

# Provider-Based Deterministic Packet Marking against Distributed DoS Attacks

Vasilios A. Siris\* and Ilias Stavrakis

Institute of Computer Science  
Foundation for Research and Technology - Hellas (FORTH)  
P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece  
Tel.: +30 2810 391726, fax: +30 2810 391601  
Email: {vsiris,stavraki}@ics.forth.gr

## Abstract

One of the most serious security threats in the Internet are Distributed Denial of Service (*DDoS*) attacks, due to the significant service disruption they can create and the difficulty to prevent them. In this paper, we propose new deterministic packet marking models in order to characterize *DDoS* attack streams. Such common characterization can be used to make filtering near the victim more effective. In this direction we propose a rate control scheme that protects destination domains by limiting the amount of traffic during an attack, while leaving a large percentage of legitimate traffic unaffected. The above features enable providers to offer enhanced security protection against such attacks as a value-added service to their customers, hence offer positive incentives for them to deploy the proposed models. We evaluate the proposed marking models using a snapshot of the actual Internet topology, in terms of how well they differentiate attack traffic from legitimate traffic in cases of full and partial deployment.

**Keywords** Distributed Denial of Service (DDoS), defense models, filtering

## 1 Introduction

Distributed Denial of Service (*DDoS*) attacks are one of the most serious security threats in the Internet today. The main aim of *DDoS* attacks is the disruption of services by attempting to limit access to a machine or service instead of subverting the service itself [4]. *DDoS* attacks achieve their goal either by consuming network bandwidth in the path close to the victim by sending huge amounts of traffic (*bandwidth attacks*), or by consuming the victim's memory and computational resources by exploiting an inherent protocol or implementation vulnerability (*protocol attacks*). The usual procedure of a *DDoS* attack involves the perpetrators compromising vulnerable hosts at which they install slave programs, and in a coordinated manner instruct thousands of these slave programs to attack a particular destination host or destination network.

The Internet's security vulnerability is mainly due to its open resource access model design, emphasizing on functionality and simplicity, but not on security. Furthermore, Internet's current routing protocols and forwarding procedures are

---

\*The authors are also with the Dept. of Computer Science, Univ. of Crete. A preliminary version of this paper appeared in [14].

largely based on destination addresses, and no entity is responsible for ensuring that source addresses are truthfully set. The above architectural drawbacks in combination with a large number of Internet hosts that have poor or no security make the Internet susceptible to a wide range of attacks, which are made possible with attacking tools that have been developed for this purpose.

## 1.1 Desirable properties of a defense system

A powerful defense model should have several properties in order to be characterized as effective and secure. In particular, see also [21], a defense model

- should prevent only attack traffic from reaching the victim. This requires that the defense model differentiates the legitimate traffic from the malicious traffic, and limit the disruption caused to legitimate users.
- should not be itself a target for new attacks. Thus, it should avoid direct communication between different entities, avoid single points of failure, and be stateless, i.e. not keep per-flow information in intermediate routers.
- should be simple and easily deployable. Thus, it should not require major changes or additions to the existing infrastructure.
- should not create extra traffic, thus increasing the load during attack periods, and should involve procedures that are invoked only during attacks, avoiding overhead during periods with no attacks.
- should offer positive incentives to domains that need to implement the corresponding procedures. For example, a domain has little or no incentive to allow control of its resources by an external entity.
- should have a fast response time not only in the detection of the attack but also in the establishment of the appropriate actions to counteract the attack, and should be able to adapt to changes of the attack traffic pattern.

Achieving the above objectives simultaneously is difficult, and involves tradeoffs in the degree to which each is achieved.

## 1.2 Motivation and contributions

The dominant reason that *DDoS* attacks comprise a hard security problem is that *DDoS* attack streams may have no common characteristics that can be used for detection and filtering. Furthermore, most ISPs rely on manual detection of *DDoS* attacks and perform offline fine-grain traffic analysis to identify the attacking stream features, based on packet attributes such as traffic type, size, and source address. Based on this analysis, administrators can manually install filtering rules or access control lists. Such human intervention results in poor response time and lacks adaptability to changes of the attack traffic pattern. Finally, the expressiveness of existing rule-based filtering mechanisms is too limited and as the difference between legitimate and attack packets becomes increasingly subtle, the number of required filtering rules as well as the number of packet attributes included in each rule explodes, creating scalability problems for high-speed implementations of rule-based filtering [10].

In this paper, we propose and evaluate two provider-based packet marking models: *Source-End Provider Marking* and *Source and Destination-End Provider Marking*. Both models are based on deterministic packet marking, and aim to give the victim's provider stable and secure information about the path incoming traffic

streams follow. These markings can be used for detection of attacking or suspicious streams independently of the variability the attacker gives to those streams, and provide a common attribute to perform filtering. In this direction we propose a rate control scheme that protects destination domains by limiting the amount of traffic during an attack, while leaving a large percentage of legitimate traffic unaffected. Hence, providers can offer increased protection to their customers as a value-added service, improving the available throughput for legitimate users during such attacks. We evaluate the performance of the proposed models in terms of the achieved differentiation between legitimate and attack streams using Burch and Cheswick’s traceroute map of real Internet topology [2], and in terms of the properties identified in Section 1.1. The results show that the proposed models provide better performance than the *Pi* marking scheme [21], using order of magnitude fewer routers and giving providers deployment incentives to adopt such a defense system.

The rest of this paper is organized as follows. In Section 2 we present the two provider-based packet marking models. In Section 3 we evaluate the performance of the proposed models in terms of the achieved differentiation between legitimate and attack traffic streams. In Section 4 we discuss the deployment incentives for the proposed models. In Section 5 we present related work identifying the differences with our proposed approach, and finally in Section 6 we conclude the paper identifying ongoing and future research directions.

## 2 Provider-based deterministic packet marking

In this section we present the two provider-based deterministic packet marking models: *Source-End Provider Marking* and *Source and Destination-End Provider Marking*. In both models marking is performed only at the edge routers that connect a provider to its customers or to the Internet. Following the approach presented in [21], marks are placed in the 16-bit identification field used for IP packet fragmentation. This results in losing the information that is necessary for packet reassembly. Fortunately, recent measurements [16] indicate that the percentage of fragmented packets is very small (less than 0.25%), and most modern TCP implementations set the do not fragment (DFT) bit by default [19], as specified by the Path MTU Discovery standard in RFC 1191. Moreover, as suggested in [22], compatibility with IP fragmentation can be achieved by avoiding to mark packets that will get fragmented or are fragments themselves.

### 2.1 *Source-End Provider Marking*

In this model, marking is performed by a provider’s edge router that connects its customers to the Internet, as packets enter the provider’s network from the source domain, Figure 1. Marking is deterministic in that all packets are marked with the same value, which consists of the last two bytes of the MD5 hash of the IP address belonging to the interface that connects the router to the source domain; this is performed in order to achieve uniform distribution of mark values [21]. With such an approach, all packets originating from a particular source domain have the same mark. Of course, due to the limited number of mark values ( $= 2^{16}$ ), it is possible that different domains have the same mark value. Indeed, this is the reason for not achieving perfect differentiation between legitimate traffic and attack traffic, as we investigate in Section 3.




Figure 1: *Source-End Provider Marking* model

### 2.1.1 Packet filtering

On the destination side, the provider can implement detection and filtering on the edge router that is connected to the destination domain, Figure 1, based on the marks placed by the source-end provider. One simple scheme would be to drop all packets containing a mark that has been identified as belonging to attack traffic.

An alternative to packet dropping, which behaves less drastically to traffic identified, possibly wrongly, as attack traffic and avoids starvation of such traffic, is to perform rate-limiting in a manner that ensures that all traffic that is identified as non-attack traffic is not affected. Assume that  $l_i$  is the rate of packets with mark  $i$  before an attack, and  $I$  is the set of marks. Now consider that there is a *DDoS* attack, and let  $A$  be the set of marks identified to correspond to attack traffic. Also, let  $L$  be the set of marks corresponding to non-attack traffic; hence,  $I = A \cup L$ . If  $C$  is the total capacity connecting a provider's edge router to the victim, then the provider can allocate an amount of bandwidth  $C_{legit}$  to packets containing marks in the set  $L$ . To ensure that legitimate traffic is not affected,  $C_{legit}$  must be  $C_{legit} = \frac{\sum_{j \in L} l_j}{\sum_{i \in I} l_i} C$ . The last equation ensures that the average amount of capacity for legitimate traffic is the same before and after the attack. Packets with marks identified to belong to attack traffic will be allocated capacity  $C_{attack} = \frac{\sum_{j \in A} l_j}{\sum_{i \in I} l_i} C$ . The above rate control scheme can be implemented using weighted or class-based queueing, which is supported in current routers. If  $a_i$ , for  $i \in A$ , is the rate of attack traffic with mark  $i$  and  $\sum_{i \in A} (l_i + a_i) > C_{attack}$ , then limiting traffic with mark  $i \in A$ , to rate  $C_{attack}$  will result in dropping packets identified as attack traffic with percentage  $1 - \frac{\sum_{j \in A} l_j}{\sum_{i \in A} (l_i + a_i)} \frac{C}{\sum_{i \in I} l_i}$ .

Rather than handle all packets containing a mark identified to belong to attack traffic in the same way, we can set different rate-limits  $C_j$  for each mark  $j \in A$  given by  $C_j = \frac{l_j}{\sum_{i \in I} l_i} C$  for  $j \in A$ . This rate-limiting scheme results in dropping a percentage of packets with mark  $j \in A$  equal to  $1 - \frac{l_j}{l_j + a_j} \frac{C}{\sum_{i \in I} l_i}$ . Hence, the percentage of packets with mark  $j$  that are dropped is an increasing function of the amount of attack traffic  $a_j$ , i.e. the intensity of the attack. One can show that this multiple rate-limiting approach allows a larger percentage of legitimate packets that contain a mark corresponding to attack traffic, to enter the destination domain, compared to the approach where there is a single rate-limiter for all packets containing a mark corresponding to attack traffic; this is achieved at the cost of implementing a larger number of rate-limiters.

### 2.1.2 Advantages and limitation

The degree of differentiation between legitimate and attack traffic achieved using this model will be investigated in Section 3. Here we discuss the basic features of

the model in terms of the properties identified in Section 1.1, and its basic limitation that leads us to the design of the *Source and Destination-End Provider Marking* model.

The *Source-End Provider Marking* model gives an indirect way of communicating information about the originating source domain of packets, without explicit communication between providers. The model is completely decentralized and stateless, avoiding central coordinators and single points of failure. It has fast response time, because filtering can be applied at the same point where attack detection is performed. It does not demand changes of the existing infrastructure. The marking procedure is quite straightforward and simple, and can be implemented with current router capabilities. Furthermore, the location of the detection and filtering module at the edge router connecting the provider with the destination domain enables the provider to protect a customer’s access link and servers. Hence, providers can offer enhanced protection against DoS attacks as a value-added service to their customers.

The main disadvantage of the *Source-End Provider Marking* model is its inability to handle false marking attacks in an environment of partial deployment. In particular, if a source-end provider does not implement marking, then an attacker that has compromised hosts in domains that are connected to that provider can instruct these hosts to mark packets using a value that corresponds to some other source domain. If the destination-end provider applies filtering actions after detecting an attack, then using the above method an attacker can harm the legitimate traffic that originates at the actual source domain. Note, nevertheless, that false marking does not influence other source domains that have a different mark value, nor does marking using random values.

## 2.2 *Source and Destination-End Provider Marking*

In the *Source and Destination-End Provider Marking* model the mark value of each packet is produced in two phases. The first phase is identical to the *Source-End Provider Marking* model, where marking is performed at the provider’s edge router to which the source domain is connected. The second phase involves marking at the edge routers that connect the destination-end provider to the Internet, Figure 2. These edge routers mark  $n$  (for  $n < 16$ ) of the 16 bits in the IP identification field, with a mark value that is different for different edge routers. The remaining  $16 - n$  bits maintain the value placed by the source-end provider. For example, if the destination-end provider is connected to the Internet through 4 edge routers, two bits are enough to differentiate these routers. Note that hashing is used only for producing the first phase mark at the source-end provider, and not for the mark at the destination-end provider. The reason for this is that a destination can communicate with a potentially large number of source domains, whereas the number of edge routers in a destination-end provider are much smaller, hence a few bits can be enough to identify them. Nevertheless, in Section 3.4 we consider the alternative of using hashing at the destination-end provider, when the number of bit is not enough to differentiate all its edge routers.

Note that the *Source and Destination-End Provider Marking* model maintains all the benefits of the *Source-End Provider Marking* model that were discussed in Section 2.1.2, and limits the impact of false marking in the case of partial deployment, as we discuss next.

### 2.2.1 Limiting impact of false marking in case of partial deployment

We now discuss how the above model can reduce the impact of false marking attacks, which is a limitation of the *Source-End Provider Marking* model. The basic intuition




Figure 2: *Source and Destination-End Provider Marking* model

is that under normal conditions we expect the legitimate traffic from each source domain to enter the last provider from one or a few edge routers. In case of a false marking attack however, the destination provider would receive the same mark from a larger number of edge routers. With the *Source and Destination-End Provider Marking* model, because different edge routers at the destination-end provider mark packets differently, only the attack traffic, containing a spoofed mark corresponding to the source domain under attack, that enters through the same edge routers as the legitimate traffic from that source domain will maintain the same mark value as the legitimate traffic.

To make the above more precise, consider a provider with  $E$  edge routers that faces a *DDoS* attack from  $N$  attacker hosts in an environment of partial deployment of the proposed marking models. Suppose that  $P$  is the percentage of providers that adopt the model, and  $U$  is the sending rate of each attacker. If all attackers mark packets with a value belonging to a legitimate domain, then the amount of traffic  $T$  that enters the target domain's provider with that mark value is  $T = N \cdot (1 - P) \cdot U$ . Note that  $T$  is the amount of excess (attack) traffic with the specific mark that would reach the detection and filtering module when the *Source-End Provider Marking* model is used.

If we assume that the attack packets are uniformly distributed over all edge routers, then the amount of traffic that enters each edge router is  $T' = \frac{N \cdot (1 - P) \cdot U}{E}$ . Note that  $T'$  is the amount of excess (attack) traffic with the specific mark value that would reach the detection and filtering module when the *Source and Destination-End Provider Marking* model is used, if we assume that the traffic from a source-end provider enters the destination-end provider from a single edge router. Hence, for an attacker to produce the same aggregate amount of attack traffic as in the case of the *Source-End Provider Marking* model, he would need to increase the number of attacker hosts or the rate of each attacker host by a factor  $E$ . From the above analysis, we understand that a large destination-end provider, which has a large number of edge routers  $E$ , can offer better protection using the *Source and Destination-End Provider Marking* model compared to a small destination-end provider.

### 3 Performance evaluation

In this section we evaluate, using a real snapshot of the Internet's topology, the *Source-End Provider Marking* and the *Source and Destination-End Provider Marking* models in terms of the achieved differentiation between legitimate and attack traffic during simulated *DDoS* attacks. Furthermore, we investigate how this differentiation is affected by the number of attackers, the number of bits required for marking at the destination-end provider, and the percentage of providers that

implement the approach. Since our focus is on the performance of the proposed marking models, we assume that the attack detectors have optimal performance, i.e. they have 100% detection probability and 0% false alarm probability. Finally, because we focus on evaluating the differentiation achieved by the proposed marking schemes, our performance metric considers the legitimate traffic that is not affected by the filtering scheme employed. In the case of complete dropping, where all packets containing a mark identified as belonging to attack traffic are dropped, our results refer to the legitimate traffic that reaches the victim, whereas in the case of rate-limiting, our results refer to the legitimate traffic that is not rate-limited.

### 3.1 Experiment scenario and metrics

The topology used in our experiments was Burch and Cheswick’s Internet Map [2], which was created using traceroute messages from a single host to destination hosts throughout the Internet, thus producing a tree with thousands of paths. We assume the victim of the *DDoS* attacks to be the root host of the tree and the legitimate and attack hosts to be specific leaves of the tree.

In our experiments, similar to [21], we choose 5000 leaves at random to act as legitimate users that send 10 packets each, and a variable number of leaves to act as attackers that send 100 packets each during an attack; these two sets are disjoint. As we discuss later, our comparison metric considers only the percentage of accepted traffic, hence does not depend on the absolute values of the packet rate or on the relative rate of legitimate and attack traffic. Finally, unless otherwise noted, we apply the first marking phase at the third hop away from the source. The results we present are the average of 5 runs of each experiment with the same parameters.

The performance metrics we consider, for comparison reasons, are identical to the ones used for evaluating the *Pi* marking scheme in [21]. The basic performance metric is the *acceptance ratio gap*, which is the difference between the *user acceptance ratio* and the *attacker acceptance ratio*. The *user acceptance ratio* is the ratio of user packets that are not affected by filtering to the total number of user packets, and the *attacker acceptance ratio* is the ratio of attack packets that are not affected by filtering to the total number of attack packets sent to the victim during the attack. Hence, the *acceptance ratio gap* gives the degree of differentiation between legitimate traffic and attack traffic. In a real environment with no protection the *acceptance ratio gap* would be zero, since we have no information to differentiate the legitimate traffic from attack traffic. On the other hand, in the case of perfect differentiation, the *acceptance ratio gap* would be 1.

### 3.2 Attack and legitimate traffic differentiation

The performance of *Source-End Provider Marking* is shown in Figure 3. In this experiment we consider 100% deployment, hence the *attacker acceptance ratio* is zero. Thus, the *acceptance ratio gap* coincides with the *user acceptance ratio*. From this graph we see that, e.g. in the case of 2000 attackers, the acceptance ratio gap, which is equal to the user acceptance ratio, is 70%; this means that 70% of the legitimate users will not be affected by filtering. The decrease of the *user acceptance ratio* when the number of attackers increases is due to the increase of the number of collisions of legitimate traffic marks with attack traffic marks.

Figure 4 shows the performance of the *Source and Destination-End Provider Marking* model for a different number of bits required by the destination-end provider. Note that more bits are required by a larger provider, since such a provider has a larger number of edge routers connecting it to the Internet. Providing more bits for marking at the destination-end provider gives rise to two opposite effects: First,

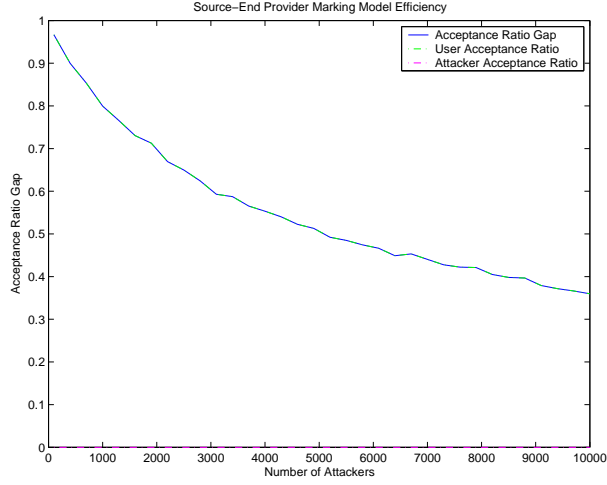


Figure 3: *Source-End Provider Marking* model performance




Figure 4: *Source and Destination-End Provider Marking* model performance

decreasing the number of bits for marking at the source-end provider tends to decrease the differentiation achieved by the source-end marking side, as shown in Figure 5, whereas increasing the number of bits for marking at the destination-end provider tends to increase the differentiation achieved by the destination-end marking side. Which of the two effects is dominant, hence the net increase or decrease of the achieved differentiation depends on the number of bits, as shown in Figure 4. In particular, this figure shows that giving 2 or 5 bits for marking at the destination-end provider, which leaves 14 or 11 bits for marking at the source-end provider, results in an overall increase of the performance compared to when all 16 bits are used for marking at the source-end provider. The opposite is true when 4 or 7 bits are used for marking at the source-end provider. We anticipate that the above tradeoff depends on the topology and the length (number of hops) of the path between the source and the destination.

Note that increasing the number of bits used for marking at the destination-end provider offers protection against false marking attacks in the case of partial

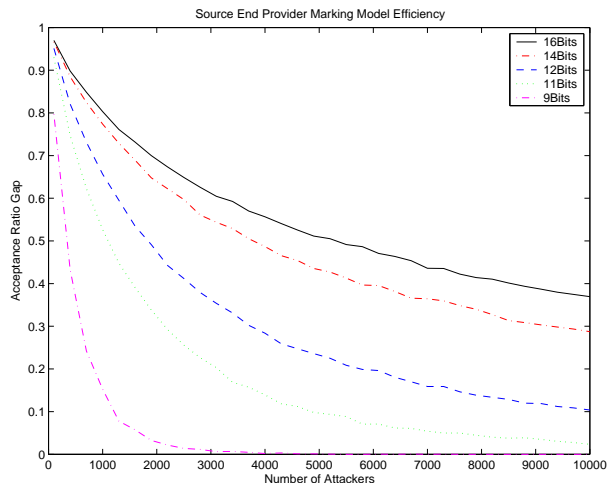


Figure 5: *Source-End Provider Marking* performance with different marking field size

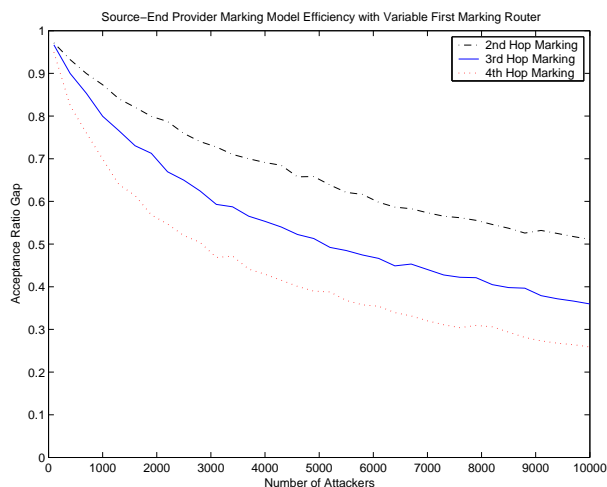



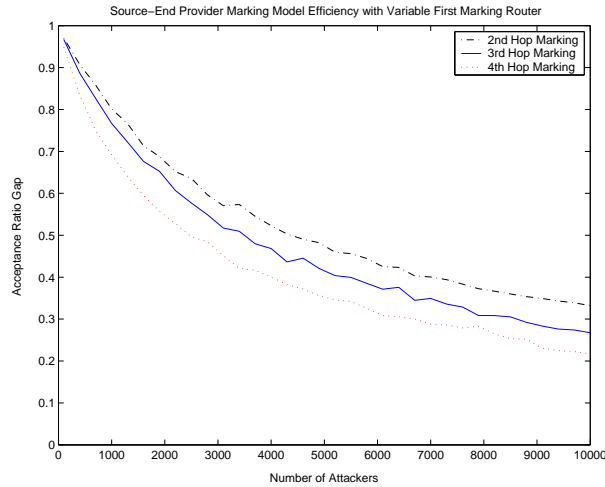
Figure 6: *Source-End Provider Marking* with variable first marking router

deployment, as discussed in Section 2.2.1.

Figures 6 and 7 show the performance of the *Source-End Provider Marking* and *Source and Destination-End Provider Marking* models, respectively, for different first marking routers. Different first marking routers effectively correspond to different sizes of the source domain, since we assume that the source-end provider marks packets at the edge router that connects it to the source domain. The results show that the acceptance ratio is higher when marking is performed closer to the source. Also observe that the *Source and Destination-End Provider Marking* model is less affected by the first marking router, compared to the *Source-End Provider Marking* model.



(a) 2 bits needed for last provider

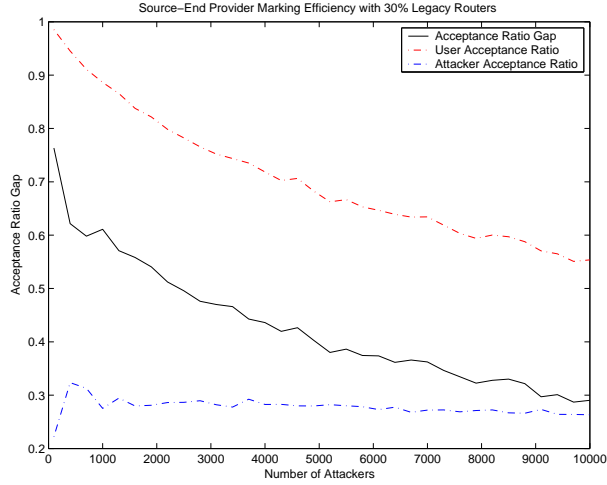


(b) 7 bits needed for last provider

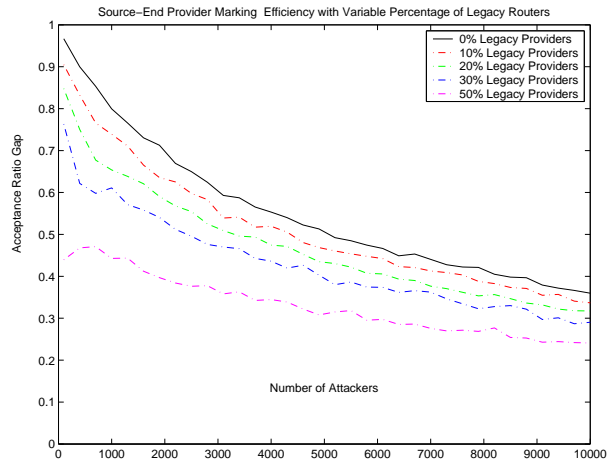
Figure 7: *Source and Destination-End Provider Marking* with variable first marking router

### 3.3 Partial deployment

Next we investigate the performance of the two models in an environment of partial deployment. We assume that some percentage of providers do not implement our marking model, hence their edge routers are legacy routers. In our experiments legacy routers are chosen randomly from the set of leaves representing legitimate users and attackers. Figure 8(a) shows that the *attacker acceptance ratio* is no longer zero, since due to partial deployment not all attack packets will be marked, and those not marked will avoid filtering. Figure 8(b) shows the performance of the *Source-End Provider Marking* model for different percentages of legacy routers. In this experiment the marking field of packets coming from legacy providers has a random value. Figure 9 shows the performance of the *Source and Destination-End Provider Marking* model for different sizes of the last provider. In this experiment the marking field has a random value only in the part that corresponds to the source-end provider mark. The results in Figures 8 and 9 show that there are substantial



(a) User and attacker acceptance ratio (30% legacy routers)



(b) Different percentage of legacy routers

Figure 8: *Source-End Provider Marking* with partial deployment

gains even under partial deployment of the proposed models.

### 3.4 Performance with IPv6

Figure 10 shows the performance of the *Source-End Provider Marking* model when the 20 bit flow label field of the IPv6 header is used for marking, which gives us 4 more bits than the IPv4 identification field. Figure 10 shows that by using the larger flow label field we improve the performance of the *Source-End Provider Marking* model by approximately 2% for a small number (1000) of attackers and 15% for a large number (10000) of attackers.

Figure 11 shows the performance of the *Source and Destination-End Provider Marking* model, when a different percentage of the 20 bits are assigned to the source-end and to the destination-end provider. In this experiment we assume that the edge routers of the destination-end provider are five hops away from the destination, and marking at the destination-end provider uses hashing. The results show that by assigning 16 bits to the source-end provider and the remaining 4 bits to




Figure 9: *Source and Destination-End Provider Marking* with partial deployment

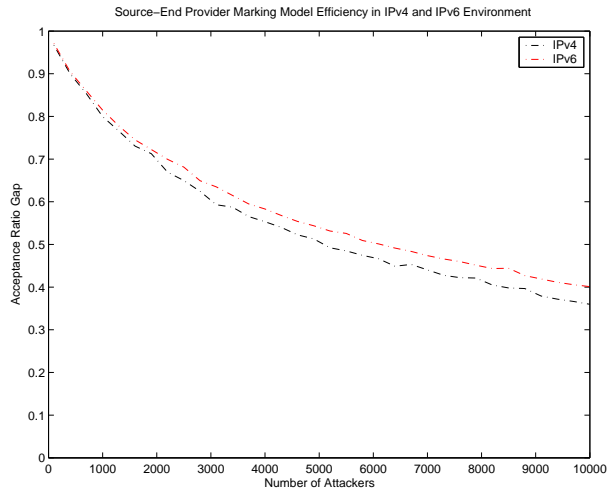


Figure 10: *Source-End Provider Marking* with 20 bit IPv6 flow label

the destination-end provider, we improve performance by approximately 10% for a small number (1000) of attackers, and 30% for a large number (10000) of attackers. Figure 11 also shows that increasing further the number of bits assigned to the destination-end provider reduces performance; this is due to the significant reduction of differentiation that is achieved with source-end marking when the number of bits used is less than 14, as shown in Figure 5.

## 4 Deployment incentives

One of the most important features of the proposed models are the positive economic incentives they give to a provider to deploy them. If increased traffic volume due to *DDoS* attacks results in increased demand, hence increased revenue, a provider has no incentives to deploy a *DDoS* defense model. However, if he can use the defense models to offer better protection as an added value service to his customers, hence

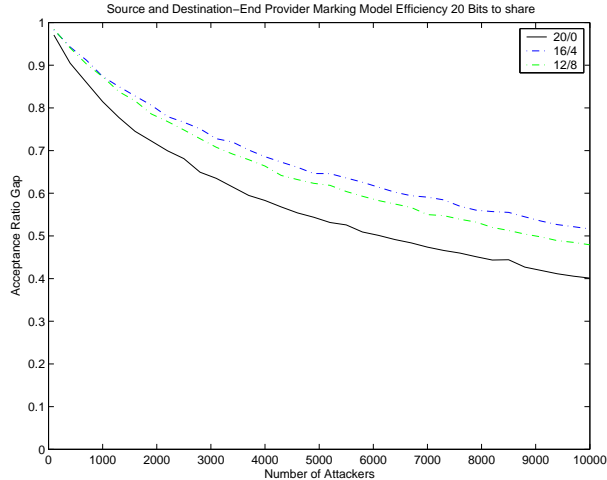


Figure 11: *Source and Destination-End Provider Marking* in IPv6

increase his revenue stream, then he does have a major incentive to adopt such a defense model.

Furthermore, all necessary countermeasures (detection and filtering) belong to the administration of the last provider, which is the stakeholder that gains from the defense model. Thus, there are no security threats for a provider to deny to apply filters, as happens in many direct cooperation schemes.

The above discussion referred to the incentives for a destination-end provider. For the source-end provider there are also deployment incentives, when the source-end provider simultaneously acts as a destination-end provider, e.g. when traffic originates and is destined for customers of this provider. In this case, the provider can offer increased protection to its customers from attacks that originate from domains it is directly connected to, in addition to attacks that originate in source domains connected to other providers.

## 5 Related work

In this section we discuss other *DDoS* defense systems that use deterministic packet marking policies, identifying the similarities and differences with our work. Furthermore, we present different approaches addressing the problem of *DDoS* attacks.

According to the *Pi* marking scheme [21], every router in the Internet marks packets with one or two bits that are produced by hashing the IP addresses of the marking router and its previous hop. Compared to *Pi* marking, our approach achieves from 10% to 20% better *acceptance ratio gap*, and even more with the *Source and Destination-End Marking* model, using order of magnitude fewer marking routers, since we assume that marking is performed only by edge routers belonging to the provider network. Furthermore, the *Pi* scheme suffers from short paths false marking attacks, which can arise when there are unmarked bits due to short paths. The *Stack Pi* marking scheme [22] is an improvement of the *Pi* scheme, the main difference being how it handles the existence of legacy routers. In particular, this model uses the identification field as a stack of marking bits; this is achieved if each marking router shifts the mark value before adding its own mark. The main improvements of *Stack Pi* over *Pi* arise in cases of partial deployment. The evaluation in [22] is for the combined operation of the *Stack Pi* marking scheme and an

optimal threshold-based filtering mechanisms, hence the results are not comparable with the results in this paper, which consider a simple filtering scheme.

Similar to *Pi* marking is the approach of [9], which assumes that the marking field is initialized to 0's by the first marking router. Each router along the path to the destination marks a bit of the marking field, the position of which is chosen randomly and remains the same for minutes or hours. The marking value is produced by simply changing the bit's previous value from 0 to 1 or the opposite. This scheme faces problems in an environment of partial deployment. Moreover, if the participating domains are separated by non-participating domains all initial marks can be lost. In the work of [6], the first four and the last four routers along a packet's path each mark two bits in the packet's header. The two marking bits are determined based on the four color theorem and the Internet's hierarchy.

The work in [20] presents a *DDoS* defense scheme that utilizes IP traceback to perform packet filtering. The approach considers two types of marks: one for performing IP traceback and one for performing filtering. Packets are marked with one of the two mark types with some percentage. The mark corresponding to IP traceback is used to measure the traffic rate received from a particular path, which is then used to compute a drop probability. The packet dropping scheme gives priority to packets containing an IP traceback mark or a mark identified as belonging to non-attack traffic. An issue with this approach is how to correlate marks used for IP traceback and marks used for filtering in an environment where IP source addresses can be spoofed. Our work differs from the above approach in the marking scheme, where we do not rely on IP traceback, and in the filtering scheme, where we ensure that traffic identified as non-attack traffic receives the same average throughput that it received before the attack, while not starving traffic containing a mark identified to belong to attack traffic.

Another approach for DoS protection is the controller-agent model [17]. According to this approach, edge routers connecting an ISP to the Internet mark packets with id's determined by a controller. After detecting an attack, the victim communicates with the controller, which in turn establishes filters with the signature of the attacks at the edge routers. The work in [18] is an extension of the controller-agent model for TCP SYN flooding attacks. Unlike the controller-agent model, our approach does not involve any direct communication between different domains, and there is no single point of failure. Furthermore, our approach can differentiate traffic based on source domain information, in addition to destination domain information.

Next we present approaches for Distributed DoS detection and protection that differ from deterministic packet marking. The approach in [8, 3] uses an overlay network to forward packets from nodes that have proven their legitimacy, whereas traffic from all other nodes are filtered.

Source-end models [11] consider the application of detection mechanisms or countermeasures at the source networks. Hence, possible attacks are detected and filtered before they enter the Internet core. An issue with source-end models is that there is a lack of incentives for their deployment.

Probabilistic Packet Marking models [13, 15] encode partial route path information, like hash-based information, for traceback purposes by probabilistically marking IP packets during or after the attack. With such probabilistic marking schemes, the victim needs to receive a large amount of marked packets in order to trace the attack back to its source. Moreover, after tracing the source of an attack, some communication between the destination and the source is required, in order for the latter to take some countermeasures.

Filtering models such as egress and ingress filtering [5], use topology based address information in order to filter attack or suspicious traffic. Egress filtering focuses on protecting other domains from outgoing attack or suspicious traffic,

whereas ingress filtering focuses on protecting a domain from incoming attack or suspicious traffic. Another filtering approach is route-based filtering [12], which uses route information to filter packets containing spoofed IP addresses. Pushback [7] is an approach according to which a router notifies upstream routers to apply filtering rules, thus pushing filtering away from congestion points near the victim. In the same direction, the work in [1] uses IP's record route option to detect the source of an attack, and a communication model for the targets of the attack to inform domains closest to the source of the attack to apply filters.

## 6 Conclusion and future work

We have presented two provider-based deterministic packet marking models which aim to provide a common attribute for attack streams, which can be used by providers to establish filters, hence offer their customers increased protection against *DDoS* attacks. Our experiments demonstrate that there are significant gains in using the proposed models even under partial deployment. Moreover, by offering their customers increased protection against *DDoS* attacks as a value-added service, providers can increase their revenue stream; this provides positive incentives for deploying the proposed models.

We are currently evaluating parameterized filtering and rate-limiting mechanisms, such as the ones presented in this paper, that can give providers the flexibility to define different levels of protection against *DDoS* attacks. Another interesting issue is how the marking information added by the proposed models can be used to improve the performance of attack detection algorithms. Finally, in this paper we have assumed that the gain in accepting some percentage of legitimate traffic is equal to the cost of accepting the same percentage of attack traffic, and for this reason we subtract the attack acceptance ratio from the user acceptance ratio to obtain the acceptance ratio gap. We are investigating a more general metric, which considers the relative cost of accepting legitimate traffic and the cost of accepting attack traffic.

## References

- [1] A. Argyraki and D.R. Cheriton. Active Internet Traffic Filtering: Real-Time Response to Denial-of-Service Attacks. In *Proc. of USENIX Annual Technical Conference*, 2005.
- [2] H. Burch and B. Cheswick. Internet watch: Mapping the Internet. *Computer*, 32(4):97–98, April 1999.
- [3] D. Cook, W. Morein, A. Keromytis, V. Misra, and D. Rubenstein. WebSOS: Protecting Web Servers From DDoS Attacks. In *Proc. of IEEE International Conference on Networks (ICON)*, 2003.
- [4] C. Douligeris and A. Mitrokotsa. DDoS attacks and defense mechanisms: classification and state-of-the-art. *Computer Networks*, 44(5):643–666, April 2004.
- [5] P. Ferguson and D. Senie. Network ingress filtering: Defeating Denial of Service attacks which employ IP source address spoofing. *Internet Engineering Task Force, RFC 2827*, May 2000.
- [6] Z. Gao, N. Ansari, and K. Anantharam. A New Marking Scheme to Defend against Distributed Denial of Service Attacks. In *Proc. of IEEE Globecom'04*.

- [7] J. Ioannidis and S. M. Bellovin. Implementing Pushback: Router-Based Defense Against DDoS Attacks. In *Proc. of Network and Distributed System Security Symposium, California*, 2002.
- [8] A. Keromytis, V. Misra, and D. Rubenstein. SOS: Secure Overlay Services. In *Proc. of ACM SIGCOMM'02*.
- [9] Y. Kim, J. Jo, and F. Merat. Defeating Distributed Denial-of-Service Attack with Deterministic Bit Marking. In *Proc. of IEEE Globecom'03*, 2003.
- [10] Y. Kim, W. C. Lau, M. C. Chuah, and H. J. Chao. PacketScore: Statistics-based Overload Control against Distributed Denial-of-Service Attacks. In *Proc. of IEEE INFOCOM'04*.
- [11] J. Mirkovic, G. Prier, and P. L. Reiher. Attacking DDoS at the Source. In *Proc. of IEEE International Conference on Network Protocols (ICNP)*, 2002.
- [12] K. Park and H. Lee. On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack Prevention in Power-Law Internets. In *Proc. of ACM SIGCOMM'01*.
- [13] S. Savage, A. R. Karlin D. Wetherall, and T. Anderson. Practical Network Support for IP Traceback. In *Proc. of ACM SIGCOMM'00*.
- [14] V. A. Siris and I. Stavrakis. Provider-Based Deterministic Marking against Distributed DoS Attacks. In *Proc. of 19th Int'l Parallel and Distributed Processing Symposium (IPDPS), 1st Int'l Workshop on Security in Systems and Networks (SSN), 2005*.
- [15] D. X. Song and A. Perrig. Advanced and Authenticated Marking Schemes for IP Traceback. In *Proc. of IEEE INFOCOM'01*.
- [16] I. Stoica and H. Zhang. Providing Guaranteed Services Without Per Flow Management. In *Proc. of ACM SIGCOMM '99*.
- [17] U. K. Tupakula and V. Varadharajan. A Practical Method to Counteract Denial of Service Stacks. In *Proc. of 16th Australasian Computer Science Conference on Research and Practice in Information Technology (CRIPTS)*, 2003.
- [18] U. K. Tupakula, V. Varadharajan, and A. K. Gajam. Counteracting TCP SYN DDoS Attacks using Automated Model. In *Proc. of IEEE Globecom'04*.
- [19] R. van den Berg and P. Dibowitz. Over-zealous Security Administrators are Breaking the Internet. In *Proc. of LISA Conference*, 2002.
- [20] J. Xu and M. Sung. IP Traceback-based Intelligent Packet Filtering: A Novel Technique for Defending Against Internet DDoS Attacks. In *Proc. of IEEE International Conference on Network Protocols (ICNP)*, 2002.
- [21] A. Yaar, A. Perrig, and D. Song. Pi: A Path Identification Mechanism to Defend against DDoS Attacks. In *Proc. of IEEE Symposium on Security and Privacy*, 2003.
- [22] A. Yaar, A. Perrig, and D. Song. StackPi: New Packet Marking and Filtering Mechanism for DDoS and IP Spoofing Defense. Technical Report CMU-CS-02-208, Carnegie Mellon University, February 2003.