Cooperative Defence against DDoS Attacks

Guangsen Zhang and Manish Parashar
628 CoRE, Department of Electrical and Computer Engineering
Rutgers, The State University of New Jersey
94 Brett Road, Piscataway, NJ 08854-8058
{gszhang, parashar}@caip.rutgers.edu

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) cannot 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 defence 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 defence 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.

ACM Classification: C.2(Computer-Communication Networks), D.2(Software Engineering)

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 (Mirkovic, Prier and Reiher, 2002; Papadopoulos, Lindell, Mehringer, Hussain and Govindan, 2003).

DDoS attacks are likely to become an increasing threat to the Internet due to the easy availability of userfriendly 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, and Stacheldraht and have been used in DDoS attacks against well-known commercial web-sites, such as Yahoo, Amazon, Ebay (Dittrich, 2004).

Manuscript received: 12 April 2005
Communicating Editor: Julio Cesar Hernandez

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 defence mechanism that will detect the attack and respond to it by dropping the excess traffic. Generally it is easy to detect the abnormal behaviour 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 (Chang, 2002). 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 legitimate traffic. To improve the defence efficiency and accuracy, we propose a dynamic defence infrastructure composed of a diverse collection of independent defence 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 defence rather than to improve upon current available DDoS detection methods. The primary contribution of this paper is a global defence 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 20 October, 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 Internet user, it caused seven of the 13 root servers to shut down, demonstrating the vulnerability of the Internet to DDoS attacks (Peng, Leckie and Kotagiri, 2003). 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.

2.1 DDoS Attack Classification

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 (Huang, Kobayashi and Liu, 2003).

2.2 DDoS Characteristics

There are several features of DDoS attacks that hinder their successful detection and defence:

- DDoS attacks generate a large volume flow to overwhelm the target host. The victim can not protect itself even if it detects this event. So the detection and defence of DDoS should ideally be near the source of the attack or somewhere in the network.
- It is difficult to distinguish attack packets from legitimate packets. Attack packets can be identical to legitimate packets, since the attacker only needs volume, not content, to inflict damage. Furthermore, the volume of packets from individual sources can be low enough to escape notice by local administrators. Thus, a detection system based on a single site will have either high positive or high negative rates.
- DDoS traffic generated by available tools often has identifying characteristics, making the detection based on statistics analysis possible. However, given the inherently busty nature of Internet, detecting DDoS attacks is error prone.

2.3 A Taxonomy of the DDoS Detection and Defence

Based on the underlying strategies, we can categorize current DDoS detection and defence approaches into three categories: Proactive Mechanisms, Reactive Mechanisms and Post Attack Analysis.

Proactive defence mechanisms. (Keromytis, Misra and Rubenstein, 2002) The motivation for these approaches is based on the observation that it is hard to detect DDoS attacks. So instead of detecting the attacks by using signatures (attack pattern) or anomaly behaviour, these approaches try to improve the reliability of the global Internet infrastructure by adding extra functionality to Internet components to prevent attacks and vulnerability exploitation. The primary goal is to make the infrastructure immune to the attacks and to continue to provide service to normal users under extreme conditions.

Reactive defence mechanisms using available IDS. (Ioannidis and Bellovin, 2002) These mechanisms typically deploy third-party Intrusion Detection Systems (IDS) to obtain attack information and take action based on this information. Consequently their usefulness depends on the capability of the IDS systems. Different strategies are used based on the assumptions made by the IDS systems. If the IDS system can detect the DDoS attack packets accurately, filtering mechanism are used, which can filter out the attack stream completely, even at the source network. If the IDS can not detect the attack stream accurately, rate limiting is used. This mechanism imposes a rate limit on the stream that is characterized as malicious by the IDS.

Post attack analysis. (Song and Perrig, 2001) The purpose of post attack analysis is to either look for attack patterns that will be used by IDS or identify attackers using packet tracing. The goal of packet tracing is to trace Internet traffic back to the true source (not spoofed IP address). As attackers change their strategy frequently, analyzing huge amounts of traffic logs is time consuming and useless in detecting new attacks. Trace back mechanism can help to identify zombies in some situations, however, it is impractical to defend against DDoS attacks for the following reasons. First, during a DDoS attack, the attacker will control thousands of zombies (numbers will increase in the future) to launch an attack. As a result, identifying these zombies is expensive and infeasible. Second, since different network administrators control different sections of the global Internet, it would be difficult to determine who would be responsible for providing trace back information.

3. DISTRIBUTED COOPERATIVE MITIGATION APPROACH

The mitigation mechanism presented in this paper consists of two key stages. In the first stage, each defence node detects traffic anomalies locally using a variety of existing IDS tools such as Snort (Roesch, 2002). According to its local defence policy, each local defence node exerts a rate limit to the traffic identified as attack traffic. Due to the dynamic nature of the Internet, defence based on local detection mechanism alone will have high false positives. In the second stage, we enhance the accuracy of the defence mechanism by using gossip based communication mechanism to share information among the defence nodes. As the information sharing proceeds, we dynamically adjust the rate limit at each individual defence node. Finally, when this gossip based information aggregation mechanism converges, the rate limit mechanism of each individual defence node will have approximate global information about the attack behaviour, 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 purposes (Renesse, Birman and Vogels, 2003).

The p2p computing model offers a radically different and appealing alternative to the traditional client-server model for many large scale applications in distributed settings. In this model, end user processes share resources in a peer style, potentially acting as both client and server. The p2p approach removes central points of failure and associated performance bottleneck; it also balances the load such as forwarding messages or storing data among all system processes, each of which requires only local knowledge of system state. Our proposed p2p defence network deploys gossip based protocol as described in Eugster, Guerraoui, Kermarrec and Massoulie (2004).

Traditional communication techniques have absolute guarantees, but are unreliable or fail to make progress during periods of even modest disruption. Gossip based protocol usually do not require error recovery mechanisms, and thus enjoy a large advantage in simplicity, while often incurring only moderate overhead compared to deterministic protocols, such as the construction of data dissemination trees. The guarantees obtained from gossip are usually probabilistic in nature; they achieve high stability under stress and disruptions, and scale gracefully to a huge number of nodes. In this research, we propose to use gossip based communication mechanism for attack information sharing purpose in our p2p defence overlay network.

3.1 Architecture Overview

We assume that the Internet is composed of a set of Autonomous Systems (AS). Individual defence 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 (Karp, Schindelhauer, Shenker and Vocking, 2000) across the Internet. This is illustrated in the Figure 1.

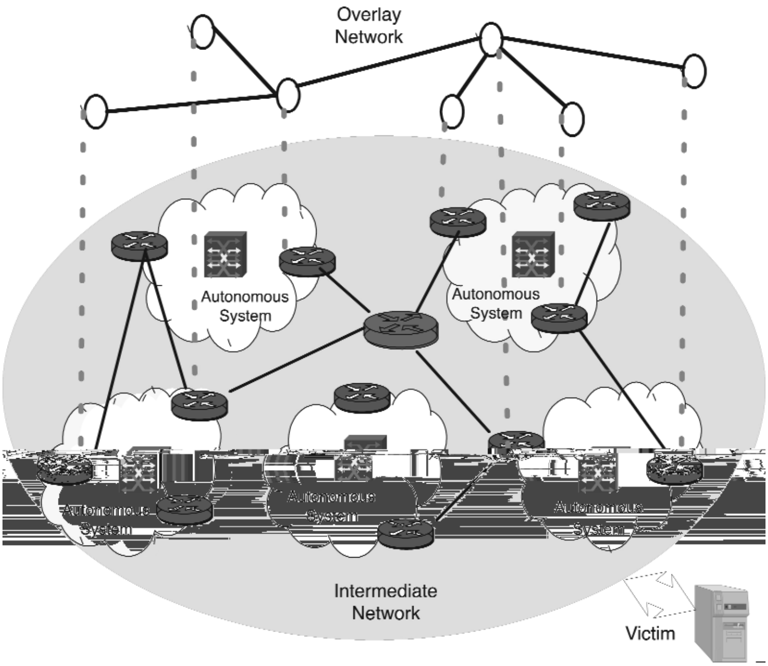

Figure 1: System Architecture

The internals of an individual defence node can be fairly complex, but conceptually it can be structured into six components, as shown in the Figure 2. 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 neighbouring nodes, using the message

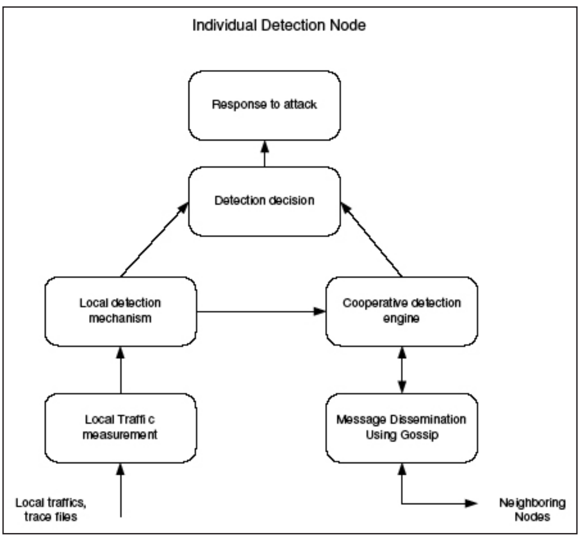

Figure 2: A conceptual architecture for an individual defence node

dissemination module, to make a global detection decision. Finally, the detection decision module will inform the local response module to take action to defend against an attack.

In our approach, the egress routers of the intermediate network coordinate with each other to provide the information necessary to detect and respond to the attack. This mechanism can improve the accuracy of rate limiting. The operations on each defence nodes are described below:

- The local detection node uses various detection mechanism to detect attacks. If a target is suspected to be a victim, the local detection node advertises its suspicions of the attack target to its peer group. If the peer group already exists, the local detection node just joins that peer group.
- Individual nodes deploy countermeasures to prevent continuance of the attack that it has identified based on the local policies.
- As current available attack detection mechanisms have high false detection rates, actions based on the local detection may drop lots of legitimate traffic and let malicious traffic pass unnoticed. Each local defence node shares information about an attack with other nodes using a gossip mechanism to improve accuracy of defence.

3.1.1 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* (Estan and Varghese, 2002; Akella, Bharambe, Reiter and Seshan, 2003) 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 Mi 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 (Mirkovic *et al*, 2002) and MULTOPS (Gil and Poleto, 2001).
- ICMP and UDP packets are mainly used by bandwidth consumption attacks and as these traffic types generally utilize small amounts of bandwidth, sudden change in the transferred ICMP or UDP byte/sec are a good indication of attacks.

Other candidate packet attributes considered for traffic profiling include the marginal distributions of the fraction of recently arrived packets having various characteristics as shown in Table 1.

For each of these attributes A_i, we use corresponding metrics M_i to measure them. Let conf_i 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): $\delta(M_i) \alpha \frac{1}{err(M_i)}$ where $err(M_i)$ is the sum of the false positive and negative rates for M_i. The appropriate ± can be configured from measurements. When a local detection node detects an attack, it will exert a rate

No.	Attributes
1	IP protocol type values
2	packet size
3	server port numbers
4	source/destination IP prefixes
5	Time-to-Live (TTL) values
6	IP/TCP header length
7	TCP flag patterns
8	IP/TCP/UDP checksums

Table 1: Attributes to identify the attack flows

limit on the traffic with identified attributes and send the ($\text{conf}_i;A_i;$ *dest*) tuples to its neighbour nodes in the overlay network infrastructure for correlation.

3.1.2 Rate Limit Mechanism

Attack detection itself is not the final goal of the defence 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 an identified attack signature will be rate-limited according to the formula below:

$$rate_{out}(A_i) = rate_{in}(A_i) * \lambda(\text{conf}_i)$$

Where $\lambda(\text{conf}_i) \leq 1$ is a factor defined by the confidence level of the attack signature identified. When the value of conf_i is 0, $\lambda(\text{conf}_i) = 1$. If each local defence 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 behaviour, reducing the affect on legitimate traffic while dropping malicious traffic.

3.2 Global Defence 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 (Gupta, Birman and van Renesse, 2002). 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

```
    when  ( node n builds a new  (conf, attribute, dest) tuple )
{
        while  ( node n believes that not enough of its
        neighbors have received  (const, attribute, dest) tuple)
        {
                m =a neighbor node of p;
                send  (conf, attribute, dest) tuple to m;
        }
}
```

Figure 3: Gossip protocol for Our Approach

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 (Kempe, Dobra and Gehrke, 2003), and thus enjoy a large advantage in simplicity, while often incurring only moderate overhead compared to optimal deterministic protocols. The gossip protocol running at each node *n* has the structure shown in Figure 3.

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 to maintain the speedy information delivery provided by multicast or broadcast. A possible variant is directional gossip (Lin and Marzullo, 1999). 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 neighbours in the network. Our gossiping protocol is described as the following: An individual node sends the (confi; A_i; *dest*) tuples to the node on its path to the destination target node with probability 1. It forwards the (confi; A_i; *dest*) tuple to all other nodes at random. At anytime t, each node *i* maintains a list of (confi; A_i; *dest*) tuples. Each node will compute the aggregated information about the attack behaviour. Every time the aggregate information is computed, the defence 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 defence network will get the approximate global information about the network behaviour 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 Kempe *et al* (2003). The algorithm we use to get aggregated information about the DDoS attacks is described as follows:

1. *Let* ($\mathrm{conf}_{r,k}$; $attribute_{r,k}$; $dest_{r,k}$) *be all pairs sent to node i in round t-1.*
2. *For each* $attribute_{r,k}$, *compute* $d_{r,k} = \frac{\sum_r conf_{r,k}}{m}$ *,where m is the number of messages received.*
3. *Based on this* $d_{r,k}$, *adjust the rate limit of the traffic with attribute* $attribute_{r,k}$.
4. *Query the routing table, find out the next hop to* $dest_{r,k}$, send the tuple ($\mathrm{conf}_{r,k}$; $attribute_{r,k}$; $dest_{r,k}$) *to that node with probability 1. Send the pair to other neighbours 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. The advantage of the optimized directional gossip strategy is illustrated in Figure 4.

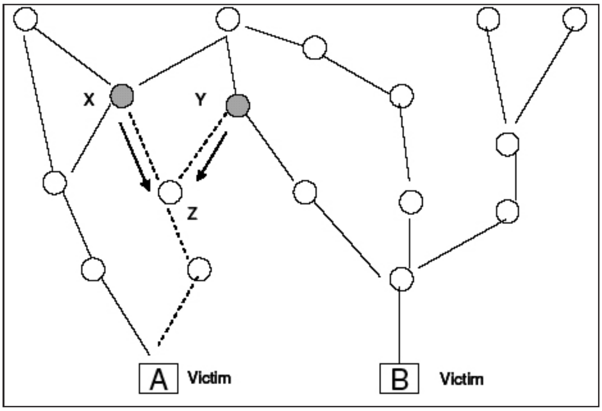

Figure 4: Gossip strategy in our approach

Suppose node X and node Y suspect that the destination host A is under attack, and both of them use node Z to forward packets to destination A. Obviously, it is better to send the ($conf_{r,k}$; $attribute_{r,k}$; $dest_{r,k}$) tuple with a higher priority to Z than to other neighbours. The rationale behind this scheme is as follows. Sending the attack information with higher probability to critical nodes allows them to make a decision early, allowing the DDoS attacks to be mitigated earlier. For each destination with $conf > 0$; each individual node in the overlay network sends the ($conf_{r,k}$; $attribute_{r,k}$; $dest_{r,k}$) tuple to its neighbours. On receiving such a message, the neighbours discard duplicates, compute the aggregate (*Aggr*) of the *conf* values received per destination, and forward non-duplicate values to their neighbours. If, for any destination, *Aggr* exceeds a pre-defined threshold, the individual node concludes that the destination is under attack. This cooperation stage helps reduce the errors in the identification of attacks.

3.3 Security of the Infrastructure

As a widely distributed defence infrastructure, the overlay network itself must be considered a target for attacks. To be effective, this infrastructure must be resilient to a variety of attacks. Relevant security issues are discussed below:

Trust: Trust is an important issue in this system, more so in the absence of a centralized trusted authority to provide digital certificates. The usual decentralized alternate to central CAs is the web-of-trust model, where certifying happens among peers rather than from a central authority. We believe, the overlay nodes can build trust relationships based on this model. Ideally, every overlay node should digitally sign their messages sent to other nodes in a manner that allows other nodes to validate the authenticity of the sending nodes. Current available technologies should suffice for this purpose.

Attack against infrastructure: Another issue that must be addressed is how to protect the communications of the detection nodes when the links are completely saturated during a DDoS attack. In the face of standard packet flood attacks, it is certainly possible that some set of nodes could be effectively removed from the infrastructure. Yet, if any connectivity remains at all, the gossip exchange of data will eventually prevail, and data stored within the infrastructure will reach all sites in the system. Also, the distributed and coordinated nature of the infrastructure makes it robust to the removal of nodes through failures or attacks. Thus, the infrastructure is relatively tolerant of denial of service attacks. In the case that an overlay node is compromised and sends large amounts of data to flood other overlay nodes, overlay nodes can apply filters to incoming data such that data sent by any nodes or set of nodes can not exceed a specified threshold. We are still working on these aspects.

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 and with reasonable overheads.

4.1 Performance Metrics

In our model, we assume that it is not easy to distinguish DDoS attack traffic from legitimate traffic. Therefore, our rate limiting mechanism will block legitimate traffic as well. To evaluate performance of the proposed defence mechanism, we adopt the following measurement:

1. Measure the legitimate traffic drop rate and the malicious traffic drop rate under the different pattern of DDoS attacks. Since the algorithm dynamically adjusts the rate limiting to the suspicious traffic based on the global information, the legitimate user should be adequately served.
2. Partial placement of the defence node will be effective. One advantage of distributed cooperative defence mechanism lies in the fact that with partial converge or deployment a synergistic defence effect is achieved. Since not every router or gateway in the Internet will be a defence node, the cooperative defence mechanism is designed to be effective in partial deployment. This feature is supported by an overlay network topology in which only nodes that have established direct peering relationship are aware of each other. The system provides a significant level of defence for potential targets with only a few defence nodes deployed, and becomes more effective as more defence nodes are added, protecting a larger community.
3. Gossip mechanism is reliable and scalable, we will analyze message dissemination converge rate and the overhead introduced by this distributed cooperative mechanism.

4.2 Results

We implemented our distributed cooperative defence 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. In our experiment, we use snort (Roesch, 2002) for this purpose. 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 defence 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. Figure 5 shows the experimental topology with 100 nodes. 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.

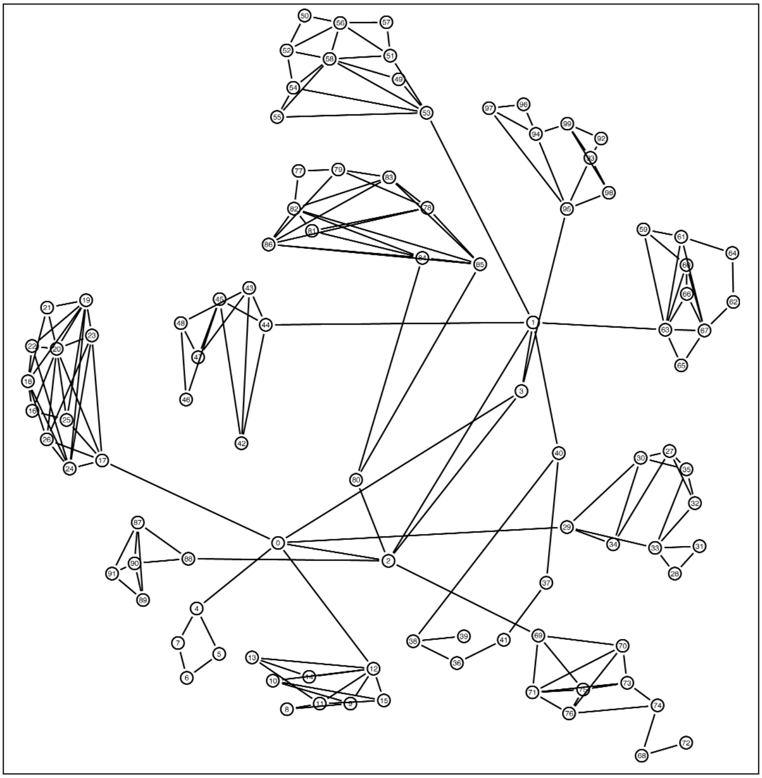

Figure 5: Simulated network topology

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 6 shows the result from the experiment runs. The x axis represents time intervals in seconds; the y axis represents the number of packets

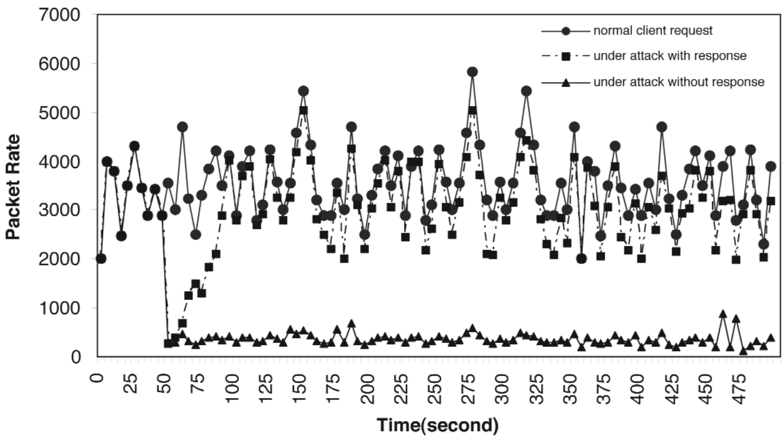

Figure 6: Legitimate user packet rate under different test conditions

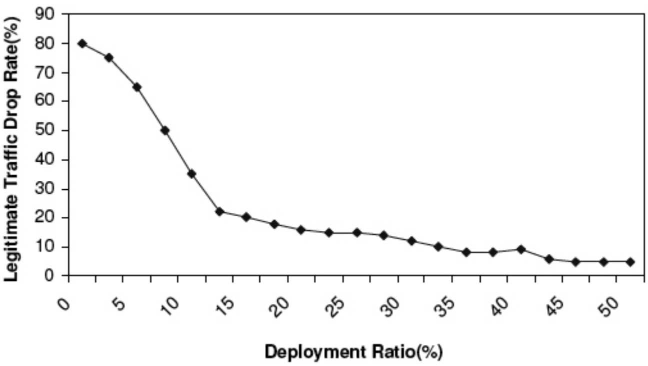

Figure 7: Reduction in false rate with increased deployment

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 defence mechanism enabled, we can notice a gradual increase of the legitimate packet rate. The ramp-up behaviour is due to the false detection of local defence node. As a result, some legitimate traffic will be dropped by the rate limiting mechanism as well. As the algorithm converge, each defence node gets more precise information about the global attack information and thus can rate-limit attack traffic with more accuracy.

In the second set of experiments, we investigate the benefit of increased deployment of defence node. Figure 7 displays the average number of false alarm rates, which decrease gradually with more nodes joined in the peer to peer defence overlay network. By adding sufficient nodes to the defence overlay, attack traffic can be dropped efficiently and the amount of the attack packets reaching the victim server will decrease. The decrease of legitimate packet drop rate will stabilize when the deployment ratio is greater than 20%. This is because as we add more routers as local defence nodes into the cooperative defence network, more of them will be on the same path from the attack traffic source to the victim destination. As a result, the attack information sharing among them will not increase the overall knowledge about the network attack that much.

In the third set of experiments, we vary the parameters of the gossip mechanism to investigate the relationship between the overhead of information sharing and defence efficiency. Let p represent the probability that each detection node in the detection overlay network sends the local attack information to its neighbour 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 2. The

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%

Table 2: Cooperative defense performance

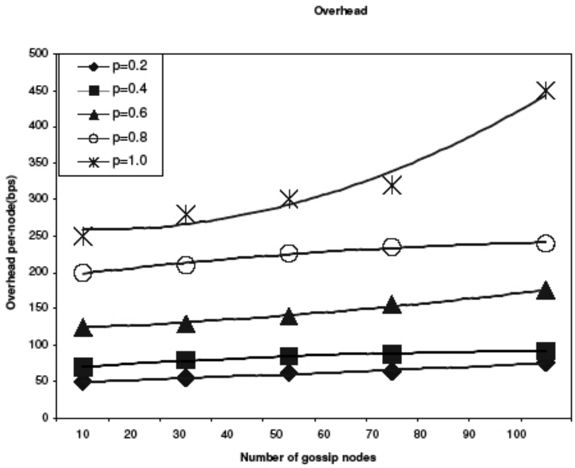

Figure 8: Information sharing overhead

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

As we can see from the simulation results, our algorithm can detect and defend 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 defence node detects an attack, as a result legitimate packets will be dropped dramatically in the case of false detection.

Defences 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 8 shows the per-node overhead with different number of nodes in the system. The packets processed by each node for the cooperative defence purpose do not increase much as we add more nodes into the defence overlay. So the gossip based information sharing mechanism is scalable to be used in larger and higher speed network situations. When the gossip probability is increased, the overhead will increase as well. This parameter can be tuned to adapt different applications to achieve optimum performance.

5. RELATED WORK.

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

Distributed Packet Filtering(DPF) (Park and Lee, 2001) explores the power-law of Internet topology in source address validation. It can be distributed in Internet core routers to probatively stop packet flows with obviously wrong source addresses, and meanwhile to reactively trace back the attack sources. Empirical experiments show that DPF can efficiently identify spoofed IP addresses outside autonomous system(AS) where the attacker resides. In particular for 1997–1999 Internet topology, enabling DPF on 18:9% vertex cover of Internet AS topology can effectively stop 88% traffic with a spoofed source address. However, though the vertex cover ratio is relatively stable according to Internet AS topology study, the membership of the vertex cover may vary over time. Besides, computing appropriate filtering tables based on existing inter domain routing protocol is a non-trivial problem. Inconsistent filter table will lead DPF to regular traffic.

Pushback (Ioannidis and Bellovin, 2002) and Aggregate-Based Congestion Control(ACC) are projects at AT&TxCenter for Internet Research. The routers in 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 for 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 neighbours who carry the aggregates 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 requests all the routers on the path of aggregate traffic be augmented with the pushback capability.

In Xuan, Bettati and Zhao (2001) a collaborative DDoS defence system is proposed in which routers act as gateways, detecting DDoS attacks locally and identifying and dropping packets from misbehaving flows. Gateways are installed and communicate only within the source and the victim domains, thus providing cooperative defence of a limited scope. Similarly, COSSACK (Papadopoulos *et al*, 2003) forms a multicast group of defence nodes which are deployed at source and victim networks. Each defence node can autonomously detect the attack and issue an attack alert to the group. Sources involved in the attack cooperate with the victim to suppress it. Since intermediate networks do not participate in defence, systems described in Xuan *et al* (2001) and Papadopoulos *et al* (2003) cannot control attack traffic from networks that do not deploy proposed defence.

6. CONCLUSION AND FUTUREWORK

Distributed denial of service is a major threat that cannot be addressed through isolated actions of sparsely deployed defence nodes. Instead, various defence systems must organize into a framework and inter-operate, exchanging information and service, and acting together, against the threat (Mirkovic, Prier and Reihe, 2003; Papadopoulos *et al*, 2003). In this paper we proposed a global defence 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 defend DDoS attacks with higher accuracy. Compared to the existing solutions, our contribution is to provide a distributed proactive DDoS detection and defence mechanism. Our approach continuously monitors the network. When an attack begins, individual defence nodes drop attack traffic identified according to the local information and mitigate load to the target victim. However, as local detection has a high false alarm rate, the legitimate traffic will drop 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 defend against DDoS attacks more effectively. Future work will fold in more topology information and vulnerability information gleaned from automated scanning and mapping tools. When the nodes know more topology information of the global Internet, it can utilize more intelligent gossip strategy to reduce the information sharing overhead while trying to detect attacks with speed. Armed with these more sophisticated methods, our approach can defend attacks more efficiently. We are investigating several important questions that still need to be addressed. These include the consensus algorithm and optimal gossip period. We also plan to validate this scheme by running them on real attack data sets.

REFERENCES

AKELLA, A., BHARAMBE, A., REITER, M. and SESHAN, S. (2003): Detecting DDoS attacks on ISP networks, in *ACM SIGMOD Workshop on Management and Processing of Data Streams*, San Diego, CA, 20–23.

CHANG, R. K. C. (2002): Defending against flooding-based, distributed denial-of-service attacks: A tutorial, *IEEE Communications Magazine* 40(10): 42–51.

DITTRICH, D. (2004): Distributed denial of service (DDoS) attacks/tools. http://staff.washington.edu/dittrich/misc/ddos/.

ESTAN, C. and VARGHESE, G. (2002): New directions in traffic measurement and accounting, in *Proceedings of SIGCOMM 2002*, Pittsburgh, PA, USA, 270–313.

EUGSTER, P., GUERRAOUI, R., KERMARREC, A. and MASSOULIE, L. (2004): Epidemic information dissemination in distributed systems, *IEEE Computer* 37(50): 60–67.

GIL, T. M. and POLETO, M. (2001): Multops: a data-structure for bandwidth attack detection, in *Proceedings of 10th Usenix Security Symposium*, Washington, D.C., USA, 23–28.

GUPTA, I., BIRMAN, K. P. and VAN RENESSE, R. (2002): 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.

HUANG, Q., KOBAYASHI, H. and LIU, B. (2003): 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.

IOANNIDIS, J. and BELLOVIN, S. M. (2002): Implementing pushback: Router-based defence against DDoS attacks, in *Proceedings of Network and Distributed System Security Symposium*, NDSS '02, Reston, VA, USA, 100–108.

KARP, R. M., SCHINDELHAUER, C., SHENKER, S. and VOCKING, B. (2000): Randomized rumor spreading, in *IEEE Symposium on Foundations of Computer Science*, 565–574.

KEMPE, D., DOBRA, A. and GEHRKE, J. (2003): Computing aggregate information using gossip, in *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, Cambridge, MA.

KEROMYTIS, A., MISRA, V. and RUBENSTEIN, D. (2002): Using overlays to improve network security, in *Proceedings of the ITCom Conference, special track on Scalability and Traffic Control in IP Networks*, Boston, MA, 245–254.

LIN, M. and MARZULLO, K. (1999): Directional gossip: gossip in a wide area network, in *Proceedings of Dependable Computing – Third European Dependable Computing Conference*, Berlin, Germany, 364–379.

MIRKOVIC, J., PRIER, G. and REIHE, P. (2003): Alliance formation for DDoS defence, in *Proceedings of the New Security Paradigms Workshop, ACM SIGSAC*, Ascona, Switzerland, 11–18.

MIRKOVIC, J., PRIER, G. and REIHER, P. (2002): Attacking DDoS at the source, in *Proceedings of ICNP 2002*, Paris, France, 312–321.

PAPADOPOULOS, C., LINDELL, R., MEHRINGER, J., HUSSAIN, A. and GOVINDAN, R. (2003): Cossack: Coordinated suppression of simultaneous attacks, in *DARPA Information Survivability Conference and Exposition*, Washington, DC, 1: 2–13.

PARK, K. and LEE, H. (2001): On the effectiveness of route-based packet filtering for distributed DoS attack preventation in power-law internets, in *Proceedings of ACM SIGCOMM*, San Diego, CA, USA, 15–26.

PENG, T., LECKIE, C. and KOTAGIRI, R. (2003): Protection from distributed denial of service attack using history-based IP filtering, in *Proceedings of IEEE International Conference on Communications*, Anchorage, Alaska, USA, 1: 482–486.

RENESSE, R., BIRMAN, K. and VOGELS, W. (2003): Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining, *ACM Transactions on Computer Systems* 21(2): 164–206.

ROESCH, M. (2002): The snort network intrusion detection system. http://www.snort.org.

SONG, D. X. and PERRIG, A. (2001): Advanced and authenticated marking schemes for IP traceback, in *Proceedings of IEEE Infocomm*, Anchorage, Alaska, USA, 2: 878–886.

XUAN, D., BETTATI, R. and ZHAO, W. (2001): A gateway-based defence system for distributed dos attacks in high-speed networks, in *Proceedings of 2001 IEEE Workshop on Information Assurance and Security*.

BIOGRAPHICAL NOTES

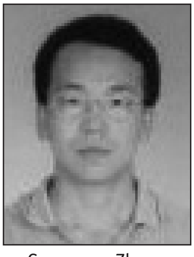
Guangsen Zhang

Guangsen Zhang is a PhD student in the Department of Electrical and Computer Engineering at Rutgers University. He received a BE degree in Electrical Engineering from Xi'an Jiaotong University, China in 1994 and a ME degree in Information Engineering from Beijing University of Posts and Telecommunication, China in 1997. His research interests include autonomic and grid computing, networking and computer security.

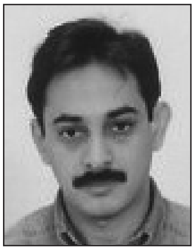

Manish Parashar

Manish Parashar is Associate Professor of Electrical and Computer Engineering at Rutgers University, where he is also director of the Applied Software Systems Laboratory. He received a BE degree in Electronics and Telecommunications from Bombay University, India, and MS and PhD degrees in Computer Engineering from Syracuse University. He has received the NSF CAREER Award and the Enrico Fermi Scholarship from Argonne National Laboratory. His current research interests include parallel and distributed computing (including autonomic, Grid and peer-to-peer computing) with applications to computational science and engineering, networking and software engineering. Manish is a member of the executive committee of the IEEE Computer Society Technical Committee on Parallel Processing (TCPP), part of the IEEE Computer Society Distinguished Visitor Program (2004–2006), and a member of ACM. He is also the cofounder of the IEEE International Conference on Autonomic Computing (ICAC). Manish has co-authored over 140 technical papers in international journals and conferences, has co-authored/edited six books/proceedings, and has contributed to several others in the area of computational science and parallel and distributed computing. For more information please visit http://www.caip.rutgers.edu/~parashar