CDC02- REG1097

Optimal Filtering for Denial of Service Mitigation

Stephan Bohacek
University of Southern California
bohacek@math.usc.edu

Abstract

An optimal approach to mitigation of flooding denial of service attacks is presented. The objective is to minimize effect of the mitigation while protecting the server. The approach relies on routers filtering enough packets so that the server is not overwhelmed while ensuring that as little filtering is performed as possible. The optimal solution is to filter packets at routers through which the "attack packets" are passing. The identification of which router the packets are passing is carried out by routers filtering a small but time varying fraction of the packets. The arrival of packets at the server is correlated to router filtering providing an indication through which routers the attack packets are passing. Once sufficient confidence in the identification is achieved, the routers that forward more attack packets filter more packets than router that forward less attack packets.

1 Introduction

There have been major attacks on important Internet servers such as the February 7, 2000 attack of the Yahoo servers and the following days attacks on ZDNet.com, eBay, Amazon.com, Buy.com and CNN.com. These attacks were classified as SYN floods. An attacker carries out a SYN flood [4], [6] by sending packets that request a TCP connection. The server response by sending an acknowledgment packet and reserves system resources for this connection. The attacker does not respond to the acknowledgment, but instead continues to send SYN packets. The server reserves system resources for each SYN and continues to reserve the resources for 3 minutes. Thus, if the SYN packets are sent sufficiently fast, the servers resources will be entirely allocated to these malicious SYNs.

Other attacks that are similar to the SYN attack include UDP and ICMP echo request floods. Like the SYN attack, the attacker or attackers send many packets to the host under attack. Depending on the details of the attack, there are two possible effects of the attack. Either the server resources can be exhausted by responding to the attack packets or the data link that carries the packets will become congested. In either case, the presence of a large number of packets results in the disruption of intended services. This paper will focus on the SYN attack. However, since the objective is to stop malicious packets from arriving at the server, this defense is also useful for stopping other types of flooding attacks.

Since all Internet packets have return addresses, one defensive strategy is to ignore packets from addresses that send too may SYNs. However, since this return addresses are easily falsified (return address spoofing), this defense has a limited effect. One approach to counteract spoofed return addresses is to require that all gateways ensure that the return address is correct. This approach is known as egress filtering [5]. While this approach is very effective, it requires that each subnet implement such filtering. In the current heterogeneous Internet, with no central authority, there is little hope of all subnets would implement such filtering. Even in the case where a large fraction of the subnets have implemented egress filtering address spoofing would work for senders in the unregulated subnet. Hence, the attacker only needs to gain control of hosts in these unregulated subnets. It is this scenario that the present effort is most useful. In particular, the attacks is not extremely widely spread since only a small fraction of subnets are unregulated. In this setting, locating and isolating these relatively few subnets is a reasonable objective.

Another approach to locating and isolating the hosts or subnets responsible for a DoS attack is for a router to randomly tag packets. The server under attack then correlates the attack packets with the tags and determine through which routers the packets are passing through. This approach is known as IP traceback [8] and is similar to the approach taken here one difference is that the approach here does not require any modification to the IP format to enabling a field for the router tag. Furthermore, the approach presented here directly addresses the problem of server overflow[1]. In particular, the tagging approach provides some mechanism for locating through which router the attack is passing through. Other, similar approaches include controlled flooding [3] and input debugging [9]. However, there has been no effort focused on determining the optimal approach to utilizing this localizing information. Here we present an optimal approach that both eliminates the attack and minimizes the effect on other non-attackers. It is conceivable that the approach presented here could be extended to the tagging approach.

Note that is very difficult to differentiate between an attack and naturally occurring extreme traffic. However, from the server's point of view, the effect is the same; the link is heavily congested and/or the server is overburdened. In either case, the server is not able to process all the packets and many packets are neglected. Furthermore, the extreme traffic may lead to many connections being prematurely terminated. The approach here filters packets based on whether the router through which they pass also forwards many other packets of this type. Thus, if a large number of SYN packets pass through a router, then this router will drop a large fraction of these packets. In the case of an attack, this action is likely the best thing to do. In the case of extreme traffic (not an attack), then this is not necessarily the best thing to do. However, since the server and/or link is overburdened and hence packets must be dropped, there is no compelling reason to drop packets from one router or another. Hence, there is no reason against dropping packet as described here. The follows refers to attack packets while these may actually be non-malicious packets but merely part of a traffic extreme.

The paper proceeds as follows. The next section details the approach. Section 3 provides an example. Section 4 provides some concluding remarks. The proofs of the theorems of Section 2 are withheld until Section 4.1.

2 Method

There are multiple parts of the defense system. This paper focuses on the workings of the mitigation module. This module is activated by a detection module that determines that an attack is underway. Detection is a active field. Optimal sequential approached have been developed in [2]. The mitigation module discussed here need not run on the same system as the host under attack (HUA). All that is required is that the mitigation module can record all the packets that arrive at the HUA. Hence, the mitigation module can merely be on the same LAN as the HUA.

The objective of this defense is for the routers to filter (drop) malicious packet before they arrive at the HUA. However, this filtering must be carried out in such a way that the effect on non-malicious packets is minimized. In order to accommodate these two objectives, it is necessary to identify the router through which the attack is coming and, if there is sufficient confidence in the identification, filter packets at this router.

After the attack is detected, the mitigation module commands the set of routers that are one hop away to randomly drop SYNs destined for the host under attack (HUA). The probability of a router dropping a packet varies with time in such a way that the drops on one router are uncorrelated to the drops in other routers. The mitigation module correlates the arrival of SYNs with the router drop probabilities and decides through which of the routers the attack is coming[2]. Based on the confidence of the decision, the mitigation module informs some routers to drop more SYNs and others to drop less. Once enough confidence has been gained that the correct router

[1] In the case of UDP or ICMP flooding, it is link congestion, not server overflow that is of concern. This work can easily be extended to the case where link congestion is the objective of teh attack.

[2] The rate at which the drop rate varies can be chosen to be much larger than the packet latency, so that latency can be neglected.

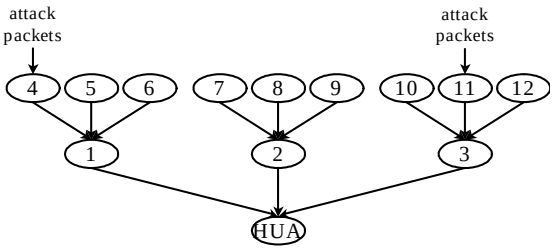

Figure 1: Once the attack is detected, the mitigation module informs the routers one hop away to randomly drop SYNs destined for the HUA, that is, the routers labeled 1,2 and 3 perform this random dropping. After some time, confidence is gained that the attack is coming though routers 1 and 3. At this point the filtering is begun by routers two hops away from the HUA, and which lead to routers 1 and 3, that is 4, 5, 6, 10, 11 and 12. This process is repeated until the attackers are isolated.

or routers have been identified, the filter process is repeated but for the router which are two hops away and lead to the routers which the attack passes through. This process is repeated until the attackers are isolated. This process is depicted in Figure 1. The objective of this paper is to determine a filter strategy. Issues regarding which sequence of routers that carry out the filtering are not discussed. All that is required is that at any time, the set of routers that perform filtering are such that they *exactly* isolate the HUA. For example, referring to Figure 1, $\{1, 2, 3\}$, $\{1, 2, 7, 8, 9\}$ exactly isolate the HUA, while $\{1, 2, 3, 7\}$ does not exactly isolate the HUA.

One approach to filtering is to occasionally filter all SYN packets passing through a router. Then, if the SYN attack ceases, it can be concluded that the attack is coming through the router. However, if the attack is not coming through that router, then some "good" SYNs are needlessly dropped. Furthermore, it is not clear how long the router must drop SYNs to ensure that SYN arrivals at the HUA vary enough to detect the attack. For example, a stealthy attack could send the SYNs at a moderate rate so as to overflow the HUA, but not so many SYNs that the stream of SYNs is easily detected. Note that, it does not take that many half-open connections to disable a server. For example, an older version of BSD Unix server may crash with as little as 3000 half-open connections. Since each connection remains open for up to 3 minutes, the attacked must send only 3000 SYNs in three minutes. Thus, if the router simply drops all SYNs, then it may take on the order of a minute to determine if this is the correct router or not. In the case that it is not the correct router, or in the case where the attack has stopped all together, many good SYNs will be dropped.

The approach presented here is similar to the brut-force approach discussed above, but attention is paid to protecting both the router and well-behaved users. Instead of dropping all SYNs, only a small fraction of SYNs are dropped. Then, based on the arrival of SYNs a filter is designed that guarantees that with some probability near one, the server will not overflow and that minimize the effect on good users. Note that in the case of an attack that repeatedly stops and starts, this method will repeatedly readjust the filter to track the attackers behavior. Let f_i be the fraction of packets filtered by router i. Then, the optimal filter satisfies

$$\begin{aligned} &\min \sum_{i=1}^{n} f_i \\ &\text{subject to}: \ \text{Prob}\left(\text{Server overflow}|\ \text{Observations}\right) < 1 - \alpha \\ &0 \le f_i \le 1 \text{ for } i = 1, \cdots, n \end{aligned}$$

where $0 < \alpha < 1$ and there are n routers that are carrying out the filtering. The observations are SYN arrivals. This optimization problem is to be solved at all times.

In solving this optimization, the probability of server overflow given the observations of arrivals

must be determined. We assume that the SYNs arrivals at router i can be modeled as a Poisson process with unknown rate λ_i. That is, for small Δt, the probability of a packet arriving during the time interval Δt, is $\Delta t \lambda_i$. This modeling assumption is in accordance with measurements [7]. With the filtering in place, SYNs leave router i at rate $f_i \lambda_i$. We assume that the filtering ratios varying around some mean filtering ratio $\bar{f}_i$. Hence, router i filters a fraction $\tilde{f}_i + \bar{f}_i$ of SYNs. Thus the average total ratio of SYN arrivals at the HUA is given by $\bar{F}'\Lambda$, where $\bar{F}$ is the column vector of $\bar{f}_i$'s and Λ is the column vector of λ_i's. The optimization problem above is given by

$$\begin{aligned} &\min \sum_{i=1}^{n} \bar{f}_i \\ &\text{subject to : } \operatorname{Prob}\left(\bar{F}'\Lambda < L_{\max} | O\right) < 1 - \alpha, \\ &0 \le f_i \le 1, \end{aligned} \tag{1}$$

where $L_{\max}$ is the maximum rate at which the host can process SYNs and O is the set of observations so far.

The unknown arrival rates λ_i can be estimated with least squares. Define

$$H := \begin{bmatrix} \tilde{f}_1(1) + \bar{f}_1 & \cdots & \tilde{f}_1(M) + \bar{f}_1 \\ \vdots & & \vdots \\ \tilde{f}_n(1) + \bar{f}_n & \cdots & \tilde{f}_n(M) + \bar{f}_n \end{bmatrix}^T \begin{bmatrix} \tilde{f}_1(1) + \bar{f}_1 & \cdots & \tilde{f}_1(M) + \bar{f}_1 \\ \vdots & & \cdots \\ \tilde{f}_n(1) + \bar{f}_n & \cdots & \tilde{f}_n(M) + \bar{f}_n \end{bmatrix},$$

where there are M number of observations. Define

$$\left\langle \tilde{f}_i, \tilde{f}_j \right\rangle_M := \frac{1}{M} \sum_{k=0}^{M} \tilde{f}_i(k) \tilde{f}_j(k)$$

$$\left\langle 1, \tilde{f}_i \right\rangle_M := \frac{1}{M} \sum_{k=0}^{M} \tilde{f}_i(k).$$

If $\left\langle 1, \tilde{f}_j \right\rangle_M = 0$ and $\left\langle \tilde{f}_i, \tilde{f}_j \right\rangle_M = 0$ for $i \neq j$, which is true asymptotically when $\tilde{f}_i$ have zero mean and $\tilde{f}_i$ and $\tilde{f}_j$ are uncorrelated for $i \neq j$, then

$$H = M \begin{bmatrix} \left\langle \tilde{f}_1, \tilde{f}_1 \right\rangle_M & \bar{f}_1 \bar{f}_2 & \cdots \\ \bar{f}_2 \bar{f}_1 & \left\langle \tilde{f}_2, \tilde{f}_2 \right\rangle_M & \\ \vdots & & \ddots \end{bmatrix}.$$

Define

$$B := M \begin{bmatrix} \frac{1}{M} \sum_{k=1}^{M} \left(\left(\tilde{f}_1(k) + \bar{f}_1 \right) 1_{\{SYN \text{ arrived at tine } k\}} \right) \\ \vdots \\ \frac{1}{M} \sum_{k=1}^{M} \left(\left(\tilde{f}_n(k) + \bar{f}_n \right) 1_{\{SYN \text{ arrived at tine } k\}} \right) \end{bmatrix}.$$

Then the least-squares estimate of actual SYN arrival rates, Λ^*, is $\hat{\Lambda} = H^{-1}B$. With this estimate of the SYN arrivals, the optimal filter can be determined as follows.

The optimal filter for (1) is given by

$$F = G^{-1}\left(L_{\max} \hat{\Lambda} + \mu \vec{1} \right),$$

where $G = \left(\hat{\Lambda}\hat{\Lambda}' - \gamma s^2 H^{-1}\right)$, μ solves

$$0 = \left(\frac{1}{2}\right)^2 \mu^2 \vec{1}^T G \vec{1} + \mu \left(\frac{1}{2}\right)^2 (L - 2L) \left(\hat{\Lambda}\frac{2}{\gamma s^2}\right)^T G\vec{1} \tag{2}$$
$$- \left(\left(\frac{1}{2}\right) L - \left(\frac{1}{2}\right)^2 L\right) \left(\frac{2}{\gamma s^2}\hat{\Lambda}\right)^T G \left(\hat{\Lambda}\frac{L2}{\gamma s^2}\right) - L^2 \frac{1}{\gamma s^2},$$

γ is such that $P\left(t_{M-n} \geq \gamma\right) = 1 - a$, where t_{M-m} is a random variable distributed according to the student's t distribution with $M - m$ degrees of freedom, and s is given by

$$s^2 = \frac{1}{M-n} \sum_{k=0}^{M} \left(\left[\begin{array}{ccc} f_1(k) & \cdots & f_n(k) \end{array} \right] \hat{\Lambda} - 1_{\{SYN \text{ arrived at time } k\}} \right)^2. \tag{3}$$

The above theorem assumes that each router can and perhaps should filter packets. It is possible that some prior information leads to the conclusion that some routers should not or cannot implement any filtering. For generality we simply assume that some of the filtering ratios are predetermined, i.e. F is restricted so that $F_i = f_i$ for $i \in I$. The optimization problem becomes

$$\begin{aligned} &\min_{\bar{F}} -\bar{F}\vec{1} \\ &\text{subject to Prob}\left(F'\Lambda < L_{\max} | O\right) < 1 - \alpha, \\ &0 \leq F \leq 1,\ F_i = f_i \text{ for } i \in I, \end{aligned} \tag{4}$$

where, $F = S\tilde{F} + \bar{F}$, where $\bar{F}_i = f_i$ for $i \in I$ and $\bar{F}_i = 0$ for $i \notin I$. and $\tilde{F} \in \mathbb{R}^{n-|I|}$, where $|I|$ is the number of elements in I and $S \in \mathbb{R}^{n \times (n-|I|)}$ is a matrix of 0's and 1's

The optimal filter for (4) is given by

$$F = \left(S^T G S\right)^{-1} \left(L_{\max} S^T \hat{\Lambda} + \mu\vec{1} - S^T G \bar{F}\right),$$

where

$$\mu = \left(\frac{1}{\vec{1}^T \left(S^T G S\right)^{-1} \vec{1}}\right) \left(L_{\max}^2 \left(\hat{\Lambda}^T S \left(S^T G S\right)^{-1} S^T \hat{\Lambda} - 1\right)\right. \tag{5}$$
$$\left. - \left(2L_{\max}\hat{\Lambda}^T S - \bar{F}^T G S\right) \left(S^T G S\right)^{-1} S^T G \bar{F} + 2L_{\max}\hat{\Lambda}^T \bar{F} - \bar{F}^T G \bar{F}\right)$$

with G and s the same as in Theorem 2.

An application of this theorem arises when a set of filtering indicates that only a subset of the routers are forwarding attack packets. For example, referring to Figure 1, suppose that it has been determined that attack packets are not passing through router 2 and are passing through routers 1 and 3. The next set of filtering can focus on filtering in routers 4, 5, 6, 10, 11, and 12. Since packets still arrive through router 2, the arrival rate at router 2 must be determined. However, since router 2 should perform not perform any filtering, the above provides the optimal filter. Another important application of this theorem is in the case when some upstream does not perform any filtering. Then these routers can be treated as one router with $f = 0$.

Note that for γ large $G \to -\gamma s^2 H^{-1}$ and $\left(\Lambda^T G^{-1} \Lambda - 1\right) \to -1$ and $\vec{1}^T G^{-1} \vec{1} \to < 0$. In this case, the real solutions to (2) and (5) exist. However, for γ small enough, μ may be imaginary. That is, there may not be an optimal solution. This is the case when the constraints $0 \leq F \leq 1$ come into play. There are two options. One is to increase γ and slowly decrease until some F_i cross the boundary. Once the i is determined, a similar problem is solved but with F_i set to the constraint 0 or 1. For example, if, as γ is decreased, $F_1 < 0$, then F_1 is set to zero, and the optimization is carried out with the remaining F_i's. Alternatively, one can try different each combinations of F_i to 0 or 1.

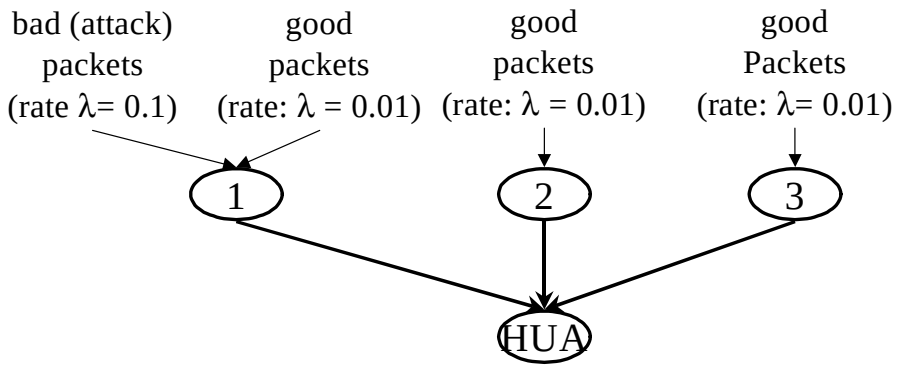

Figure 2: The filtering experiment was performed on the above topology. Both bad SYNs and good SYNs pass through router 1. The rate of bad SYNs is 0.1 while the rate of good SYNs is 0.01. Only good SYNs pass through routers 2 and 3, and they pass at a rate of 0.01.

3 Example

Next we present results of an example of this approach. We concentrate on the simple topology shown in Figure 2. The results are shown in Figures 3 - 6. Figure 3 shows the effect of the filtering on the good SYNs, while Figure 4 shows the effect of the filtering on the bad SYNs. Figure 6 shows how the filtering ratios vary with time. Initially, each router filters the same amount. This is because it is not known through which router the attack is coming, so all routers filter equally so as to protect the HUA. Later, it becomes clear that the attack is coming through router 1. Hence, there is no need to filter SYNs passing through routers 2 and 3. Note that even after time step 200, when there is much certainty that the attack is coming through router 1, router 1 does not greatly increase its filter ratio. The reason for this is that the server can withstand some bad SYNs. So there is no need to filter out all bad SYNs. Furthermore, allowing some bad SYNs to pass also allows the good SYNs that pass through router 1 to also pass. In this way as many good SYNs reach the router as possible. Note that the filter ratio constantly varies. This allows the filter to quickly react to an attack that stops. Figure 5 shows the estimate of the arrival rates at each time step. Note that initially, router 1 is estimated to have a high arrival rate, but shortly later router 2 is estimated to have a high arrival rate. However, the filter ratios do not reflect these estimates. The reason for this is that there is little confidence in estimates, hence no action is taken. It is not until time step 75 that it is estimated that router 1 has a high arrival rate and there is substantial confidence in this estimate. As the confidence increases, the filter takes further action and decreases the filtering on routers 2 and 3. Note that as confidence is gained, and the filters are increased or decreased, the effect of the filtering is stronger leading to greater confidence in the identification of the router through which the attack is passing. This effect can be seen in Figure 6, where the filter rapidly vary after the 50^{th} time step.

4 Conclusions

While there has been extensive work focused on detecting the origin of denial of service attacks, there has been little work on real-time DoS mitigation. This paper presents an approach to DoS mitigation that balances the need to protect the server and minimize the effect on non-attack packet. In most cases, the computation involved in determining the filtering is computationally straight forward and the mitigation occurs rapidly. A drawback to the approach is that the routers must filter packets. Furthermore, in order to minimize the effect on non-attack packets, the routers should filter according to packet type, e.g., TCP-SYN, UDP, etc. Thus, the routers must examine the contents of the packet and make a decision as to whether to drop the packet. This requires

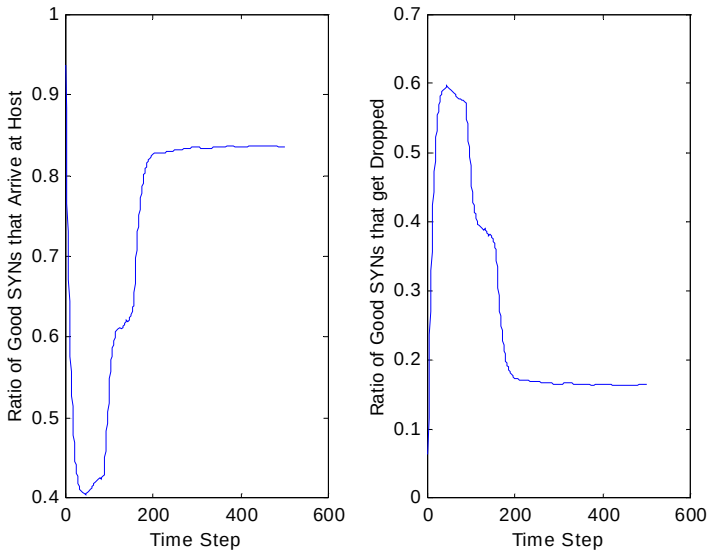

Figure 3: Above shows the effects of the filtering on the good SYNs. The plot on the left shows the ratio of SYNs that arrive at the host, while the plot on the right shows the ratio of good SYNs that where dropped. Note that at time step 300, the filtering changes to permit most of the good SYNs to pass to the router. However, some good SYNs are still filetred. This is due to the fact that some good SYNs pass through router 1 and are filtered out along with the bad SYNs.

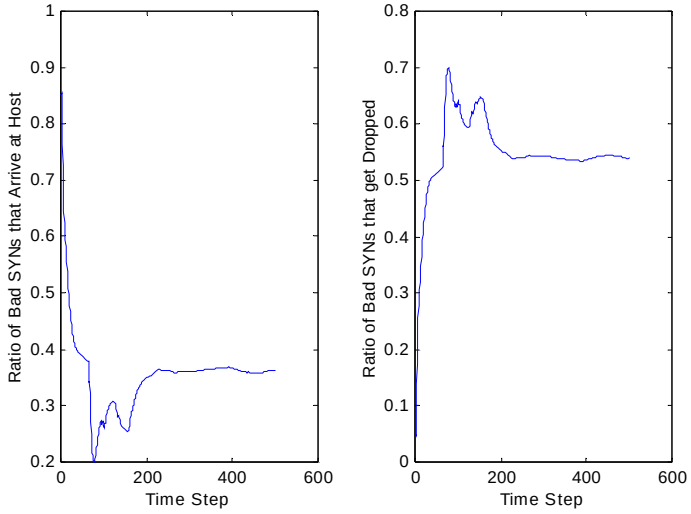

Figure 4: Above shows the effect of the filtering on the bad SYNs. Note that shortly after the attack is detected, extensive filtering takes places and filter a large fraction of the bad SYNs.

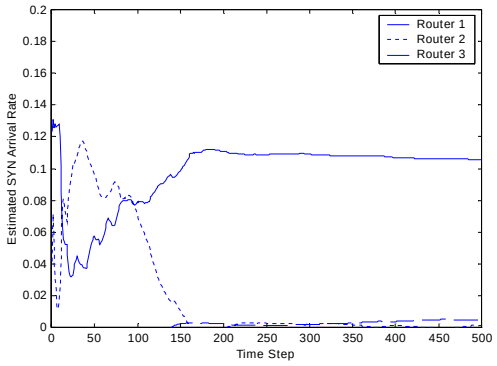

Figure 5: The above shows the estimated rate of SYN arrivals at each router.

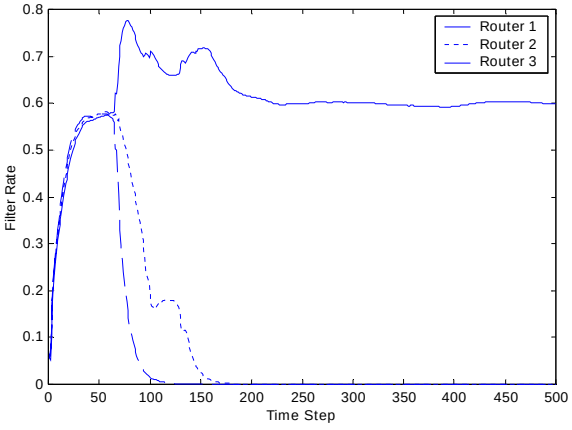

Figure 6: Above shows the time series of the average filter rate for each router.

considerable processing by the routers. However, routers today are beginning to focus on valued added features rather than strictly on bandwidth. Hence, such DoS mitigation could be carried out by future generation routers. Another drawback is that this approach requires that the mitigation module know the topology and which routers to filter packets. While this is not realistic for the common user, it is not unreasonable to large web server farms. Another consideration is the communication between the host and the router must be secure and authenticated. Clearly, if an attacker could use this approach to convince a router to filter all SYN packets destined for a server, then this approach could be used to implement a DoS attack.

The work presented here is an initial step toward DoS mitigation. Extensive work remains before its implementation. For example, since it is unlikely that all routers would be capable of implementing this approach, the effectiveness of the approach when not all routers are capable of filtering packets must be addressed. Future work will also develop a mitigation approach that uses IP traceback in combination with router filtering. Lastly, the method above focused on the average filter ratio with the constraint that the time-varying component of the filter was predetermined. Future work will focus on the better design of the time-varying component of the filters.

4.1 Proofs of Theorems

4.1.1 Proof of Theorem 2.

Once the estimate for Λ is determined, the variance of the error in this estimate can be estimated via

$$s^2 = \frac{1}{M-n}\sum_{k=0}^{M}\left(\left[\begin{array}{ccc} f_1(k) & \cdots & f_n(k) \end{array}\right]\hat{\Lambda} - 1_{\{SYN \text{ arrived at time } k\}}\right)^2.$$

Once Λ has been estimated, one can estimate which λ_i is largest, i.e., through which router are the most SYNs arriving. Of course, there may be some error in this conclusion. The probability of not making an error can be found with use of confidence intervals. Suppose that $\hat{\lambda}_1 > \hat{\lambda}_i$ for $i \neq j$. Define $\hat{d}_i = \lambda_1 - \lambda_i$. Since $\left[\begin{array}{ccc} 1 & -1 & 0 & \cdots \end{array}\right]\Lambda^* > 0$ implies that $\lambda_1^* > \lambda_2^*$, if $\left[\begin{array}{ccc} 1 & -1 & 0 & \cdots \end{array}\right]\left(\hat{\Lambda} - \Lambda^*\right) < \hat{d}_2$ then $\lambda_1^* > \lambda_2^*$. Similarly, if

$$\begin{array}{c} \dfrac{\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]\left(\hat{\Lambda} - \Lambda^*\right)\left(\hat{\Lambda} - \Lambda^*\right)^T\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]^T}{s^2\left(\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]H^{-1}\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]^T\right)^{-1}} \\ < \hat{d}_2^2 s^2\left(\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]H^{-1}\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]^T\right)^{-1}, \end{array} \tag{6}$$

then $\lambda_1^* > \lambda_2^*$. The left hand side of (6) has the Student's t distribution with $M-n$ degrees of freedom. Therefore, the probability of making an incorrect conclusion is bounded by

$$P\left(t_{M-n} \geq \hat{d}_2^2 s^2\left(\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]H^{-1}\left[\begin{array}{ccc} 1 & -1 & 0 \end{array}\right]^T\right)\right),$$

where t_{M-n} has the Student's t distribution with $M-n$ degrees of freedom.

The average rate at which SYN's arrive is $\sum_{i=1}^{n}\bar{f}_i\lambda_i$. It is required that $\sum_{i=1}^{n}\bar{f}_i\lambda_i < L_{\max}$. Once Λ is estimated, one can choose $\bar{f}$ such that $\sum_{i=1}^{n}\bar{f}_i\hat{\lambda}_i < L_{\max} - d$. In this case $\sum_{i=1}^{n}\bar{f}_i\lambda_i^* < L_{\max}$ whenever $\sum_{i=1}^{n}\bar{f}_i\left(\hat{\lambda}_i - \lambda_i^*\right) < d$, or $\bar{F}\left(\hat{\Lambda} - \Lambda^*\right)\left(\hat{\Lambda} - \Lambda^*\right)^T\bar{F}^T < d^2$. Similarly, $\sum_{i=1}^{n}\bar{f}_i\lambda_i^* < L_{\max}$ whenever

$$\frac{\left(\hat{\Lambda} - \Lambda^*\right)\bar{F}^T\bar{F}\left(\hat{\Lambda} - \Lambda^*\right)^T}{s^2\left(\bar{F}H^{-1}\bar{F}^T\right)} < \frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1}. \tag{7}$$

The left hand side of (7) has the Student's t distribution with $M-n$ degrees of freedom. Hence the probability of $\sum_{i=1}^{n}\bar{f}_i\lambda_i^* > L_{\max}$ is bounded by

$$P\left(t_{M-n} \geq \frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1}\right),$$

where t_{M-n} has the Student's t distribution with $M-n$ degrees of freedom. Clearly this probability depends on the choice of $\bar{F}$. Since a small $\bar{f}_i$ implies that the rate through router i is small, the larger $\bar{f}_i$, the less the obtrusive. Hence the objective is

$$\begin{aligned} &\max \sum \bar{f}_i \\ &\text{subject to } P\left(t_{M-n} \geq \frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1}\right) = 1-\alpha, \\ &0 \leq \bar{F} \leq 1 \\ &\bar{F}\hat{\Lambda} = L_{\max} - d \end{aligned}$$

where α is a design parameter. Given α and $M-n$, $P\left(t_{M-n} \geq \frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1}\right) = 1-\alpha$ when $\frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1} = \gamma$, where γ is given by tables etc. Hence the above maximization problem is

$$\begin{aligned} &\max \sum \bar{f}_i \\ &\text{subject to } \frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1} = \gamma \end{aligned}$$

or

$$\begin{aligned} &\max \bar{F}\vec{1} \\ &\text{subject to } \frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1} = \gamma \end{aligned}$$

where $\vec{1} = \begin{bmatrix} 1 & \cdots & 1 \end{bmatrix}^T$. Or

$$\begin{aligned} &\min -\bar{F}\vec{1} \\ &\text{subject to } \left(\bar{F}H^{-1}\bar{F}\right)^{-1} - \gamma\frac{s^2}{d^2} = 0. \end{aligned}$$

Lagrange multiplier theory implies that the optimal $\bar{F}$ occurs when

$$\begin{aligned} -\vec{1} + \mu H^{-1}\bar{F}^T &= 0 \\ \left(\bar{F}H^{-1}\bar{F}^T\right)^{-1} - \gamma\frac{s^2}{d^2} &= 0. \end{aligned}$$

The first equation implies that

$$\bar{F}^T = \frac{1}{\mu}H\vec{1}.$$

Hence$\frac{1}{\mu^2}\vec{1}^T H H^{-1} H\vec{1} = \gamma\frac{s^2}{d^2}$,or $\sqrt{\frac{d^2}{\gamma s^2}\vec{1}^T H\vec{1}} = \mu$.Therefore

$$\bar{F}^T = \frac{1}{\sqrt{\frac{d^2}{\gamma s^2}\vec{1}^T H\vec{1}}}H\vec{1}.$$

The objective

$$\begin{aligned} &\max_{F,d} \sum \bar{f}_i \\ &\text{subject to } P\left(t_{M-n} \geq \frac{d^2}{s^2}\left(\bar{F}H^{-1}\bar{F}^T\right)^{-1}\right) = 1-\alpha, \\ &0 \leq \bar{F} \leq 1 \\ &\bar{F}\hat{\Lambda} = L_{\max} - d \end{aligned}$$

or

$$\begin{aligned}
&\max_{F,d} \sum \bar{f}_i \\
&\text{subject to } \frac{d^2}{\gamma s^2} = \left(\bar{F} H^{-1} \bar{F}^T\right), \\
&0 \leq \bar{F} \leq 1,\, d \geq 0 \\
&\bar{F}\hat{\Lambda} = L_{\max} - d
\end{aligned}$$

Hence $d = L_{\max} - \bar{F}\hat{\Lambda}$ and $d^2 = L_{\max}^2 + \bar{F}\hat{\Lambda}\hat{\Lambda}^T\bar{F}^T - 2L_{\max}\bar{F}\hat{\Lambda}$. Thus the first condition becomes

$$\begin{aligned}
0 &= \bar{F} H^{-1} \bar{F}^T - \left(L_{\max}^2 + \bar{F}\hat{\Lambda}\hat{\Lambda}^T\bar{F}^T - 2L_{\max}\bar{F}\hat{\Lambda}\right)\frac{1}{\gamma s^2} \\
&= \bar{F}\left(H^{-1} - \frac{s^2}{\gamma}\hat{\Lambda}\hat{\Lambda}^T\right)\bar{F}^T - L_{\max}^2\frac{1}{\gamma s^2} + 2L_{\max}\bar{F}\hat{\Lambda}\frac{1}{\gamma s^2}
\end{aligned}$$

Therefore the objective is

$$\begin{aligned}
&\min_{\bar{F}} -\bar{F}\vec{1} \\
&\text{subject to } \bar{F}\left(H^{-1} - \frac{s^2}{\gamma}\hat{\Lambda}\hat{\Lambda}^T\right)\bar{F}^T - L_{\max}^2\frac{1}{\gamma s^2} + 2L_{\max}\bar{F}\hat{\Lambda}\frac{1}{\gamma s^2} = 0 \\
&L_{\max} - \bar{F}\hat{\Lambda} \geq 0, 0 \leq F \leq 1
\end{aligned}$$

Note that the set $\left\{F : \bar{F}\left(H^{-1} - \frac{s^2}{\gamma}\hat{\Lambda}\hat{\Lambda}^T\right)\bar{F}^T - L_{\max}^2\frac{1}{\gamma s^2} + 2L_{\max}\bar{F}\hat{\Lambda}\frac{1}{\gamma s^2} = 0\right\}$ is not convex. Lagrange multiplier theory [1] implies that the optimal solution obeys

$$\begin{aligned}
\left(2\left(H^{-1} - \frac{1}{\gamma s^2}\hat{\Lambda}\hat{\Lambda}^T\right)\bar{F}^T + 2L_{\max}\hat{\Lambda}\frac{1}{\gamma s^2}\right) - \frac{1}{\mu}\vec{1} = 0 \qquad (8) \\
\bar{F}\left(H^{-1} - \frac{1}{\gamma s^2}\hat{\Lambda}\hat{\Lambda}^T\right)\bar{F}^T - L_{\max}^2\frac{1}{\gamma s^2} + \left(2L_{\max}\frac{1}{\gamma s^2}\right)\hat{\Lambda}^T\bar{F}^T = 0
\end{aligned}$$

The first equation implies that

$$\bar{F}^T = \frac{1}{2}\left(H^{-1} - \frac{s^2}{\gamma}\hat{\Lambda}\hat{\Lambda}^T\right)^{-1}\left(\frac{1}{\mu}\vec{1} - 2L_{\max}\hat{\Lambda}\frac{s^2}{\gamma}\right) \qquad (9)$$

Then the second implies that

$$\begin{aligned}
0 = &\left(\frac{1}{2}\right)^2\left(\frac{1}{\mu}\right)^2\vec{1}^T\left(H^{-1} - \frac{1}{\gamma s^2}\hat{\Lambda}\hat{\Lambda}^T\right)^{-1}\vec{1} + \frac{1}{\mu}\left(\frac{1}{2}\right)^2(L - 2L)\left(\hat{\Lambda}\frac{2}{\gamma s^2}\right)^T\left(H^{-1} - \frac{1}{\gamma s^2}\hat{\Lambda}\hat{\Lambda}^T\right)^{-1}\vec{1} \\
&- \left(\left(\frac{1}{2}\right)L - \left(\frac{1}{2}\right)^2 L\right)\left(\frac{2}{\gamma s^2}\hat{\Lambda}\right)^T\left(H^{-1} - \frac{1}{\gamma s^2}\hat{\Lambda}\hat{\Lambda}^T\right)^{-1}\left(\hat{\Lambda}\frac{L2}{\gamma s^2}\right) - L^2\frac{1}{\gamma s^2}
\end{aligned}$$

This equation easily gives a solution for μ which can then be substituted into (9) yielding $\bar{F}$.

$$
\begin{aligned}
0 &= \left(\frac{1}{2}\right)^2 \left(\frac{1}{\mu}\vec{1} - 2L\hat{\Lambda}\frac{s^2}{\gamma}\right)^T G \left(\frac{1}{\mu}\vec{1} - 2L^2\hat{\Lambda}\frac{s^2}{\gamma}\right) - L^2\frac{s^2}{\gamma} + \frac{1}{2}\left(2L\frac{s^2}{\gamma}\hat{\Lambda}\right)^T G \left(\frac{1}{\mu}\vec{1} - 2L\hat{\Lambda}\frac{s^2}{\gamma}\right) \\
&= \left(\frac{1}{2}\right)^2 \left(\frac{1}{\mu}\right)^2 \vec{1}^T G\vec{1} + \left(\frac{1}{2}\right)^2 \left(2L\hat{\Lambda}\frac{s^2}{\gamma}\right)^T G \left(2L\hat{\Lambda}\frac{s^2}{\gamma}\right) - 2\left(\frac{1}{2}\right)^2 \frac{1}{\mu}\left(2L\hat{\Lambda}\frac{s^2}{\gamma}\right)^T G\vec{1} \\
&- L\frac{s^2}{\gamma} - \left(\frac{1}{2}\right)\left(2L\frac{s^2}{\gamma}\hat{\Lambda}\right)^T G \left(2L^2\hat{\Lambda}\frac{s^2}{\gamma}\right) + \frac{1}{\mu}2\left(\frac{1}{2}\right) L\frac{s^2}{\gamma}\hat{\Lambda}G\vec{1} \\
&= \left(\frac{1}{2}\right)^2 \left(\frac{1}{\mu}\right)^2 \vec{1}^T G\vec{1} + \frac{1}{\mu}\left(\frac{1}{2}\right)^2 (L - 2L)\left(\hat{\Lambda}\frac{2}{\gamma s^2}\right)^T G\vec{1} \\
&- \left(\left(\frac{1}{2}\right)L - \left(\frac{1}{2}\right)^2 L\right)\left(\frac{2}{\gamma s^2}\hat{\Lambda}\right)^T G\left(\hat{\Lambda}\frac{L2}{\gamma s^2}\right) - L^2\frac{1}{\gamma s^2},
\end{aligned}
$$

yielding the desired result.

4.1.2 Proof of Theorem 2

Now suppose that F is restricted so that $F_i = f_i$ for $i \in I$. The the problem becomes

$$
\begin{aligned}
&\min_F -F\vec{1} \\
&\text{subject to } F\left(H^{-1} - \frac{s^2}{\gamma}\hat{\Lambda}\hat{\Lambda}^T\right)F^T - L_{\max}^2\frac{1}{\gamma s^2} + 2L_{\max}F\hat{\Lambda}\frac{1}{\gamma s^2} = 0 \\
&L_{\max} - F\hat{\Lambda} \geq 0, 0 \leq F \leq 1, \text{ and } F_i = f_i \text{ for } i \in I
\end{aligned}
$$

Define $F := S\tilde{F} + \bar{F}$, where $\bar{F}_i = f_i$ for $i \in I$ and $\bar{F}_i = 0$ for $i \notin I$. and $\tilde{F} \in \mathbb{R}^{n-|I|}$, where $|I|$ is the number of elements in I and $S \in R^{n\times(n-|I|)}$ is a matrix of 0's and 1's. Then the above becomes

$$
\begin{aligned}
&\min_{\tilde{F}} -\tilde{F}\vec{1} \\
&\text{subject to } \left(S\tilde{F} + \bar{F}\right)\left(H^{-1} - \frac{s^2}{\gamma}\hat{\Lambda}\hat{\Lambda}^T\right)\left(S\tilde{F} + \bar{F}\right)^T \\
&- L_{\max}^2\frac{1}{\gamma s^2} + 2L_{\max}\left(S\tilde{F} + \bar{F}\right)\hat{\Lambda}\frac{1}{\gamma s^2} = 0 \\
&L_{\max} - \left(S\tilde{F} + \bar{F}\right)\hat{\Lambda} \geq 0, 0 \leq \tilde{F} \leq 1.
\end{aligned}
$$

Now define $G = \left(H^{-1} - \frac{s^2}{\gamma}\hat{\Lambda}\hat{\Lambda}^T\right)$

$$
\begin{aligned}
&\min_{\tilde{F}} -\tilde{F}\vec{1} \\
&\text{subject to } \left(S\tilde{F} + \bar{F}\right)^T G\left(S\tilde{F} + \bar{F}\right) \\
&- L_{\max}^2\frac{1}{\gamma s^2} + 2L_{\max}\left(S\tilde{F} + \bar{F}\right)^T\hat{\Lambda}\frac{1}{\gamma s^2} = 0 \\
&L_{\max} - \left(S\tilde{F} + \bar{F}\right)^T\hat{\Lambda} \geq 0, 0 \leq F \leq 1.
\end{aligned}
$$

The Largrange multiplier conditions become

$$
\begin{aligned}
2S^T GS\tilde{F} + 2S^T G\bar{F} + 2L_{\max}S^T\hat{\Lambda}\frac{1}{\gamma s^2} - \frac{1}{\mu}\vec{1} = 0 \\
\left(S\tilde{F} + \bar{F}\right)^T G\left(S\tilde{F} + \bar{F}\right) - L_{\max}^2\frac{1}{\gamma s^2} + 2L_{\max}\left(S\tilde{F} + \bar{F}\right)^T\hat{\Lambda}\frac{1}{\gamma s^2} = 0.
\end{aligned}
$$

where $\vec{1} \in R^{n-|I|}$. After a simply redefinition of the Lagrange multiplier, the first equation becomes

$$S^T G S \tilde{F} + S^T G \bar{F} + L_{\max} S^T \hat{\Lambda} \frac{1}{\gamma s^2} - \mu \vec{1} = 0$$

so

$$\tilde{F} = \left(S^T G S\right)^{-1} \left(L_{\max} S^T \Lambda + \mu \vec{1} - S^T G \bar{F}\right). \tag{10}$$

Expanding the second equation yields,

$$0 = F^T S^T G S F + \bar{F}^T G \bar{F} + 2 F^T S^T G \bar{F} + L_{\max}^2 - 2 L_{\max} \Lambda^T S F - 2 L_{\max} \Lambda^T \bar{F}.$$

Substituting (10) yields

$$\begin{aligned}
0 = & \left(\left(S^T G S\right)^{-1} \left(L_{\max} S^T \Lambda + \mu \vec{1} - S^T G \bar{F}\right)\right)^T \\
& \times S^T G S \left(\left(S^T G S\right)^{-1} \left(L_{\max} S^T \Lambda + \mu \vec{1} - S^T G \bar{F}\right)\right) \\
& + \bar{F}^T G \left(\bar{F}\right) \\
& + 2 \left(\left(S^T G S\right)^{-1} \left(L_{\max} S^T \Lambda + \mu \vec{1} - S^T G \bar{F}\right)\right)^T S^T G \bar{F} \\
& + L_{\max}^2 - 2 L_{\max} \Lambda^T S \left(\left(S^T G S\right)^{-1} \left(L_{\max} S^T \Lambda + \mu \vec{1} - S^T G \bar{F}\right)\right) - 2 L_{\max} \Lambda^T \bar{F}
\end{aligned}$$

Solving for μ yields 5.

References

[1] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1995.

[2] R. B. Blazek, H. Kim, B. Rozovskii, and A. Tartakovsky. A novel approach to detection of ŞDenial-of-serviceŤ attacks via adaptive sequential and batch-sequential change-point detection methods. In *2001 IEEE Man Systems and Cybernetics Information Assurance Workshop*, 2001.

[3] H. Burch and B. Cheswick. Tracing anonymous packets to their approximate source. In *Proc. Usenix LISA '00*, 2000.

[4] CERT. http://www.cert.org/advisories/CA-1996-21.html, 2001.

[5] S. Institute. http://www.sans.org/y2k/egress.htm, 2001.

[6] S. McClure, J. Scambray, and G. Kurtz. *Hacking Exposed*. Osborne, 1999.

[7] V. Paxson and S. Floyd. Wide area traffic: The failure of Poisson modeling. *IEEE/ACM Transactions on networking*, 1995.

[8] S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Practical network support for IP traceback. In *Proceedings of the 2000 ACM SIGCOMM Conference*, pages 295–306, Stockholm, Sweden, 2000.

[9] R. Stone. CenterTrack: An IP overlay network for tracking DoS floods. In *USENIX Security Symposium*, 2000.