Adjusted Probabilistic Packet Marking for IP Traceback

Tao Peng¹, Christopher Leckie¹, and Kotagiri Ramamohanarao²

¹ ARC Special Research Center for Ultra-Broadband Information Networks
Department of Electrical and Electronic Engineering

The University of Melbourne

Victoria 3010, Australia

{t.peng, c.leckie}@ee.mu.oz.au

<http://www.ee.mu.oz.au/cubin>

² Department of Computer Science and Software Engineering

The University of Melbourne

Victoria 3010, Australia

rao@cs.mu.oz.au

Abstract. Distributed denial-of-service attack is one of the greatest threats to the Internet today. One of the biggest difficulties in defending against this attack is that attackers always use incorrect, or “spoofed” IP source addresses to disguise their true origin. In this paper, we present a packet marking algorithm which allows the victim to traceback the approximate origin of spoofed IP packets. The difference between this proposal and previous proposals lies in two points. First, we develop three techniques to adjust the packet marking probability, which significantly reduces the number of packets needed by the victim to reconstruct the attack path. Second, we give a detailed analysis of the vulnerabilities of probabilistic packet marking, and describe a version of our adjusted probabilistic packet marking scheme whose performance is not affected by spoofed marking fields.

1 Introduction

Distributed denial-of-service (DDoS) attacks have become a major threat to the Internet [10]. At the same time, DDoS is extremely difficult to defend [6]. The reason lies in the fact that the attackers use incorrect (“spoofed”) IP addresses in the attacking packets and therefore disguise the real origin of the attacks. This has made it very difficult or impossible to traceback the source of attacking IP packets.

A number of recent studies have approached the problem of IP packet traceback by Probabilistic Packet Marking (PPM) [15] [17]. It is assumed that the attacking packets are much more frequent than the normal packets. The main idea is to let every router mark packets probabilistically and let the victim reconstruct the attack path from the marked packet. All of the probabilistic marking algorithms try to overload the marking information into the 16 bit identification field in the IP packet header, which is seldom used [5] [18]. A major issue with

existing probabilistic marking schemes is that they use a fixed marking probability, which means that there is a greatly reduced probability of getting packets from routers which are far away from the victim. Consequently the number of packets needed to reconstruct the attack path depends on the number of packets which are marked by the furthest router in the attack path. If we can increase the marking probability for the routers which are far away from the victim, then we need less packets to reconstruct the attack path.

A potential problem with packet marking is that the attacker can forge the marking field. The authenticity of the marking field is the biggest challenge for Probabilistic Packet Marking, which is discussed in [13]. Although Song and Perrig [17] have proposed a scheme for router authentication, it is still hard to implement and there are still some chances for the attacker to spoof the marking field. However, if we can let the routers mark all the packets when they first enter the network, then there is no way for the attacker to use the spoofed marking field to decoy the victim.

In this paper, we make two contributions to the technique of Probabilistic Packet Marking. First, we have developed three techniques for adjusting the probability used by routers to mark packets, in order to reduce the number of packets needed by the victim to reconstruct the attack path. Second, we give a detailed analysis of the vulnerabilities of PPM, and describe a version of our adjusted probabilistic packet marking scheme whose performance is not affected by spoofed marking fields. We demonstrate the benefits of our approach with an analytical model as well as providing an experimental evaluation using simulated packet traces.

The paper is organized as follows. We present a brief background to this problem and highlight the main challenges of IP marking in Section 2. Section 3 introduces our Adjusted Probabilistic Marking Algorithm and shows a theoretical analysis. Simulation results of all these three techniques are provided in Section 4. From the analysis and simulation results, we can see that our Adjusted Probabilistic Marking Algorithm is more efficient and secure than the previous marking schemes. We discuss some practical issues in Section 5 and the related work is given in Section 6. Finally we conclude in Section 7.

2 Background on Probabilistic Packet Marking (PPM)

Once an attack has been detected, an ideal response would be to block the attack traffic at its source. Unfortunately, there is no easy way to track IP traffic to its source. This is due to two features of the IP protocol. The first feature is the ease with which IP source addresses can be forged. The second feature is the stateless nature of IP routing, where routers normally know only the next hop for forwarding a packet, rather than the complete end-to-end route taken by each packet. This design decision has given the Internet enormous efficiency and scalability, albeit at the cost of traceability. In order to address this limitation, Probabilistic Packet Marking (PPM) has been proposed to support IP traceability.

2.1 Definitions

The main idea of PPM is to let routers mark the packets with path information probabilistically and let the victim reconstruct the attack path using the marked packets.

Denial-of-service attacks are only effective so long as they occupy the resources of the victim. As a result, most denial-of-service attacks are comprised of thousands or millions of packets. PPM is based on the assumption that when we mark each packet with only a small probability then the victim will receive sufficient packets to reconstruct the attack path.

The network can be viewed as a directed graph $G = (V, E)$ where V is the set of nodes and E is the set of edges. V can be further partitioned into end systems (leaf nodes) and routers (internal nodes). The edges denote physical links between elements in V . Let $S \subset V$ denote the set of attackers and let $t \in V/S$ denote the victim. We will first consider the case when $|S| = 1$ (single-source attack) and treat the distributed DoS attack case separately. We assume that routes are fixed, and that the attack path $A = (s, v_1, v_2, \dots, v_d, t)$ is comprised of d routers (or hops) and has path length d [13].

Let N denote the number of packets sent from s to t . A packet x is assumed to have a marking field where the identity of a link $(v, v') \in E$ traversed can be inscribed. A packet travels on the attack path sequentially. At a hop $v_i \in \{v_1, \dots, v_d\}$, packet x is marked with the edge value (v_{i-1}, v_i) , $i = 1, \dots, d$, with probability p . As we seen in Fig.1, packet 1 is marked with edge value (v_1, v_2) and distance 2; packet 2 is marked with edge value (v_2, v_3) and distance 1. When t receives two packets it can reconstruct the attack path (v_1, v_2, v_3) .

Each router marks a packet with probability p . When the router decides to mark a packet, it writes its own IP address into the edge field and zero into the distance field. Otherwise, if the distance field is already zero, which means this packet has been marked by the previous router, it processes the packet as follows: (1) It combines its IP address and the existing value in the edge field and writes the combined value into the edge field. (2) It increases the distance value by 1. Thus, the edge value contains both information from the previous router and the current router. Finally if the router does not mark the packet, then it always increments the distance field. This distance field indicates the number of hops between the victim and the router that has marked the packet. The distance field should be updated using saturating addition, meaning the distance field is not allowed to wrap. When using this scheme, any packet written by the attacker will have a distance field greater than or equal to the real attack path. In contrast, a packet which is marked by the router should have a distance field which is less than the length of the path traversed from that router.

Savage et al. propose a method called Fragment Marking Scheme (FMS) [15] to compress the IP addresses and reconstruct the attack path. It is later improved by Song and Perrig[17]. Unless otherwise stated, when we talk about PPM in the rest of this paper, we are referring to Song and Perrig's version of PPM.

Fig. 1. Probabilistic Packet Marking

2.2 Limitation of Previous PPM Schemes

Our aim is to minimize the time required to reconstruct the attack path. This depends on the time it takes to receive packets that have been marked by each router on the attack path. This in turn depends on the choice of the marking probability p . In this section, we model the performance of PPM in terms of p , and highlight the limitation of using a fixed marking probability.

Definition 1. Let α_i denote the probability that packet arriving at the victim is lastly marked at node v_i but nowhere after v_i . For a uniform marking probability, $\alpha_i = \Pr\{x_d = (v_{i-1}, v_i)\} = p(1-p)^{d-i}$ ($i = 1, 2, \dots, d$).

Definition 2. Let α_0 denote the probability that a packet sent from the attacker reaches the victim without being marked at any of the routers. For a uniform marking probability, $\alpha_0 = (1-p)^d$. In order to reconstruct the attack path as quickly as possible, the victim needs to receive a sample of packets marked by each router in the path. An unmarked packet provides no information to the victim. In fact, there is a risk that unmarked packets may contain misleading information that has been spoofed by the attacker. Consequently, we want as many packets to be marked as possible. This implies that p should be large, so that α_0 is as small as possible. However, there is a penalty for making p too large. As p increases, there is a greater likelihood that packets marked by routers close to the source will be overwritten by routers close to the victim. Note that $\alpha_d \geq \dots \geq \alpha_2 \geq \alpha_1$, so α_1 is the smallest value. This is worst for packets marked by the first router after the source. So we need to choose p such that α_0 is minimized and α_1 is maximized.

According to the *coupon collecting problem* [8], for each attack path with d routers (excluding the victim), and with marking probability p , the expected number of packets needed to reconstruct the attack path is $N(d) = \frac{\ln(d)+O(1)}{p(1-p)^{d-1}}$ [15].

We can show that $N(d)$ is minimised when $p = \frac{1}{d}$. Consequently, the number of packets needed to get one sample from each router is $N_f(d) \simeq \frac{\ln(d)+O(1)}{\frac{1}{d}(1-\frac{1}{d})^{d-1}}$. This is the best result we can achieve for marking algorithms with a fixed probability. Our proposal is to reduce the total number of packets required $N(d)$ by using a higher marking probability for routers close to the source. Ideally, we want to receive an equal number of packets marked by each router on the attack path, i.e. $\alpha_i = 1/d$. In this case, the number of packets needed for reconstruction is $N_a(d) = d \ln(d)$. The savings of this approach are $\frac{N_f(d)}{N_a(d)} = (1 - \frac{1}{d})^{1-d}$, which is greater than 2 for $d \geq 2$. Our aim has been to develop a technique for adjusting the marking probability so that we can achieve the performance of $N_a(d)$.

3 Adjusted Probabilistic Packet Marking Schemes

According to the analysis in section 2, we propose that a router should adjust its packet marking probability based on its position in the attack path. However, the position of the router in the attack path is not known, since the position of the attacker is unknown. We need to estimate this distance based on the available information. In this section, we propose 3 different schemes for adjusting the marking probability based on the different distance measures d_1 , d_2 and d_3 . The definition of these distances is shown in Fig. 2

Fig. 2. Definitions of different distance measures

3.1 Number of Hops Traversed by Packet d_1

Let every router mark the packet with probability $p_1(d_1) = 1/d_1$. The ideal case for packet marking is to receive packets marked by each router with equal probability $\alpha_i = 1/d$, if the path length is d . Let $p_1(d_1)$ represent the marking probability of the router at distance d_1 from the source, where $d_1 = 1, 2, \dots, d$. Then we obtain the following equations:

$$\alpha_d = p_1(d) = 1/d \quad (1)$$

$$\alpha_{d-1} = p_1(d-1)[1 - p_1(d)] = 1/d \quad (2)$$

$$\alpha_{d-2} = p_1(d-2)[1 - p_1(d-1)][1 - p_1(d)] = 1/d \quad (3)$$

From equation 1 we can get $p_1(d) = 1/d$; from equation 2 we can get $p_1(d-1) = 1/(d-1)$; from equation 3 we can get $p_1(d-2) = 1/(d-2)$. Accordingly, we can summarize the marking probability formula as $p_1(d_1) = 1/d_1$. Then for the router at distance d_1 , $\alpha_{d_1} = \frac{1}{d_1} \times (1 - \frac{1}{d_1+1}) \times (1 - \frac{1}{d_1+2}) \times \dots \times (1 - \frac{1}{d})$. This equation can be simplified as $\alpha_{d_1} = \frac{1}{d_1} \times \frac{d_1}{d_1+1} \times \frac{d_1+1}{d_1+2} \times \dots \times \frac{d-1}{d} = \frac{1}{d}$. This means if each router marks the packet with the probability $p_1(d_1) = 1/d_1$, we can receive the packets marked by each router with equal probability $1/d$, given the path length is d .

In order to implement this marking scheme, we need to know the distance measure d_1 . We propose to add an extra field in the IP option field. This field can be used to record the number of hops (d_1) traversed by the packet. The default value for this field is 0, and the router increases this value by 1 every

time it forwards the packet. Every time the router gets the packet, it extracts the information d_1 from the option field and marks the packet with probability $1/d_1$. In order to prevent the attacker from spoofing this field, we can use the encryption schemes which are discussed in [17].

3.2 Number of Hops Traversed Since the Packet Was Last Marked (d_2)

In the original Probabilistic Packet Marking (PPM) scheme [15], there are three parts in the marking field. One part is called the distance field (d_2), which is used to hold the distance information from last router to mark the packet to the current router. We denote $d_2 = 0$ for routers next to each other. Let each router mark the packet according to the formula: $\frac{1}{2(d_2+1)}$. Since the larger the d_2 value, the higher the likelihood that it will be overwritten. Thus, we believe we should use a low marking probability for the packets with high d_2 value. Let us now illustrate the derivation of this formula by considering an example when the attack path length is 3.

The router marks the packet which has a distance value d_2 in the marking field with a probability $p_2(d_2)$. We assume the routers mark each packet when it first enters the network. So when the packet passes the first router, the d_2 value will be set to 0. By analyzing all the possibilities of the d_2 value when the packets traverse the attack path, we can derive expression for $\alpha_i, i = 1, 2, \dots, d$. Using these equations, we can find optimal marking probabilities for $\alpha_1, \alpha_2, \alpha_3$. However, the equations become more complicated as the path length increases, we consequently propose that the general marking probability should be $p_2(d_2) = \frac{1}{2(d_2+1)}$, which has been shown through experiments to have the best performance.

Since there are 5 bits in the marking field to hold the information in the existing probabilistic marking scheme [15] [17], we only need to extract this information from the marking field and mark the packet according to the formula $p_2(d_2) = \frac{1}{2(d_2+1)}$.

3.3 Number of Hops from Current Router to Destination (d_3)

If we can get the distance of the current router to the destination (d_3), we can mark each packet with a probability $p_3(d_3) = 1/(c+1-d_3)$ where c is a constant, and then we can receive packets marked by each router with a probability of $1/c$.

According to the marking scheme, we can have $\alpha_{d_3} = \frac{1}{c+1-d_3}(1 - \frac{1}{c-d_3+2}) \dots (1 - \frac{1}{c-1})(1 - \frac{1}{c}) = \frac{1}{c+1-d_3} \times \frac{c-d_3+1}{c-d_3+2} \dots \frac{c-2}{c-1} \times \frac{c-1}{c} = \frac{1}{c}$. In order to make this scheme work, we have to make sure $c+1-d_3 > 0$. Since most path lengths in the Internet are bounded by 30 [4] [1] [19], we can take $c = 30$ for safety. So if we mark with probability $p_3(d_3) = 1/(31-d_3)$, we can make sure we can receive the packets marked by each router with probability $1/30$.

We rely on the routing protocol to provide us with the distance measure d_3 . Current Internet routing protocols are destination-based and every time the router forwards the packet, it will look at the routing table to find the next

hop to the destination. Internet protocols provide us with a measure of the number of hops to each destination, which can be stored in the routing table as a measure of distance d_3 . When the router starts to route the packet, it can extract the distance information d_3 from the routing table and then mark the packet according to the formula $p_3(d_3) = 1/(31 - d_3)$.

3.4 Summary

We can summarize each marking scheme in term of its performance and practicality.

Marking scheme 1: $p_1(d_1) = 1/d_1$ can achieve the ideal marking performance. With this marking scheme, we can receive the packets marked by each router with equal probability for path length. Furthermore, every packet is marked under this scheme, and the attacker has no chance to spoof the marking field. However, this scheme requires a special hop count field and there is a risk that this field can be spoofed by the attacker. In order to make this scheme work, we need a strong authentication scheme which can stop the attacker from spoofing, e.g. [17].

Marking scheme 2: $p_2(d_2) = \frac{1}{2(d_2+1)}$ uses the distance field that is part of the packet marking scheme. This scheme can achieve a performance which is close to the optimal performance. In order to make this scheme work, we need to make sure the distance value in the marking field is trustable. One possibility is to let the routers mark all the packets when they first enter the network, then the attackers have no way to spoof the distance value. However, this is only practical if we control the ingress routers to our network, and thus is effectively the same as a technique called ingress filtering [9].

Marking scheme 3: uses information from the routing protocol and can achieve better results than using the uniform marking probability. Since the information is from the routing protocol, it can not be manipulated by an attacker. So scheme 3 is the safest and most practical scheme.

4 Evaluation

Our aim is to compare the performance of each scheme to PPM. Our comparison is based on the number of packets needed to reconstruct an attack path for a range of simulated attacks.

4.1 Methodology

We simulate attacks from different distances using the methodology in [17]. The network topology is based on a real traceroute dataset obtained from Lucent Bell Labs [11]. In our simulation, we vary the attack path from 1 to 30 hops and conduct 1000 random trials at each path length value. We measured the

number of packets required to reconstruct the attack path using our schemes, and compared this to the number of packets required by PPM [17], where our implementation of PPM used a threshold of $M=5$ as defined in [17]. We varied the uniform marking probability of PPM using the values $p = 0.01, 0.04$, and 0.1 . Note that $p = 0.04$ is recommended as the optimum choice for PPM [15].

4.2 Results

The performance of schemes 1 to 3 are shown in Fig. 3.

Schemes 1 and 2 perform the best, outperforming PPM for all values of p tested. However, these results assume that the distance field has not been tampered with. *Scheme 3* is the most practical, since its distance measure cannot be tampered with by the attacker.

Scheme 3 outperformed PPM with $p = 0.01$ and 0.04 . Although PPM with $p = 0.1$ outperforms *scheme 3* for small hop counts, *scheme 3* performs far better when the attack path is large.

Scheme 3 outperforms *scheme 2* when path length is 20 or higher as shown in Fig. 3. This is because as the path length increases, *scheme 3* approaches optimum performance while *scheme 2* cannot achieve the optimum performance as we discussed in Section 3.2. Furthermore, *scheme 1* and *scheme 3* converge when the path length equals 30 because c equals the path length, which makes $p_3(d_3)$ equivalent to $p_1(d_1)$.

Fig. 3. Scheme 1,2,3 compared with uniform marking probability

5 Discussion

5.1 Distributed Denial-of-Service Attacks

During a distributed denial-of-service attack, there are many attacking sources. We have found that the number of packets needed for reconstruction increases linearly with the number of attackers. So it will become very hard to verify all the attacking sources during a DDoS attack. Thus, our method to reduce the number of packets needed for reconstruction becomes extremely important to improve the reconstruction efficiency.

5.2 Spoofing the Marking Field

By spoofing the marking field, it is possible for attackers to make the attack appear as though it has come from a more distant source, e.g. a false source s_f as shown in Fig. 4. However, the attacker cannot change the marking of routers between it and the victim, e.g., v_1 to v_3 . Consequently, we can always reconstruct the path to the attacker, although we may also reconstruct a false sub-path at the start of the true path, e.g., v_{f_1} to v_{f_3} .

Fig. 4. Effect of Spoofing the Marking Field (Fake sub-path: v_{f_1} to v_{f_3} , true path: v_1 to v_3)

If we are unable to authenticate the marking field, then this false sub-path can affect the performance of our first two schemes. This is because distance measures d_1 and d_2 will be inflated by the false sub-path, thus decreasing the packet marking probability of routers in the true attack path.

However, our third scheme is unaffected by the actions of the attacker. This is because d_3 is derived from information in the routing table of each router, and the destination field. The attacker cannot fake the destination field without defeating the purpose of the attack, and the attacker cannot manipulate the contents of the routing tables in the routers. Thus, the performance of our third scheme is secure against manipulation by the attacker.

6 Related Work

Burch and Cheswick [3] propose a link-testing traceback technique. It infers the attack path by flooding the links with large bursts of traffic and observing how this perturbs the attack traffic. This scheme requires considerable knowledge of network topology and the ability to generate huge traffic in any network

links. Mahajan et al. [12] provide a scheme in which routers learn a congestion signature to tell good traffic from bad traffic. The router then filters the bad traffic according to this signature. Furthermore, a pushback scheme is given to let the router ask its adjacent routers to filter the bad traffic at an earlier stage. This scheme is effective for some types of DDoS attacks but it needs a narrow and accurate congestion signature to make sure the bad traffic is filtered while the good traffic is not affected.

Bellovin [2] proposed an ICMP "traceback" scheme to let router generate ICMP packets to the destination containing the address of the router with a low probability. For a significant traffic flow, the destination can gradually reconstruct the route that was taken by the packets in the flow. ICMP packets are often treated with a low priority by routers to reduce the additional traffic, which undermines the effectiveness of the scheme. This scheme is later extended by Wu et al. [20]. An alternative approach is to mark the packets themselves. Savage et al. [15] describe a scheme for routers to probabilistically mark packets. They propose using the identification field of the IP header, which is normally used to control fragmentation. They point out that IP fragmentation is seldom used in practice. While their approach overcomes many of the limitations of the ICMP traceback proposal, there are some security problems when the attackers fake the marking field. Song et al. [17] propose an enhanced scheme of probabilistic packet marking and also set up a scheme for router authentication. However, the authentication scheme is complex to implement. Dean et al. [7] propose an alternative marking scheme using noisy polynomial reconstruction. This scheme is backwards compatible, and incrementally deployable compared with the former proposals. Unfortunately their scheme is very vulnerable to fake markings put in the packets by the attackers. Furthermore, the number of packets needed to reconstruct the attack path is quadratic to the number of attackers. Snoeren et al. [16] propose a scheme to let routers store a record of every packet passing through the router, so that the router can then trace back the origin of the packet by using the history in the router. Although they describe a smart scheme to compress the storage, it is still a huge overhead for the router to implement this scheme, especially with the increasing network speed. Park and Lee [14] propose to put distributed filters in the routers and filter the packets according to the network topology. This scheme can stop the spoofed traffic at an early stage. However, in order to place the filters effectively, it needs to know the topology of the Internet and routing policy between Autonomous Systems, which is hard to achieve in the expanding Internet.

In summary, every marking scheme uses a fixed marking probability which will result in a small number of packets marked by the more distant routers when all the packets arrive at the victim. In contrast, we have developed several schemes that solve this problem by adjusting the marking probability in each router, which significantly reduces the number of packets required to reconstruct the attack path. Furthermore, no one has set up a scheme to completely solve the security problem that the attacker can fake the marking field. However, our third marking scheme does not use the contents of the marking field to adjust the marking probability, and thus cannot be manipulated by the attacker while at the same time requiring fewer packets to trace the packet.

7 Conclusion

In this paper, we make the following two contributions to Probabilistic Packet Marking (PPM). First, we developed three techniques to adjust the marking probability used by each router so that the victim receives packets marked by each router with equal probability. *Scheme 1* is to let the IP packet carry a message to inform the router how far the packet has traveled. *Scheme 2* is to use the distance value of the marking field in the IP packet. *Scheme 3* is to get the distance between the router and destination from the routing table. Both *scheme 1* and *scheme 2* need authentication to prevent the attacker from spoofing the required information. *Scheme 3* is the most practical one and can improve the reconstruction efficiency compared with the optimal uniform marking probability ($p = 0.04$). By implementing this scheme, we can substantially reduce the number of packets needed to reconstruct the attack path in comparison to PPM. Our second contribution is that we give a detailed analysis of the vulnerability of PPM, and describe a version of our adjusted probabilistic packet marking scheme whose performance is not affected by the vulnerability caused by spoofed marking fields.

Acknowledgment. We would like to thank the AT&T Internet Mapping Project for making available their traceroute data and the anonymous reviewers for their helpful comments.

References

1. Skitter Analysis. Cooperative association for internet data analysis, 2000. <http://www.caida.org/Tools/Skitter/Summary/>.
2. S. Bellovin. *The icmp traceback message*. Internet Draft,IETF, March 2000. draft-bellovin-itrace-05.txt (work in progress).<http://www.research.att.com/~smb>.
3. Hal Burch and Bill Cheswick. Tracing anonymous packets to their approximate source. In *Proceedings of the 14th Systems Administration Conference*, New Orleans, Louisiana, U.S.A., December 2000.
4. R.L. Carter and M.E. Crovella. Dynamic server selection using dynamic path characterization in wide-area networks. In *Proceedings of the 1997 IEEE INFOCOM Conference*, Kobe,Japan, April 1997.
5. K. Claffy and S. McCreary. Sampled measurements from june 1999 to december 1999 at the ames inter-exchange point. *Personal Communication*, January 2000.
6. Computer emergency response team. *cert advisory ca-2000-01: Denial-of-service developments*, 2000. <http://www.cert.org/advisories/CA-2000-01.html>.
7. Drew Dean, Matt Franklin, and Adam Stubblefield. An algebraic approach to ip traceback. In *Network and Distributed System Security Symposium, NDSS '01*, Feburary 2001.
8. W. Feller. *An Introduction to Probability Theory and Its Applications(2nd edition)*, volume 1. Wiley and Sons, 1966.
9. P. Ferguson and D. Senie. *Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing*. RFC2267,IETF, January 1998.
10. John D. Howard. *An Analysis of Security Incidents on the Internet*. PhD thesis, Carnegie Mellon University, 1998.

11. Lucent Lab. Internet mapping, 1999.
<http://cm.bell-labs.com/who/ches/map/dbs-/index.html>.
12. Ratul Mahajan, Steven M. Bellovin, Sally Floyd, John Ioannidis, Vern Paxson, and Scott Shenker. Controlling high bandwidth aggregates in the network. Technical report, AT&T Center for Internet Research at ICSI (ACIRI) and AT&T Labs Research, February 2001.
13. K. Park and H. Lee. On the effectiveness of probabilistic packet marking for ip traceback under denial of service attack. In *Proceedings of IEEE INFOCOM 2001*, 2001.
14. Kihong Park and Heejo Lee. On the effectiveness of router-based packet filtering for distributed dos attack prevention in power-law internets. In *Proceedings of the 2001 ACM SIGCOMM Conference*, San Diego, California, U.S.A., August 2001.
15. Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson. Practical network support for ip traceback. In *Proceedings of the 2000 ACM SIGCOMM Conference*, August 2000. <http://www.cs.washington.edu/homes/savage/traceback.html>.
16. Alex C. Snoeren, Craig Partridge, Luis A. Sanchez, Christine E. Jones, Fabrice Tchakountio, Stephen T. Kent, and W. Timothy Strayer. Hash-based ip traceback. In *Proceedings of the 2001 ACM SIGCOMM Conference*, San Diego, California, U.S.A., August 2001.
17. Dawn X. Song and Adrian Perrig. Advanced and authenticated marking schemes for ip traceback. In *Proceedings of IEEE INFOCOM 2001*, 2001. <http://paris.cs.berkeley.edu/perrig/projects/iptraceback/tr-iptrace.ps.gz>.
18. I. Stoica and H. Zhang. Providing guaranteed services without per flow management. In *Proceedings of the 1999 ACM SIGCOMM Conference*, Boston, MA, August 1999.
19. W. Theilmann and K. Rothermel. Dynamic distance maps of the internet. In *Proceedings of the 2000 IEEE INFOCOM Conference*, Tel Aviv, Israel, March 2000.
20. S. Felix Wu, Lixia Zhang, Dan Massey, and Allison Mankin. *Intension-Driven ICMP Trace-Back*. Interner Draft, IETF, February 2001. draft-wu-itrace-intension-00.txt(work in progress).