Packet-Marking Scheme for DDoS Attack Prevention

K. Stefanidis and D. N. Serpanos

{stefanid, serpanos}@ee.upatras.gr

Electrical and Computer Engineering Department
University of Patras
Patras, Greece

Abstract

One of the main difficulties in the detection and prevention of Distributed Denial of Service (DDoS) attacks is that the incoming packets cannot be traced back to the source of the attack, because (typically) they contain invalid or spoofed source IP address. For that reason, a victim system cannot determine whether an incoming packet is part of a DDoS attack or belongs to a legitimate user. Various methods have been proposed to solve the problem of IP traceback for large packet flows. These methods rely on the assumption that they can gather a sufficient number of packets from the same source, in order to reconstruct the traversed path or to determine the source address. In this paper we introduce a packet marking scheme which enables the unique identification of the path that each incoming packet has traversed, relying only on the information inside that packet. We show how the proposed scheme enables real time identification and filtering of the DDoS attack traffic. The proposed scheme is simple to implement, introduces no bandwidth overhead, low computational overhead and has low fault probability. Using the above metrics, we compare our proposed scheme with existing marking schemes and demonstrate its advantages over them. Finally, we introduce a method that can be used post mortem, in order to determine the source IP address of the attacking systems (up to the nearest router to the source).

Keywords: Distributed Denial of Service Attacks, Packet Marking, IP Traceback.

1 Introduction

One of the major open problems in network security today is Distributed Denial of Service (DDoS) Attacks. In a DDoS attack, the attacker sends vast amounts of traffic from a large number of systems that are controlled by him/her to a victim network or system. The result is that the victim's resources become overloaded and it cannot process the requests of legitimate users, thus any services that this system provides become unusable. One of the main difficulties in the detection, and prevention of DDoS attacks is that the incoming packets cannot be traced back to the source of the attack, because (typically) they contain invalid or spoofed source IP address. For that reason, a victim system cannot determine whether an incoming packet is part of a DDoS attack or belongs to a legitimate user.

This paper addresses the problem of identifying the sources of a DDoS attack even if the source IP address of the incoming packets is spoofed. In section 2 we will present the various traceback methods that have been proposed. In section 3 we will discuss our approach towards the solution of the traceback problem. In section 4 we will present the basic assumptions the have been made. In section 5 we will present our traceback method in detail. In section 6 we will present an analysis of our proposed method as well as comparisons with other related methods. Further work and conclusions can be found in sections 6 and 7 respectively.

2 Related Work

Over the last years, various methods have been proposed to provide an effective solution to the IP spoofing problem. Those methods use different approaches such as link testing, packet logging, packet marking and ingress filtering. In this section we will present the basic characteristics of some of these methods.

One of the earliest solutions that have been proposed regarding the IP spoofing problem is ingress filtering. According to this scheme, the routers drop packets that have illegitimate source IP address [15]. This requires that the router has sufficient information on legitimate addresses for each ingress interface. The required information and the lookup procedure for each packet, introduces a large computational and memory overhead.

Moreover, the administrative burden and possible complications with protocols that depend on source IP address spoofing are some of the factors that prevent universal deployment of ingress filtering.

A method suggested in [9] [12] is to log the packets in the routers and use this information for traceback purposes. This scheme can trace an attack long after the attack has completed. The obvious drawback of packet logging is the amount of memory required by the routers. The memory overhead can be reduced by storing only a digest of the packet's header instead of the whole packet [9].

Another proposed method involves the input debugging feature of some routers which enables the operator to filter some packets on the egress port and determine which ingress port they arrived on. This feature can be used in a hop by hop fashion to determine the source of packets that bear a common characteristic or signature. This method introduces high management overhead because it requires the collaboration of all the ISPs along the packet's path. There are some tools that have been developed and can automate the procedure to some point [12].

One last method that has been proposed is packet marking. In packet marking, routers alter the IP header of the traversing packets (they mark them) in order to notify the end host of their presence on the route. The end host can gather those markings and rebuild the traversed path for large packet flows. In probabilistic packet marking [7] [10], the marking procedure is performed once every N packets. This reduces the computational overhead of the marking but increases the number of packets needed to reconstruct the path. In deterministic packet marking [1] [6], the marking procedure is performed for each packet at edge routers only. This reduces the number of packets needed for path reconstruction. The traceback method that we present in this paper is based on packet marking.

3 Purpose of this Scheme

The basic notion of this marking scheme is that in order to stop an ongoing DDoS attack; we primarily need the information that enables us to distinguish the packets that belong to the attack from the packets that belong to normal users of the service. The source IP address of the packet is not reliable during a DDoS attack, as discussed in section 1.

What is of most importance in the procedure of preventing an ongoing DDoS attack is that we need this information to be part of the packet itself. This information must be reliable and enable us to identify the true source of the packet as accurately as possible. Thus a victim of a DDoS attack can use this information together with a DDoS detection system in order to correctly identify and filter in real time the attack traffic. After the end of the DDoS attack, we would like to be able to trace the sources of the attack. We need to be able to use the gathered information in order to trace as accurately as possible the source of the packets that have been classified as part of the DDoS attack.

Most existing packet-marking schemes need more than one packet to determine the source of an incoming packet and the traceback procedure is often computational prohibiting to be done for each packet in real time. Other traceback schemes rely on the fact that traceback of a packet is an infrequent procedure. This marking scheme provides the means to filter DDoS attack traffic in real time and to trace the sources of the attack in a post-mortem fashion. The compromise that has to be done is that due to the limited available space for the traceback information (see section 5.3) and the fact that exact host tracing is very difficult on the IP level; this marking scheme as well as similar traceback schemes, can trace the source of the packet up to the closest router. Another compromise is that there will be some false positives, meaning that some packets will be falsely considered part of the attack traffic. However "An ideal traceback system produces no false negatives while attempting to minimize false positives" [9].

4 Basic assumptions

The assumptions that will be used in this paper are largely borrowed from [7] and are the following:

- The attacker may generate any packet
- The attacker knows that he is being traced
- The attacker knows the traceback scheme
- Routing is stable most of the time
- Routers are not compromised

- Routers are both CPU and memory limited

The first three assumptions mean that the proposed marking scheme cannot contain any weaknesses that could be exploited by the attacker. The attacker can craft any kind of packet, even packets that bear such markings that could circumvent traceback or filtering of his/her packets. The fourth assumption dictates that we expect most of the packets from a specific source that have the same destination, to follow the same path. Efficiency of this marking scheme can be degraded if the assumption is not true, but success of the scheme is not compromised. The fifth assumption has already been thoroughly discussed in [7] [9] and the last assumption dictates that the overhead that this marking scheme poses to the routers should be limited.

5 Packet Marking Scheme

5.1 Overview

The basic notion of our marking scheme is that when a packet enters the network, it is marked along its path in such a way that when it arrives at its destination it carries a distinctive mark that can be used for filtering purposes and the detection of its origin.

In figure 1 we show an example network as seen by the victim. The routers that forward and mark the packets are noted as Ri and the nodes that produce the traffic are noted as Ni. Each node can be part of the DDoS attack.

The routers mark each incoming packet and leave outgoing packets unchanged. We use the existing IP header to place the packet marking. The marking consists of two fields. The first field is a 12 bit digest of the path that the packet has traversed and we will call it the path field. The routers along the packet's path inject their router signature into the path field. The router signature is a 12 bit representation of the router's IP address. In our scheme we will use bits 0-2, 9-11, 17-19 and 25-28 of the router's 32 bit IP address as the router signature. Ideally, we would like to pass the whole 32 bit address into the path field but this is not possible due to restrictions discussed in section 5.3.

The second field of the marking is the 5 bit distance field. The distance field shows the distance between the source of the packet and the receiver (counting in hops). The distance field has a maximum value of 31 which is a sufficient maximum distance value for the internet today [13]. Nevertheless if the distance reaches the maximum value, it refolds back to 0. This poses no harm to the filtering or traceback procedure as we will show in section 5.5.

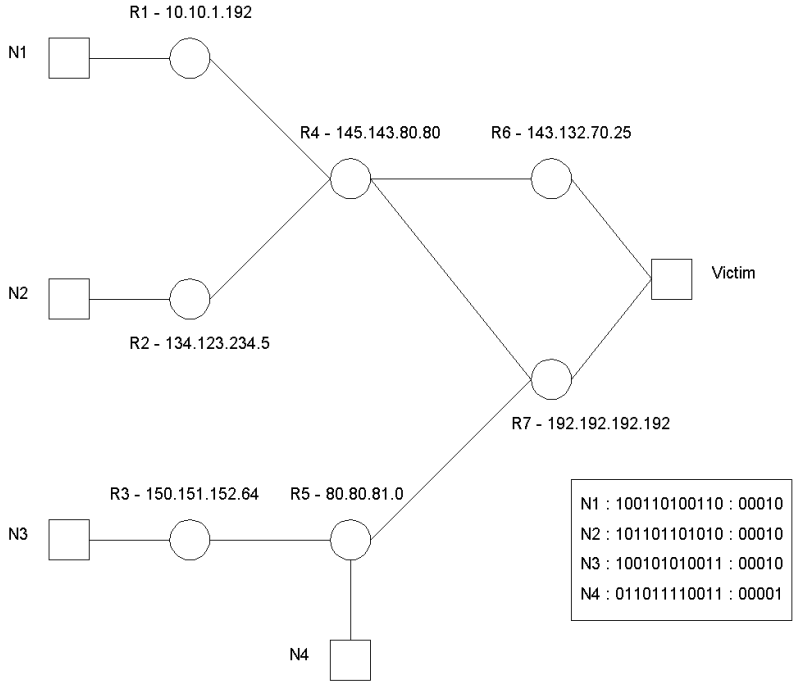

Figure 1: Example network topology, were Ri denotes the routers and their IP addresses and Ni denotes the hosts of the network. We can see how the packets from the various hosts carry unique markings upon their arrival to the victim.

5.2 Marking procedure

As mentioned above, all packets are marked by the routers along the traffic path. The marking procedure is the same for all routers except the edge router that the packet meets first. More specifically, when a packet enters the network, it gets marked by the ingress interface of the closest router. The router marks the path field of the packet with its router signature, overwriting any marking information that the packet has and sets the distance field to 0. All the other routers along the packet's path alter the existing marking by putting the result of the XOR of the existing path field and their router signature. They also increase the distance field by one.

This procedure ensures that every packet in the network gets a marking and possible false markings from the attacker are overwritten. Thus our marking scheme is robust against false markings. The marking procedure in pseudo code can be found below:

Ingress Interface of Edge Router:

```
for each incoming packet p
  p.path = router_signature
  p.distance = 0
```

Other Routers:

```
for each incoming packet p
  p.path = p.path <XOR> router_signature
  if p.distance == 31
    p.distance = 0
  else
    p.distance ++
```

5.3 Coding issues

Figure 2: The IP Header. We use the 16 bit Identification field and 1 bit from the Flags field for the marking procedure. From those 17 bits, we use 12 bits for the path field and 5 bits for the distance field.

One of the main problems of packet marking schemes is that they have to use the existing IP header as a placeholder for the markings. This ensures that there will be no bandwidth overhead from the marking procedure and that the information required for packet traceback will be present inside the packet. In our scheme, as well as in related packet marking schemes [1] [6], we will use the 16 bit identification field and 1 bit of the flags field that is reserved. This means that by overloading the ID field all fragmented packets will be corrupt, thus there are backward compatibility problems that have to be addressed. There are measurements that suggest that only 0.25% of internet traffic is fragmented [11] and it is known that modern network stacks implement automatic MTU discovery to prevent fragmentation [7].

One possible solution is to turn off marking in fragmented traffic but this could be quickly exploited by the attacker. We would like to remind that one of the assumptions taken is that the attacker has full knowledge of the used traceback scheme. For that reason and because we don't want to add any bandwidth overhead by using external signals to identify fragmented traffic, we will accept the trade-off that out marking scheme will not be backward compatible with fragmented traffic. One possible solution would be to set the "don't fragment" flag to 1. This could degrade performance in networks that do not use automatic MTU discovery but it would ensure that no packet gets corrupted.

In figure 2 we show part of the IP header. The 17 bit shaded part is the part that will be used for packet marking. Those 17 bits will be divided into two fields, the 12bit path field and the 5bit distance field.

5.4 Filtering procedure

From the victim's point of view, all incoming packets hold in their IP header a distinctive marking; that is an encoded form of the path that this packet traversed. We will assume that a detection system handles the task of discovering an ongoing DDoS attack. This detection system can be one of the various systems that have been proposed and use statistical methods, neural networks or other methods to identify the existence of a DDoS attack.

When a DDoS attack occurs, the detection system can use the markings of the incoming packets instead of the source IP address to identify the attacking packets and order the border firewall of the victim's network to drop all packets bearing the specified mark. This means of course that the detection system and the firewall have to be reconfigured in order to scan the 17 bits packet marking in addition to the source IP address of a packet. This can be done fairly easily on software firewalls based on iptables or similar technologies and on most of the proposed DDoS detection mechanisms.

Further discussion should be done about the mapping between the marking and the source network of a packet. When the detection system decides that packets bearing a certain marking are part of an ongoing attack and therefore should be dropped, all packets that followed the same exact path will be considered part of the attack. While it is true that all those packets originate from the same source network, the inverse is not always true.

There exists the possibility that packets from the same network will bear different marks since each packet on the internet is individually routed and packets with the same source and destination can follow different routes. On the other hand, it is safe to assume that for the duration of an attack, the routing on the Internet is stable. Nevertheless, if during an attack, the traffic from one network is split in two or more paths; all those traffic flows will have the same characteristics and will be classified by the detection mechanism as part of the ongoing attack. Therefore, the only impact that this routing behaviour will have is that the distinct sources of a DDoS attack will be virtually increased.

We would also like to remind that this marking scheme, as most of the current schemes, is able to provide information up to the closest router to the source. Thus when one or more hosts of a network participate in a DDoS attack, all the hosts of this network will be placed in the victim's drop list.

5.5 Traceback procedure

The information that each packet carries is sufficient for traceback purposes as long as the victim has a map that shows the victim's upstream routers. Later we will show why this assumption is reasonable and practical. In order to trace the possible sources of a packet, we start the traceback procedure from the victim's own edge router towards the leaves of the tree-like map of upstream routers. The procedure starts with the path and distance information that the packet carries in the appropriate fields. For each upstream router we check if the distance is 0 and if the path information is identical to the router signature of the upstream router. If the previous probe is successful, we have identified one possible source. If the router in question is not a possible source then we calculate the new path field as the result of the <XOR> between the current path field and the router signature of the current router. We also decrease the distance field by one. Then we initiate the same traceback procedure starting from the current router and using the new path and distance fields.

The above is a recursive procedure that ends when all the nodes of the upstream map are traversed and returns a list of all the possible sources of the packet. The traceback procedure is computationally intensive and has to be initiated once for each distinct DDoS attack source. Ideally the traceback will produce one result for each DDoS attack source, but there exists the possibility of false positives. This means that there can be more than one source for a single marking. False positives will be discussed in detail in section 6.2.2.

Due to the computational intensiveness of the traceback procedure, it can only be initiated after the end of the DDoS attack. Thus the markings of the packets that have been characterised as part of the DDoS attack must be logged. Further discussion about the memory requirements of this marking scheme can be found in section 6.1. The traceback procedure in pseudo code can be found below:

```
function trace(path, distance, router)
  for each upstream_router of router
    if (distance == 0 <AND> path == router_signature(upstream_router)
      return upstream_router;
    if (distance == 0 <AND> path != router_signature(upstream_router)
      new_distance = 32;
    new_path = path <XOR> router_signature(upstream_router);
    new_distance --;
    call trace(new_path, new_distance, upstream_router);
```

As mentioned above, the victim needs a map of upstream routers in order to initiate the traceback procedure. That kind of map can be obtained by using standard traceroute tools such as Skitter [14] that can map thousands of nodes every day in a non-intrusive manner. Furthermore, such a map can be obtained after the DDoS attack since the amount of data needed to be maintained for traceback is limited as discussed in the following section.

6 Analysis

6.1 Overheads

In this section we will discuss the computational overhead (at router level), the memory requirements and bandwidth overhead that this marking scheme introduces.

The marking procedure dictates that a router has to read the path field of the packet and <XOR> it with its router signature. The router signature is a constant value based on the IP address of the router and needs no recalculation. Furthermore the router has to increment the distance field in a similar way that it decrements the

TTL field of the IP header. This marking procedure is very simple, comparing to existing marking schemes [1] [6] [7] [10], and involves only one write and one increment operation per packet. Furthermore there are no requirements that involve the router's memory.

During the DDoS attack, the victim has to keep track of all the packet markings that have been classified as part of the attack and store them for traceback purposes. The amount of information that has to be stored is one marking for each attacking source. Furthermore the victim has to store a map of all upstream routers. Such a map has doubles of 32 bit IP addresses that constitute the edges of the graph. Compressed forms of similar maps exist [14] and do not exceed the amount of 10Mbyte.

We would also like to note that since this marking scheme produces no additional control traffic during the marking procedure and no network traffic is required during the filtering or traceback procedure; no bandwidth overhead is introduced by this marking scheme.

6.2 Faults

In every detection system exists the probability to classify a host that is not part of a DDoS attack as an attacking host (false positive) or an attacking host as legitimate (false negative). In this section we will discuss the false positive and false negative probability of this marking scheme.

6.2.1 False negatives

The presented marking scheme is supposed to be deployed along with a DDoS detection mechanism. The detection mechanism carries the burden of identifying the sources of the DDoS attack as long as the information about the source of each packet provided by the marking scheme is correct. By design the packets that come from a host that participates on the attack will bear the same marking. In the rare exception that the routing changes during the attack, either this host will disappear, from the victim's point of view, and be replaced by another host (the marking of the host will change) or other hosts will appear as parts of the DDoS attack (some of the host's traffic will follow another route). Nevertheless the marking scheme produces no false negatives and the false negative probability of the combination of the two systems depends only on the detection system.

6.2.2 False positives

The false positive probability depends on the number of the attackers, the total number of edge routers and the length of the marking field. Further discussion on the length of the marking field can be found in section 7. In our analysis we will concentrate on the marking field as a whole and we will not take into account the distance factor because it makes no difference to the result.

Let R be the number of edge routers and A the number of attacking hosts. Let n be the length of the marking field. The number of edge router IP addresses that match a specific marking is $\mu=R/2^n$. The number of distinct markings that the attacking packets will bear is the number of the resulting faces of a 2^n faced die after A throws. The later is a special case of the occupancy problem [3]. Thus the expected number of distinct markings of the attack packets is:

$$M = 2^n - 2^n(1-1/2^n)^A$$

The number of false positives F can be calculated as follows:

$$F = M * \mu - A$$

We would like to note that as the number of attackers increases very much, the collisions between sources that belong to the attack also increase. Thus the percentage of false positives lowers for large numbers of attacking hosts. In figure 3 we can see the aforementioned effect.

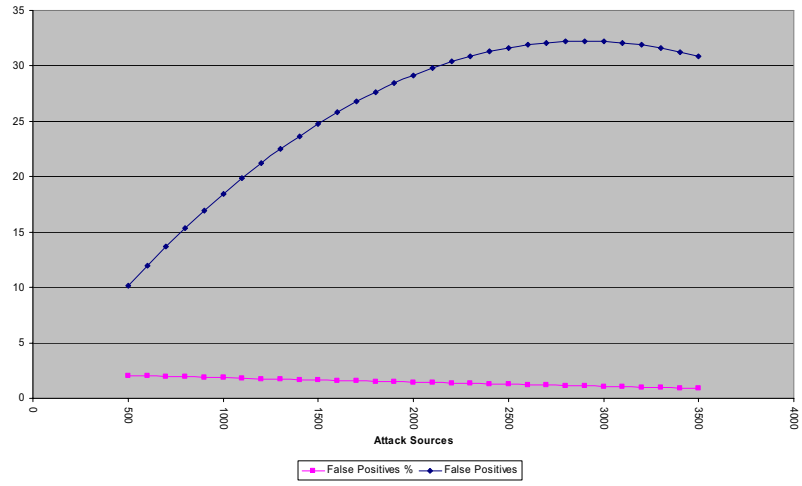

Figure 3: Chart showing the false positive probability against the number of attacking hosts for R=133000

6.3 Comparisons

6.3.1 Probabilistic Packet Marking

The basic advantage of this marking scheme against the various probabilistic marking schemes such as PPM [7] is that with this marking scheme the victim can successfully filter the incoming packets that belong to the DDoS attack in real time. The number of packets the victim needs in order to identify the source is only one. In PPM the number of packets needed is 500 – 4000 depending on the path length [7].

One more advantage of this marking scheme over PPM and others is that the attacker in this marking scheme cannot inject false markings into the network. This is called mark spoofing and is one of PPM's disadvantages. We would like to note that the effect of false markings injected by compromised routers is existent in both schemes and has been solved in [10].

Furthermore, as shown in [10] the PPM scheme cannot withstand distributed DoS attacks. The mere existence of 25 attacking hosts can result in thousands of false positives. On the other hand, this marking scheme scales perfectly as the number of attacking hosts rises and we even detected a small decrease of the false positive percentage.

6.3.2 Deterministic Packet Marking

As with probabilistic packet marking, the basic advantage of this marking scheme over DPM [1] is the number of packets needed for source identification. DPM needs about 7 – 10 packets in order to identify the source. On the other hand, DERM [6] shows some filtering capabilities but with very high false positive probability compared to this marking scheme since it uses 16 bit for the marking. The "multiple hash DERM" which can perform efficient traceback has an even larger false positive probability by a factor of 16.

7 Further work

In the previous sections we showed that the proposed marking scheme is capable of tracing back attack sources up to their nearest router as well as filtering the traffic that those sources produce. Further work should be done on limiting the inherit problem of false positives. This can be done in two ways.

First, the router signature that the scheme uses can be improved in order to ensure fewer collisions. This means either to identify those bits from the router IP address that have the largest entropy or construct a hash function that minimises possible collisions. Second, the fact that 17 bits are used in order to identify the source of a packet and according to the analysis in section 6.2, the false positive probability can be severely increased as the

number of edge routers rises. Thus, research must be done on the possibility of exploiting more bits of the IP header for traceback reasons in order to minimise false positives.

8 Conclusions

Identifying the true source of incoming packets is the key problem that has to be solved in order to effectively prevent DDoS attacks. In this paper we have presented a packet marking scheme that enables not only effective traceback of incoming packets but also real time filtering of packets that are part of a DDoS attack. We have showed that this scheme introduces no bandwidth overhead and low processing overhead to the Internet infrastructure. We also have presented the advantages of this scheme over current marking schemes. This marking scheme cannot be considered as a perfect filtering and traceback solution mainly because of its, low but not insignificant, false positive probability. We believe that this scheme is a step towards an effective prevention of DDoS attacks.

References

[1] Belenky, A., and Ansari, N.: IP Traceback With Deterministic Packet Marking, in *IEEE Communications Letters*, pp. 162-164, Vol. 7, No. 4, 2002

[2] Burch, H., and Cheswick, B.: Tracing Anonymous Packets to their Approximate Source, in *Proc. Of USENIX LISA Conf.*, pp 319-327, 2000

[3] Feller, W.: An Introduction to Probability Theory and its Applications (2nd Edition), Vol 1, 1966

[4] Ioannidis, J., and Bellovin, S. M.: Implementing Pushback: Router-based Defense against DDoS Attacks, in *Proc. Of Symposium of NDSS*, 2001

[5] Lee, C. J. H., Thing, L. L. V., Xu, Y., and Ma, M.: ICMP Traceback with Cumulative Path, an Efficient Solution for IP Traceback, in *Proc. Of ICICS*, pp. 124-135, 2003

[6] Rayanchu, K. S., and Barua, G.: Tracing Attackers with Deterministic Edge Router Marking (DERM), in *Proc. Of ICDCIT*, pp. 400-409, 2004

[7] Savage, S., Wetherall, D., Karlin, A., and Anderson, T.: Network Support for IP Traceback, in *IEEE/ACM Transactions on Networking*, pp. 226-237, Vol. 7, No. 3, 2001

[8] Shannon, C., Moore, D., and Claffy, K.: Characteristics of Fragmented IP Traffic on Internet Links, in *Proc. Of SIGCOMM,* pp. 83-97, 2001

[9] Snoeren, C. A., Partrige, C., Sanchez, A. L., Jones, E. C., Tchakountio, F., Schwartz, B., Kent, T. S., and Strayer, T. W.: Single-Packet IP Traceback, in *IEEE/ACM Transactions on Networking*, Vol. 10, No. 6, 2002

[10] Song, X. D., and Perrig. A.: Advanced and Authenticated Marking Schemes for IP traceback, in *Proc. Of IEEE INFOCOMM*, 2001

[11] Stoica, I., and Zhang, H.: Providing Guaranteed Services without Per-flow Management, in *Proc. Of ACM SIGCOMM*, pp. 81-94, 1999

[12] Stone, R.: CenterTrack: An IP Overlay Network for Tracking DoS Floods, in *Proc. Of 9th USENIX Security Symposium*, 2000

[13] Theilmann, W., and Rothermel, K.: Dynamic Distance Maps of the Internet, in *Proc. Of IEEE INFOCOM*, pp. 275-284, Vol. 1, 2000

[14] CAIDA, http://www.caida.org

[15] Ferguson, P., and Senie, D.: Network Ingress Filtering: Defeating Denial-of-Service Attacks which Employ IP Source Address Spoofing, RFC 2827, 2000