

Cooperative Defense against Network Attacks^{*}

Guangsen Zhang, Manish Parashar
The Applied Software Systems Laboratory
Department of Electrical and Computer Engineering
Rutgers University
{gszhang,parashar}@caip.rutgers.edu

Abstract. Distributed denial of service (DDoS) attacks on the Internet have become an immediate problem. As DDoS streams do not have common characteristics, currently available intrusion detection systems (IDS) can not detect them accurately. As a result, defend DDoS attacks based on current available IDS will dramatically affect legitimate traffic. In this paper, we propose a distributed approach to defend against distributed denial of service attacks by coordinating across the Internet. Unlike traditional IDS, we detect and stop DDoS attacks within the intermediate network. In the proposed approach, DDoS defense systems are deployed in the network to detect DDoS attacks independently. A gossip based communication mechanism is used to exchange information about network attacks between these independent detection nodes to aggregate information about the overall network attacks observed. Using the aggregated information, the individual defense nodes have approximate information about global network attacks and can stop them more effectively and accurately. To provide reliable, rapid and widespread dissemination of attack information, the system is built as a peer to peer overlay network on top of the internet.

1 Introduction

A Distributed Denial of Service (DDoS) attack is a large-scale, coordinated attack on the availability of services at a victim system or network resource. The DDoS attack is launched by sending an extremely large volume of packets to a target machine through the simultaneous cooperation of a large number of hosts that are distributed throughout the Internet. The attack traffic consumes the bandwidth resources of the network or the computing resource at the target host, so that legitimate requests will be discarded. The impact of these attacks can vary from minor inconvenience to the users of a web site, to serious financial losses to companies that rely on their on-line availability to do business [11, 12].

DDoS attacks are likely to become an increasing threat to the Internet due to the easy availability of user-friendly attack tools, which help to coordinate and execute a large scale DDoS attack. Even an unsophisticated individual can launch a devastating attack with the help of these tools. Available tools include Trinoo, TFN, TFN2K, Shaft,

^{*} The research presented in this paper is supported in part by the National Science Foundation via grants numbers ACI 9984357, EIA 0103674, EIA 0120934, ANI 0335244, CNS 0305495, CNS 0426354 and IIS 0430826.

and Stacheldraht and have been used in DDoS attacks against well-known commercial web-sites, such as Yahoo, Amazon, Ebay [2].

The only way to completely eliminate the DDoS threat is to secure all machines on the Internet against misuse, which is unrealistic. Most large web sites currently handle the problem by equipping critical systems with abundant resources. While this raises the bar for the attacker, any amount of resources can be exhausted with a sufficiently strong attack. The only viable approach is to design defense mechanism that will detect the attack and respond to it by dropping the excess traffic. Generally it is easy to detect the abnormal behavior of attack near the victim. However, it is also often too late to detect the DDoS attack at the victim network. The attack should ideally be stopped as close to the sources as possible, saving network resources and reducing congestion. However, there are no common characteristics of DDoS streams that can be used to detect the attacks near the source [12]. To balance this tradeoff, in this paper we try to detect the DDoS attacks in the intermediate network. As the traffic is not aggregated enough in the intermediate network, current single deployment detection systems can not detect DDoS attacks with high accuracy. As a result, the reported false alarms will lead to dramatically affect on legitimate traffic. To improve the defense efficiency and accuracy, we propose a dynamic defense infrastructure composed of a diverse collection of independent defense nodes located in the intermediate network of the Internet. We make the assumption that in the intermediate network, the aggregated attack flows toward the victim consume more bandwidth than aggregated normal flows to the victim. This is reasonable because if every attacker sends at a rate comparable to a good user, then an attacker must recruit or compromise a large number of hosts to launch an attack with sufficient traffic volume.

The focus of this research is to develop methods to efficiently share the information provided by existing DDoS attack detection systems to improve the accuracy of defense rather than to improve upon current available DDoS detection methods. The primary contribution of this paper is a global defense infrastructure built as an overlay network on top of the Internet. This infrastructure provides reliable, rapid and widespread cooperation among individual detection nodes to improve the accuracy of DDoS detection in the intermediate network. Given the large scale of the internet and purpose of this infrastructure, we need resilient and scalable communication mechanism to exchange the attack information. We design directional gossip mechanisms to fulfill this need while reducing the overhead of information sharing. Initial results using a simulation illustrate that the proposed approach is both efficient and feasible.

The rest of the paper is organized as follows. Section 2 gives an overview of DDoS. Section 3 explains our approach. Section 4 presents an experimental evaluation. Section 5 discusses related work. Section 6 concludes the paper.

2 DDoS Background

Distributed denial of service attacks (DDoS) pose a great threat to the Internet. A recent DDoS attack occurred on October 20, 2002 against the 13 root servers that provide the Domain Name System (DNS) service to Internet users around the world. Although the attack only lasted for an hour and the effects were hardly noticeable to the average Inter-

net user, it caused 7 of the 13 root servers to shut down, demonstrating the vulnerability of the Internet to DDoS attacks [11]. Distributed denial of service attacks occur when numerous subverted machines (zombies) generate a large volume of coordinated traffic toward a target, overwhelming its resources. DDoS attacks are advanced methods of attacking a network system to make it unavailable to legitimate network users. These attacks are likely to become an increasing threat to the Internet due to the convenience offered by many freely available user-friendly attack tools. Furthermore, attackers need not fear punishment, as it is extremely difficult to trace back the attack and locate even the agent machines, let alone the culprits who infected them.

There are two main classes of DDoS attacks: bandwidth depletion and resource depletion attacks. A bandwidth depletion attack is designed to flood the victim network with unwanted traffic that prevent legitimate traffic from reaching the victim system. A resource depletion attack is an attack that is designed to tie up the resources of a victim system. This type of attack targets a server or process at the victim making it unable to legitimate requests for service [6].

3 Distributed Cooperative Mitigation Approach

The mitigation mechanism presented in this paper consists of two key stages. In the first stage, each defense node detects traffic anomalies locally using a variety of existing IDS tools such as Snort [14]. According to its local defense policy, each local defense node exerts a rate limit to the traffic identified as attack traffic. Due to the dynamic nature of the Internet, defense based on local detection mechanism alone will have high false positives. In the second stage, we enhance the accuracy of the defense mechanism by using gossip based communication mechanism to share information among the defense nodes. As the information sharing proceeds, we dynamically adjust the rate limit at each individual defense node. Finally, when this gossip based information aggregation mechanism converges, the rate limit mechanism of each individual defense node will have approximate global information about the attack behavior, and will be able to defend against attack traffic more efficiently by dropping the traffic with higher accuracy.

To enhance the security and reliability of information sharing, our system is built on a peer-to-peer overlay network composed of local detection nodes, which may be routers with DDoS detection and attack packets filtering functionality. The peer-to-peer overlay, which we will reference as p2p networks, have been shown to be highly resilient to disruption and are reliable and scalable for information dissemination purpose [13].

3.1 Architecture Overview

We assume that the Internet is composed of a set of Autonomous Systems (AS). Individual defense nodes are located at the egress routers of an Autonomous System, which collect meaningful information and detect DDoS attacks locally. The system then uses the overlay network to share the attack information using a gossip protocol based on epidemic algorithm [8] across the Internet.

The internals of an individual defense node can be fairly complex, but conceptually it can be structured into six components, as shown in the Figure 1. The traffic measurement module is responsible for measuring local traffic. Next, the local detection mechanism will use this data to detect any local anomaly. This local decision will be sent to the cooperative detection engine, which will combine this local decision with the decisions from neighboring nodes, using the message dissemination module, to make a global detection decisions. Finally, the detection decision module will inform the local response module to take action to defend against an attack.

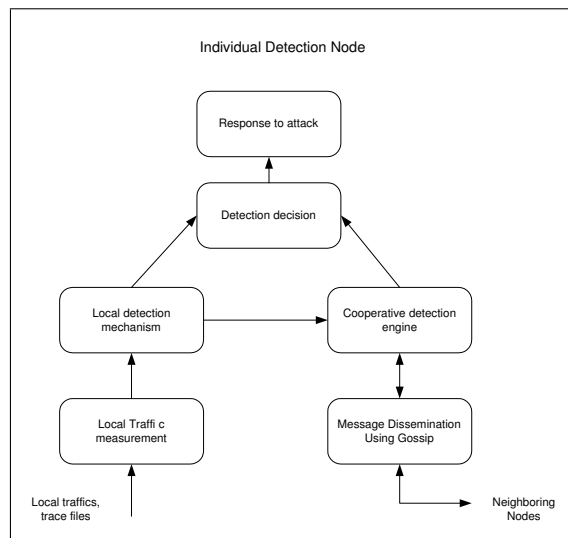


Fig. 1. A conceptual architecture for an individual defense node

3.2 Local Defense Mechanism

As we have mentioned, the local defense node can utilize heterogeneous attack detection mechanisms to monitor local network traffic. The main function of the local defense nodes include two aspects: local attack signature generation and rate limiting of identified attack traffic.

Attack Signature Generation Generally, the attack signature of the DDoS attacks can be acquired using the network monitoring capability of the IDS. Current IDS have the capability to produce traffic statistics based on captured packet data. As the high-traffic destinations are most likely to be under attack, it is reasonable to keep traffic statistics only for those high traffic flows that have the same destination IP addresses. We can use a *sample-and-hold* [3, 1] algorithm to let the local detection nodes keep track of

destinations whose traffic occupies greater than a fraction r of the capacity C of the outgoing link. We call these destinations popular and destinations not in this list as unpopular. Traffic profiles at each router are essentially a set of metrics M_i for the traffic to popular destinations. An effective choice of such metrics is key to characterizing traffic streams. However, computing arbitrary fingerprints might require excessive memory and computation. Several metrics have been proposed by the research community. Some of them are:

- The ratio of TCP traffic between the two directions. Due to the nature of the TCP protocol we expect a loose symmetry on the incoming versus outgoing packet rates. This principle has been used by local detection mechanisms such as D-WARD [11] and MULTOPS [4].
- ICMP and UDP packets are mainly used by bandwidth consumption attacks and as these traffic types generally utilize small amounts of bandwidth, suddenly change in the transferred ICMP or UDP byte/sec are good indication of attacks.

For each of these attributes A_i , we use corresponding metrics M_i to measure them. Let $conf$ denote the confidence with which the individual detection node suspects an attack with attributes discussed above. We set $conf_i = \delta(M_i) * d_N(M_i)$. δ assigns “weights” to a metric, depending on the extent to which the metric contributes to errors (false positive or negatives): $\delta(M_i) \propto \frac{1}{err(M_i)}$ where $err(M_i)$ is the sum of the false positive and negative rates for M_i . The appropriate δ can be configured from measurements.

When a local detection node detects an attack, it will exert a rate limit on the traffic with identified attributes and send the $(conf_i, A_i, dest)$ tuples to its neighbor nodes in the overlay network infrastructure for correlation.

Rate Limit Mechanism Attack detection itself is not the final goal of the defense system. Once a DDoS attack signature is detected, the next step is to rate-limit the traffic with the identified attack signature. The objective is to maximize friendly traffic throughput while reducing attack traffic as much as possible. According to the confidence of the attack signature, the traffic with identified attack signature will be rate-limited according to the formula below:

$$rate_{out}(A_i) = rate_{in}(A_i) * \lambda(conf_i)$$

Where $\lambda(conf_i) \leq 1$ is a factor defined by the confidence level of the attack signature identified. When the value of $conf_i$ is 0, $\lambda(conf_i) = 1$. If each local defense node rate-limits traffic based on local information only, legitimate traffic will usually be wrongly dropped as well. In the next section, we will discuss how to share the information of the attack signature so that each individual detection node has more accurate information about the attack behavior, reducing the affect on legitimate traffic while dropping malicious traffic.

3.3 Global Defense Using Aggregated Information

A key requirement of an anomaly detection model is low false positive rates, calculated as the percentage of normalcy variations detected as anomalies, and high positive rate,

calculated as the percentage of anomalies detected. In our approach, there are two factors which will affect the system performance: the overhead of the information sharing mechanism, and the delay for the decision making. Communication bandwidth is often a scarce resource during the DDoS attack, so the attack information sharing should involve only small messages. In particular, any protocol collecting all local data at a single node will create communication bottlenecks, or a message implosion at that node. Recently, gossip-based protocols have been developed to reduce control message overhead while still providing high reliability and scalability of message delivery [5]. Gossip protocols are scalable because they don't require as much synchronization as traditional reliable multicast protocols. In gossip-based protocols, each node contacts one or a few nodes in each round (usually chosen at random), and exchanges information with these nodes. The dynamics of information spread bears a resemblance to the spread of an epidemic, and leads to high fault tolerance. Gossip-based protocols usually do not require error recovery mechanisms [9], and thus enjoy a large advantage in simplicity, while often incurring only moderate overhead compared to optimal deterministic protocols.

Compared with reliable multicast or broadcast protocols, the gossip protocol has a smaller overhead. However, it requires a longer time for each node get the message. While reducing message dissemination overhead, we still want maintain the speedy information delivery provided by multicast or broadcast. A possible variant is directional gossip [10]. Directional gossip is primarily aimed at reducing the communication overhead of traditional gossip protocols. In our approach, we use a modified directional gossip strategy. We assume that the individual node knows its immediate neighbors in the network. Our gossiping protocol is described as the following: An individual node sends the $(conf, attribute, dest)$ tuples to the node on its path to the destination target node with probability 1. It forwards the $(conf, attributed, dest)$ tuple to all other nodes at random.

At anytime t , each node i maintains a list of $(conf_k, attribute_k, dest_k)$ tuples. Each node will compute the aggregated information about the attack behavior. Every time the aggregate information is computed, the defense node will adjust the rate-limit to the identified attack traffic (traffic with attributes monitored) according to this new information. As this process converges exponentially, all the nodes in the peer to peer defense network will get the approximate global information about the network behavior quickly. Thus we can have a more accurate rate limit on the attack traffic. The convergence of information aggregation using epidemic algorithm has been discussed in [9]. The algorithm we use to get aggregated information about the DDoS attacks is described as follows:

1. Let $(conf_{r,k}, attribute_{r,k}, dest_{r,k})$ be all pairs sent to node i in round $t-1$.
2. For each $attribute_{r,k}$, compute $d_{i,k} = \frac{\sum_r conf_{r,k}}{m}$, where m is the number of messages received.
3. Based on this $d_{i,k}$, adjust the rate limit of the traffic with attribute $attribute_{r,k}$.
4. Query the routing table, find out the next hop to $dest_{i,k}$, send the pair $(conf_{i,k}, dest_{i,k})$ to that node with probability 1. Send the pair to other neighbors with probability p .

Based on the aggregated information of the attack signature, each individual detection node dynamically adjusts the rate limit factor for the identified attack traffic.

4 Simulation Results and Analysis

To further examine system performance, under detailed network models, we conduct experiments using the Emulab testbed. The objective of the emulation is to illustrate that our approach can effectively defend against DDoS attack with high accuracy with reasonable overheads.

4.1 Results

We implemented our distributed cooperative defense mechanism in a Linux router and tested it with live traffic in the Emulab testbed. As mentioned earlier, we rely on existing intrusion detection systems to detect attacks at each individual detection node. We implemented dynamic coordination mechanism based on gossip in a Linux router which will dynamically adjust the rate limiting parameters according to the information aggregated from the detection nodes of peer to peer defense overlay network.

We use a simple HTTP client-server as the model of the simulated application. We use the GT-ITM topology generator to generate the Internet topology. Which can generate a random transit-stub graph based on input parameters. This graph closely resembles the Internet topology. The attack is simulated using a given number of compromised nodes in different sub networks. Detection agents are deployed at selected nodes and execute the algorithm described in Section 3. The communication agents use gossip to share information. In these experiments, there are 10 attackers, each of them send out 1.3Mbps UDP traffic to the victim. The good user makes request with traffic rates chosen randomly and uniformly from the range [2Kbps, 6Kbps]. If a request arrives at the server successfully, the server will return the requested document after a random processing time, chosen according to collected empirical distributions.

In the first set of experiments, we performed test runs for normal use, under attack without response, and under attack with distributed cooperative response. In each case, we measured the packets rate of a selected client at the HTTP server. Figure 2 shows the result from the experiment runs. The x axis represents time intervals in seconds; the y axis represents the number of packets received at the server. The attack starts 50 seconds after the start of legitimate traffic and last for 500 seconds. Compared with the packet rate of normal run, the selected legitimate client's packet rate at server drop dramatically under attack without response. For the experiment that we ran attacks with cooperative defense mechanism enabled, we can notice a gradual increase of the legitimate packet rate. The ramp-up behavior is due to the false detection of local defense node. As a result, some legitimate traffic will be dropped by the rate limiting mechanism as well. As the algorithm converge, each defense node get more precise information about the global attack information thus can rate-limit attack traffic with more accuracy.

In the second set of experiments, we vary the parameters of the gossip mechanism to investigate the relationship between the overhead of information sharing and defense efficiency. Let p represent the probability that each detection node in the detection overlay network sends the local attack information to its neighbor nodes. We vary the Gossip probability p between 0.2, 0.4, 0.6, 0.8, 1.0. The performance of the approach with different gossip probability p used are shown in Table 1. The *false positive rate* measures




Fig. 2. Legitimate user packet rate under different test conditions

the percentage of legitimate packets dropped by the rate limiting mechanism, and *false negative rate* measures the percentage of attack traffic pass the defense node.

Table 1. Cooperative defense performance

Gossip Prob.	False Positive	False Negative
0.2	12.12%	5.2%
0.4	10.03%	4.13%
0.6	8.32%	4.32%
0.8	8.15%	3.56%
1.0	7.67%	3.12%

As we can see from the simulation results, our algorithm can detect and defense DDoS attacks with high accuracy. With $p = 0.4$ we have low false positive and low false negative packet drop rate respectively. The false positive rate is relatively higher than the false negative rate. This is because we adopt high initial drop rate when the local defense node detects an attack, as a result legitimate packets will be dropped dramatically in the case of false detection.

Defenses mitigate the impact of the attack traffic on the victim network but may impose an additional overhead on the networks that implements them. We measure the overhead introduced by distributed cooperative information sharing in this experiment as well. Figure 3 shows the per-node overhead with different number of nodes in the system. The packets processed by each node for the cooperative defense purpose do not increase much as we add more node into the defense overlay. So the gossip based

information sharing mechanism is scalable to be used in larger and higher speed network situation. When the gossip probability is increased, the overhead will increased as well. This parameter can be tuned to adapt different application to achieve optimum performance.




Fig. 3. Information sharing overhead

5 Related Work

The idea of cooperative defense against network attacks has been proposed in a number of projects. Projects closely related to this paper are discussed below.

Pushback [7] and Aggregate-Based Congestion Control(ACC) are project at AT&T Center for Internet Research. The routers int the system assume that the congestion of local packet queue is the sign of DDoS attack and take action to rate limit the identified aggregates which are responsible of queue congestion according to local policy. If the congested router cannot control the aggregate itself, it issues a rate limit request to its immediate upstream neighbors who carry the aggregated traffic to apply rate limiting to specified excessive flows. These requests will be propagated upstream as far as the identified aggregates have been effectively controlled. This approach request all the routers on the path of aggregate traffic be augmented with the pushback capability.

6 Conclusion and Future Work

In this paper we proposed a global defense infrastructure by building an overlay network on top of the internet. A gossip-based scheme is used to get global information about

distributed denial of service attacks by information sharing. We assume with global information, we can defense DDoS attacks with higher accuracy. Compared to the existing solutions, our contribution is to provide a distributed proactive DDoS detection and defense mechanism. Our approach continuously monitors the network. When an attack begins, individual defense nodes drop attack traffic identified according to the local information and mitigate load to the target victim. However, as local detection has high false alarm rate, the legitimate traffic will dropped as well with high rate. By correlating the attack information of each individual nodes, our scheme can get more information about the network attack thus can defense against DDoS attacks more effectively.

References

1. A. Akella, A. Bhambe, M. Reiter, and S. Seshan. Detecting DDoS attacks on ISP networks. In *ACM SIGMOD Workshop on Management and Processing of Data Streams*, pages 20–23, San Diego, CA, 2003.
2. D. Dittrich. Distributed denial of service (DDoS) attacks/tools, 2004. <http://staff.washington.edu/dittrich/misc/ddos/>.
3. C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proceedings of SIGCOMM 2002*, pages 270–313, Pittsburgh, PA, USA, 2002.
4. T. M. Gil and M. Poletto. Multops: a data-structure for bandwidth attack detection. In *Proceedings of 10th Usenix Security Symposium*, pages 23–28, Washington, D.C., USA, August 2001.
5. I. Gupta, K. P. Birman, and R. van Renesse. Fighting fire with fire: using randomized gossip to combat stochastic scalability limits. *Special Issue Journal Quality and Reliability Engineering International: Secure, Reliable Computer and Network Systems*, 18(3):165–184, May 2002.
6. Q. Huang, H. Kobayashi, and B. Liu. Analysis of a new form of distributed denial of service attack. In *Proceedings of CISS03, the 37th Annual Conference on Information Science and Systems*, Johns Hopkins University, Baltimore, Maryland, March 2003.
7. J. Ioannidis and S. M. Bellovin. Implementing pushback: Router-based defense against DDoS attacks. In *Proceedings of Network and Distributed System Security Symposium, NDSS '02*, pages 100–108, Reston, VA, USA, February 2002.
8. R. M. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. Randomized rumor spreading. In *IEEE Symposium on Foundations of Computer Science*, pages 565–574, 2000.
9. D. Kempe, A. Dobra, and J. Gehrke. Computing aggregate information using gossip. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, Cambridge, MA, October 2003.
10. M. Lin and K. Marzullo. Directional gossip: gossip in a wide area network. In *Proceedings of Dependable Computing - Third European Dependable Computing Conference*, pages 364–379, Berlin, Germany, 1999.
11. J. Mirkovic, G. Prier, and P. Reiher. Attacking DDoS at the source. In *Proceedings of ICNP 2002*, pages 312–321, Paris, France, November 2002.
12. C. Papadopoulos, R. Lindell, J. Mehringer, A. Hussain, and R. Govindan. Cossack: Coordinated suppression of simultaneous attacks. In *DARPA Information Survivability Conference and Exposition*, volume 1, pages 2–13, Washington, DC, April 2003.
13. R. Renesse, K. Birman, and W. Vogels. Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining. *ACM Transactions on Computer Systems*, 21(2):164–206, May 2003.
14. M. Roesch. The snort network intrusion detection system, 2002. <http://www.snort.org>.