

# 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.*

## 1. Introduction

Distributed Denial of Service (*DDoS*) attacks are one of the most serious security threats in the Internet today, undermining the further deployment of new services and limiting the usage of existing services, such as e-commerce. 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 [1]. *DDoS* attacks achieve their goal either by con-

suming 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 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 largely based on destination addresses, and no entity is responsible for ensuring that source addresses are correct. Thus, anyone could generate attack traffic that appears to have originated from almost anywhere, by simply forging the source address in the IP header. This process is called *spoofing* and is widely adopted in flooding attacks.

### 1.1. Desirable properties of a defense system

In the last several years many *DDoS* countermeasures have been deployed. Most of the work in this area has focused on tolerating attacks by mitigating their effects on the victim. A powerful defense model should have several properties in order to be characterized as effective and secure. In particular, see also [2], 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 of 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.

---

\* The authors are also with the Dept. of Computer Science, Univ. of Crete.

- 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. Such an approach 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.

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 [3], 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 [2], using order of magnitude fewer routers and giving providers deployment incentives for its adoption.


The rest of this paper is organized as follows. In Section 2 we present the two provider-based 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 the 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 of providers that connect a provider to its customers or the Internet. Following the approach presented in [2], 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 [4] indicate that the percentage of fragmented packets is very small (less than 0.25%). Moreover, as suggested in [4], 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



**Figure 1. Source-End Provider Marking model**

that all packets are marked with the same value, which consists of the last two bytes of the MD5 hash of the IP address of the interface that connects the router to the source domain; this is performed in order to achieve uniform distribution of mark values [2]. 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.

**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  $a_i, i \in A$ , is the rate of attack traffic with mark  $i$ , then limiting attack traffic to rate  $C_{attack}$  will result in dropping packets identified as attack traffic with percentage

$$\frac{\sum_{j \in A} a_j}{\sum_{i \in A} (l_i + a_i)} \frac{C}{\sum_{i \in I} l_i}.$$

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

$$C_j = \frac{l_j}{\sum_{i \in I} l_i} C \text{ for } j \in A.$$

This rate-limiting scheme results in dropping a percentage of packets with mark  $j \in A$  equal to  $\frac{a_j}{l_j + a_j} \frac{C}{\sum_{i \in I} l_i}$ . Hence, the percentage of dropping for a mark is an increasing function of the amount of actual attack traffic with this mark, i.e. the intensity of the attack. One can show that this multiple rate-limiting approach allows a larger percentage of legitimate packets, which 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 provides 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 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 attacks, 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.

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.




Figure 2. Source and Destination-End Provider Marking model

**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 is that in normal circumstances we expect the legitimate traffic to enter the last provider from one or few edge routers from each source domain. Thus, in large *DDoS* attacks we expect that attack traffic to enter the destination-end provider from all of its edge routers that connect it to the Internet. Because different edge routers at the destination-end provider mark packets differently, only the attack traffic that enters through the same edge routers as the traffic from a particular source domain will maintain the same mark value as the legitimate traffic from that source domain. Indeed, a situation where a destination-end provider receives packets with the same mark from a large number of edge routers that connect it to the Internet can be used as an indication of a false marking attack; development of such a detection mechanism is part of our future work.

## 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 probabil-

ity 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; we leave the comparative evaluation of different filtering schemes for future work.

### 3.1. Experiment scenario and metrics

The topology used in our experiments was Burch and Cheswick’s Internet Map [3], which was created using traceroute messages from a single host to destination hosts throughout the Internet, producing a tree with thousands of paths. We assume the victim of the *DDoS* attacks is the root of the tree and the legitimate and attack hosts to be specific leaves of the tree.

In our experiments, similar to [2], 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; 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 [2]. 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.

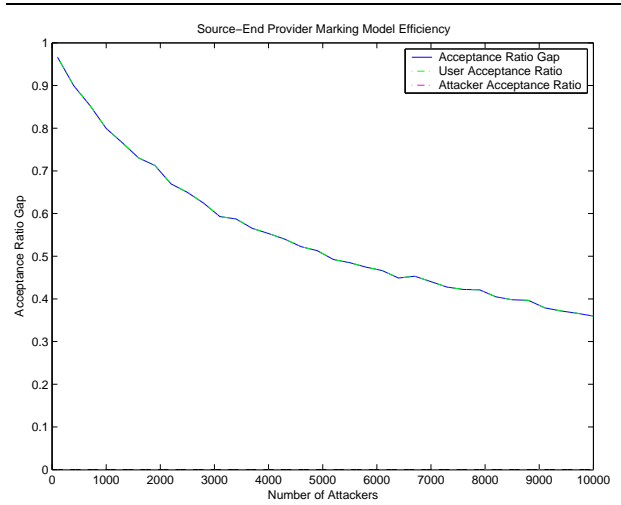



Figure 3. Source-End Provider Marking model performance

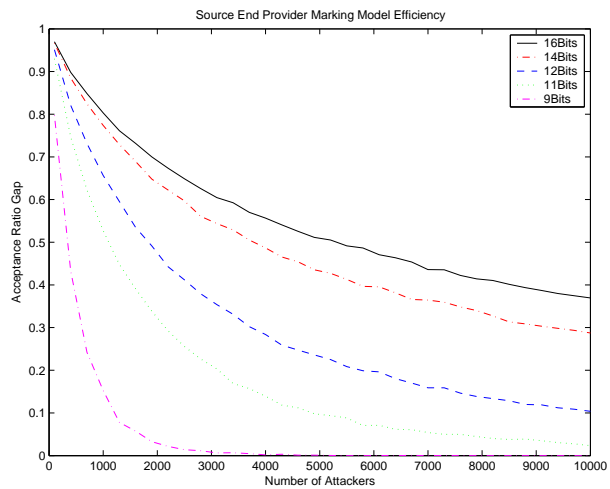
### 3.2. Attack and legitimate traffic differentiation

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 a larger number of bits is 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, 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



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



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


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 are investigating how the topology and the length (number of hops) of the path between the source and destination influences this tradeoff.

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 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.

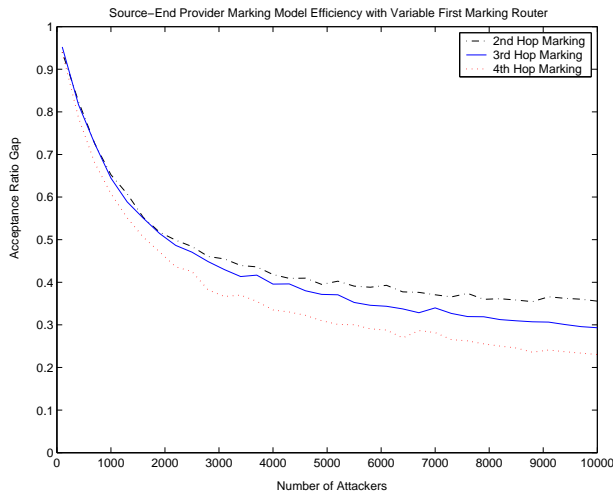
### 3.3. Partial deployment

Next we investigate the performance of the two models in an environment of partial deployment. We as-




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

sume 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 shows that the *Source and Destination-End Provider Marking* model exhibits substantial gains even under partial deployment. In this experiment the marking field has a random value only in the part that corresponds to the source-end provider mark.



(a) 2 bits needed for last provider

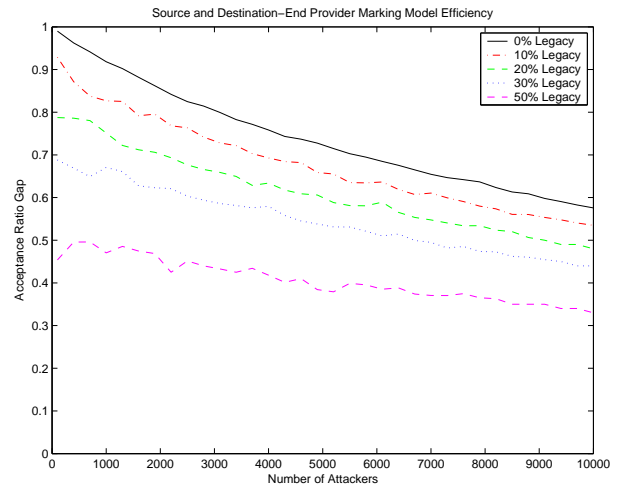


(b) 7 bits needed for last provider


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

### 3.4. Performance with IPv6

Figure 9 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 9 shows that by using the larger flow label field we improve the performance by approximately 2% for a small number (1000) of attackers and 15% for a large number of attackers. Due to space limitations, we do not show the increase in performance for the *Source and Destination-End Marking* model. However, based on the results in Figure 9 and the comparison of Fig-



**Figure 8. Source and Destination-End Provider Marking with partial deployment**



**Figure 9. Source-End Provider Marking with 20 bit IPv6 flow label**

ures 4 and 5, we expect that the performance will increase more by allocating the extra 4 bits for marking at the destination-end provider.

## 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. Since increased traffic volume due to *DDoS* attacks results in increased de-

mand, 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 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 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 summarize *DDoS* defense systems that use deterministic packet marking policies. One approach considers that every router along a path adds a mark to all packets. The basic representative of this class is the *Pi* marking scheme [2], with several descendants [4, 5, 9]. 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.

The work in [6] presents a *DDoS* defense scheme that utilizes IP traceback to perform filtering. Our work differs 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 [7, 8]. According to this approach, edge routers connecting an ISP to the Internet mark packets with id's determined by a controller. Unlike the controller-agent model, our approach does not involve any communication between different entities, 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.

## 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 offering their customers increased protection against *DDoS* attacks. Our experiments demonstrate that there are significant gains in using the proposed models even under partial deployment.

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 *DDoS* detection algorithms. Finally, 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] C.Douligeris, A.Mitrokotsa, DDoS attacks and defense mechanisms: classification and state-of-the-art, *Computer Networks* 44 (2004) 643-666.
- [2] A. Yaar, A. Perrig, D. Song, Pi: A path identification mechanism to defend against DDoS attacks, *IEEE Symposium on Security and Privacy*, May 2003.
- [3] H. Burch, B. Cheswick. Internet watch: Mapping the Internet. *Computer*,32(4):97-98, Apr. 1999.
- [4] A. Yaar, A. Perrig, 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.
- [5] Y. Kim, J.-Y. Jo, F. L. Merat, Defeating Distributed Denial-of-Service Attack with Deterministic Bit Marking, *IEEE Globecom* 2003, .
- [6] M.Sung, J.Xu, IP Traceback-based Intelligent Packet Filtering: A Novel Technique for Defending Against Internet DDoS Attacks, in *Proceedings of the 2002 IEEE International Conference on Network Protocols (ICNP'02)*.
- [7] U. K. Tupakula, V. Varadharajan, A. K. Gajam, A practical method to counteract Denial of Service Attacks, *26th ACCRPIT, ACM ICPS*, 2003, pp. 204-275.
- [8] U. K. Tupakula, V. Varadharajan, A. K. Gajam, Counteracting TCP SYN DDoS Attacks using Automated Model, *IEEE Globecom* 2004.
- [9] Z. Gao, N. Ansari, K. Anantharam, A New Marking Scheme to Defend against Distributed Denial of Service Attacks, *IEEE Globecom* 2004.