provider should not trust its customers for protection. Furthermore, their solution relies heavily on the operator and customers knowing *a priori* who are the good and bad network entities, and their solution has a scalability issue in that it is not scalable to maintain detailed per-customer state for all customers within the network.

Of course, preferential dropping in the network is not new, as it has been proposed as a mechanism for providing Quality-of-Service (QoS) [8, 9]. However, for providing QoS, the service-level-agreements that dictate the bandwidth allocation are assumed to be either specified by customers or decided by the operator for the purpose of traffic engineering. There is also a body of work on measurement-based admission control for determining whether or not to admit new traffic into the network, e.g. [12, 14]. With both service-level-agreement-based and admission-control-based bandwidth reservation schemes, rate limits are enforced. Our work here is different in that we use preferential dropping for a different purpose to provide bandwidth isolation between traffic flows to minimize the damage that attack traffic can cause to regular traffic. Our solution is based on a combination of traffic forecasting, proportional bandwidth allocation, soft admission control at the network perimeter, and lazy dropping of traffic inside the network only when needed. As the mechanisms of differential tagging and preferential dropping are already available in modern routers, our solution is readily deployable.

3. PROACTIVE SURGE PROTECTION


In this section, we present a high-level architectural overview of our DDoS defense solution called *proactive-surge-protection* (PSP). To illustrate the basic concept, we will depict an example scenario on the Abilene network. It consists of core routers that are interconnected by OC192 (10 Gb/s) links, as shown in Figure 1(a). For the purpose of depiction, we will zoom in on a portion of the Abilene network, as shown in Figure 1(b). Consider the situation where there is a sudden bandwidth-based attack along the origin-destination (OD) pair Chicago/NY. Suppose the magnitude of the attack traffic is 10 Gb/s. This attack traffic, when combined with the regular traffic for the OD pairs Sunnyvale/NY and Denver/NY, will cause the total load along the Chicago/NY link ($3 + 3 + 10 = 16$ Gb/s) to greatly exceed its 10 Gb/s link capacity, which will result in a high percentage of indiscriminate packet drops. Although the OD pairs Sunnyvale/NY and Denver/NY are not under *direct* attack, these flows will also suffer substantial packet loss because these flows pass through links that are common to the attack route, resulting in significant *collateral damage*. The flows between Sunnyvale/NY and Denver/NY are said to be caught in the *crossfire* of the Chicago/NY attack.

3.1 Approach


Our approach is based on providing *bandwidth isolation* between different traffic flows so that the collateral damage to crossfire traffic flows is minimized. This bandwidth isolation is achieved by using a form of *soft* admission control at the perimeter of a provider



(a) Abilene network



(b) 10 Gb/s attack traffic Chicago/NY



(c) Shielded Sunnyvale/NY and Denver/NY traffic from collateral damage

Figure 1: Attack scenario on the Abilene network.

network. In particular, to avoid saturation of network links, we impose *rate limits* on the amount of traffic that gets injected into the network for each OD pair. However, rather than imposing a *hard* rate limit, where packets are *blocked* from entering the network, we classify packets into two priority classes, *high* and *low*. Metering is performed at the perimeter of the network, and packets are tagged *high* if the arrival rate is below a certain threshold. But when a certain threshold is exceeded, packets will get tagged as *low* priority. Then, when a network link gets saturated, e.g. when an attack occurs, packets tagged with a low priority will be dropped preferentially. This ensures that our solution does not drop traffic unless a network link capacity has indeed been exceeded. Under normal network conditions, in the absence of sustained congestion, packets will get forwarded in the same manner as without our solution.

Consider again the above example, now depicted in Figure 1(c). Suppose we set the high priority rate limit for the OD pairs Sunnyvale/NY, Denver/NY, and Chicago/NY to 3.5 Gb/s, 3.5 Gb/s, and




Figure 2: Proactive-Surge-Protection (PSP) architecture.

3 Gb/s, respectively. This will ensure that the total traffic admitted as high priority on the Chicago/NY link is limited to 10 Gb/s¹. If the limit set for a particular OD pair is *above* the *actual* amount of traffic along that flow, then all packets for that flow will get tagged as high priority. Consider the OD pair Chicago/NY. Suppose the actual traffic under an attack is 10 Gb/s, which is above the 3 Gb/s limit. Then, only 3 Gb/s of traffic will get tagged as high priority, and 7 Gb/s will get tagged as low priority. Since the total demand on the Chicago link exceeds the 10 Gb/s link capacity, considerable packets would get dropped. However, the packets drop will come from the OD pair Chicago/NY since all packets from Sunnyvale/NY and Denver/NY would have been tagged as high priority. Therefore, the packets for the OD pairs Sunnyvale/NY and Denver/NY would be shielded from collateral damage and get through to their final destination.

3.2 Architecture

Figure 2 depicts the architecture of our PSP solution.

Forecaster: It determines the expected amount of regular traffic at different times throughout a day, a week, a month, etc. In particular, for each time period t , the forecaster will determine a *forecast matrix*, $F = [f(i, j)]$, where $f(i, j)$ is the expected rate of regular traffic for the OD pair with ingress node i and egress node j .

For this forecasting component, a number of sophisticated forecasting methods could be envisioned. For example, Bayesian-based methods used for traffic predictions in transportation highways could potentially be adapted for network traffic predictions [13, 23]. Such forecasting methods take into consideration both dynamic shifts in past traffic patterns as well as current traffic measurements. However, rather than relying on the availability of sophisticated forecasting methods, we derive a coarse-grained forecast matrix by taking the traffic average over the same time of day from the previous months. Earlier study [20] has shown that traffic over the same time of day indeed exhibits strong similarity. Since the derivation is based on historical measurements and determined *offline*, the operator can filter out any known attack traffic during the measurement periods and prevent from history poison. Although this forecasting method can have potential inaccuracies, the simulation results show that the bandwidth allocation algorithm based on this coarse-grained forecasting method is quite robust.

¹Operators can also set admission limits to some factor below the link capacity to provide the desired headroom.

Bandwidth Allocator: It determines the *rate limits* at different times throughout a day, a week, a month, etc. For each time period t , the bandwidth allocator will determine a *bandwidth allocation matrix*, $B = [b(i, j)]$, where $b(i, j)$ is the rate limit for the OD pair with ingress node i and egress node j . One method is to simply use the *forecast matrix* as the bandwidth allocation. However, under normal operating conditions, network links are typically underutilized. Therefore, our forecasting based on recent historical measurements will reflect this underutilization. If we based our bandwidth allocation of high priority packets on just the forecasted traffic, then the amount of traffic admitted into the high priority class is expected to be less than what the network can handle. Alternatively, since there is *room* for admitting more traffic into the high priority class, we can *proportionally scale up* the forecasted traffic to admit as much high priority traffic as the network capacity permits. This way, if our *forecasted* traffic for a particular OD pair underestimates the *actual* traffic, then the proportionally scaled bandwidth allocation will likely be higher than the actual traffic, thus ensuring the actual traffic will get through. This proportional scaling algorithm is described in more details in Section 4.


Differentiated Tagging: Given the rate limits determined by the bandwidth allocator, the metering and tagging of packets in accordance with the determined rate limits can be implemented at the perimeter of the network. In particular, packets arriving at ingress node i and destined to egress node j are tagged as high priority if their metered rates are below the threshold given by $b(i, j)$; otherwise, they are tagged as low priority. These traffic management mechanisms for metering and tagging are commonly available in modern routers at linespeeds.

Preferential Drops: Finally, at each network router, in the case of sustained congestion where a link capacity has been exceeded, low priority packets are preferentially dropped over high priority packets. Again, this preferential dropping mechanism [8] is commonly available in modern routers at linespeeds. By using this mechanism, our solution ensures that it does not drop traffic unless a network link capacity has indeed been exceeded. Under normal network conditions, in the absence of sustained congestion, packets should get forwarded in the same manner as without our surge protection scheme.

4. PROPORTIONAL SCALING

The basic idea of the proportional scaling algorithm is to *fairly* allocate all available network capacity by proportionally increasing the amount of traffic that we admit into the high priority class, starting with the forecast traffic matrix. By raising the threshold of how much traffic we admit to the high priority class, we can better accommodate bursty traffic and forecast inaccuracies. Intuitively, our scaling algorithm works by allocating bandwidth iteratively in a “water-filling” manner. Each iteration increases the bandwidth allocation to flows proportional to some weight. Once a link is fully allocated, we remove the link from consideration and fix the rates of the flows that cross this link to the bandwidth allocated to them so far. We repeat this procedure until no more residual capacity can be allocated to non-fixed flows.

Consider the example shown in Figure 3. A simple network topology is shown in Figure 3(a) with 3 nodes that are interconnected by 10 Gb/s links. The forecast matrix is shown in Figure 3(b). Consider the links $A \rightarrow B$ and $B \rightarrow C$. $A \rightarrow B$ is shared by




(a) Network

(b) Forecast

	A	B	C
A	∞	6 Gb/s	4 Gb/s
B	4 Gb/s	∞	6 Gb/s
C	6 Gb/s	4 Gb/s	∞

(c) Bandwidth allocation



(d) After 1st round

(e) After 2nd round

Figure 3: Proportional scaling algorithm illustration.

the OD pairs (A, B) at 1.5 Gb/s and (A, C) at 1 Gb/s, and $B \rightarrow C$ is shared by the OD pairs (A, C) at 1 Gb/s and (B, C) at 0.5 Gb/s. Given the aggregated traffic forecast on the 10 Gb/s link $A \rightarrow B$ is only 2.5 Gb/s, we can potentially increase the bandwidth allocation to the two flows crossing this link, namely (A, B) and (A, C) , by a factor of $10/2.5 = 4$. Similarly, given the aggregated traffic forecast on $B \rightarrow C$ is only 1.5 Gb/s, we can potentially increase the bandwidth allocation to (A, C) and (B, C) by a factor of $10/1.5 = 6.67$. Since (A, C) crosses both $A \rightarrow B$ and $B \rightarrow C$, we are constrained to increase its allocation by a factor of 4, whereas we can safely increase the allocation for (B, C) by a factor of 6.67. The same analysis applies to the proportional increase of the other OD pairs. The result after the first iteration is shown in Figure 3(d). After the first iteration, the links $A \rightarrow B$ and $C \rightarrow B$ are fully allocated. We then remove these links from consideration and fix the rates of the corresponding flows, namely (A, B) , (A, C) , (C, A) , and (C, B) , to the allocation after the first iteration. In the next iteration, we take the residual capacity available on the links $B \rightarrow C$ and $B \rightarrow A$ and again proportionally increase the previous allocations of the non-fixed flows. The result after the second iteration is shown in Figure 3(e). The procedure then ends because the network capacity has been fully allocated. Figure 3(c) shows the final. Note that, for OD pairs (i, j) where $i = j$, the corresponding traffic does not actually traverse the network. Thus, we set their bandwidth allocation to infinity.

5. EXPERIMENTAL SETUP

We employed NS2 based simulations to evaluate our PSP scheme on three large real networks. While the results from the individual networks cannot be directly compared to each other because the difference of their network characteristics and traffic behavior, multiple network environments allow us to explore and understand the performance of our PSP scheme for a range of scenarios.

Abilene (AB): This is a public academic network in the US [2] with has 11 nodes² interconnected by OC192, 10Gb/s, links. The full topology is shown in Figure 1(a).

US: This is the backbone of a large service provider in the US, and consists of around 700 routers and thousands of links ranging from T1 to OC192 speeds.

EU: This is the backbone of a large service provider in Europe. It has a similar, but larger network structure as the US backbone, about 150 routers and 500 links more.

5.1 Normal Traffic Demand

For each networks, using the methods outlined in [10], we build ingress router to egress router traffic matrices from several weeks worth of sampled Netflow data that record the traffic for that network : AB (10/01/06 - 12/06/06), US (09/01/06 - 11/17/06) and EU (11/18/06 - 12/18/06). For each time interval τ , the corresponding OD flows are represented by a $N \times N$ traffic matrix where N is the number of access routers providing ingress or egress to the backbone, and each entry contains the average demand between the corresponding routers within that interval.

The above traffic data are used both for creating the normal traffic demand for the simulator as well as for computing the forecast and corresponding bandwidth allocation matrices. One goal of our study was to evaluate the effectiveness of using a single representative bandwidth allocation matrix for an extended period of time. An implicit hypothesis is that the bandwidth allocation matrix does not need to be computed and updated on a fine timescale. To this end, in the simulations, we use a finer timescale traffic matrix with $\tau = 1$ min for determining the normal traffic demand, and a coarser timescale 1 hour interval for computing the forecast and bandwidth allocation matrix from historical data.

5.2 DDoS Attack Traffic

We explored a number of different DDoS scenarios:

AB DDoS: We identified the bottleneck links and picked the corresponding attacking ingress and egress routers to saturate the bottleneck links. The weathermap on the Abilene website [1] clearly shows the bottleneck links are from Denver to Chicago which have around 2 Gb/s, while the other links normally are under 1 Gb/s. Our DDoS attack consists of three 5 Gb/s OD flows from Denver, Kansas City and Indianapolis, and all go to Chicago.

US DDoS: The attack matrix was based on large DDoS alarms for the past year from 6/1/05 to 7/1/06 generated by a commercial DDoS detection system deployed at key locations in that network. The attacks themselves caused no overloads on any backbone link and needed to be scaled up. For the presented evaluations involving the US backbone in Section 6, we selected the attack traffic matrix with the largest number (several thousands) of OD flows, and scaled

²Abilene actually has another secondary core router at Atlanta. It only connects to the primary Atlanta core router and has much less traffic. So, we ignore it from the topology for simplicity.

up the attack demand by a factor of 1,000.

EU DDoS: We created a synthetic DDoS attack matrix as follow: First, we randomly select 5% of access routers to be the victims. For each selected victim, 2% of the access routers are randomly chosen as the attackers. Then for each attack OD flow, its demand is 10% of the total access capacity on the corresponding ingress router. The resulting attack scenario consists of attack flows along only 0.1% of the OD pairs. This is in keeping with the very limited distributedness of DDoS attacks in a core network as discovered by earlier study [18].

5.3 NS2 Simulation Details

Our experiments are implemented on the NS2. This involved implementing the 2-class bandwidth allocation, and simulating both the normal and DDoS traffic flows.

Bandwidth Allocation and Policing: The metering and class differentiation of packets are implemented at the perimeter of each network using the differentiated service module in NS2, which allows users to set rate limits for each individual OD pair. Our simulation updates the rate limits hourly by pre-computing the bandwidth allocation matrix based on the hourly traffic matrix of several weeks prior to the attack date: AB (10/01/06 - 11/30/06), US (09/01/06 - 11/16/06) and EU (11/18/06 - 12/17/06).

The differentiated service module marks incoming packets into different priorities based on the configured rate limits and the estimated incoming traffic rate of the OD pair. Specifically, we implemented differentiated service using TSW2CM (Time Sliding Window with 2 Color Marking), an NS2 provided policer. As its name implies, the TSW2CM policer uses a sliding time window to estimate the traffic rate. If the estimated traffic exceeds the given threshold, the incoming packet is marked in low priority class, otherwise in high priority class. The packets for either class would be stored in the same physical router queue but separate virtual queues. We configure the drop policy of two virtual queues differently. The drop policy for the high priority class queue is drop tail, such that the packets are dropped only when the physical queue is completely full. Whereas the drop policy for low priority class queue is RED [11] which early drops packets proportional to the physical queue length.

Traffic Simulation: For the evaluations in the paper, we select a different weekday at random for each network: 12/06/06, 11/17/06, and 12/18/06 for the AB, US, and EU networks, respectively, and consider 96 1-minute time intervals (one every 15-minutes) across that entire day³. Recall that for a given time interval τ , we compute normal and DDoS traffic matrices that give average traffic rates across that interval. These matrices are used to generate the traffic flows for that time interval. Both DDoS and network traffic are simulated as constant bandwidth UDP streams with fixed packet sizes.

6. EXPERIMENTAL RESULTS

We first evaluate the collateral damage in each network in the absence of any protection, then explore the utility of using our surge protection scheme. For simplicity, we use the terms Forecast-PSP and Proportional-PSP to refer to the PSP schemes using forecast and proportional scaling bandwidth allocation matrices respectively, and refer to the baseline unprotected scenario as No-PSP. Recall that an OD pair is considered as (i) *attacked* if there is attack

³Results for other days were similar and therefore omitted.

traffic along that pair (ii) *crossfire* if it shares at least one path link with an OD pair containing attack traffic and (iii) *non-crossfire* if it neither *attacked* nor *crossfire*.

6.1 Potential for Collateral Damage




Figure 4: The percentage of the number of the three OD pair types classified under an attack traffic.




Figure 5: The proportion of normal traffic demand corresponding to the three types of OD pairs.

We first explore the extent to which OD pairs and offered traffic demand are placed in potential harm's way, because of paths or subpaths shared with a given set of attack flows. For a given network and attack matrix, by tracing the routing paths, we count the number of each category (*crossfire*, *non-crossfire* or *attacked*) of OD pairs, and report the relative proportions in Figure 4. Because the attack traffic involves a small number of interfaces, the percentage of attacked OD pairs is fairly low, less than 3%. However, for the US and EU backbones, the percentage of crossfire OD pairs is large, 95.5% and 83.5%, even though the attacks were only directly over 1.2% and 0.1% of the OD pairs, respectively. This is because, for the US and EU backbones, the attack flows went over common links that were shared by the vast majority of other OD pairs. On the other hand, the percentage of crossfire OD pairs for Abilene is much lower, mainly because the attack was much more concentrated (on three links only). In general, the percentage of OD pairs that are in the crossfire of attack flows depends on where the attacks occurred and how traffic is routed over a particular network.

We next depict the relative proportions of the overall normal traffic demand corresponding to each type of OD pairs. While the classification of the OD pairs into the 3 categories is fixed for a given network and attack matrix, the relative traffic demand for the different classes is time-varying, depending on the actual normal traffic demand in a given time interval. Figure 5 presents a breakdown of the total normal traffic demands for the 3 classes across the 96 time intervals. For all 3 networks, crossfire OD pairs account for a significant proportion of the traffic demand. Figures 4 and 5 together suggest that, even with an attack coming from limit amount of interfaces, it could be routed around the network in a manner that can impact a significant portion of OD pairs and overall network traffic. The numbers above give us an indication of the potential "worst case" impact that an attack can unleash, if its strength is sufficiently scaled up. This is because a crossfire OD pair will only suffer collateral damage and packet losses if some link on its path gets congested. While the above results do not provide any measure of actual damage impact, they do nevertheless suggest that there is

	Impacted OD Pairs(%)	Impacted Demand(%)	Packet Drop(%)
AB	78.4 [70.6, 93.2]	92.8 [89.6, 96.9]	31.9 [30.2, 35.5]
US	50.6 [47.7, 52.1]	39.9 [36.5, 42.1]	47.7 [46.5, 48.8]
EU	47.1 [41.8, 51.5]	47.0 [42.8, 49.3]	68.0 [64.8, 71.2]

Table 1: Collateral damage in the absence of PSP.

	No-PSP	Forecast-PSP		Proportional-PSP	
	Packet drop(%)	Packet drop(%)	Reduction(%)	Packet drop(%)	Reduction(%)
AB	22.8	3.1	86.4	0.5	97.8
US	25.0	11.8	52.7	9.9	60.5
EU	34.7	14.2	59.0	6.8	80.3

Table 2: The time-averaged mean crossfire OD-pair packet drop rate in Figure 6.

a real potential for widespread collateral damage, and underline the importance and urgency of developing techniques to minimize such damage.

We next evaluate the actual collateral damage induced by the specified attacks in the absence of any protection scheme. We define a crossfire OD pair to be *impacted* in a given time interval, if it suffered some packet loss in that time interval under No-PSP. Table 1 presents the total number of and traffic demand for impacted OD pairs as a percentage of the corresponding values for all crossfire OD pairs, and the average packet drop rate across the impacted OD pairs. To account for time variability, we present the average (with the 10 and 90 percentile indicated in the brackets) for the three metrics across the 96 attacked time intervals.

Overall, the results show that not only can the attacks impact most of the crossfire OD pairs, but they can cause severe packet drops in many of them. The percentage of impacted OD pairs and corresponding demand in the Abilene network is larger than other two networks because the attack flows are concentrated and overlaid on a single path from Kansas City to Indianapolis to Chicago, such that all attacked links are congested, except the link from Denver to Kansas City. Traffic belonging to the crossfire OD pairs therefore has a high possibility of suffering packet loss.

6.2 Surge Protection Evaluation

We next evaluate the performance of the proposed PSP techniques from following three perspectives:

(A) Time series packet drop: For each of the 96 attack time intervals, for each crossfire OD pair that had some traffic demand in that interval, we compute the packet drop rate, i.e., the number of packets dropped as a percentage of the total offered load for that OD pair. The mean crossfire drop rate is obtained by averaging across these per-OD pair drop rates for that interval. Figure 6 shows the time series of the mean crossfire drop rate across all the 96 attacked time intervals, for Abilene, US and EU. Even though the attack remains the same at each time interval, the normal traffic demand matrix is time-varying, hence the observed variability in the drop rate time series. In particular, we found slightly greater packet drop rates at the network traffic peak time, such as 12PM (GMT) in the Europe backbone and 6PM (GMT) in the US backbone. However, our PSP scheme consistently and significantly re-

	Forecast-PSP			Proportional-PSP		
	0%	10%	50%	0%	10%	50%
AB	0.380	0.737	0.899	0.367	0.922	0.996
US	0.198	0.256	0.535	0.254	0.331	0.613
EU	0.218	0.315	0.644	0.438	0.704	0.856

Table 3: Distribution of per-OD pair drop rates.


	Scale factor		1	2	3
	No-PSP	Forecast-PSP	23.05	59.87	85.24
AB	No-PSP	Forecast-PSP	3.48	8.22	10.21
	No-PSP	Proportional-PSP	0.58	2.70	4.09
	Forecast-PSP	Proportional-PSP	24.78	37.85	45.17
US	No-PSP	Forecast-PSP	11.71	15.72	17.56
	No-PSP	Proportional-PSP	9.74	13.23	14.56
	Forecast-PSP	Proportional-PSP	33.81	46.84	57.71
EU	No-PSP	Forecast-PSP	14.01	17.14	19.70
	No-PSP	Proportional-PSP	6.40	7.28	7.94
	Forecast-PSP	Proportional-PSP	6.40	7.28	7.94

Table 4: The time-averaged mean crossfire OD-pair packet drop rate as the attack volume scaled up to 3.


duce packet drop rate across all 3 networks. Table 2 summarizes Figure 6 and presents, for each PSP scheme, (i) the time-averaged mean crossfire drop rate, and (ii) the reduction in drop rate from the corresponding No-PSP drop rate as a percentage of the No-PSP value. Under Forecast-PSP, the time-average drop rate reduction with respect to the No-PSP case is 86.4%, 52.7% and 59.0% for the AB, US and EU networks, respectively. Under Proportional-PSP, the corresponding reduction percentages are 97.8%, 60.5% and 80.3%, respectively. In the US network, Proportional-PSP provides less improvement over Forecast-PSP because the network utilization is higher. On the other hand, both Abilene and the EU network have more residual capacity, which allow Proportional-PSP to offer greater headroom for unexpected normal traffic from either forecast inaccuracies or sudden traffic bursts.

(B) Packet drop distribution: The above results provide network-wide average performance measures. We next analyze the drop rates for individual crossfire OD pairs. For each time interval, for each crossfire OD pair that had some loss under No-PSP in that interval, we compute its drop rate under the 2 PSP schemes as a percentage of the drop rate it suffered under No-PSP. Table 3 summarizes some interesting points from the CDF across all the time intervals, of the per-OD-pair values. Each table entry with column header $x\%$ represents the total percentage of crossfire OD pairs that experienced a drop rate of at most $x\%$ of its corresponding drop rate under No-PSP. The results indicate substantial reductions in drop rate for many OD pairs across all 3 networks, with Proportional-PSP exhibiting the best performance. Among all the OD pairs that had some packet loss under No-PSP, a significant proportion had no loss under Proportional-PSP: 36.7%, 25.4% and 43.8% for Abilene, US and EU respectively. Also for a vast majority of the OD pairs (99.6%, 61.3% and 85.6% for Abilene, US and EU), the loss rate under Proportional-PSP was 50% or less than under No-PSP.


(C) Behavior under Scaled attacks: Given the growing presence of broadband connections and availability of large armies of botnets “for hire”, it is important to understand the effectiveness of the PSP techniques as these attacks scale. For each network, we next increase the intensity of the attack matrix by scaling the demand of every attack flow by a factor ranging from 1 to 3. For each scaling factor, we measure the time-averaged drop rate of crossfire OD pairs for eight 1-min time intervals, equally spaces across 24



(a) AB



(b) US



(c) EU

Figure 6: The mean crossfire OD-pair packet drop rate across 24 hours.

hours. Table 4 shows that the drop rate under No-PSP increases much faster than under Forecast-PSP and Proportional-PSP, as the attack intensity increases. This is because under No-PSP, all the normal traffic packets have to compete for limited bandwidth resources with the attack traffic, while with our protection scheme only normal traffic marked in low priority class is affected by the increasing attack. Therefore, even in the extreme case when the attack traffic demand is sufficient to clog all the links, our protection scheme can still guarantee that the normal traffic marked in the high priority class goes through the network. Consequently, our protection scheme is much less sensitive to the congestion degree, as evidenced by the much slower growth of the drop rate. For example, in the EU network, under Proportional-PSP, as the scale factor goes from 1 to 3, the mean drop rate increases very little from 6.40% to 7.94%. Whereas under No-PSP, the mean drop rate jumps from 33.81% to 57.71% in the same attack intensity range.

7. CONCLUSION

In this paper, we presented our proactive surge protection (PSP) scheme, which significantly reduces the impact of DDoS attacks on regional and core networks. Extensive simulations using real traffic, topology, and DDoS data shows that the packet loss rates caused by DDoS traffic can be reduced by 60.5-97.8%. Importantly, with our protection in place, the damage caused by a larger DDoS attack increases only slightly as the attack volume increases by several factors. In addition, our solution is resilient to IP spoofing as well as changes in the underlying traffic characteristics such as the number of TCP connections. This is due to the fact that we focus on protecting traffic between different ingress-egress interface pairs in a provider network and both the ingress and egress interface of an IP datagram can be directly determined by the network operator. Therefore, the network operator does not have to rely on unauthenticated information such as a source or destination IP address to tag a packet. As most network operators collect coarse-grained network statistics today, and most routers allow for traffic prioritization, the solution presented in this paper can be readily deployed. While our current approach has shown surprisingly good results, there is clearly room for improvement. In particular, more accurate bandwidth allocation algorithms might further improve our results. We are currently exploring the practical aspects of deploying PSP in large provider networks as well as providing enhancements to our algorithms.

8. REFERENCES

- [1] Abilene NOC weathermap. <http://weathermap.grnoc.iu.edu/abilene.jpg.html>.
- [2] Advanced networking for leading-edge research and education. <http://abilene.internet2.edu>.
- [3] Arbor peakflow. www.arbor.net.
- [4] CERT CA-1996-21 TCP SYN Flooding and IP Spoofing Attacks.
- [5] Cisco guard. <http://www.cisco.com/en/US/products/ps5888/index.html>.
- [6] Washington Post, The Botnet Trackers, Tuesday, February 16, 2006.
- [7] H. Ballani, Y. Chawathe, S. Ratnasamy, T. Roscoe, and S. Shenker. Off by default! In *ACM HotNets Workshop*, November 2005.
- [8] D. Clark and W. Fang. Explicit allocation of best-effort packet delivery service. *IEEE/ACM ToN*, August 1998.
- [9] M. A. El-Gendy, A. Bose, and K. G. Shin. Evolution of the Internet QoS and support for soft real-time applications. *Proceedings of the IEEE*, 91(7):1086–1104, July 2003.
- [10] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True. Deriving traffic demands for operational IP networks: Methodology and experience. In *ACM SIGCOMM*, June 2000.
- [11] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM ToN*, August 1993.
- [12] M. Grossglauser and D. N. C. Tse. A framework for robust measurement-based admission control. In *IEEE/ACM ToN*, 1999.
- [13] E. Horvitz, J. Apacible, R. Sarin, and L. Liao. Prediction, expectation, and surprise: Methods, designs, and study of a deployed traffic forecasting service. In *Conference on UAI*, July 2005.
- [14] J. Jamin, P. B. Danzig, S. Shenker, and L. Zhang. A measurement-based admission control algorithm for integrated services packet networks. *IEEE/ACM ToN*, February 1996.
- [15] S. Kandula, D. Katabi, M. Jacob, and A. Berger. Botz-4-Sale: Surviving organized DDoS attacks that mimic flash crowds. In *ACM/USENIX NSDI*, May 2005.
- [16] K. Lakshminarayanan, D. Adkins, A. Perrig, and I. Stoica. Taming IP packet flooding attacks. In *ACM HotNets Workshop*, 2003.
- [17] X. Li, F. Bian, M. Crovella, C. Diot, R. Govindan, G. Iannaccone, and A. Lakhina. Detection and identification of network anomalies using sketch subspaces. In *ACM/USENIX IMC*, October 2006.
- [18] Z. M. Mao, V. Sekar, O. Spatscheck, J. Van der Merwe, and R. Vasudevan. Analyzing large ddos attacks using multiple data sources. In *ACM LSAD Workshop*, pages 161–168, November 2006.
- [19] B. Raghavan and A. C. Snoeren. A system for authenticated policy-compliant routing. In *ACM SIGCOMM*, October 2004.
- [20] M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates, and Y. Zhang. Experience in measuring internet backbone traffic variability: models, metrics, measurements and meaning. In *International Teletraffic Congress (ITC)*, November 2003.
- [21] S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Network support for IP traceback. *IEEE/ACM ToN*, 9(3), June 2001.
- [22] A. C. Snoeren, C. Partridge, L. A. Sanchez, C. E. Jones, F. Tchakountio, B. Schwartz, S. T. Kent, and W. T. Strayer. Single-packet IP traceback. *IEEE/ACM ToN*, 10(6):721–734, December 2002.
- [23] S. Sun, C. Zhang, and Y. Zhang. Traffic flow forecasting using a spatio-temporal bayesian network predictor. In *International Conference on Artificial Neural Networks*, September 2005.
- [24] P. Verkaik, O. Spatscheck, J. V. der Merwe, and A. C. Snoeren. Primed: community-of-interest-based ddos mitigation. In *ACM LSAD Workshop*, pages 147–154, November 2006.
- [25] A. Yaar, A. Perrig, and D. Song. Pi: A path identification mechanism to defend against DDoS attacks. In *IEEE Security and Privacy Symposium*, pages 93–107, May 2003.
- [26] A. Yaar, A. Perrig, and D. Song. An endhost capability mechanism to mitigate DDoS flooding attacks. In *IEEE Security and Privacy Symposium*, May 2004.