*Integrated Design and Process Technology*, **IDPT-2006**
Printed in the United States of America, June, 2006


# An Effective Defence Mechanism against Massively Distributed Denial-of-Service Attacks

**Y. Chen, S. Das, P. Dhar, A. El Saddik and A. Nayak**
**School of Information Technology and Engineering,**
**University of Ottawa, 800 King Edward Avenue, Ottawa, ON K1N 6N5, Canada.**
{***yaochen, shantdas, elsaddik, anayak***}***@site.uottawa.ca***
**Cistech Limited,**
**30 Concourse Gate, Unit 35, Ottawa ON K2E 7V7, Canada**
***pulak@cistech.ca***

**Abstract.** The Denial of Service(DoS) attacks, and more specifically Distributed Denial of Service(DDoS) attacks, have become one of the major threats to the operation of the Internet today. The victim (typically a web server) is taken off the internet by exhausting its resources under a massive flow of attack packets, keeping it too busy to continue its normal services. Most often the attackers use spoofed IP addresses to disguise the source of packets which makes it very difficult to track and stop these attacks.

In this paper, we investigate existing and possible defense mechanisms against DDoS attacks and then propose a novel Marking-based DDoS Attack Detection and Filtering(MDADF) scheme. Our scheme is based on a firewall that can distinguish the attack packets (containing spoofed source addresses) from the packets sent by legitimate users, and thus filters out most of the attack packets before they reach the victim. Unlike previously proposed solutions, our scheme has a low deployment cost requiring the cooperation of only about 20% of the Internet routers in the marking process. (The scheme is also designed to quickly detect the occurrence of an attack and warn the victim, so that appropriate action can be taken.) We have extensively tested our scheme on simulated DDoS attacks with up to several thousand attackers and the simulation results show that the MDADF scheme can effectively filter-out more than 90% of attack packets without much affecting the flow of legitimate packets to the victim web-server.

**Keywords:** Denial-of-Service attack, Distributed DoS, packet marking, firewall, filtering scheme

## I. Introduction

In recent years, the Internet has experienced rapid growth—more than 353 million computers are connected to the Internet [8] and the number of users has increased to nearly 1 billion in 2005 [9]. The Internet has penetrated into many aspects of our life and many important services such as power, transportation, banking, and medicine are now dependent on the Internet, making it crucial that there are no disruptions in the operation of the Internet. However, the Internet was designed for openness and scalability without much concern for security and this point has been exploited by attackers and worm writers to unleash a flurry of disruptive activities sometimes causing severe financial damages to users of the Internet. The number of reported Internet security incidents has increased from 252 in 1990 to 137,529 in 2003 and a lot more in last two years due to the use of automated attack tools [6]. Email viruses, computer worms, and Denial of Service(DoS) attacks have severely threatened the normal operation of the Internet, with the DoS attacks and its distributed version (DDoS) being the most difficult to prevent or control.

A Denial of Service (DoS) attack has the sole purpose of reducing or eliminating the availability of a service provided over the Internet, to its legitimate users. This is achieved either by exploiting the vulnerabilities in the software, network protocols, or operation systems, or by exhausting the consumable resources such as the bandwidth, computational time and memory of the victim. In the distributed DoS attacks, the attacker first takes control of a large number of vulnerable hosts on the internet, and then uses them to simultaneously send a huge flood of packets to the victim, exhausting all of its resources.

There are a large number of exploitable machines on the internet, which have weak security measures, for attackers to launch DDoS attacks, so that the attacks can be executed by an attacker with limited resources against the large, sophisticated sites, including Yahoo, Amazon, CNN, eBay, and Microsoft. In October 2002, some of the Internet root servers, Domain Name System (DNS), also became the target of DDoS attack.

The devastating effects of the DoS and DDoS attacks have caused people's attention, and many different mechanisms have been proposed to deal with them. However, most of them are ineffective against massively distributed DoS attacks involving thousands of compromised machines. In the following, we analysis the existing solutions for DoS attack and propose a Marking-based Detection and Filtering (MDADF) scheme to defend massively distributed DoS attacks.

## II. Approaches for Defending DoS/DDoS Attacks

Current DoS/DDoS defenses can be classified into three categories: preventive mechanisms, reactive mechanisms, and source-tracking mechanisms. The preventive schemes aim at improving the security level of a computer system or network; thus preventing the attacks from happening, or enhancing the resistance to attacks (see for example [11]). Such solutions are generally costly and difficult to implement.

The source-tracking schemes on the other hand, aim to track-down the sources of attacks, so that punitive action can be taken against them and further attacks can be avoided. The existing solutions fall into four groups: packet marking [1], [16], [17], [19], message traceback[3], [12], logging [18], [21], and traffic observation [5]. Many different packet marking schemes have been proposed, for encoding path information inside IP packets, as these are routed through the internet. The idea is to use the encoded information to reconstruct path through which a packet has travelled and thus track down the sources of offending packets. In the message *traceback* method, routers send *traceback* messages in separate packets to the destination. When enough messages from the routers along the path have been received by the victim, the attack path can be determined. Another method called *logging* is to record packet information at routers. The path to attackers can be determined by the routers exchanging information with each-other. The traffic-observation method is based on observing the changes in the traffic flow to detect the attacks and determine the attack paths. A common problem existing in these four solutions is that the reconstruction of attack path becomes quite complex and expensive when there are a large number of attackers (i.e. for highly distributed DoS attacks). Also, these type of solutions are designed to take corrective action after an attack has happened and cannot be used to stop an ongoing DDoS attack.

The reactive measures for DDoS defence are designed to detect an ongoing attack and react to it by controlling the flow of packets to mitigate the effects of the attack. One of the proposed reactive schemes, given by Yaar et al. [22] uses the idea of packet marking for filtering out the attack packets instead of trying to find the source of such packets. This scheme uses a path identifier (called *Pi*) to mark the packets; the *Pi* field in the packet is separated into several sections and each router inserts its marking to one of these. Once the victim has identified the marking corresponding to attack packets, it can filter out all such packets coming through the same path.

Other reactive schemes proposed to control DDoS attacks include D-WARD [14] and *Pushback* [10]. D-WARD is designed to be deployed at the source network; it monitors the communication rate and imposes a rate-limit on any suspicious outgoing flow according to its objective function. The *Pushback* method generates an attack signature after detecting a congestion, and applies a rate limit on corresponding incoming traffic, then this information is propagated to upstream routers so that the attack flow can be pushed-back. The success of both these methods depend on the precision of generating a signature for the attack flow.

## III. Designing an Effective Protection Scheme

Generalizing from the various defense mechanisms, a good protection scheme against DDoS attacks should be based on continuous monitoring, precise detection and timely reaction to attacks. The following characteristics are desirable:

- The scheme should be able to control or stop the flow of attack packets before it can overwhelm the victim. The timely detection and immediate reaction to an attack is essential, to prevent the depletion of resources at the victim location. The suitable place to deploy defense scheme are the perimeter routers or the firewall of a network.
- In stopping the flow of attack packets to the victim, the scheme must ensure that packets from legitimate users are successfully received so that the service to the legitimate users is not denied or degraded. Any degradation in service would signify a partial success for the denial of service attack.
- The implementation cost should be low. Unless most internet users fully recognize the threats posed by DoS/DDoS attacks, it is difficult to get cooperation from them in defending such attacks, especially when the investment required is costly. Therefore, any viable DDoS defence scheme should require minimal participation of third party networks or intermediate routers on the internet.

A good defence mechanism should be able to precisely distinguish the attack packets from the legitimate packets. What makes it difficult to control or stop the DDoS attacks is the use of spoofed IP address. Spoofed packets are commonly used in DoS/DDoS attacks to hide the location of attackers and the compromised machines, so that the paths to them are concealed. If we can distinguish the packets which have spoofed IP addresses, then these packets can be selectively filtered out by a firewall to stop most attacks.

## IV. Distinguishing the Attack Packets

In this section, we present a new packet marking method which will help us to distinguish DDoS attack packets from packets sent by legitimate users.

In order to make the marking scheme fast and efficient we use part of the header in an IP packet, as the marking field. The 16-bit Identification field in IP header has been commonly employed as the marking space (see [1], [2], [16], [17], [19], [22]). The Identification(ID) field is currently used to indicate IP fragments belonging to different packets, but only less than 0.25% of the packets on the Internet actually use this feature [20]. Therefore, employment of ID-field as the marking space will not much affect

the normal transmission of IP packets. In our scheme each cooperating router on the path of an IP packet would insert a mark on the ID-field of the packet. The genrated marking should be such that two packets reaching the victim through diffrent routes are guaranteed to have distinct markings.

### A. Computing the Packet Marking

The mark made by a router would be a function of its IP-address. To fit the 32-bit IP address $A$ of a router into the ID-field, we employ a hash function $h$ that converts $A$ to a 16-bit value.

As the IP addresses of routers can be easily known, each router $R$ also uses a 16-bit (randomly generated) key $K_R$ when computing its marking, so that the markings cannot be spoofed. The marking for a router *R* is thus calculated as $M_R$ = h(A) XOR $K_R$, where $A$ is the IP address of the router. To make the marking scheme more effective, we let each router perform a Cyclic Shift Left(CSL) operation on the old marking $M_{old}$ and compute the new marking as $M = CSL(M_{old}) \oplus M_R$. In this way, the order of routers on the path of a packet influences the final marking on the packet.

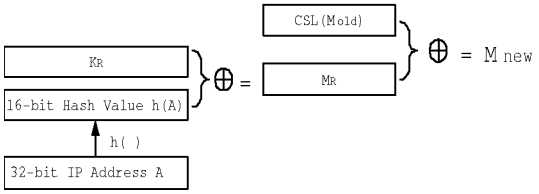


Fig. 1. The Marking Scheme

## V. Filtering Scheme

The MDADF scheme employs a firewall at each of the perimeter routers of the network to be protected and the firewall scans the marking field of all incoming packets to selectively filter-out the attack packets (see Figure 2).

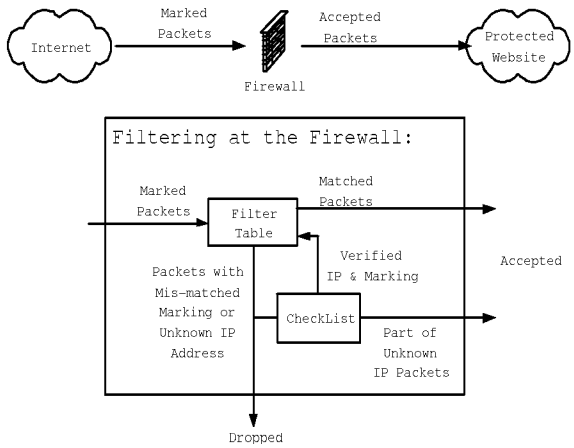


Fig. 2. The System Structure

On employing our marking scheme, when a packet arrives at its destination, its marking depends only on the path it has traversed. If the source IP address of a packet is spoofed, this packet must have a marking that is different from that of a genuine packet coming from the same address. The spoofed packets can thus be easily identified and dropped by the filter, while the legitimate packets containing the correct markings are accepted.

### A. Learning Phase

To distinguish the spoofed packets, the firewall needs to keep a record of the genuine markings. During normal time that no attacks are happening, the firewall can learn about the correct markings for packets sent from specific IP addresses. The (IP-address, Marking) pairs are stored in a *Filter Table*[1], which are later used to verify each incoming packet and filter-out the spoofed ones. The learning phase continues for a sufficient time to allow most of the filter table to be filled up. If the Filter Table gets full, any new entry to be added replaces the oldest one.

### B. Normal Filtering Procedure

After the learning phase, the firewall begins to perform its normal filtering operations. To the packet from an IP address recorded in the Filter Table, it is accepted if it has a consistent marking; otherwise, it is dropped. For the packet from a new IP address, we accept it with probability $p$ and put the (IP-address, Marking) pair to a *Check List*, so that the marking can be verified. The value of $p$ is set to high (close to 1) initially. When an attack is detected, the value of $p$ is decreased according to the packet arrival rate and the victim's capability for handling the incoming traffic.

### C. Marking Verification

To verify the markings in the Check-List, a random *echo* message is sent periodically to the source address for each (IP-address, Marking) pair in the Check-List, and a counter is used to record the number of echo messages have been sent for it. To avoid the reply being imitated by the attacker, the content of the echo message is recorded in the Check-List and compared with the content of reply received.

On receiving an echo reply from the source, the marking can be verified and the (IP-address, Marking) pair is moved to the Filter Table; otherwise, it indicates the previously received packet was spoofed, then this pair is deleted from the Check List. If the counter in the Check List shows that more than $d(= 10)$ echo messages have been sent to an IP address $x$, then the entry for this IP address is removed from the Check List and the pair (x,$\phi$) is added to the filter table, where $\phi$ is a special symbol denoting that all packets having source IP address x should be discarded. Since in this

[1]The filter table can be implemented as a content-addressable memory to speed up the filtering process.

situation, this source IP must be either non-existent or inactive, so that the packets received with this source address are coming from the attacker and need to be rejected.

### D. Attack Detection

To detect the start of a DDoS attack, we use a counter called Total-Mismatches-Counter ($TMC$), which counts the number of packets whose marking cannot be matched at the firewall. This includes both packets with incorrect markings as well as packets from unknown source addresses that are not recorded in the Filter Table. When the $TMC$ value becomes greater than a threshold $\theta$, it is considered as a signal of DoS/DDoS attack. The value of $TMC$ is reset to zero after fixed intervals to ensure that the cumulative results over a long duration is not considered as the indication of attack by mistake.

### E. Route Change Consideration

Though routes on the Internet are relatively stable, they are not invariable. Once the route between two hosts has changed, the packet received by the destination will have a different marking with the one stored in the Filter Table, so that it may be dropped according to our basic filtering scheme.

Taking route changes into consideration, we introduce another counter called $SMC$, to count the number of mismatching packets for any IP address $A$. When the value of $SMC_A$ reaches a threshold $\delta$, the entry ($A$, $Marking_A$) is copied to the Check List to test whether the route from this source has changed and $SMC_A$ is reset to zero. If the new marking is verified by the Check List verification process, the marking for this IP address is updated in the Filter Table. Otherwise, the original marking is preserved. Unless the route change has been verified, the original marking is still used to filter packets.

## VI. Experimental Results

We have evaluated the performance of the MDADF scheme under various parameter settings by simulating DDoS attacks of different magnitudes.

### A. Simulation of Internet Traffic

In our simulation, we have used the topological data obtained from the Internet Mapping Project [4] of Lumeta Corporation. This data was generated by using *traceroute* to probe the paths in Internet from a single host (netmapper.research.lumeta.com, 65.198.68.56). Since all the paths in the database congregate at the single node, this node is quite suitable to act as the victim in the simulations of DDoS attacks.

We use a packet generator process to simulate the normal Internet traffic, which periodically sends packets from a randomly selected internet user. Then the packet marking process is simulated, by computing the markings for each cooperating router on the route for this particular user. Finally, the marked packet is inserted into a packet-queue at the firewall of the victim. The rate, at which packets are added to the packet-queue, mimics the normal traffic flow for a typical web-server on the Internet.

Attackers usually have two methods to disguise the source locations: spoofing a genuine host's IP address, or inserting a randomly generated IP address into source address field. We simulated both types of attacks, called *Spoofed* attack and *Randomized* attack respectively. Packets are generated from each attacker to simulate the attack traffic. So, higher the number of attackers, more will be the volume of the attack flow. In the simulation of Spoofed attack, for each attack packet, one of the legitimate user is randomly selected and its IP address is used as the spoofed value of the source address. The marking field is initially filled with a random value and the marking process is simulated, as before.

### B. Parameter Selection

The choice of values for different parameters affects the performance results of the MDADF scheme. In our experiments, we have come up with some suitable values for these parameters by trail and error and we have tested the effects of changing the values of these parameters. The data-set used in the experiments contained 10,000 hosts and 50,000 intermediate routers. The size of the filter table was varied from 5000 to 10,000. The participation rate of routers was varied from 100% to 0%. For the parameter $p$, the most suitable value were found to 0.75 and 0.1 respectively for the pre-attack and post-attack scenarios. A learning phase of 10 minutes gave good results in most scenarios. In the following, we discuss the results obtained in the simulations.

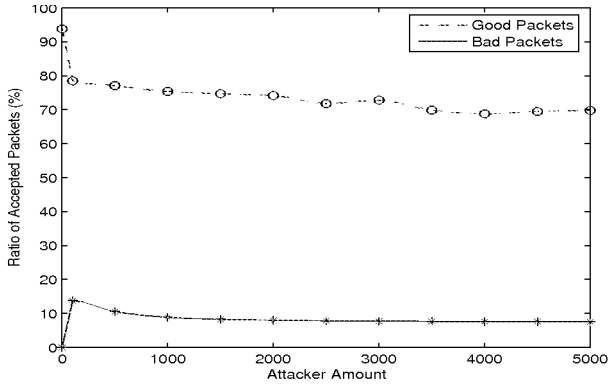


Fig. 3. Ratio of Accepted Packets vs. Number of Attackers under Spoofed Attack

### C. Performance under Spoofed Attacks

Figure 3 shows the ratio of packets that were accepted at the firewall under different magnitudes of attack. As can be

seen, more than 70% good packet are accepted even in the most severe condition under spoofed attack, in which the attack traffic is almost 10 times of the normal traffic. As the number of attackers is increased, there is a slight decrease in the acceptance rate of legitimate packets, as due to the heavy congestion some packets are dropped before the firewall can handle them. Though the good packets acceptance ratio decreases a little with the increasing number of attackers, the bad packets accepting ratio stays at a very low level.

The participation of routers in the marking scheme is important to our MDADF scheme; However we cannot expect all the routers to be willing to cooperate. Therefore, we have tested the effect of different participation rates on our scheme when the attackers are 500, 2000, and 5000 respectively. We show the difference between acceptance ratio of good and bad packets, which is called "Acceptance Ratio Gap". Obviously, when the participation rate is zero and no markings are applied, the acceptance ratio of good and bad packets should be equal and the gap is zero. As shown in Figure 4, many more good packets are accepted than the bad ones even when the participation rate is only 20% under all three conditions. This means that the MDADF scheme can efficiently distinguish between good and bad packets when just 20% of routers in the Internet deploy our marking scheme.

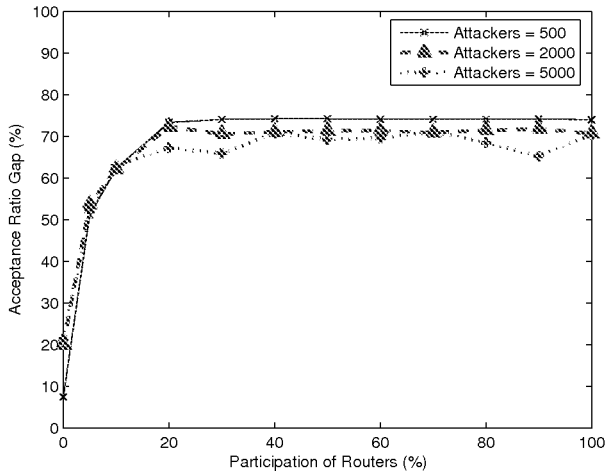


Fig. 4. Acceptance Ratio Gap vs. Router Participation under Spoofed Attacks

The Filter Table is a critical part of our system, which stores the (IP-address, Marking) pairs and more the number of records in the Filter Table, less good packets will be dropped by mistake. We have tested the performance of our scheme with different sizes of Filter Table under different attack environments and Figure 5 shows that keeping 80% of the legitimate user records is sufficient to filter packets, and even 70% is good enough.

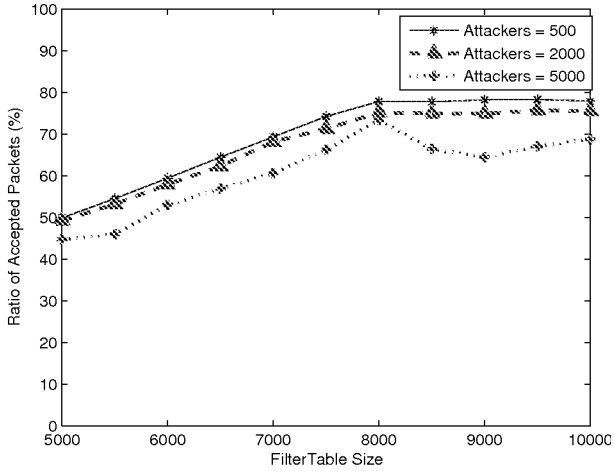


Fig. 5. Acceptance Ratio Gap vs. Filter-Table Size

### D. *Performance under Randomized Attacks*

We tested our scheme's performance under randomized attack in which attackers use randomly generated IP addresses. From Figure 6, we can see that our scheme is efficient under aggregate attacks.

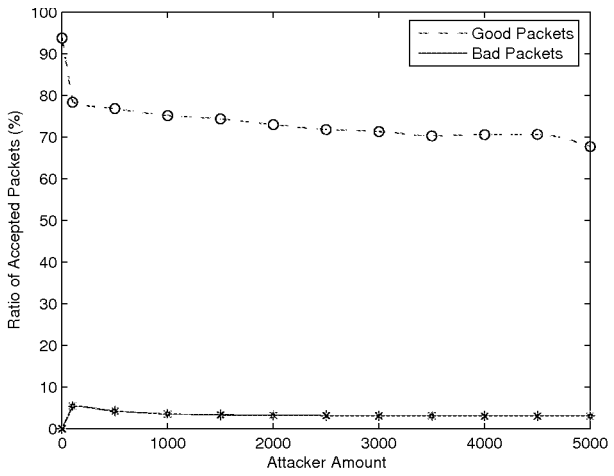


Fig. 6. Ratio of Accepted Packets vs. Number of Attackers under Randomized Attacks

### E. *Attack Detection Time*

Under both types of attacks, our scheme can detect the occurrence of an attack in 3 - 4 seconds in most conditions as shown in Table I However if the number of attackers is too small (100), it takes a little more time for the system to notice the effects of the attack.

| # of Attackers | Attack Detection Time (sec) | |
|---|---|---|
| | Spoofed Attack | Randomized Attack |
| 100 | 6.86 | 6.85 |
| 500 | 3.61 | 3.65 |
| 1000 | 3.53 | 3.56 |
| 1500 | 3.50 | 3.53 |
| 2000 | 3.50 | 3.53 |
| 2500 | 3.49 | 3.53 |
| 3000 | 3.47 | 3.50 |
| 3500 | 3.47 | 3.49 |
| 4000 | 3.47 | 3.50 |
| 4500 | 3.47 | 3.50 |
| 5000 | 3.47 | 3.50 |

TABLE I
ATTACK DETECTION TIME FOR DIFFERENT NUMBER OF ATTACKERS

## VII. CONCLUSIONS AND DISCUSSION

In this paper, we have proposed a low-cost and efficient scheme called MDADF, for defending against DDoS attacks, The MDADF scheme is composed of two parts: marking process and filtering process. The marking process requires the participation of routers in the Internet to encode path information into packets. We suggest the use of a hash function and secret key to reduce collisions among packet-markings. The scheme also includes mechanisms for detecting and reporting DDoS in a timely manner.

The evaluation of the scheme under simulations, show that our scheme can effectively and efficiently differentiate between good and bad packets under spoofed attack when the routers' participation rate is as low as 20%, so the deployment cost of our scheme is very low. Also, most good packets are accepted even under the most severe attack, whose traffic is about 10 times of normal traffic. At the same time, the bad packet acceptance ratio is maintained at a low level. Our scheme performs well even under massively distributed DoS attacks involving upto 5000 attackers.

Under both spoofed and randomized DDoS attacks, the MDADF scheme detected the occurrence of attack precisely within 3 - 4 seconds. The quick detection is valuable to the victim so that appropriate actions can be taken to minimize the damage caused by a DDoS attack.

### *Acknowledgements*

The authors would like to thank Bill Cheswick and Hal Burch for providing the internet dataset used in our simulations.

## REFERENCES

[1] Andrey Belenky and Nirwan Ansari. IP Traceback With Deterministic Packet Marking. *IEEE Communications Letters*, vol. 7, no. 4, pp. 162–164, April 2003.

[2] Andrey Belenky and Nirwan Ansari. Tracing Multiple Attackers with Deterministic Packet Marking (DPM). *2003 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing (PACRIM 03)*, pp. 49-52, August 2003.

[3] Steve Bellovin. ICMP Traceback Messages. *Internet draft, work in progress*, March 2000.

[4] Bill Cheswick and Hal Burch. Internet Mapping Project. http://research.lumeta.com/ches/map/

[5] Hal Burch and Bill Cheswick. Tracing Anonymous Packets to Their Approximate Source". *Proceedings of the 14th Systems Administration Conference (LISA 2000)*, pp. 319-327, December 2000.

[6] CERT® Coordination Center. CERT/CC Statistics 1988-2005. http://www.cert.org/stats/cert_stats.html#incidents

[7] Yao Chen. A Novel Marking-based Detection and Filtering Scheme Against Distributed Denial of Service Attack. Masters Thesis, University of Ottawa, 2006.

[8] Internet System Consortium. ISC Domain Survey: Number of Internet Hosts. http://www.isc.org/index.pl?/ops/ds/host-count-history.php

[9] Internet World Stats. Internet User Statistics – The Big Picture: World Internet Users and Population Stats. http://www.internetworldstats.com/stats.htm

[10] John Ioannidis and Steven M. Bellovin. Implementing Pushback: Router-Based Defense Against DDoS Attacks. *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, pp.6-8, February 2002.

[11] Sherif M. Khattab, Chatree Sangpachatanaruk, Rami Melhem, Daniel Mosse, and Taieb Znati. Proactive Server Roaming for Mitigating Denial-of-Service Attacks. *Proceedings of the 1st International Conference on International Technology: Research and Education (ITRE 03)*, pp. 500-504, August 2003.

[12] Allison Mankin, Dan Massey, Chien-Lung Wu, S. Felix Wu, Lixia Zhang. On Design and Evaluation of "Intention-Driven" ICMP Traceback. *IEEE International Conference on Computer Communication and Networks (ICCCN'01)*, pp. 159-165, October 2001.

[13] Jelena Mirkovic, Sven Dietrich, David Dittrich, and Peter Reiher. *Internet Denial of Service – Attack and Defense Mechanisms*. The Radia Perlman Series in Computer Networking and Security. December 2004.

[14] Jelena Mirkovic, Gregory Prier, and Peter Reiher. Attacking DDoS at the Source. *Proceedings of the IEEE International Conference on Network Protocols 2002*, pp. 312-321, November 2002.

[15] Kihong Park and Heejo Lee. On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack Prevention in Power-Law Internets. *Proceedings of ACM SIGCOMM 2001*, August 2001.

[16] Tao Peng, Christopher Leckie, and Kotagiri Ramamohanarao. Adjusted Probabilistic Packet Marking for IP Traceback. *Networking 2002*, pp. 697-708, May 2002.

[17] Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson. Practical Network Support for IP Traceback. *Proceedings of ACM SIGCOMM 2000*, August 2000.

[18] Alex C. Snoeren, Craig Partridge, Luis A. Sanchez, Christine E. Jones, Fabrice Tchakountio, Stephen T. Kent, and W. Timothy Strayer. Hash-Based IP Traceback. *Proceedings of ACM SIGCOMM 2001*, pp. 3-14, August 2001.

[19] Dawn Xiaodong Song and Adrian Perrig. Advanced and Authenticated Marking Schemes for IP Traceback. *Proceedings of IEEE INFOCOM*, pp.878-886, April 2001.

[20] Ion Stoica and Hui Zhang. Providing Guaranteed Services Without Per Flow Management. *Proceedings of ACM SIGCOMM 1999*, pp. 81-94, April 1999.

[21] Robert Stone. CenterTrack: An IP Overlay Network for Tracking DoS Floods. *Proceedings of USENIX Security Symposium '00*, pp.199-212, August 2000.

[22] Abraham Yaar, Adrian Perrig, and Dawn Song. Pi: A Path Identification Mechanism to Defend against DDoS Attacks. *Proceedings of the IEEE Symposium on Security and Privacy*, pp. 93-109, May 2003.