

Spooing Prevention Method

Anat Bremler-Barr
 Interdisciplinary Center Herzliya
 bremler@idc.ac.il

Hanoch Levy
 Tel-Aviv University
 hanoch@cs.tau.ac.il

Abstract—A new approach for filtering spoofed IP packets, called Spooing Prevention Method (SPM), is proposed. The method enables routers closer to the destination of a packet to verify the authenticity of the source address of the packet. This stands in contrast to standard ingress filtering which is effective mostly at routers next to the source and is ineffective otherwise. In the proposed method a unique temporal key is associated with each ordered pair of source destination networks (AS's, autonomous systems). Each packet leaving a source network S is tagged with the key $K(S, D)$, associated with (S, D) , where D is the destination network. Upon arrival at the destination network the key is verified and removed. Thus the method verifies the authenticity of packets carrying the address s which belongs to network S . An efficient implementation of the method, ensuring not to overload the routers, is presented. The major benefits of the method are the strong incentive it provides to network operators to implement it, and the fact that the method lends itself to step-wise deployment, since it benefits networks deploying the method even if it is implemented only on parts of the Internet. These two properties, not shared by alternative approaches, make it an attractive and viable solution to the packet spoofing problem.

Index Terms—System design, Security, Distributed Denial of Service, Source IP spoofing

I. INTRODUCTION

Spooing the source IP address of packets on the Internet is one of the major tools used by hackers to mount denial of service (DoS) attacks. In such attacks the attackers forge the source IP of packets that are used in the attack. Instead of carrying the source IP of the machine the packet came from, it contains an arbitrary IP address which is selected either randomly or intentionally. The ease with which such attacks are generated made them very popular. According to a study [1] there are at least four thousand such attacks every week in the Internet.

Aside from being very effective in generating denial of service on the victim, the spoofed attacks give hackers two additional advantages: First, it weakens the ability to mitigate the attack since the malicious traffic cannot be categorized (by source) and hence much harder to filter

out. Second, it makes law enforcement harder, since it is much harder to trace back the source of the attack and hence to trace the attacker.

There are very few and not very effective mechanisms that network operators may use today to detect and filter out spoofed packets. These mechanisms are briefly reviewed in the next subsection. The most prominent of them is the ingress and egress filtering (sometimes relying on the uRPF [2] feature in routers). In ingress filtering an ISP prohibits receiving from its stub connected networks, packets whose source address does not belong to the corresponding stub network address space [3]. In egress filtering a router or a firewall, which is the gateway of a stub network, filters out any packet leaving the network whose source address does not belong to the network address space [2], [4]. Thus, both mechanisms ensure that the traffic leaving out of a stub network may only spoof addresses that belong to the same stub network. The former does it at the ISP router while the latter at the stub edge equipment¹.

Ingress and egress filtering are "good-will" preventive and *not* self-defensive methods. Cooperative and "good-netizen" network operators deploy the method to avoid being the source of such attacks, however these methods do not provide any remedy to victims while being attacked. Only if implemented by all (or most) networks, ingress and egress filtering could significantly reduce packet spoofing in the Internet. Furthermore, the deployment of ingress/egress filtering inflicts significant costs on the ISP implementing it, both in equipment and labor (administration, management, configuration) without being assured significant benefit or protection. Thus the incentive for an ISP to deploy these mechanisms is relatively low. In addition, whenever multi homing is used, egress filtering is problematic since it may block legitimate traffic and requires very careful network administration.

In this paper we present an alternative solution, the Spooing Prevention Method (SPM), which offers an

¹Note that these mechanisms do not deal with forging of source addresses to another address that belongs to the stub IP address domain. Such spoofing may be easier to swart by filtering on source-IP subnet.

efficient and defensive method by which routers on destination networks can detect and filter out spoofed packets. Together with the ingress/egress methods SPM is both more effective in stopping spoofed attacks and provides an incentive to the network owners to implement the methods. Thus becoming a defensive method that overcomes the deficiencies of the ingress/egress filtering approach. An ISP that joins SPM marks all the packets originating in its domain with a special key that is known only to the participants of SPM. The key placed on each packet is a function of the source network (AS) and the destination network (AS). For each packet that arrives at its destination network, routers in the destination network check the source network and verify that the key on the packet equals to the key that corresponds to that source network. If not, then the packet is marked spoofed and may be discarded. The assignment of a key to an outgoing packet, and the verification of the key on an incoming packet, each requires one IP lookup operation. The keys are changed periodically (every few hours) and are a fixed string of 32 bits. All the "online" operations related to the keys are thus simple memory look-ups (one look-up per packet), no cryptographic computations are involved. Keys are selected and distributed by the origin networks as discussed in Subsection III-B. We propose two methods for the distribution of key information: The first uses a key distribution protocol, while the second does not require such a protocol. In the latter method the key information is learned passively by observing the keys used in the regular non-spoofed traffic, where non-spoofed traffic is identified (and verified) via other techniques, such as TCP intercept. These techniques resemble techniques that were proposed in order to secure BGP [5], [6], [7], [8], and can use the same infrastructure.

As described, the SPM method alone does not stop a source from spoofing another source IP address from the same source AS network. While the method greatly limits the amount of spoofing a source can do, we suggest to further limit the spoofing range by combining the SPM method with the ingress/egress filtering method described above. That is, participating SPM ASes are also required to enforce ingress/egress filtering with their stub networks. The combination imposes a nearly tight anti-spoofing mechanism that enables trace back and enforcement of non-spoofed traffic in the Internet.

The important characteristic of the SPM method is that packets that come from an SPM member are marked and are authenticated. Hence, if an ISP detects an attack on itself or on one of its customers, it protects itself from spoofed packets by allowing in only packets that originate from SPM member ASes. Moreover, the clients

of an SPM member get proper service from the attacked site, in case that the attacked site is an SPM member. In addition to these two direct incentives, we argue that the SPM architecture introduces high visibility of the networks deploying SPM and hence is likely to be attractive for ISP's to join.

In addition to providing its servers and clients a better service, the method enjoys from the benefits of step-wise deployment. That is, a subset of the ASes that deploys the method enjoys the benefits and is able to provide its member ASes spoof-less traffic between their customers, even if other ASes have not yet joined SPM. Moreover, as stated above, if and when the members detect a spoofed attack, they can guarantee clean service between their customers while blocking any suspicious traffic coming from ASes that do not participate in the SPM system. Notice that there are today methods that use authentication key, like the authentication header in IPv6 (RFC 2460 [9]). However, these methods use the authentication in order to secure sessions between the clients. They do not provide any defense against DDoS in general and spoofed packets in particular. To the opposite, they can be even used to perform a hard DDoS attack on the clients. The attacker would send many packets with invalid authentication field. The clients would perform computational heavy cryptographic check. Moreover, these methods secure the validity of the session only after the session has been established, since the key exchange is part of the session initiation. Hence, these methods are also vulnerable to spoofed SYN attacks. The SPM, on the other hand, uses a very lightweight authentication mechanism on the source address, and can authenticate the source address, from the first (SYN) packet. This is due to the fact that SPM uses a key per AS pair, to verify the source address. Where as the key is known in advance to the two endpoints.

The structure of this paper is as follows: In section III we give a detailed description of the architecture of SPM. In section IV, we provide an analytic model evaluating and quantifying the added motivation of ISP's to implement SPM. The model shows that SPM provides significant protection to its participants, and relative benefit (in comparison to not joining SPM). The benefits are in the reduction in the amount of spoofed DDoS attacks and in higher availability of services to its clients. This benefit increases with the number of participating ISP's. Nonetheless, the model also shows that even when SPM consists only of part of the ISP's, a participant of SPM derives significant benefits from the system, and thus has strong incentive to joining it.

II. RELATED WORK

There are three mechanisms used today in order to stop spoofed packets: ingress/egress filtering at the origin, trace back, and attempts to mitigate the packets at the destination.

In addition to the ingress/egress filtering described above, some network operators configure their routers to drop any packet that is obviously spoofed, i.e., packets that use private source addresses that try to leak to the Internet [10], or packets whose source IP address is from an unallocated address space (bogons addresses [11]). This can be done, by configuring explicit filters, or by using the rpf (Reverse Path Forwarding) feature [2], that does not allow receiving packets with source addresses that are not in the forwarding table. Notice, that these methods require an extra lookup at the router forwarding table, on the source address.

Traceback is the ability to determine the path(s) an attack flow traverses in the Internet and by this to discover the sources of the malicious traffic. While this feature is not currently supported by the routers, there are many new suggestions on how to add this capability to routers [12], [13], [14], [15], [16], [17], [18]. Those suggestions are divided into two major classes of methods: suggestions that involve stamping the traffic packets with a signature of the routers they pass through [12], or suggestions that involve sending samples of the traffic packets, to a special collector that analyzes the path [13].

The capability to mitigate DDOS attacks at the destination network is very limited and is done today by the TCP intercept [19] feature. In TCP intercept the router checks if there is a real host behind the source address by completing the 3-way handshake of the TCP protocol on behalf of the destination server. In case that the connection with the client is successfully established, the client address is not spoofed, and the router establishes the connection with the server on behalf of the client and then acts as a TCP slicer, i.e., carefully and transparently stitching the two connections together. The drawback of this solution is that it is applicable only to TCP, and hence cannot protect from spoofed attacks on UDP traffic (or any other connection less traffic). Moreover, in many router vendors the TCP intercept poses a serious performance penalty on the router. Other issues with this method are that it either relies on symmetric and consistent routing or hides the true source IP of a client.

A few recent studies [20], [21], [22] proposed solutions that resemble, in some aspects, to SPM, but where the key is a function of the route traversed by the packet traverses and not a function of the source destination network (as it is in SPM). The main advantage of SPM

on these techniques is that SPM directly benefits the ASes implementing it and not the other ASes, thus providing ASes strong incentive to implement SPM. In contrast, those techniques do not provide extra-benefits to the organizations adopting them and thus may seriously suffer from lack of incentives to adopt them. One of those studies [20], [21] proposed to use the TTL field in the IP header, to play the role played by the key in SPM. The disadvantage of this technique is that it may allow an attacker to spoof up to 10% of the packets [21] while SPM can filter out 99.9999999% of the spoofed traffic.

A second approach proposed by these recent studies [22] suggested to add a key field to each IP header, that would contain a signature of the routers on the path traversed by the packet. Choosing keys as a function of the route, may be problematic since the route (and consequently the key information) changes very dynamically. Moreover, it is possible that there exist multiple routes per single source/destination network pair, making aggregation of the key information very problematic. This difficulty is overcome by SPM. Partial results of this work was presented at [23].

III. THE SPOOFING PREVENTION METHOD

In the SPM architecture a key is added to each packet, to validate that the packet is not spoofed. The key, is a constant number, that is chosen for marking all the traffic between a source AS and destination AS. The fact that the key is a function of the source and destination makes it hard to spoof.

To enable the SPM some routers at participating ASes are required to: 1) Mark the outgoing packets with the appropriate key, and 2) Verify the authenticity of the key on incoming packets.

Similar to the BGP architecture, the elementary players in the SPM are the ASes². Every AS chooses independently the set of keys to mark traffic that originated from its AS. This set of keys is distributed to other participants in SPM. The distribution of the key can be achieved either by designing a special distribution protocol, or by passive label distribution protocol, where the key assignments, is derived from the normal traffic.

Under SPM keys are placed on the packets by routers at the source AS, their authenticity is checked by routes at the destination AS, and the keys are removed from the packets after the authenticity check has been performed.

²Autonomous system - is a unit that is controlled by a common network administrator. In many cases AS is an ISP or a big Enterprise, in some cases due to administrator convince, ISP's/Enterprises are split to several AS's or an AS is an aggregation of several small Enterprises/ISPs.

Therefore, attackers that have access only to edge devices, such as standard computers and servers cannot see or affect the keys or the method. Note that SPM can be used with or without the ingress/egress filtering.

In the following section, we describe the three basic building blocks of the architecture: The key, the key distribution protocol, and the routers new tasks.

A. The key

Two basic issues regarding the key are: where is the key placed in the packet, i.e., in which layer? and what is the key?

Since the role of the key is to verify that the source address of a packet is not spoofed its natural and most effective placement is in the IP header, where the packet source address appears. Note that the IP layer is the largest common denominator protocol of the Internet, in the sense that all Internet protocols are running over it. Thus, adding the key to the IP header allows SPM to capture any spoofing attack carried out over any Internet protocol, such as, TCP, UDP, or other protocols.

For the ease of deployment and implementation, the key can be placed in two places in the IP header. The first place is in the IP option field, and second is to use the ID field in the IP header. The disadvantage of adding a new IP option is that most of the routers today do not process the IP options, due to the performance penalty involved. Using the ID header, for a DDoS related solution was previously suggested [12]. Two disadvantages in using the ID header are, first, that the packet ID may be used in all of the packets that are not fragmented. However, today, most of the traffic in the Internet is not fragmented [24]. And second, that the ID field length is only 16 bit, which in turn constrains the key to be 16 bits long.

The main guideline for the key handling procedures is that it should be very lightweight. One should be careful, not to use heavy calculations in marking or verifying the key, otherwise those functions may be themselves a target for a DDoS attack on the routers, since the routers use these functions on high volume flows.

In general, a different key is selected for each source destination AS pair. Specifically, for source AS and destination AS, that participate in SPM, a random constant number is chosen, so that all the traffic flowing from the source AS to the destination AS is marked by the same key. The key is renewed periodically, every few hours. Traffic, to AS's that do not participate in SPM, does not have to carry a specific key.

Other key selection methodologies that one may consider are:

- 1) **A key for each source address:** In this case it is easy for an attacker inside the SPM, to acquire the

key of the source address she wants to spoof: The attacker sends a request (e.g., SYN packet) to the source address she wants to spoof³ and deduces the source key from the reply.

- 2) **A key for each source-destination address pair:** To eliminate the above drawbacks one may suggest having a unique key for each source-destination address pair. This solution is not practical due to the number of keys and other data structures each router would have to maintain.
- 3) **A key for each source-destination network pair:** Aggregation is applied in order to eliminate the complexities of the per address pair suggestion. One option is to store a key per source-destination subnet (prefix) pair. While this method may provide a more fine-grained anti spoofing by preventing a source from spoofing any other subnet than its own, it has two drawbacks: the amount of key information to be stored in each router is still huge (squaring the routing table size, which is already 120k today), and secondly, special care is required in the design of the protocol, since routers in different ASes store different lists of networks (prefixes), mainly due to the CIDR [25] protocol.
- 4) **A key for each source-destination AS pair:** The selected methodology this paper focuses on associates a key with the source AS and destination AS of each packet⁴ that participate in SPM. When a router in an AS that participate in the SPM, receives a packet, it can verify that the source address is not spoofed by checking that the key on the packet matches the key that is associated with the source AS and the current AS.

Even though attackers cannot effect this anti spoofing method without sniffing or hijacking traffic on backbone or peering links we use 32 bits keys to secure the mechanism.

Unlike standard cryptography, in which the attacker should not be able to guess the correct key even once, here the requirement is that the attacker guesses the key correct only with low probability as long as she does not know when the guess was right. This is because stopping 99.9% of the malicious packets in a DDoS attack is good enough, unlike in cryptography where the attacker should

³Notice, that this would require the spoofer to find an address that would answer him. This can be achieved for example by scanning the address space for servers, by sending TCP/SYN requests, and waiting for the SYN-ACK

⁴AS stands for Autonomous system, which is a portion of the Internet controlled by a common network administrator. In many cases an AS is an ISP or a big Enterprise, in some cases due to administrator convince, ISP's/Enterprises are divided into several ASes or an AS is an aggregation of several small Enterprises/ISPs.

not be able to guess the key even once. Therefore, we use a simple string of 32 bits as the key, and require a periodic key change. In such a case the attacker may guess the right key in only one of every four billion (2^{32}) packets.

While guessing the key does not pose a threat to the SPM method, acquiring a correct set of keys by a malicious attacker could pose a real danger to the method. However, learning the correct key requires from an attacker to sniff real traffic on some backbone or peering links in the Internet. A hacker that has these capabilities is by far more dangerous to the Internet and any server on the Internet than spoofed DDoS attacks.

While, we are not concerned with real time sniffing, one may be concerned that eventually the key value would some how be disclosed and published. Therefore, in order to make the key more secure, the key should be changed periodically.

B. The key distribution protocol

An important property of the SPM architecture is that all the peering routers (or peering interfaces) in an AS have the key information to verify incoming packets and mark the key on outgoing packets.

The key label information is sum up to two small tables: (1) **The AS-out table**, that maintains keys for marking flows, that originate in this AS, and destined to another AS in the SPM. (2) **The AS-in table**, maintains keys for verification of flows that are destined to networks attached to the local AS.

The AS-out and AS-in tables take around 120KB of memory each, which is very moderate size: there are 16,000 ASes, the AS number is coded by two bytes, and we consider key fields that are bounded by four bytes. Each entry in the table holds two keys, the old key and the new key (see Subsection III-B.4). Hence total of 120KB. Notice, that the maximum total number of possible ASes is bounded by 2^{16} , hence the maximal size that the *AS-in* and *AS-out* tables may reach in the future is bounded by 480KB.

In this subsection we overview two methods for key distribution: First, a passive key information distribution method which does not require a distribution protocol, and, second, an active key distribution protocol.

1) Passive key information distribution method:

This method avoids the use of a dedicated key distribution protocol. The verification key table **AS-in Table**, is learned passively from the tagged keys in the traffic that comes from non-spoofed addresses. One can identify that that traffic is not spoofed, if it is a TCP traffic (TCP connection), that completed the 3-way handshake.

In the case of symmetric routing, we can identify such connections by passively monitoring the incoming and outgoing traffic and verifying that the connection was closed. The situation in asymmetric routing, is more problematic. In such a case, the attacker may mislead the identification mechanism and cause it to conclude that the connection is closed, by sending a SYN packet followed by an ACK packet. Without knowing the sequence number that was used in the response from the server, it would be impossible to verify the connection. In these asymmetric cases the router can use the TCP intercept [19], in order to verify the connection and then retrieve the key information. The overhead of the passive learning key information, is small, since it must be operated only when there is a need for learning the new key values.⁵

The idea of retrieving control information, from the data information, is not new, and was proposed earlier in order to detect hijacking of BGP [8] and also in anti-spoofing mechanisms which use TTL [21], [20].

Up to this point we dealt with how the receiving domain learns the key information for the incoming traffic arriving from other domains. Nonetheless, the key information must be also learned by the sending domain devices in order to conduct the traffic marking.

In this case all the routers residing inside the local AS, need to decide on one way to mark the traffic. When a route reflector is used in the AS, the route reflector can choose the key information, and distribute the key information to the routers inside that AS, by piggybacking the key information, on the route reflector BGP updates.

An alternative approach is to use a configuration that would synchronize the tables of all the routers in the AS. This method requires minimal configuration, i.e, only setting the same secret password at each router. The main idea is to use one way hash function with the parameters of the destination AS, and the configurable password. Using this function, we produce in each router the same AS-out table. The use of a one-way hash-function makes the password cracking task hard.⁶

2) **Active Distribution Protocol:** Distributing key information between routes as required in our method is not new. Due to various BGP security issues such as

⁵The worst case may occur when the retrieving of new key information, is needed while there is a denial of service attack. However, it is sufficient to verify one good connection per source AS, which makes the method reasonable, even in this extreme scenario.

⁶To change the AS-out table from time to time, we can add the time as a parameter.

prefix hijacking⁷ there are suggestions to secure BGP by adding certificate keys to BGP announcements, in order to validate them. The cryptographic keys suggested in secure BGP are not suitable for our anti-spoofing method, since in our case a light function is required in order to be to validate the actual traffic. However, the problem and method of distributing the (BGP) keys are applicable to SPM. Two basic approaches of keys distribution have been suggested in secure BGP and are under consideration of the IETF. One approach is the central approach such as suggested in S-BGP [5] and in So-BGP [6], where the organization of IP registries (ARIN, APNIC, RIPE) are in charge of distributing the keys. The second is the IRV (Inter-domain Routing Validator) approach [7] that uses a distributed approach, where a central server in each AS manages the key distribution and selection. We concentrate on the latter approach, IRV, since we believe it provides better flexibility in partial deployments, which is important for the success of SPM (see IV)

Following the IRV (Inter-domain Routing Validator) architecture and protocol suggested in [7] we use one central server in each AS to manage the key distribution and selection. IRV, is a new architecture and protocol designed to solve security issues in BGP. The basic architecture uses a dedicated server in each AS, called IRV server, that stores information about the integrity of the BGP information of the corresponding AS. The receiver of BGP information, can corroborate the information it received by querying the IRV servers of other AS's. Each AS server, maintains a list of the other AS servers. This list is obtained from information that is piggybacked on the BGP announcement, or from a central registry. The queries between the IRV servers are carried out on a secure channel such as IPsec.

The AS server, performs the following tasks: (1) choosing the keys for the AS-out table, (2) distributing the AS-out table to the routers in this AS, (3) announcing the corresponding keys from the AS-out table to each of the other AS servers that participate in SPM, (4) building the AS-in table, from the announcements of the other AS servers (5) updating the AS-in table in the routers in its AS.

The traffic that an AS server needs to send to other AS servers when it updates its out keys, totals the size of the AS-out table (i.e., a total of 120KB). This is equal also to the total traffic the AS server receives from all the other AS servers, when they update their keys. The total

⁷Prefix hijacking, happens when a router at an ISP announces maliciously or erroneously an IP address space that does not belong to it, and all or some of its neighbors and other networks in the Internet start to route traffic to these destination IP addresses, incorrectly.

inbound and outbound traffic to the AS server during the keys update phase, is quite moderate, 120KB. The bandwidth requirement for such a transfer which is the bandwidth requirement of SPM, is negligible, comparing to the BGP information. This amount of traffic may be transmitted over a T1 line in less than half a second.

3) *Protecting the AS server from DDoS*: It is rather easy to protect the AS server itself from DDoS attacks. The SPM system, after the first key exchange, protect the servers from spoofed attacks. The key exchange itself is over TCP over IPsec (or simply a pair of public and private cryptographic keys for each server). Thus, any packet destined to the AS-server other than TCP, may be discarded. TCP packets to the AS-server are allowed only from the other AS-servers IP addresses, i.e., a white list may be used to enhance its protection.

4) *Changing the SPM keys periodically*: To increase the system security we suggest periodical key updates. E.g., every few hours. Due to the lack of space, the two alternative ways to achieve the goal is given in the appendix VI-A.

C. The SPM routers

Here we discuss which routers should tag outgoing packets with the appropriate key and which routers should perform the authentication on AS incoming packets. In Subsections III-C.1 and III-C.2 we describe the algorithms and data structures used in each of these routers to carry out these operations.

Since the routers that tag the packets need to tag only packets that originate in the local AS, we place the tagging task at the edge routers at the ISP. Since these are the routers that can distinguish between traffic originated in the AS and that should be labeled, and traffic that comes from outside of the AS. These are the same routers that should implement the ingress filtering, and hence a natural option is to combine the two mechanisms, and require that each SPM member will also implement ingress filtering.

The packet authentication, can be best done at the peering routers. This way the packets are verified as early as possible, so the attack would not congest the links of the AS network. However, in most of the cases the edge routers, are less overloaded, and hence could be an attractive alternative to preform the packet authentication.

1) *Tagging packets with keys*: Each router that carry out the tagging operation maintains a network-in table. In order to tag a packet, a lookup on the destination address is required (and in order to authenticate a packet, a lookup on the source address is required). Hence,

one can combine the process of tagging a packet with the regular IP-lookup. The information of the network-out table and network-in table are stored as additional fields in the FIB (forwarding information table). This solution requires that the routers are BGP routers since it requires a detailed forwarding table in which networks from different ASes have different entries. This might be another consideration in selecting the routers that would carry out this task. The process of inflating the network-in and network-out tables given the AS-in and AS-out tables is described in Appendix VI-B.

2) *Dynamic authentication process*: Notice that the cost of tagging a packet is minimal as it is piggybacked on the IP lookup process. The cost of authenticating a packet is higher, since it requires an additional IP-lookup on the source address of the packet. However, the extra IP lookup is also required in other technique such as the uRPF [2] method (see the Introduction). Moreover, one may choose to activate the authentication only during attack time as described below.

During the authentication process packets are classified into one of three categories:

- 1) **SPM-recognized spoofed traffic**: There are two types of spoofed traffic that an SPM member can filter : 1. Spoofed traffic from some AS, where the spoofed address is an SPM address. In this case the traffic is tagged incorrectly. 2. Spoofed traffic that originates from an SPM domain - in this case the source AS, which is a member of SPM, will tag the traffic. However, the source address will not match the SPM key. This can happen since the spoofed source address belongs to another SPM member, and hence it was tagged incorrectly, or that the spoofed source address is of a non SPM member, and hence, the existence of a key, is a sign that the source is spoofed. In case that SPM members also use ingress/egress filtering, the traffic will be filtered at the source AS. In case that SPM members do not use ingress/egress filtering, the traffic will be filtered using the SPM mechanism at the destination AS, only if the destination address is an SPM member. The latter implies that when SPM does not conduct ingress/egress filtering it provides relative advantage to its members.
- 2) **SPM certified non-spoofed traffic**: This is traffic that originates at SPM domains, gets certified by the sending domain and thus can be recognized as authentic traffic by the receiving domain. Traffic of this category passes untouched to the destination.
- 3) **All other traffic**: Traffic that comes from non SPM domains, and may be either non-spoofed or spoofed (that spoofs to non SPM addresses).

The receiving end thus may have to use additional means to sort out this traffic.

In authenticating incoming packets one may either apply the above method continuously on all the packets or distinguish between two types of verification and discard modes.

Peace time (conservative) verification and discard: Only traffic of the first category is completely discarded. This elimination comes with SPM and does not require any additional mechanism on behalf of the receiving domain/server. Dropping this type of traffic, can be done on regular basis, even if there is no detection of attack. Moreover, an increase in the amount of traffic of this category can be used as an indication of a spoof DDoS attack.

Attack time (aggressive) verification and discard: If and when a DDoS attack is detected, the traffic of the third category may be completely discarded as to provide further advantage to traffic coming from SPM domains. Alternatively, additional detection and recognition mechanisms may have to be deployed in order to sort out this traffic and categorize it into spoofed and non-spoofed traffic.

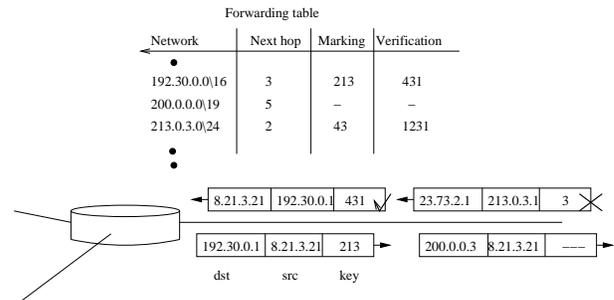


Fig. 1. Illustration of the the routers tasks

IV. ANALYSIS OF THE BENEFITS AND INCENTIVES OF SPM

The objective of this section is to evaluate the benefits that SPM provides to its participants. It is important to evaluate these benefits in a relative-comparative manner, that is, the benefit one gains from joining SPM (as opposed to not joining). The importance of the relative benefit is in evaluating the *incentive* one has to join SPM.

In the context of the analysis below we will take a conservative approach and account only for the benefit from SPM when operating in the peace time verification mode. This will allow us to focus only on the amount of spoofed traffic *directly eliminated* by SPM.

Accounting the additional benefit of SPM when operating in the attack time verification mode requires further

assumptions on the additional protection mechanisms taken by the receiving domain and is outside the scope of this paper. Roughly speaking, in the attack verification mode, the traffic that cannot be verified since it comes from a non-SPM member address and without a key, can be completely blocked as to provide further advantage to traffic coming from SPM domains. Note, that after the establishment of the SPM, if the core of the big ISP's decide to participate in SPM, this approach would give other ISP's high incentive to join SPM. This is true since by allowing only SPM member traffic to reach a server, the server still gets high share of the legitimate traffic, despite filtering out all the traffic of clients whose ISP is not an SPM member.

A. Model and Formulation

Our aim is to evaluate the relative benefits of SPM, in comparison to current methods. Our interest is in the amount of damage caused to domain i as a result of attacks performed in the Internet, as a function of the defense mechanisms deployed by the various domains. In particular we will be interested in evaluating the amount this damage to the server traffic of domain i . We will conduct this evaluation under i) the *No defense* approach, ii) the *ingress/egress filtering* approach and iii) the *SPM* approach. We will use this evaluation to compare the relative merits of the different approaches. Below we use a simplistic model aiming at the comparison.

We assume that the Internet consists of N domains, indexed $1, 2, \dots, N$. Let $INT = \{1, 2, \dots, N\}$ denote this set. Each of the domains is in charge of the traffic originating at the domain and destined at the domain. Each domain is economically responsible for the servers and the clients operating at the domain and thus its objective is to provide good quality of service to these clients and servers. It should be noted that for large portion of the Internet traffic the client may reside in domain i while the server resides in domain j , in which case there are mutual interests of domains i and j to provide good quality for the traffic between the corresponding client and server.

Our focus in this analysis is on demonstrating the benefits of SPM with respect to the server traffic of the domains, which is perhaps the primary objective of the domains. Additional benefits are derived by SPM domains, due to their client traffic. We describe the properties of these benefits (see Section IV-G) but do not provide its exact analysis since it depends on the exact verification and discard algorithm used by the servers.

Let $A_{i \rightarrow j}^{(k)}$ be the rate of attacks performed from domain i to domain j where the address of i is spoofed

to an address in domain k . Let $A_{\rightarrow i} = \sum_{k=1}^N \sum_{j=1}^N A_{j \rightarrow i}^{(k)}$ denote the total attack rate directed at domain i .

We will focus on domain i and aim at deriving the amount of *damage* inflicted on servers placed in the domain. We will denote this D_i^{server} and we will evaluate it through the rate of attack traffic that *reaches* the domain. For each of the defense mechanisms we will then evaluate the *damage reduction*, denoted DR_i^{server} , measured via the *reduction in attack traffic rate*. In the comparison we will deal with the *relative damage reduction* which is $DR_i^{server} / D_i^{server}$.

B. Damage (attack rate) under No Defense

The total damage to domain i is given by the overall attack rate aiming at the domain:

$$D_i^{server} = \sum_{k=1}^N \sum_{j=1}^N A_{j \rightarrow i}^{(k)} = A_{\rightarrow i}. \quad (1)$$

C. Damage Reduction (attack rate reduction) under ingress/egress Filtering Defense

In this section we assume that a set of domains denoted $IGR \subseteq \{1, 2, \dots, N\}$ conducts ingress/egress filtering. That is, each of these domains does not allow to send out traffic whose source address is spoofed. Recall that INT denotes the set of all Internet domains.

The damage reduction of domain i , as experienced by the servers residing in the domain, denoted DR_i^{server} is the reduction in attack rate arriving at the domain. This is given by:

$$DR_i^{server} = \sum_{j \in IGR} \sum_{k \in INT} A_{j \rightarrow i}^{(k)}. \quad (2)$$

This means that the damage reduction, due to the existence of ingress/egress filtering at the set IGR , is *identical* to domains which have implemented ingress/egress filtering and to those that did not implement it. Thus, accounting to its server traffic, domain i has no incentive implementing ingress/egress filtering.

D. Damage Reduction Under ingress/egress Club Defense

One may attempt to overcome the deficiencies of the ingress or egress filtering model by creating an "ingress/egress Club". Under the ingress/egress Club model, the domains which implement ingress/egress filtering conduct ingress/egress filtering exclusively to traffic destined to domains in IGR . This is done in order to increase the relative benefit provided to members of IGR compared to non-members (as to increase the incentives of non members to become members). Let

$IGRCLUB$ denote the set of domains which participate in the ingress/egress club.

The damage reduction to server traffic is given by:

$$DR_i^{server} = \sum_{j \in IGRCLUB} \sum_{k \in INT} A_{j \rightarrow i}^{(k)} \quad i \in IGRCLUB \quad (3)$$

$$DR_i^{server} = 0 \quad i \notin IGRCLUB. \quad (4)$$

Note that it is relatively advantageous for domain i to belong to the ingress/egress club.

E. Damage Reduction under SPM Defense

As described earlier, domain i , when participating in SPM , marks all the packets leaving the domain with a special cookie. This cookie can be uniquely recognized by the SPM participants and helps in identifying the authenticity of the packet source. Thus, all partners of the SPM can recognize the authenticity of the packets generated by other partners of SPM . To analyze the benefits of the SPM method we will further assume that the partners of the SPM do treat the SPM produced and authenticated packets at higher priority.

The damage reduction to the servers of domain i is expressed in two ways. First it includes all attacks where the spoofed address belongs to SPM . Second, it includes all attacks generated by domains in SPM

(since that traffic is tagged with a wrong key).

$$DR_i^{server} = \sum_{j \in INT} \sum_{k \in SPM} A_{j \rightarrow i}^{(k)} \quad (5)$$

$$+ \sum_{j \in SPM} \sum_{k \in INT-SPM} A_{j \rightarrow i}^{(k)} \quad i \in SPM$$

In contrast, a domain that is not in SPM will not benefit. This yields:

$$DR_i^{server} = 0; \quad i \notin SPM \quad (6)$$

If SPM selects to conduct ingress/egress filtering then Equation 5 does not change while Equation 6 changes to:

$$DR_i^{server} = \sum_{j \in SPMIE} \sum_{k \in INT} A_{j \rightarrow i}^{(k)} \quad i \notin SPMIE, \quad (7)$$

where $SPMIE$ is the set of node conducting SPM (with ingress/egress filtering). However, if SPM selects to conduct ingress/egress filtering in a "club mode", then the damage reduction of non-members is identical to Equation 6.

F. The Benefits of SPM: A Comparison to Other Methods

To demonstrate the benefits of SPM , we next consider several special case scenarios.

1) *Fully Symmetric System*: First, let us consider a fully symmetric system. In this system let $A_{i \rightarrow j}^{(k)} = A/N^3$ for all $1 \leq i, j, k \leq N$. Also let us assume that the size of each of the defense sets, IGR , $IGRCLUB$, SPM and $SPMIE$, is given by K (that is, the number of domains that implement the defense is K).

Under these parameters we compare the relative damage reduction as follows:

1) Under no defense (from Equation 1):

$$D_i^{server} = \frac{A}{N^2}. \quad (8)$$

2) Under ingress/egress filtering (from Equation 2):

$$DR_i^{server} / D_i^{server} = \begin{cases} \frac{K}{N} & i \in IGR \\ \frac{K}{N} & i \notin IGR \end{cases} \quad (9)$$

3) Under Ingress/Egress Club:

$$DR_i^{server} / D_i^{server} = \begin{cases} \frac{K}{N} & i \in IGRCLUB \\ 0 & i \notin IGRCLUB \end{cases} \quad (10)$$

4) Under SPM or Under SPM plus Ingress/Egress Club:

$$DR_i^{server} / D_i^{server} = \begin{cases} \frac{2K}{N} - \frac{K^2}{N^2} & i \in SPM \\ 0 & i \notin SPM \end{cases} \quad (11)$$

5) Under SPM plus ingress/egress filtering:

$$DR_i^{server} / D_i^{server} = \begin{cases} \frac{2K}{N} - \frac{K^2}{N^2} & i \in SPMIE \\ \frac{K}{N} & i \notin SPMIE \end{cases} \quad (12)$$

2) *Discussion*: These results are depicted in Figure 2. In this figure we assume that the number of ISP's in the Internet is 10,000 and depict the relative damage reduction (evaluated via the attack rate reduction) for an ISP as a function of the number of ISP's participating in the protection system. We do this for Ingress/Egress Filtering and SPM with Ingress/Egress.

The results demonstrate clearly the following properties:

- 1) Under ingress/egress Filtering there is no added value for a domain to conduct ingress/egress filtering. The relative benefit for a participant is identical to that of a non-participant (both for the client traffic and for the server traffic).
- 2) The ingress/egress Club provides some *relative benefit* to its participants ; that is, the benefit to participants is larger than to non participants. Note however, that when the club is of small size, the benefit is relatively small and thus there is little incentive to join the club when it is small. Note also that the benefit is only relative since it

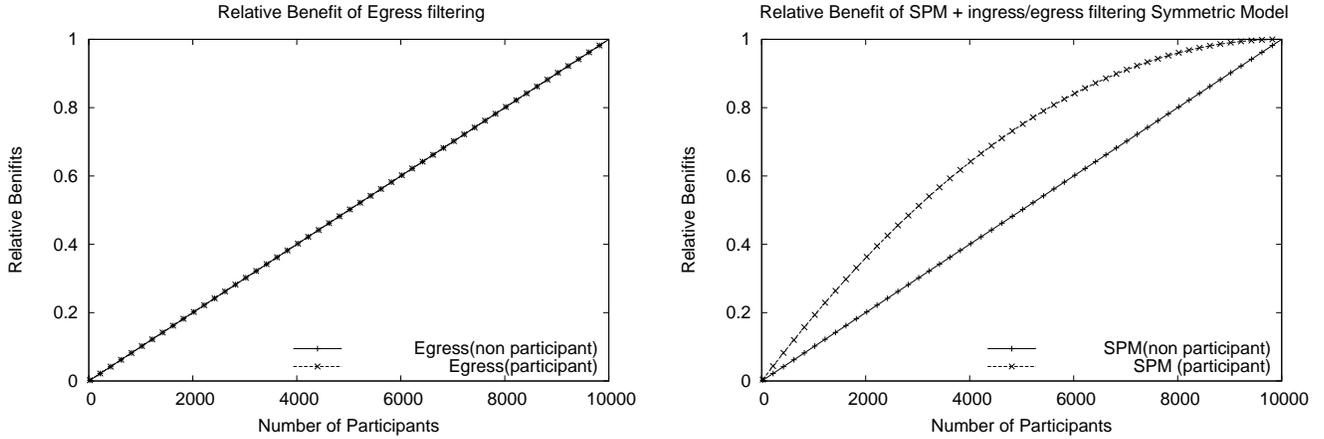


Fig. 2. The Benefits of SPM plus Ingress/Egress filtering under symmetric traffic (compared to Ingress/Egress Filtering)

is achieved by eliminating the benefit from non-participants.

- 3) SPM provides significant benefit to its participants. The benefit is always larger than that of Egress Club; further, when the number of participants is small (small club) this benefit is roughly twice as high as that of the Egress club.
- 4) The reader should recall that this comparison accounts only to parts of the benefits of SPM (see discussion at the beginning of IV). Thus the benefits of SPM are in practice higher.

3) *An Asymmetric System.* Next we consider an asymmetric system, namely where the domain sizes are not all identical and so is the traffic generated by them. To this end we follow some of the analysis conducted in the field and assume that the domain size is distributed in a Zipf like distribution. Zipf law distributions have been shown to closely approximate the relative size of populations and in analogy they may properly reflect the relative sizes of ISP's (domains) In similar contexts, a Zipf law distribution has been used in [26] to model the relative frequencies of Web pages. The use of a Zipf law distribution means that if X_i is the size of domain i , $i = 1, \dots, N$, then X_i is proportional to $1/i^\alpha$, where α is close to unity. For simplicity we assume that $X_i = X/i$ for some constant X .

For the sake of conciseness, the analysis of this system is given in the appendix VI-C. The results of this system are depicted in Figure 3 which demonstrates the following properties:

- 1) For ingress/egress filtering the benefits of joining IGR are identical to those in the symmetric system: The benefit of being on IGR grows very slowly with the size of IGR, and *there is no relative advantage* to be in IGR (as opposed to not be in it).

- 2) The benefit for participating domains grows very rapidly with the SPM size. Thus, the benefit a domain derives from joining SPM are very significant even if the SPM conducts ingress/egress filtering, and thus domains have very strong incentives to join SPM. This results from the fact that a large fraction of the attacks and the spoofing are directed to the large domains (those indexed with low indices) as they carry the major mass of Internet traffic.

G. Client Traffic

The clients of an SPM domain do also benefit from SPM, due to receiving preferred treatment at SPM server domains. Due to the lack of space, this discussion is not provided here, and is given in the appendix VI-D.

H. The Theory in Practice

Below we demonstrate the analysis developed above in the context of the practical Internet. To be specific we estimate the relative benefit that results from constructing SPM of the ten largest ISP's. To conduct this evaluation one needs to find the relative traffic volume sent by (and to) each of the domains, a data that is hard to obtain. Nonetheless, one can estimate the relative size of a domain by the number of IP addresses it holds. This number for the 10 largest networks is provided in [27] and is listed in Table 1 (where for each AS are given the Rank, the size of the IP space, the AS number, and the name of the AS). The total size of the IP space used in the Internet is estimated in [27] as 1,507,993,620. Using these numbers one may estimate the highest rank AS to occupy roughly 9.0% of the Internet, and the 10 highest rank AS's to occupy together roughly 27.8% of the Internet.

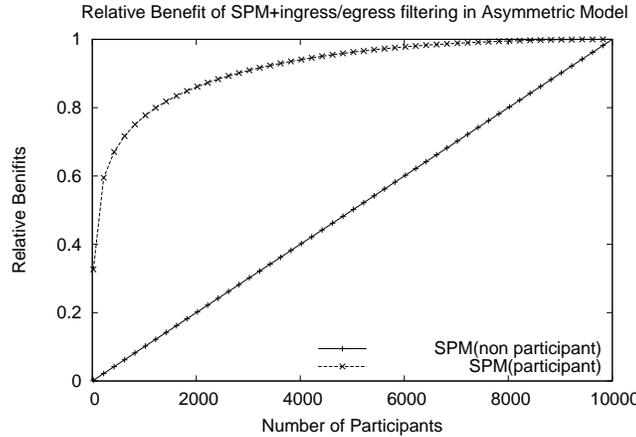


Fig. 3. The benefits of SPM plus Ingress/Egress under Asymmetric traffic

Rank	IP Space	ASN	Description
1	136,231,798	3303	Swisscom Enterprise Solutions
2	72,756,483	568	DISO-UNRRA
3	44,547,464	3356	Level 3 Communications, LLC
4	35,468,904	701	UUNET Technologies, Inc.
5	30,133,936	7132	SBC Internet Services
6	28,986,765	7018	ATT WorldNet Services
7	21,558,912	16631	Cogent Communications
8	18,576,327	237	Merit Network Inc.
9	18,512,345	2381	University of Wisconsin-Madison
10	18,047,631	7474	Optus Communications Pty

We can now compare the traffic model developed in Section IV-F.3 to these numbers. To this end, we recall the estimate that the Internet consists of roughly $N = 10,000$ AS's. Thus, the model developed in IV-F.3 is based on the assumption that the relative size (and the relative amount of traffic sent to) the largest 10 As's ($K = 10$) is given by

$$\frac{\sum_{i=1}^{10} (1/i)}{\sum_{i=1}^{10,000} (1/i)} \approx \ln(10)/\ln(10,000) = 0.25. \quad (13)$$

This is indeed very close (from below) to the relative size obtained from [27], implying that the model developed in Section IV-F.3 (asymmetric traffic) quite accurately describes today's Internet.

Thus if one wishes to estimate the relative benefit that SPM will grant to its members, when SPM consists of only the 10 largest AS's, one could apply Equation 22 with $K = 10$ and $N = 10000$, yielding relative damage reduction of roughly 25%. Further, if the 100 largest AS's cooperate in SPM, the relative damage reduction increases to roughly 50%.

V. CONCLUDING REMARKS

We examined the packet-spoofing problem of the Internet and recognized that today's technological solutions

(ingress filtering) are economically ineffective since a network, say X , that invests in such technology benefits very little from it (while the major benefit spreads over all networks other than X).

We proposed SPM as an alternative method to be used by networks routers as to eliminate or reduce spoofing attacks. The method can be implemented using a simple key mechanism to be used by the participants of SPM. Further, SPM is upper-compatible to today's Internet and can be implemented in a transparent fashion when only part of the networks deploy it. We analyzed the benefits of SPM and demonstrated that it forms an economical attractive solution since a network that elects to deploy it can derive significant relative benefits to its servers as well as to its clients, and thus has the incentive to invest in its deployment. Further, these benefits are significant even when SPM is only deployed by a fraction of the Internet networks and even if it is deployed without ingress/egress filtering.

REFERENCES

- [1] David Moore, Geoffrey M. Voelker, and Stefan Savage, "Inferring internet Denial-of-Service activity," pp. 9–22.
- [2] Cisco IOS, "Unicast reverse path forwarding," 1999.
- [3] P. Ferguson and D. Senie, "Network ingress filtering: Defeating denial of service attacks which employ ip source address spoofing," Tech. Rep., IETF, May 2000, RFC 2267.
- [4] T. Killalea, "Recommended internet service provider security services and procedures," Tech. Rep., November 2000, RFC 3013.
- [5] Stephen Kent, Charles Lynn, and TITLE = Karen Seo, , " .
- [6] Russ White, "Architecture and deployment considerations for secure origin bgp (sobgp)," Tech. Rep., IETF, <ftp://ftp-eng.cisco.com/sobgp/drafts/draft-white-sobgp-architecture-00.txt>.
- [7] Geoffrey Goodell, William Aiello, Timothy Griffin, John Ioannidis, Patrick McDaniel, and Aviel Rubin, "Working around bgp: An incremental approach to improving security and accuracy of interdomain routing," .

- [8] Lakshminarayanan Subramanian, Volker Roth, Ion Stoica, Scott Shenker, and Randy H. Katz, "Listen and whisper: Security mechanisms for bgp," in *NSDI*, 2004.
- [9] S. Deering and R. Hinden, "Rfc 2460 - internet protocol, version 6 (ipv6) specification," Tech. Rep., IETF.
- [10] Y. Rekhter, B. Moskowitz, D. Karrenberg, and G. J. de Groot and, "Address allocation for private internets," 1996.
- [11] "The bogon reference page," <http://www.cymru.com/Bogons/index.html>.
- [12] Stefan Savage, David Wetherall, Anna R. Karlin, and Tom Anderson, "Practical network support for IP traceback," in *SIGCOMM*, 2000, pp. 295–306.
- [13] Steve Bellovin, Marcus Leech, and Tom Taylor, "Icmp traceback messages," Tech. Rep., February 2003, <http://www.ietf.org/internet-drafts/draft-ietf-itrace-04.txt>.
- [14] Matt Franklin Drew Dean and Adam Stubblefield, "An algebraic approach to ip traceback," February 2001.
- [15] Dawn X. Song and Adrian Perrig, "Advanced and authenticated marking schemes for ip traceback," 2001.
- [16] S. Felix Wu, Lixia Zhang, Dan Massey, and Allison Mankin, "Intension-driven icmp trace-back," .
- [17] A. C. Snoren, C. Partridge, L. A. Sanchez, C. E. Jones, F. Tchakountio, S. T. Kent, and W. T. Strayer, "Hash-based ip traceback," .
- [18] John Ioannidis and Steven M. Bellovin., "Implementing push-back: Router-based defense against ddos attacks," February 2002.
- [19] Cisco IOS, "Configuring tcp intercept (prevent denial-of-service attacks)," 1997.
- [20] Guy Pazi, Anat Bremner-Barr, Rami Rivlin, and Dan Touitou, "Protecting against distributed denial of service attacks," 2002, Patents (PCT/IL02/00714, 10/232,993).
- [21] Cheng Jin, Haining Wang, . and Kang G. Shin, "Hop-count filtering: An effective defense against spoofed ddos traffic," in *ACM Conference on Computer and Communications Security*, 2003.
- [22] Abraham Yaar, Adrian Perrig, and Dawn Song, "Pi: A path identification mechanism to defend against ddos attacks," in *IEEE Symposium on Security and Privacy*, 2003.
- [23] Anat Bremner-Barr and Hanoch Levy, "Spoofing prevention method (brief announcement)," in *PODC 2004*, 2004.
- [24] David Moore Colleen Shannon and k claffy, "Characteristics of fragmented ip traffic on internet links," in *Internet Measurement Workshop*, 2001.
- [25] V. Fuller, T. Li, J. Yu, and K. Varadhan, "Classless iner-domain routing (cidr)," Tech. Rep. RFC 1519, IETF, 1993.
- [26] V. Almeida, A. Bestavros, M. Crovella, and A. de Oliveira, "Characterizing reference locality in the www," *International Conference in Parallel and Distributed Information Systems*, December 1996.
- [27] "Fixedorbit - the internet from the inside : General statistics," Tech. Rep., fixedorbit, <http://www.fixedorbit.com/stats.htm>.
- [28] Cisco IOS, "Netflow switching overview," 2000.
- [29] J. Hawkinson and T. Bates, "Guidelines for creation, selection, and registration of an autonomous system (as)," 1996.
- [30] Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang, "An analysis of bgp multiple origin as (moas) conflicts," in *Internet Measurement Workshop 2002*.
- [31] Xiaoliang Zhaom, Allison Mankin, Daniel Massey, Dan Pei, Lan Wang, . S. Felix Wu, and Lixia Zhang, "Validation of multiple origin ases conflicts through bgp community attribute," Tech. Rep., Internet draft, November 2001, draft-zhao-idr-moas-validation-00.txt.

The material in the Appendix is provided for the reviewer's convenience and is not part of the full paper.

A. Changing the SPM keys periodically

There are (at least) two alternative ways to achieve that:

- 1) Simply, each AS-server periodically selects a new set of random keys and distribute them to the other AS-servers as described above. To distribute the associated computation and communication overhead, each AS server chooses the time to replace its old key in an independent way. Notice that at the time of the key replacement each router should hold two keys, the old and the new one. Therefore, the key table at each router has two keys for each source AS and an arriving packet is authenticated if the key by which it is tagged equals to one of the two (i.e., each key is actually valid for two periods of time).
- 2) Another option is to associate a pseudo random number generator with each AS-server. Then, AS-server D , would compute the AS-in key for each other AS-server when D is the destination, in predefined times. E.g., every two hours each server picks the next pseudo random number for each other AS-server. One possible implementation of this, is for the AS-servers to exchange one time (could be out of band) a set of secret keys, and use the corresponding keys to encrypt the time of the day every two hours.

B. Mapping AS numbers to prefixes

The AS server distributes the key information to the routers inside its local AS. The routers need to construct tables that store for each network (prefix) its corresponding key, since a router in its forwarding table can match an address, to its network, but not to its AS. Hence, either the AS-server, or the router have to perform the mapping from a **AS-out table** and **AS-in table** into **Network-in table** and **Network-out table**.

Since the routers FIB have all the necessary information to carry out this mapping we assign this job to the routers. Some routers today have the mapping information ready since they use it for given aggregated Netflow reports, that sum up the amount of traffic between two ASes that passes through the router [28]. A router can easily deduce the mapping from prefixes to ASes from the AS-path in the BGP announcement.

The last AS in the AS path of a BGP announcement of a prefix is usually the AS of that contains (owns) the prefix. One, exception is the multi-homing case, in which the last AS in the AS-path, is one of the ISPs that multi-homed the network. This might be problematic to SPM, since last AS in the AS path indicates the ISP that is used for receiving incoming traffic which can be different than the ISP that is used for the corresponding outgoing traffic.

While, according to RFC 1930 [29], this situation, should not happen in multi-homing network, since a multi-home network should use a separate AS, in reality there are cases where a network is multi-homed, without being a separate AS. However, this is not a common scenario [30]. In the SPM architecture we allow the multi-homed prefix to be tagged by one of the keys corresponding to one of its ISP's. However, it is required, that each SPM authenticating router, in the Internet would know the multiple AS's that are entitled to announce the prefix readability as suggested in [31] ⁸.

C. Analysis of an Asymmetric System

We will assume that the amount of server traffic is proportional to the domain size, that is $S_i = S/i$ for some constant S .

The attack rates produced at the domains could be assumed also to be proportional to the domain size. However, aiming to be on the conservative side, we tend to assume that this is not the case, as larger domains seem to be more commercial and more protected than smaller domains and thus attacks are likely to be generated at smaller domains. To compensate for both factors we will assume that the attack rate, produced at domain i is the same for all domains, that is $A_i = A/N$ for some constant A . The targets of the attack are assumed to be proportional to the target size.

$$A_i = \frac{A}{N}; \quad A_{\rightarrow i} \approx \frac{A}{i \ln N}; \quad A_{j \rightarrow i} \approx \frac{A}{N} \frac{1}{i \ln N}; \quad (14)$$

Also it is assumed that the relative volume of spoofed addresses is proportional to the traffic of the spoofed address, as the attacker aims at approximating the reality on the network (to better hide its traffic). That is:

$$A_{j \rightarrow i}^{(k)} \approx \frac{A}{N} \frac{1}{i k \ln^2 N}. \quad (15)$$

Under these assumptions we get the following: **Under no defense (Equation 1)**

$$D_i^{server} \approx \frac{A}{i \ln N}; \quad (16)$$

⁸In [31] The multiple AS's that are entitled to announce the prefix, are added as BGP community to the BGP announcement that announces the prefix

Under ingress/egress filtering (Equation 2):

$$\begin{aligned} DR_i^{server} / D_i^{server} &\approx \left(\frac{A}{N \ln^2 N i} K \ln N \right) / \left(\frac{A}{i \ln N} \right) \\ &= \frac{K}{N}; \quad \forall i \in INT \end{aligned} \quad (17)$$

Under SPM with ingress/egress filtering: We have (from Equation 5, 7):

$$\begin{aligned} DR_i^{server} &\approx \frac{A}{N \ln^2 N i} \quad (18) \\ &\left(\sum_{j \in INT} \sum_{k \in SPM} \frac{1}{k} \sum_{j \in SPM} \sum_{k \in INT-SPM} \frac{1}{k} \right) \\ &\quad i \in SPMIE \end{aligned} \quad (19)$$

and

$$DR_i^{server} \approx \sum_{k \in INT} \sum_{j \in SPM} \frac{A}{N \ln^2 N i k} \quad i \notin SPMIE. \quad (20)$$

Assuming that SPMIE consists of the K largest domains, that is $SPM = \{1, 2, \dots, K\}$, we get:

$$\begin{aligned} DR_i^{server} &\approx \quad (21) \\ &\begin{cases} \frac{A}{N \ln^2 N i} (N \ln K + K \ln N - K \ln K) & i \in SPMIE \\ \frac{A}{N \ln^2 N i} (K \ln N) & i \notin SPMIE \end{cases} \end{aligned}$$

And the *relative damage reduction* is given by

$$DR_i^{server} / D_i^{server} \approx \begin{cases} \frac{N \ln K + K \ln N - K \ln K}{N \ln N} & i \in SPMIE \\ \frac{K}{N} & i \notin SPMIE \end{cases} \quad (22)$$

Note, that in the case of SPM without ingress/egress filtering the only change is that the ASes which do not participate in SPM do not enjoy any damage reduction.

D. Client Traffic

The clients of an SPM domain do also benefit from SPM, due to receiving preferred treatment at SPM server domains. To understand this consider a client from domain l that attempts reaching a server in domain i . If $i \in SPM$ then the traffic of l may benefit, depending on the discard strategy taken by i (as described in Section III-C), namely whether it is a *conservative discard* or an *aggressive discard*. Recall that in the former only SPM-recognized spoofed traffic is discarded, while in the latter, in addition some (or all) non-SPM traffic is discarded as well.

In the conservative case the client of l benefits at i due to having i performing better; however, in this case, the client does not derive any relative benefit compared to non-SPM clients, since those experience the same benefit.

The major relative benefit of a client of l is derived in the aggressive case. In that case the client traffic of l arrives uninterrupted to i (if l is in SPM). In contrast, if l is not in SPM its traffic may be discarded even if it is legal traffic.

As these benefits depend on the specific algorithm used by i and on its implementation, their exact analysis is beyond the scope of this work. Nonetheless, one can make the following observations:

- 1) When SPM consists of many members, the defense used by an attacked server (i) may be of the conservative type, and then the SPM client (l) derives little relative advantage. This relative advantage is, nonetheless, not critical since domain l already derives large advantage due to its servers (see Figures 2 3).
- 2) When SPM consists of a small number of members, the defense used by i is likely to be of the aggressive type (otherwise it may not sustain the attack), and then the clients of domain l derive large relative advantage. This extra relative benefit to l occurs exactly when it is needed, namely when the relative advantage to its servers is relatively small.

We therefore may conclude that the benefits to the domain clients complements the benefit to the domain servers, and increases the incentive of joining SPM, when it is most needed, namely at its very early formation stages (when its number of members is small).