Tracing Cyber Attacks from the Practical Perspective

Zhiqiang Gao and Nirwan Ansari

ABSTRACT

The integrity of the Internet is severely impaired by rampant denial of service and distributed DoS attacks. It is by no means trivial to devise a countermeasure to address these attacks because of their anonymous and distributed nature. This article presents a brief survey of the most promising recently proposed schemes for tracing cyber attacks: IP traceback. Since IP traceback technology is evolving rapidly, for the community to better comprehend and capture the properties of disparate traceback approaches, we first classify these schemes from multiple aspects. From the perspective of practicality and feasibility, we then analyze and explore the advantages and disadvantages of these schemes in depth so that shortcomings and possible enhancements of each scheme are highlighted. Finally, open problems and future work are discussed, and concluding remarks are drawn.

INTRODUCTION

The integrity of the Internet is seriously threatened by denial of service (DoS) and distributed DoS (DDoS) attacks, which aim to disrupt legitimate users from accessing a certain resource, such as a server or network. In a DoS/DDoS attack, a violator normally bombards the victim with a huge number of packets. Due to the stateless nature of the Internet and prevalence of attack tools, it is very easy for a hacker, or even a kid, to mount an attack with a very small chance of being caught. This may explain why DoS/DDoS attacks are rampant in the Internet.

Up to now, a vast amount of schemes have been proposed as countermeasures against DoS/DDoS attacks. These schemes can be roughly categorized into four groups: intrusion prevention, intrusion detection, intrusion mitigation, and intrusion response [1]. This article focuses on IP traceback, which belongs to the fourth group.

The objective of IP traceback is to locate the actual source of attack packets [2]. Ideally, an IP traceback scheme should be capable of identifying the real attacker. However, given the extreme complexity of the current Internet, it is difficult for the victim to ascertain the attack source in a DoS attack because the attacker routinely forges the source IP address of each attack packet. It is even harder to retrieve the sources of a DDoS attack because many attack sources are widely dispersed in the Internet and there is no apparent feature of a DDoS stream that can be directly exploited by the victim.

It is also hard to grasp the global view of traceback schemes since the research on DoS/DDoS is evolving rapidly. To facilitate a better understanding of the field, we classify traceback schemes from several dimensions. Instead of a comprehensive survey, we select the typical schemes of each group, along with the latest developments. Different from previous work [3], we focus on the issue of practicality of traceback schemes. We believe that practicality is the utmost property to be considered for eventual deployment of IP traceback. From this standpoint, we thoroughly explore the pros and cons of selected schemes. Finally, challenges to be overcome are highlighted and possible solutions are discussed.

The rest of the article is organized as follows. First, various traceback schemes are classified from multiple aspects. The metrics we select to assess each scheme are then presented. From the perspective of practicality, the benefits and potential drawbacks of existing schemes are explored in depth, and latest developments and possible further enhancements are proposed. Finally, we discuss the challenges and future work, and summarize the article.

CLASSIFICATIONS

Unruly DoS/DDoS attacks motivate the research on IP traceback. Until now, many traceback approaches have been proposed. To better understand the advantages and disadvantages of different traceback schemes, we classify existing schemes from multiple disparate standpoints. We hope this work will lay down a foundation for developing more efficient and effective traceback schemes in the future.

As shown in Fig. 1, five aspects are selected to classify existing traceback schemes into different categories. They include the basic principle, processing mode, functionality supported, location, and requirement for extra infrastructure.

The schemes illustrated in Fig. 1 include Probabilistic Packet Marking (PPM) [2], ICMP

0163-6804/05/$20.00

There are two main shortcomings in outbound marking. Outbound marking needs extra bandwidth that may further aggravate the performance of the network being attacked. A new ICMP message must be introduced into the Internet, and there shall not be any ICMP filtering.

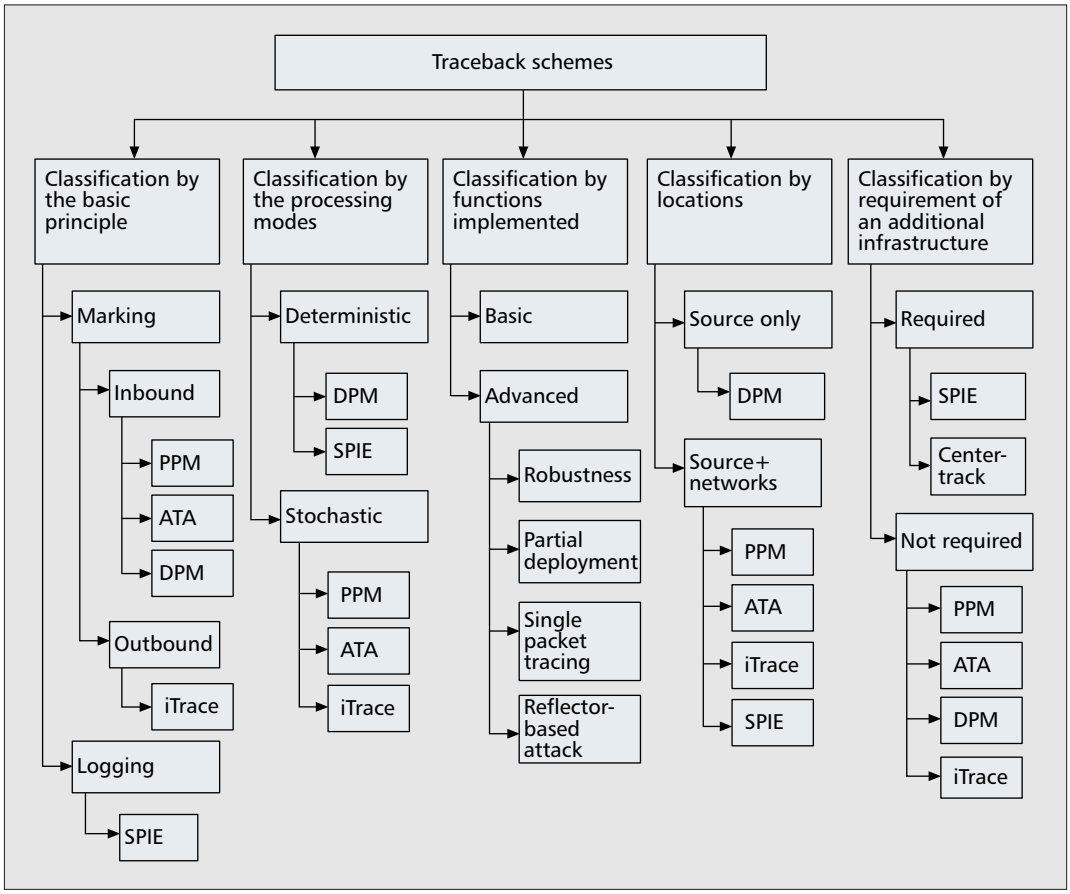

Figure 1. *Categorizing IP traceback schemes.*

traceback (iTrace) [4], Source Path Isolation Engine (SPIE, also called hash-based traceback) [5], Algebraic-Based Traceback Approach (ATA) [6], Deterministic Packet Marking (DPM) [7], and an overlay-based solution (Center-Track) [8].

Classification by the Basic Principle

According to the basic principle, most of the existing traceback schemes may be roughly categorized into two groups: marking and logging. In logging schemes, routers record some information of traversing packets for verifying whether suspected packets have been forwarded by a specific router or not. In marking schemes, a portion of or all routers along an attack path from an attack source to a victim write some information of these routers into the packets so that the attack path(s) may be recovered by the victim, even though the source IP addresses of attack packets are spoofed. The marking information may be inscribed in the same attack packets (called *inbound marking*) or extra ICMP packets (called *outbound marking*) [4]. Inbound marking does not require extra bandwidth, while the number of bits that may be used for marking is rather limited. Using the option field to store marking information is not preferable because it triggers a significant delay in processing the marked packets at routers. On the contrary, there are far more bits available for outbound marking than inbound marking, which may mitigate false positives and greatly reduce the number of marked packets required for reconstruction. There are two main shortcomings of outbound marking. First, extra bandwidth is needed, which may further aggravate the performance of the network being attacked. A new ICMP message must be introduced into the Internet, and there shall not be any ICMP filtering. Otherwise, the ICMP message may be blocked due to ICMP filtering.

Current traceback schemes based on marking include variants of PPM, ATA, DPM, and schemes that use ICMP messages, such as iTrace. Instances of logging schemes include SPIE and its variant [9].

Classification Based on the Processing Mode

From the viewpoint of the processing mode, traceback schemes may be categorized into two groups: deterministic and stochastic.

The deterministic mode implies that every packet has to be processed, either marking or logging. DPM is an example of deterministic marking, SPIE of deterministic logging. More stochastic schemes have been contrived. Three well-known examples are PPM (and many of its variants), ATA, and iTrace. Obviously, the deterministic mode incurs more processing overhead. However, it may perform single packet tracing. Furthermore, deterministic processing may be indispensable for more advanced security services, such as nonrepudiation. The probabilistic

mode is helpful to reduce bandwidth and processing overhead at the expense of increased complexity for reconstruction at the victim.

Classification by Functions Implemented

There is no panacea in IP traceback. Different tracing schemes make different assumptions and strive to solve different problems. In general, each tracing scheme has to make some tradeoffs between the performance and the overhead. In the marking schemes, factors that shall be taken into account include marking every packet or marking at a certain probability, the number of bits used for marking, the place to store the marking information, and parts of the networks (the routers, the victim, or both) that bear the incurred overheads. In the logging schemes, the issues to be addressed include the content to be logged, the frequency of logging, the place to store the logging information, and an efficient approach to communicate between the victim and the routers where the logging information is stored.

The functionalities that a tracing scheme supports may be further divided into two groups, basic functions and advanced functions. Obviously, the basic function is the ability to trace to the attack source under a DoS attack or hundreds of sources under a DDoS attack. The advanced proposals consider the following issues: the security of the scheme itself (e.g., authentication); the ability of tracing a single packet; the capability of tracing a reflector-based DDoS attack; the capability of being effective under partial deployment.

Classification by Locations

From the perspective of locations, existing traceback schemes may be divided into two types, i.e., those that inscribe information into the packets near the source, and in the network, respectively. DPM is an example that performs marking near the source (edge routers closest to the source). Most schemes work with the cooperation of the victim and the network. That is, the routers (some or all) in the network perform certain processing (marking or logging), either stochastically or deterministically, and inscribe required information into the packets. When these processed packets arrive at the victim, the victim may reconstruct the attack paths from the embedded information.

An associated issue with locations is whether the victim can reconstruct each path entirely or partially. Only recording the information of a single point is a special case of partial path information. Clearly, only single point information for each path may be provided for schemes performing marking at edge routers. In so doing, the most valuable information-the first edge router from which attack packets being mounted may be readily determined. Another benefit is that the victim is greatly relieved from the heavy computational and storage burden. However, the marking information may not be robust enough because of the lack of verifiability. If the first edge router along an attack path is compromised, no additional clue may be exploited. A tradeoff between the computational burden and the reliability is to record partial path information, e.g., storing the path information of traversing Autonomous Systems (ASs).

Classification by Requirement of an Extra Infrastructure

The current tracing schemes may also be differentiated according to whether an extra infrastructure is required. Here, we focus on additional facilities that are required for the sake of tracing rather than normal packet forwarding. An extra infrastructure refers to some additional facilities such as the Tracking Router (TR) used in CenterTrack, and SPIE Collection and Reduction Agents (SCAR) used in SPIE. In general, an extra infrastructure implies more financial investment and more management overhead; this is not attractive to the Internet Service Providers (ISPs). Note that although all traceback schemes expect certain modifications or function extensions to current facilities, especially routers, these modifications or extensions to existing devices are not considered as an extra infrastructure here. The instances that do not need an extra infrastructure include variants of PPM, iTrace, and DPM.

In the logging schemes, the issues to be addressed include the content to be logged, the frequency of logging, the place to store the logging information, and an efficient approach to communicate between the victim and the routers where the logging information is stored.

Evaluation Metrics

A number of metrics may be used to evaluate the performance of disparate traceback schemes, such as the minimum number of marked packets required for path reconstruction, processing burden, bandwidth overhead, memory overhead, robustness, scalability, ISP involvement, and so on [3]. In this article, we mainly assess disparate tracing schemes from the practical perspective. We thus consider the following criteria: the minimum number of marked packets required for path reconstruction, the computational burden, effectiveness under partial deployment, and robustness.

The Minimum Number of Marked Packets Required for Path Reconstruction — The minimum number of marked packets required for path reconstruction refers to the total number of marked packets that the victim needs to receive before it can complete the path reconstruction process. The less, the better. Note that the minimum number of marked packets is an essential and vital measure to compare different schemes. The fundamental goal of a traceback scheme is to locate each attack source. For schemes that perform marking only at the ingress edge routers, path reconstruction may be more straightforward since only single point information needs to be recovered for each path (e.g., DPM). For most other schemes, path construction is time-consuming and a huge processing overhead may be imposed on the victim.

The Computational Overhead — The computational overhead depends on several factors, such as processing packets stochastically or deterministically at the routers, inbound marking or outbound marking, and a DoS or a DDoS attack. A good design will attempt to minimize

The effectiveness of a traceback approach under partial deployment is an important factor to be considered. When a scheme is devised, issues related to partial deployment shall be taken into account.

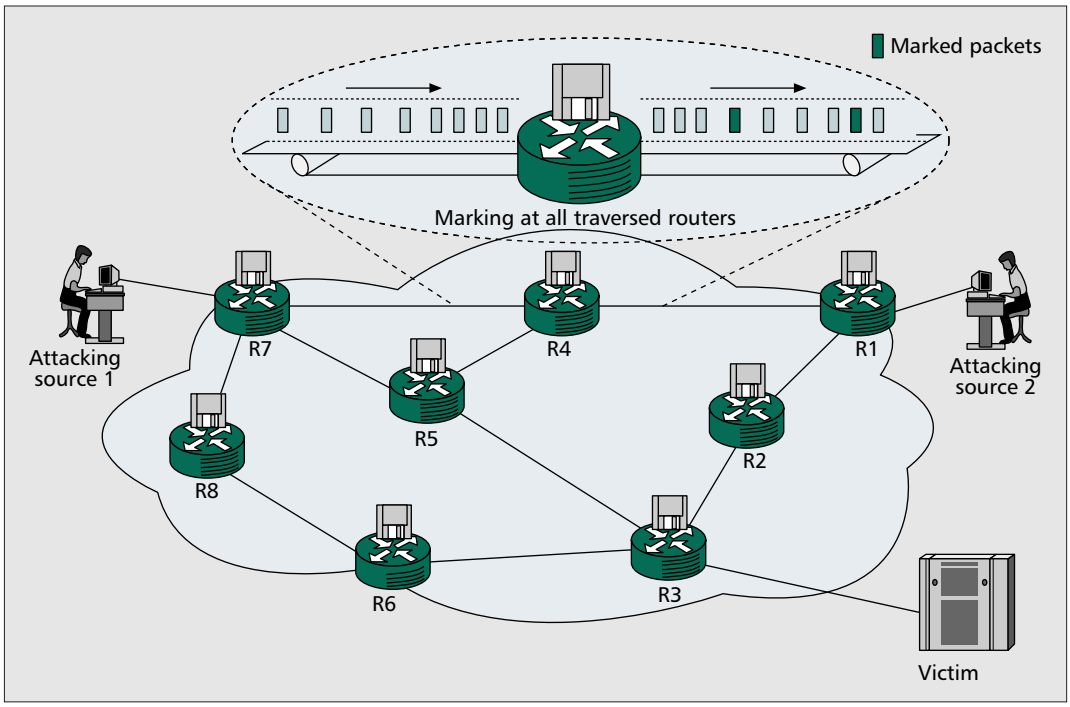

Figure 2. *Probabilistic packet marking (PPM).*

the computational burden on the victim. If an overwhelming computation is required, it may take the victim too long time to complete the path reconstruction process; this is definitely not a preferential choice.

Effectiveness under Partial Deployment — The distributed nature of the Internet renders deployment a big issue. First, the Internet Service Providers (ISPs) may lack incentives to deploy a scheme. Deploying a new scheme may imply more investment, more operational costs, and higher management complexity. Second, it may take a long time for a new scheme to be totally adopted in the Internet. Therefore, the effectiveness of a traceback approach under partial deployment is an important factor to be considered. When a scheme is devised, issues related to partial deployment shall be taken into account.

Robustness — In terms of robustness, we refer to the ability of an approach that can perform tracing reliably even under adverse conditions. When the stochastic mode is selected, it is critical to effectively process packets so that all information required for path reconstruction is reliably conveyed to the victim and that false positives are efficiently thwarted. Besides false positives incurred in processing (marking or logging) and reconstruction, it may be possible that some sophisticated attackers embed forged marking to amplify the false positives. Subverted routers are another issue to be addressed.

Evaluation of Schemes

According to the above criteria, we evaluate the following schemes: variants of PPM, iTrace, DPM, SPIE, and CenterTrack.

Variants of PPM

Basic PPM-PPM (Fig. 2) was developed by the ingenious work of Savage *et al.* [2]. The basic idea of PPM is simple. Suppose that one attack flow from an attack source to the victim traverses routers R_1, R_2, …, R_d in order. All routers employ the same marking probability which is denoted as p. For router R_i $(1 \le i \le d)$, with respect to the victim, the probability of packets marked by R_i is $p(1–p)^{d–i}$. Note that this value is different from p. The reason is that subsequent routers may "re-mark" packets already marked by previous ones, thus overriding the marking information of previous routers. Therefore, the closer a router is to the victim, the more chance its marking survives.

To handle DDoS attacks, the edge-sampling method was proposed. The detailed marking procedure at each router is depicted in Fig. 3, in which the attack traverses Routers R_1, R_2, and R_3. Each router may make the decision whether to mark the current packet or not independently. The left box in Fig. 3a shows the case that router R_1 marks packets, and the unmarked case is presented in the right box. In Fig. 3b, the left two boxes show the scenario that packets have been marked by router R_1. The upper box stands for the scenario that router R_2 "re-marks" these packets while the bottom one does not. Similarly, the right two boxes represent those packets that have not been marked by router R_1. The upper right box represents that packets have been marked by R_2, while the bottom one not marked by R_2.

Figure 3c presents the net result after combining the two boxes with the same marking. Using the similar procedure, the final result (what the victim receives) is shown in Fig. 3d.

It is easy to explain the reconstruction procedure from Fig. 3d. The victim first locates the

closest router, R_3, by looking at the packet whose dist field has a value of 0. Next, from the packets with dist=1, it can locate R_2. To save space, a new field called addr is used instead of the start field and end field as shown in Fig. 3, and its content is the result of executing the exclusive or (XOR) operation over the start and end fields. From the first step, we obtain the value of R_3; from the second step, we determine the value of $(R_2 \oplus R_3)$ [2]. Since $R_3 \oplus (R_2 \oplus R_3) = R_2$, R_2 may be located by using XOR. The procedure is repeated until the farthest router is reached.

Analysis of PPM — PPM possesses several nice features, such as low router overhead, support of incremental deployment and "post-mortem" tracing. However, several deficiencies severely impede its performance.

1) Heavy computational load for path reconstruction. When there are 25 attack sources, path reconstruction will take days and thousands of false positives may be generated [10]. Currently, an attacker may orchestrate thousands of zombies. As a result, the victim may never be able to complete the path reconstruction procedure. The daunting computational burden is caused by combinatorial explosion, which is originated from insufficient number of bits for marking.

2) High false positives. One source of false positives is the limited number of bits for marking. Another is rooted in the reconstruction algorithm. When there exists a large number of attack paths, the victim may be confused because many routers along different paths may be at the same distance to the victim.

3) Spoofed marking. The attacker may inscribe spurious marking in such a way that the victim receives more packets with forged marking information than those with the correct one. As a result, the victim will have little opportunity to resolve the attack paths.

4) Unawareness of the path length in advance. When a router decides to mark a packet, it has no idea of the path length, d. Therefore, it is incapable of setting p to the optimal value $1/d$. One possible choice is to use the recommended value, i.e., 0.04 [2]. If there are many attack paths with disparate lengths, simply using a predetermined marking probability for all paths may seriously degrade the performance.

5) Subverted routers. Few measures have been taken to defend against malfunctioned or subverted routers. Some subverted routers may be triggered by misconfigurations, and others may be resulted from internal vulnerabilities. Note that subverted routers may also generate spoofed marking. Up to now, few schemes can contain this problem.

Accordingly, PPM is a good solution for DoS attacks and small-scale DDoS attacks. It is not suitable for large-scale DDoS attacks.

Recent Development and Possible Solutions — Song *et al.* [10] proposed an advanced and authenticated PPM based on the assumption that the victim knows the mapping of the upstream routers. Their scheme can mitigate Problems 1 and 2, and effectively address Problem 3 as well. Another method to partially thwart Problem 1 is to use varying marking probability at each router [11]. The exact value of marking probability at each router depends on the hop counts between the current router and the victim. A recent work done by Tseng et al. [12] used counters to complement the loss of marking information from upstream routers. Their scheme may address Problems 1 and 3, and decrease false positives. Note that Problems 1 and 2 are related. In general, an approach that may alleviate the computational overhead is helpful to moderate false positives.

Problem 4 may not be easily resolved in the IP layer. However, it is possible for a router to know the value of d at the Autonomous System (AS) level. Schemes working at the AS level have the potential to address Problems 1 to 5 while it may only provide incomplete path information rather than hop-by-hop path information.

Problem 5 is more difficult to resolve. A good scheme shall neglect the marking information from compromised routers while a better solution is to contrive a mechanism so that the correctness of marking information embedded by the upstream routers can be verified. Unfortunately, no scheme possesses this advanced feature yet.

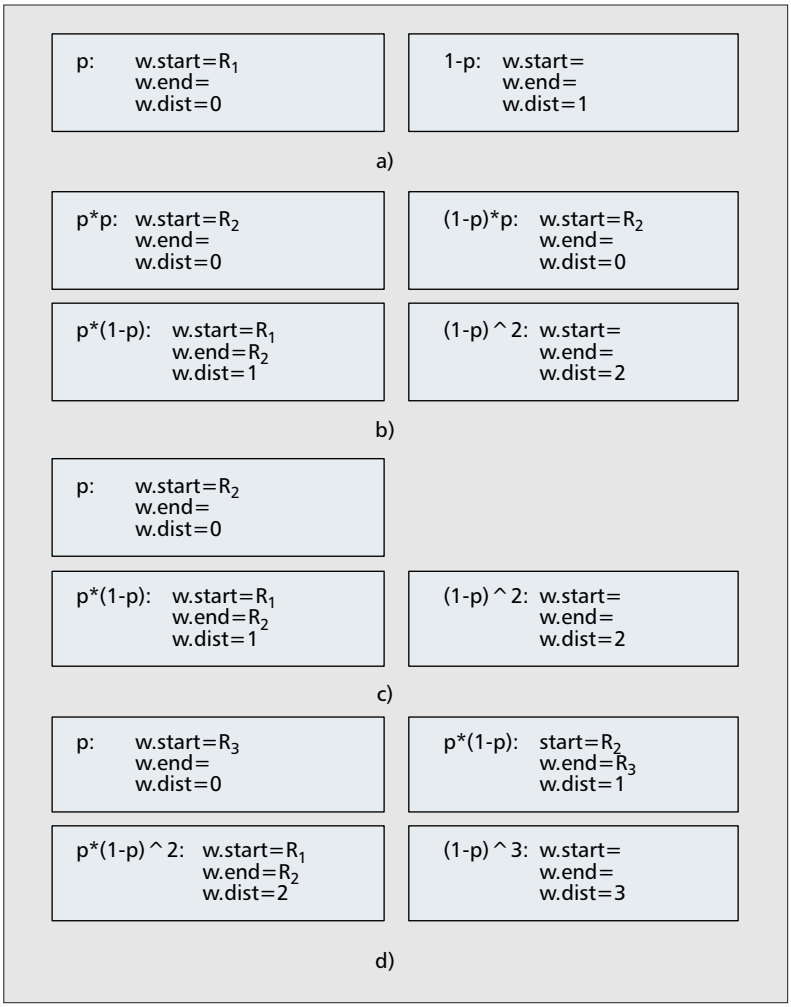

Figure 3. *PPM marking procedure.*

Unlike PPM, ICMP traceback belongs to outbound marking, which constitutes two differences. First, ICMP traceback requires additional bandwidth to convey the marking information. Second, more marking bits can be used, and thus Problems 1 and 2 (as of PPM) can be effectively solved.

Schemes	The router along an attack path	Number of packets passing by	Number of packets marked by this router	Number of marked packets from the current router received by the victim
iTrace	1st	N	Np	Np
	2nd	$N(1 + p)$	$Np(1 + p)$	$Np(1 + p)$
	3rd	$N(1 + p)^2$	$Np(1 + p)^2$	$Np(1 + p)^2$
	…	…	…	…
	dth	$N(1 + p)^{d-1}$	$Np(1 + p)^{d-1}$	$Np(1 - p)^{d-1}$
PPM	1st	N	Np	$Np(1 - p)^{d-1}$
	2nd	N	Np	$Np(1 - p)^{d-2}$
	3rd	N	Np	$Np(1 - p)^{d-3}$
	…	…	…	…
	dth	N	Np	Np

Table 1. *Comparisons of PPM and iTrace.*

ICMP Traceback

Basic Scheme — An ICMP traceback method called iTrace was proposed by Bellovin *et al.* [4]. In this scheme, each router selects one packet per 20,000 packets and then generates an ICMP message. The ICMP message has the same destination IP address as the traced packet. The ICMP message also contains the IP header of the traced packet, and the IP addresses of the incoming interface and the outgoing interface of the current router. As long as the victim receives sufficient ICMP messages, it may recover the whole attack path. In ICMP traceback, the TTL field in the IP header of the ICMP message is set to 255 so that the TTL value may be used as a clue to correctly reconstruct an attack path.

Analysis of iTrace — The marking procedure of iTrace is very similar to PPM. Therefore, it shares similar pros and cons. Unlike PPM, ICMP traceback belongs to outbound marking, which constitutes two differences. First, ICMP traceback requires additional bandwidth to convey the marking information. Second, more marking bits can be used, and thus Problems 1 and 2 (as of PPM) can be effectively solved.

Suppose that the total number of attack packets from one source is N, and the probability of generating an ICMP message at each router is p. For the first router closest to the specified source, the total number of generated ICMP packets is Np. For the second router, the total number of packets it receives (attack packets + ICMP packets) is $N(1 + p)$, and thus $Np(1 + p)$ ICMP packets are created. For a path with d routers between the attack source and the victim, the number of ICMP messages generated at the ith router ($1 \le i \le d$) is $Np(1 + p)^{i-1}$. Similar to PPM, the closer a router is to the victim, the more ICMP packets are generated. Unlike PPM, the number of ICMP messages the victim obtains from a router is the same as that generated by the router because there is no "re-marking" (Table 1). This desirable property implies a further improvement. That is, iTrace requires far fewer marked packets (ICMP packets here) than PPM for path reconstruction.

Recent Developments and Possible Solutions — Mankin *et al.* [13] proposed an "intention-driven" ICMP traceback technology. The idea is to add some intelligence to the marking procedure so that the information required for path reconstruction may be quickly gleaned by the victim. To implement "intention-driven" ICMP tracing, each router needs to modify its routing table to accommodate the intention information. This enhancement further thwarts Problems 1 and 2.

Problem 3 may be addressed by secure infrastructure such as public key infrastructure (PKI). Although PKI can tackle the issue of false marking, it imposes too high overhead on each router. Further work is required to address Problems 4 and 5 using ICMP tracing.

Hash-Based IP Traceback

Basic Scheme — Hash-based IP traceback (also called SPIE) was proposed by Snoeren *et al.* [5]. This scheme is composed of three components: data generation agents (DGAs), SPIE collection and reduction agents (SCARs), and SPIE traceback manager (STM). The function of DGA is implemented in routers using bloom filters in such a way that each router deterministically logs some information of each packet traversing the router. Each SCAR is in charge of one area of the network, and it is connected to all DGAs inside this area. STM is the central management unit that is responsible for handling the requests of the victim and assembling the path information from associated SCARs.

Whenever a server or network is under attack, the intrusion detection system (IDS) at the victim will identify the features of attack packets and report these features to STM. STM then sends an inquiry request to proper SCARs. Each SCAR collects the logging information (also called digest) of each router (or DGA) inside its area and analyzes whether the attack packets have passed through the current area or not. If this is true, the SCAR determines the routers forwarding these attack packets, and further reconstructs the attack path inside this area. All related SCARs submit their partial path reconstruction results to the STM so that the latter can reconstruct each path after gleaning these results.

The digest collected at each router is derived from the following information: the constant fields in the IP header and the first 8 bytes in the payload of the current packet. The digest table stored in a router is implemented by using the bloom filter, a specific space-efficient data structure. Whenever a bloom filter is about 70 percent full, this filter is archived for later querying and a new filter will be used.

With the help of a transform lookup table (TLT), SPIE is capable of tracing transformed packets.

Analysis of Hash-Based Traceback — SPIE is a deterministic logging scheme. It requires an additional infrastructure such as STM and SCARs, and supports advanced functions such as single packet tracing and transformed packet tracing that are especially useful in wireless networks.

Two main drawbacks exist in SPIE. It incurs very heavy computational, management, and storage overhead. Although the neat property of bloom filters mitigates the extent of the storage requirement, the deterministic nature still creates a big problem. More important, SPIE is not scalable. The current Internet is decentralized. Therefore, it is very difficult to extend this scheme from one network to the whole Internet because no STM of one network can exceed its administrative border in reality. These shortcomings seriously impede the applicability of SPIE.

Recent Developments and Possible Solutions — In terms of the problems exhibited in PPM, Problem 1 is also an issue here. Unlike PPM, however, the computational burden is distributed in the network (SCARs and STM) rather than on the victim only. Furthermore, since the logging information is distributed in each router, a high communication/bandwidth burden is incurred for SCARs and STM to recover paths. False positives depend on the performance of the selected bloom filter. An ideal bloom filter can greatly lessen false positives. Problems 3 and 4 are no longer an issue because of deterministic logging. Also, Problem 5 may be effectively thwarted with the help of the central management unit.

Li *et al.* [9] proposed a novel logging scheme that may further mitigate the storage requirement by sampling. By correlating samples, the proposal may successfully construct an attack tree. Their simulations show that the scheme may scale well to more than 5000 attack sources, a significant improvement over SPIE.

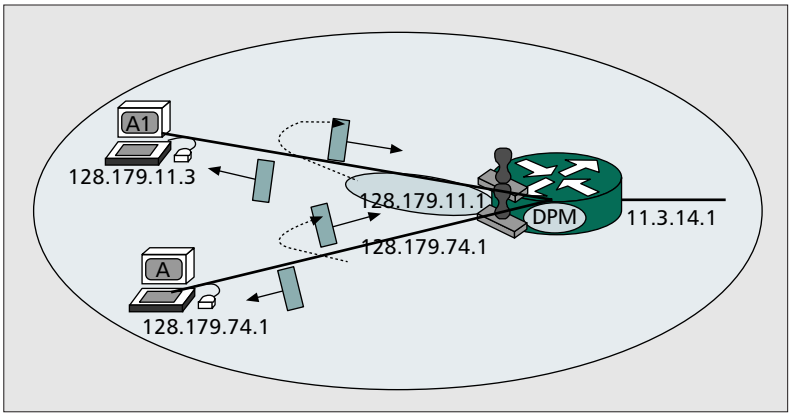

Figure 4. *Deterministic packet marking (DPM).*

Deterministic Packet Marking

Basic DPM — DPM was proposed by Belenky and Ansari [7]. In this scheme, *only* ingress edge routers perform the marking, as indicated by the DPM-enabled routers shown in Fig. 4. All other routers are exempt from the marking task.

Basic DPM uses the 16-bit ID field of the IP header and one reserved bit to record the marking information. The IP address of every ingress edge router is split into two segments with 16 bits each. One segment is randomly selected when a packet traverses this router. The idea is that the victim is capable of recovering the whole IP address of an ingress edge router once it obtains both segments from the same router. For the victim to figure out which portion of the IP address the current packet carries, one bit is used as a flag. Therefore, the marking information comprises two parts, the 16-bit partial IP address of the edge router and a 1-bit flag.

The basic scheme can effectively handle a DoS attack. For a DDoS attack, the approach may introduce high false positives. Another shortcoming is that it cannot identify the ingress edge router if the attacker uses different source IP addresses for each packet. To address these issues, they further enhanced basic DPM by using the "linkage" information [14]. That is, a hash function is used to contain the identity of the ingress edge router so that all packets traversing the same router possess the same identity. The victim can use this identity to correctly combine packets from the same source so that the whole IP address may be recovered. Thus, the marking information comprises three parts, a segment of the IP address of the current router a, the index of the current segment d (digest), and the fixed linkage information k. A good trade-off is obtained when $a = 4$, $d = 3$, and $k = 10$.

Analysis of DPM — Similar to PPM, DPM also uses the ID field to record the marking information. There are two main differences between DPM and PPM. First, PPM marks all routers along an attack path, while DPM only marks the first ingress edge router (Fig. 5). Second, PPM marks probabilistically, while DPM marks every packet at the ingress edge router.

These differences have the following implications. First, the task of ingress address recon-

There are two main differences between DPM and PPM. First, PPM marks all routers along an attack path while DPM only marks the first ingress edge router. Second, PPM marks probabilistically while DPM marks every packet at the ingress edge router.

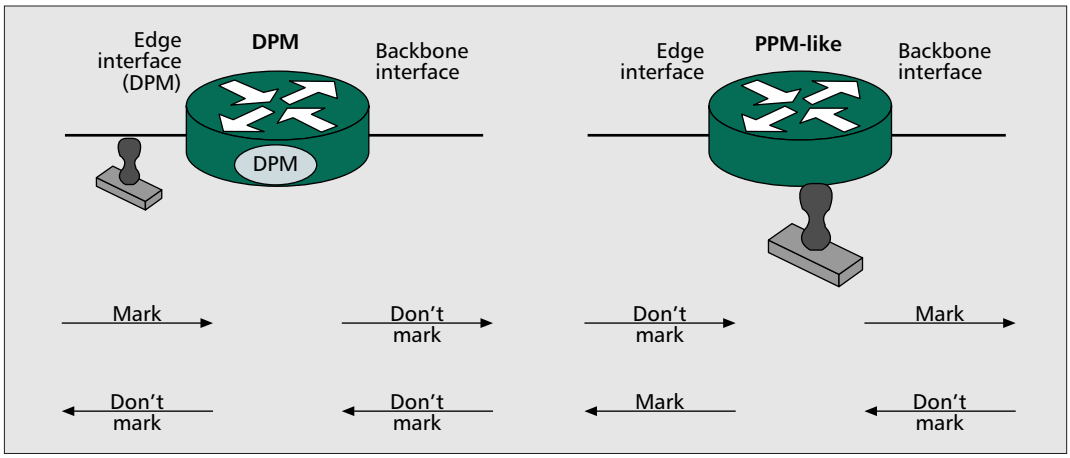

Figure 5. *Marking in PPM and DPM.*

struction in DPM is much simpler than the task of path reconstruction in PPM. As a result, DPM may handle large-scale DDoS attacks better. Second, the false positives in DPM are far less than in PPM. Third, DPM has the potential to tackle reflector-based DDoS attacks.

Recent Developments and Possible Solutions — Problems 1 and 2 are effectively thwarted in DPM. For each attack path, only the IP address of each ingress edge router needs to be recovered. Thus, the computational burden is reduced significantly. Furthermore, the linkage information may be used as a guide to effectively prevent the "combinatorial explosion" problem in PPM. This further mitigates the computational overhead. A nice collateral effect is that the false positives are decreased as well.

Problems 3 and 4 are not an issue in DPM. However, Problem 5 needs to be further addressed. One possible solution is to record partial path information rather than the whole path information in PPM and one single point in DPM (e.g., path information at the AS level).

Overlay Networks

Basic Scheme — Stone presented CenterTrack, an overlay-based solution to IP traceback [8]. In this approach, a specific router called a *tracking router* (TR) or a group of TRs is used. To trace one attack flow, dynamic routing is employed. All traffic to the victim is routed to the TR. The TR is logically directly connected to each ingress and egress edge router of the protected network through tunnels. Unlike other traceback schemes that depend on the IDS of the victim to detect invasion, IDS in CenterTrack is implemented in the TR. When an intrusion is detected, TR is capable of locating the ingress edge router of the identified attack flow because the ingress edge router may be viewed as only one hop away from the TR.

Analysis of CenterTrack — Clearly, CenterTrack enforces a heavy management burden on the network. It also wears out tremendous system resources, such as bandwidth and processing capability, due to establishment and maintenance of tunnels. Similar to SPIE, furthermore, scalability constitutes another major limitation to CenterTrack. Even though CenterTrack may determine the ingress edge router of one network with the help of TR, it cannot trace down the attack path once beyond the border of the current domain. Therefore, its applicability is rather limited.

Recent Developments and Possible Solutions — Few updates to CenterTrack have been proposed at present. Recently, an associated defensive method, Secure Overlay Service (SOS), was proposed. Unlike reactive tracing schemes, SOS is a proactive approach. By employing intensive filtering and anonymity into the forwarding structure (overlay network), SOS may effectively mitigate the impact of DDoS attacks.

In terms of the problems exhibited in PPM, Problem 1 is not a big issue. Since the ingress edge router is logically one hop away from the TR, path reconstruction in the specified network is straightforward because of the "simplified" topology. However, the usage of tunnels introduces some extra processing. Moreover, the computational burden is enforced on the TR and edge routers rather than the victim. False positives are well thwarted in this scheme. Problems 3 and 4 do not need to be considered here. Another benefit of the "simplified" topology is that the chance of routers being compromised is rather low or at least much easier to detect and diagnose.

Challenges and Future Work

Future difficult and challenging issues IP traceback should address include:

- Identifying the indirect sources of reflector-based DDoS attacks
- Identifying the attacker who conceals himself/herself with stepping stones
- Integrating IDS or defensive measures with traceback so that one mechanism may perform tracing as well as detection and/or defense
- Automatic traceback to speed up tracing and reduce human intervention

At present, all the above are still open problems. A scheme contrived to address reflector-

based DDoS attacks has to address one important issue: some kind of trust relationship must exist between the victim and the reflectors so that the reflectors may authenticate the querying requests from the victim, and the victim may obtain from the reflectors their tracing results. The trust relationship must be deliberately established and efficiently maintained. Otherwise, an attacker may exploit it to mount a DDoS attack by frequently sending bogus querying requests. Here, scalability is still a big challenge.

In addition to spoofed source IP addresses, a sophisticated attacker may use a series of stepping stones to further conceal its trail. A stepping stone is a host that is remotely logged in by a user whose physical location may be pretty far away. Employing many stepping stones can effectively hamper efforts to identify the attacker. No sound scheme has been presented to tracing through stepping stones yet.

Integrating IP traceback with other functionalities such as detection and defense is another topic of interest. Currently, a common assumption is that there exists IDS at the victim or at the TR in CenterTrack. IP traceback may identify attack sources. However, IP traceback itself is not a detection or defense scheme. A scheme that may effectively and efficiently combine detection, defense, and traceback may significantly enhance performance and mitigate false positives [15].

Instead of the current practice of human manipulation, automatic tracing is very useful, especially in a large network made up of a huge number of hosts. Automatic traceback requires more intimate coordination between IDS and traceback. To decrease the false alarm rate, the accuracy of detection needs to be significantly improved. However, improving the accuracy of DDoS detection is a daunting task given the fact that a DDoS attack may be a hybrid of different types of attacks using different protocols, ports, and attack rates [1].

Conclusions

The state of the art in IP traceback has been presented in this article, along with remaining open issues. Clearly, the current IP traceback technology is only the first step toward tackling DoS/DDoS attacks. An ideal tracing scheme has to make trade-offs among various factors. To understand the dynamics of IP traceback, we have categorized the most promising schemes from multiple aspects. From the perspective of practicality, the pros and cons of each scheme have been explored in depth, and possible future solutions have been highlighted.

References

[1] C. Douligeris and A. Mitrokotsa, "DDoS Attacks and Defense Mechanisms: Classification and State-of-the-art," *Comp. Networks*, vol. 44, 2004, pp. 643–66.

[2] S. Savage *et al.*, "Network Support for IP Traceback," *IEEE/ACM Trans. Net.*, vol. 9, Jun. 2001, pp. 226–37.

[3] A. Belenky and N. Ansari, "On IP Traceback," *IEEE Comm. Mag.*, vol. 41, July 2003, pp. 142–53.

[4] S. Bellovin, "ICMP Traceback Messages," IETF draft, Mar. 2000, http://www.research.att.com/smb/papers/draft-bellovin-itrace-00.txt

[5] A. C. Snoeren *et al.*, "Single Packet IP Traceback," *IEEE/ACM Trans. Net.*, vol. 10, Dec. 2002, pp. 721–34.

[6] D. Dean, M. Franklin, and A. Stubblefield, "An Algebraic Approach to IP Traceback," *ACM Trans. Info. and Sys. Sec.*, vol. 5, May 2002, pp. 119–37.

[7] A. Belenky and N. Ansari, "IP Traceback with Deterministic Packet Marking," *IEEE Commun. Lett.*, vol. 7, no. 4, Apr. 2003, pp. 162–64.

[8] R. Stone, "CenterTrack: an IP Overlay Network for Tracing DoS Floods," *USENIX Sec. Symp.*, July 2000, pp. 199–212.

[9] J. Li *et al.*, "Large-Scale IP Traceback in High-Speed Internet: Practical Techniques and Theoretical Foundation," *2004 IEEE Symp. Sec. and Priv.*, Oakland, CA, May 2004, pp. 115–29.

[10] D. Song and A. Perrig, "Advanced and Authenticated Marking Schemes for IP Traceback," *IEEE INFOCOM 2001*, pp. 878–86.

[11] D. Wei and N. Ansari, "Implementing IP Traceback in the Internet — An ISP Perspective," *Proc. 3rd Annual IEEE Wksp. Info. Assurance*, West Point, NY, 17–19 June 2002, pp. 326–32.

[12] Y. Tseng, H. Chen and W. Hsieh, "Probabilistic Packet Marking with Non-Preemptive Compensation," *IEEE Commun. Lett.*, vol. 8, no. 6, June 2004, pp. 359–61.

[13] A. Mankin *et al.*, "On Design and Evaluation of 'Intention-driven' ICMP Traceback," *Proc. Comp. Commun. and Net.*, Oct. 2001, pp. 159–65.

[14] A. Belenky and N. Ansari, "Tracing Multiple Attackers with Deterministic Packet Marking (DPM)," *Proc. 2003 IEEE Pacific Rim Conf. Commun., Comp. and Sig. Proc.*, Victoria, BC, Canada, Aug. 28-30, 2003, pp. 49–52.

[15] M. Sung and J. Xu. "IP Traceback-based Intelligent Packet Filtering: A Novel Technique for Defending Against Internet DDoS Attacks," *IEEE Trans. Parallel and Distrib. Sys.*, vol. 14, no. 9, Sept. 2003, pp. 861–72.

Biographies

ZHIQIANG GAO (zg4@njit.edu) received his B.S. in computer science from Zhejiang University, China, in 1989, and his M.S. from the Chinese Academy of Sciences in 1997. From 1989 to 1994 he was a senior software engineer working on Database. From 1997 to 2001 he was an assistant professor at the Department of Information Engineering, Nanjing University of Posts and Telecommunications, China, where he was primarily involved in image processing, MPEG codec, and algorithm design. He is pursuing his doctorate in computer engineering at the New Jersey Institute of Technology (NJIT), focusing on network security, in particular, IP traceback and DoS/DDoS defense.

NIRWAN ANSARI (nirwan.ansari@njit.edu) received B.S.E.E. (summa cum laude), M.S.E.E., and Ph.D. degrees from NJIT, University of Michigan, and Purdue University in 1982, 1983, and 1988, respectively. He joined the Department of Electrical and Computer Engineering, NJIT, as an assistant professor in 1988, and has been a full professor since 1997. He authored with E. S. H. Hou *Computational Intelligence for Optimization* (Kluwer, 1997, translated into Chinese in 2000), and edited with B. Yuhas *Neural Networks in Telecommunications* (Kluwer, 1994). He is a technical editor of *IEEE Communications Magazine*, *Computer Communications*, *ETRI Journal*, as well as *Journal of Computing and Information Technology*. His current research focuses on various aspects of broadband networks and multimedia communications. He organized (as General Chair) the First IEEE International Conference on Information Technology: Research and Education (ITRE 2003), was instrumental, while serving as its Chapter Chair, in rejuvenating the North Jersey Chapter of the IEEE Communications Society which received the 1996 Chapter of the Year Award and a 2003 Chapter Achievement Award, served as Chair of the IEEE North Jersey Section and on the IEEE Region 1 Board of Governors during 2001–2002, and currently serves on various IEEE committees including as TPC Co-/Vice-Chair of several conferences. He was the 1998 recipient of the NJIT Excellence Teaching Award in Graduate Instruction, and a 1999 IEEE Region 1 Award. He is frequently invited to deliver keynote addresses, tutorials, and talks.

To decrease the false alarm rate, the accuracy of detection needs to be significantly improved. However, improving the accuracy of DDoS detection is a daunting task given the fact that a DDoS attack may be a hybrid of different types of attacks using different protocols, ports, and attack rates.