

# PACKET MARKING FOR TRACEBACK OF ILLEGAL CONTENT DISTRIBUTION\*

Ihab Hamadeh

*Department of Computer Science and Engineering*

*Pennsylvania State University*

*University Park, PA 16802 USA*

hamadeh@cse.psu.edu

George Kesidis

*Department of Computer Science and Engineering*

*Department of Electrical Engineering*

*Pennsylvania State University*

*University Park, PA 16802 USA*

kesidis@engr.psu.edu

**Abstract** The continued proliferation of third-party file-sharing mechanisms (distributed peer-to-peer databases) has sustained the epidemic of copyright infringements over the Internet. An illegal exchange of copyrighted material will typically be negotiated by parties aware of each other's true IP addresses, i.e., identities may be available at the application layer. The sending party may, however, use spoofed source IP addresses in the packets thereby making up the transfer of copyrighted material difficult to trace back to the origin of the illegal transmission, assuming that the network has mechanisms to *detect* such transmissions in the first place. Alternatively, we can consider a situation where a hoard of illegally possessed copyrighted material is discovered and, based on feasibly sized logs of recent internetworking activity of the apprehended user, the sources of the copyrighted material are identified. In this paper, we describe a strategy of *network-layer* packet marking in which routers situated at the edge (access points) of the Internet mark packets with IP addresses identifying their input interfaces. The marks enable simple and rapid traceback to the edge router through which identified illegal transmissions entered into the Internet.

## 1. Introduction

The illegal exchange of copyrighted material remains an enormous problem in the Internet. Such copyright infringements are facilitated by pervasive peer-to-peer networks which are basically distributed databases. Many such databases contain vast amounts of copyrighted material illegally obtained and possessed. Notwithstanding the past punitive action taken against Napster, the first successful "third party" file distribution system, transmission of copyrighted material still constitutes a significant fraction of the Internet's total traffic burden and accounts for significant lost royalty revenues by copyright holders.

Currently, illegal exchanges of copyrighted material are "openly" negotiated and transferred typically using FTP or HTTP over TCP, i.e., the senders identity is known during the transfer of the file. Once illegal possession and distribution of copyrighted material begins to be actively investigated,

---

\*This research was funded in part by a Cisco Ltd URP grant.

we expect that transfer of copyrighted material will become more anonymous to prevent trace-back to (identification of) the unlawful transmitters of this material into the network (senders). More specifically, we expect that an illegal exchange of copyrighted material in the future will typically be negotiated through separate email, or even telephone, exchanges. The sending party will, however, then use spoofed source IP addresses in the packets making up the *UDP* transfer of copyrighted material to the receiving party. That is, the transfer itself is made anonymous in this fashion to protect the unlawful sender's identity in the network layer.

In this paper, we consider the problem of identifying the *source* of illegally distributed copyrighted material. To address this problem, we must first assume that a mechanism is in place in the network that can *detect* flows of illegally copyrighted material through the network. Once discovered, the network can attempt to trace back to the source of the flow (assuming this problem is made nontrivial by the use of spoofed source IP addresses by the unlawful sender). In a similar fashion, traceback of distributed denial-of-service attacks (DDoS) is attempted *after* an intrusion detection system (IDS) at the victim end-system has identified that an attack is occurring and which packets are participating in it. Alternatively, we can consider a situation where a hoard of illegally possessed copyrighted material is discovered and, based on *feasibly sized* logs of recent internetworking activity of the apprehended user, the sources of the copyrighted material are identified.

A variety of techniques have been proposed to deal with problems of IP spoofing. For example, under ingress filtering [7, 2] ingress network processors of all border routers check the source IP addresses of packets they forward into the Internet to verify that the addresses legitimately belong to the domain having access to the internet through them. Ingress filtering has, however, scalability concerns, c.f., Section 3. Under probabilistic packet marking (PPM) [8, 4, 5] *any* router, with specified probability, inscribes its local path information into the packet headers. The network reconstructs the transmission path of a session starting from the packets received from the closest routers moving up to the ISPs' border routers. Two prominent varieties have been proposed: the Fragment Marking Scheme (FMS) by Savage et al [4] and the Advanced Marking Scheme (AMS) by Song and Perrig [5]. FMS suffers from high computation overhead (Fig. 11 of [5]), because of the large number of combinations that need to be checked to reconstruct the routers' IP address, and large numbers of false positives (Fig. 10 of [5]). AMS requires the knowledge of a topological map of the Internet a priori to be able to reconstruct a 32-bit router IP address from the 11-bit or 8-bit hash values. Finally, an unlawful sender can insert "fake" links and distances into the overloaded fields in the packet header [3]. To overcome this "malicious false positive" problem, authentication of the packet marking using Message Authentication Codes (MAC) was proposed in [5].

In this paper, we present a new approach to the traceback problem based on marking packets with partial information reflecting IP addresses of the input interfaces of *only* the border routers through which copyrighted material enters the Internet. Such small packet markings may be stored at receiving end-systems in feasibly-sized logs to be used for future traceback operations.

## 2. Border Router Packet Marking (BRPM)

Like the PPM techniques, BRPM is based on the following assumptions: an attacker may generate any packet, multiple attackers may conspire, attackers may be aware they are being traced, packets may be lost or reordered, attackers send numerous packets, routers are both CPU and memory limited, and routers are not widely compromised.

Under BRPM, only border routers mark packets and all packets inbound into the Internet are marked. Marking every packet in this way protects against an attacking end-system inserting fake marks into this field in an attempt to compromise traceback. Traceback is made unambiguous by the deployment assumption that each packet is forwarded into the Internet by only one trustworthy marking ("border") router.

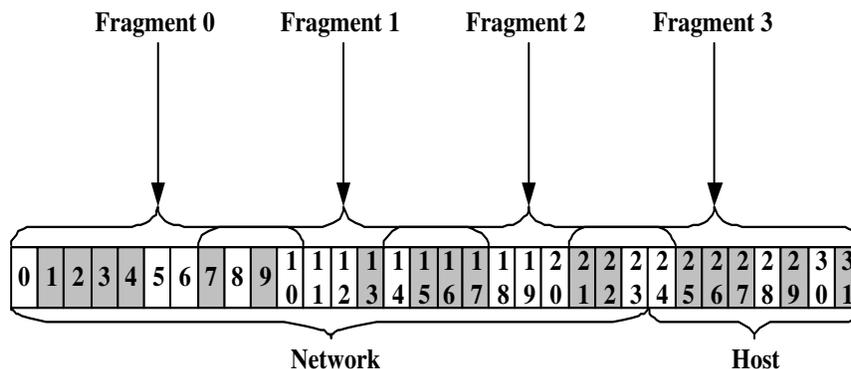


Figure 1. Example of Border Router IP address divided into overlapping 11-bit fragments (group #0).

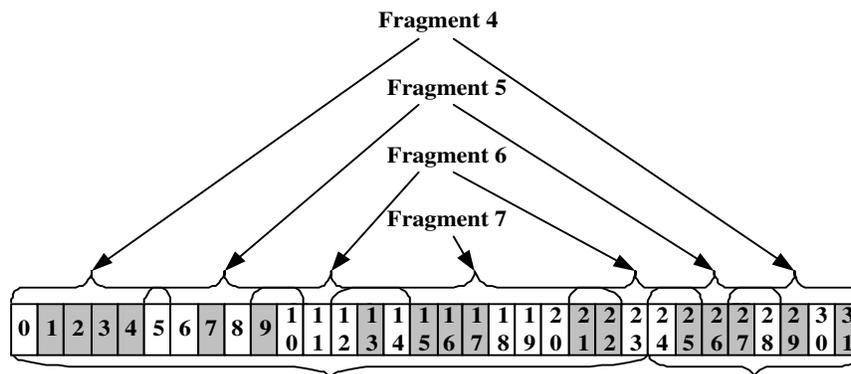


Figure 2. Example fragment group #1 overlapping in both the network and the host portion of Class-C IP address

We also associate a novel address fragmentation framework with BRPM in which a border router’s IP address<sup>1</sup> is fragmented (segmented) into several ( $k$ ) overlapping fragments where each fragment has an identifying index (IDs 0 to  $k - 1$ ). Border routers write into the header (of every packet they forward) a selected fragment and its identifier (ID). The number of bits needed for storing both a fragment and its ID is at most  $n + \lceil \log k \rceil$  where  $n$  is the fragment size and  $k$  is the total number of fragments. Example fragmentation strategies are given in Figures 1 and 2. As we will see in section 2.3, the strategy of *overlapping* fragments and grouping fragments (fragments 0-3 and 4-7 in the example) reduces false positives and address reconstruction complexity. In particular, overlapping fragments allow the victim to correlate those fragment instances from the same IP address. Note that this address fragmentation framework could also be applied to other marking strategies such as PPM.

## 2.1 Address Reconstruction Algorithm for BRPM

Under BRPM, all possible complete border router IP addresses that are consistent with the received address fragments are reconstructed. For this reason, *overlapping* fragments are used to allow correlation of those fragment instances belonging to the same 32-bit address.

<sup>1</sup>More precisely, the IP addresses of the border router’s input link interfaces.

In Figures 1 and 2, the four fragments numbered 0-3 and those numbered 4-7 are respectively *grouped* together. So, there are  $2 = k/k'$  groups of  $k' = 4$  fragments where the collection of fragments in each group span the entire 32-bit address. Suppose that an end-system has received a group of marked packets that were identified as participating in illegal content distribution. Address reconstruction works as follows. The address fragments and their identifiers are extracted from the packet headers. Only pairs of fragments with *identical* overlapping fields and with identifiers belonging to the same group are joined together to form a larger address “metafragment”. Metafragments are then made even larger, according to this same rule, by continuing the join them with other fragments whose overlapping fields agree with those of the metafragment.

Consider the simple example of two fragments ( $k = 2$ ) of  $n = 20$  bits that therefore overlap in 8 bit positions. For a given 32-bit border router address  $A$ , let

- $w(A)$  be this 8-vector of the overlapping bits
- $f_i(A)$  be the fragment with ID  $i$  for  $i = 0, 1$
- $b_i(A)$  be the 12-vector of *non*-overlapping bits of  $f_i(A)$

We therefore write  $f_i(A) = b_i(A) \oplus w(A)$ , i.e., the  $i^{\text{th}}$  fragment is composed of non-overlapping (unique to  $f_i$ ) bits  $b_i$  and the overlapping bits  $w$ . Now consider two logged fragments with different IDs,  $f_0(A_0)$  and  $f_1(A_1)$ , where  $A_0$  and  $A_1$  are the actual IP addresses of border routers that marked the corresponding packets (of course,  $A_0$  and  $A_1$  are not known a priori to the entity performing traceback). If the overlapping bits agree, i.e.,

$$w(A_0) = w(A_1) =: W, \quad (1)$$

then the following 32-bit IP address will be reconstructed given fragment instances  $f_0(A_0)$  and  $f_1(A_1)$ :

$$D_0 \equiv b_0(A_0) \oplus W \oplus b_1(A_1) = b_0(A_0) \oplus f_1(A_1) = f_0(A_0) \oplus b_1(A_1) \quad (2)$$

Note that if  $A_1 = A_0$  (i.e., the two fragments under consideration are taken from the same address) then  $D_0 = A_0$  (i.e., the address is successfully reconstructed).

In general, this process leads to a set  $S_i$  of reconstructed 32-bit addresses using only those fragments whose IDs belong to the  $i^{\text{th}}$  group. To further reduce the number of false positives, only those reconstructed IP addresses that are common to all sets, i.e., those in

$$S \equiv S_1 \cap S_2 \cap S_3 \cap \dots \cap S_n,$$

are deemed to be border routers through which originated transmission to the end-system performing traceback. Finally, those addresses in  $S$  that are invalid IP addresses are simply removed.

## 2.2 Complexity of BRPM

Unlike all PPM schemes, the BRPM scheme does **not** require any router to: decide (at random) *whether* to mark an IP datagram that has not been marked by a neighboring router and, in the case that a neighboring router has marked the packet, XOR the existing mark with its own. Note that marking *every* packet is no more complex, from either a hardware or software perspective, than marking packets at random (so that *potentially* all packets are marked).

To summarize, under BRPM:

- **only** border routers mark (interior routers may not be involved at all).
- border routers inscribe identifying information into the headers of **all** IP datagrams they forward into the Internet (hence, there is no decision making).

Table 1. The Other Two Groups of Fragments

Group #	Fragment #	Bits Used										
2	8	0	3	6	9	12	15	18	21	24	27	30
2	9	1	4	7	9	11	15	20	22	25	28	31
2	10	0	2	5	8	10	11	18	20	23	26	29
2	11	2	9	13	14	16	17	19	22	26	27	28
3	12	0	2	4	5	6	8	10	12	14	18	22
3	13	1	3	5	7	9	11	13	14	15	20	25
3	14	4	11	16	18	20	22	23	24	26	28	30
3	15	2	9	16	17	19	21	23	25	27	29	31

- marking only requires the addition of a write and a checksum update, which is already performed by routers’ network processors to update the time-to-live (TTL) field at each hop.

At the end-system performing traceback, the reconstruction algorithm requires some computation the amount of which depends on the size and number of fragments and the space of border router IP addresses to be reconstructed. Consider a single group of  $k' < k$  fragments and assume there is a common number,  $u$ , of bits in **each** fragment which are unique to that fragment (i.e., do not overlap with another). In the two-fragment example of the previous subsection,  $k' = 2$  and  $u = 12$ .

The worst-case reconstruction complexity occurs when:

- 1 The overlapping bits (shaded) are common to all ingress border routers through which content was illegally distributed, and
- 2 All non-overlapping bits are different for each border router.

We expect that this worst-case scenario is a highly unrealistic one and that, in practice, the great majority of these “possible” reconstructed addresses to be invalid, i.e., not to correspond to an actual border router.

Under the BRPM scheme, traceback requires a very small number of packets for reconstruction of a border router address. Unlike the probabilistic schemes that require several hundred to thousands of packets per border address, under BRPM only on the order of tens of packets from each border router are sufficient for address reconstruction.

### 2.3 Fragmentation Strategy and Performance

Finding the fragmentation strategy that minimizes the number of false positives is a complex task. Clearly, the best strategy will greatly depend on the *correlations* among the specific border router IP addresses to be reconstructed and the whole set of such addresses for any given collection of end-systems that wish to perform traceback. Roughly speaking, the overlapping portions of the fragments helps to prevent false positives. In Figure 2, the overlapping bits reside in both the network and host portion of the (Class C) IP address. Depending on the address space of a trust region, more bits could for example be overlapping in the host portion than in the network portion thereby reducing false positives when the host portion alone discriminates the majority of border router IP addresses under consideration.

A simulation was conducted for an  $n = 11$  bit fragment size and four groups of  $k' = 4$  fragments: fragments 0-3 of Figure 1, fragments 4-7 of Figure 2, fragments 8-11 and fragments 12-15 as shown in Table 1. This is consistent with  $15 = n + \lceil \log_2 k \rceil = 11 + \lceil 16 \rceil$  available bits for BRPM, c.f., Section 2.4. Transmitting border router addresses were generated uniformly at random and all fragments of all such addresses are assumed available for traceback. After running the traceback

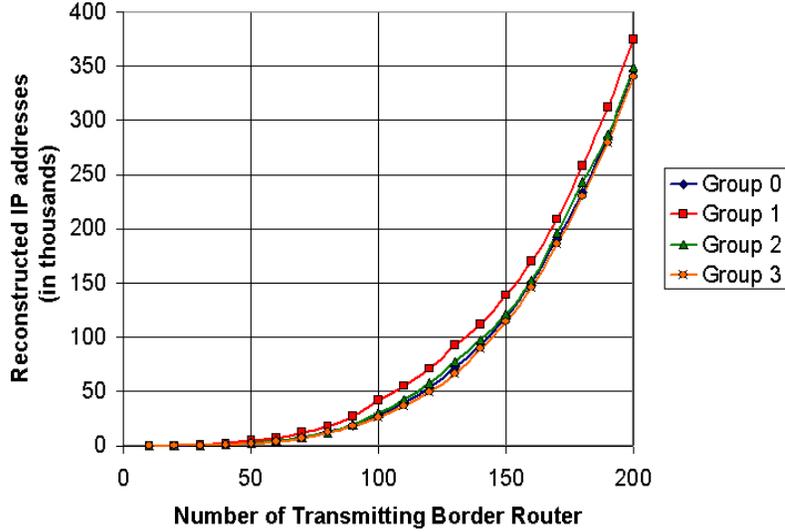


Figure 3. Number of Reconstructed Transmitting Border Router IP addresses per group.

(address reconstruction) algorithm, each group generated a set of IP addresses consisting of the true transmitting border routers (0% missed detection) in addition to false positives the number of which are depicted in Figure 3. Using the reconstruction approach of Section 2.1 above, the intersection of the four groups yielded the IP addresses of all transmitting border router (again, 0% missed detection) with no false positives. In other simulation runs we did observe negligible numbers of false positives (fewer than 5 for hundreds of transmitting border routers) depending on the degree of correlation among the transmitting border router IP addresses.

## 2.4 Overloading Issues

As mentioned previously, the number of bits occupied in a packet header by the fragment and its ID are  $n + \lceil \log_2 k \rceil$ . Security issues alone may justify the addition of a small field for BRPM in future standards for IP headers. Of course, any such additional field would, however, result in an overall reduction in transmission efficiency for the Internet.

We now consider implementation of BRPM by overwriting (or “overloading”) IPv4’s 13-bit Fragment Offset field along with the two unused bits of the ToS field. Though this is the most promising field for overloading under IPv4, as explained below, issues of backwards-compatibility for segmented IP traffic still need to be considered.

Under BRPM, when an Internet router decides to fragment a packet, it can store the 13-bit Fragment Offset field that contains the border router marking information. It can then segment the packets as required and append the border router marking information to the end of the last packet segment’s payload. When an end-system reassembles packet segments, it will retrieve the border router marking information from the end of the last segment prior to reassembly.

Clearly, without such a mechanism, packet segmentation in the Internet will cause border router markings to be lost. Since packet segmentation affects a increasingly small percentage of packets (fewer than 0.3% [6, 12]), the resulting missed detections under BRPM overloading the Fragment Offset field may not be significant.

### 3. Comparison of Other Traceback Strategies

This section briefly introduces several previously proposed techniques to trace back the origins of a DDoS attack. These methods, along with BRPM, can also be used to traceback illegal content distribution. General criteria for evaluation of traceback techniques include: false positive rates (including those maliciously caused), missed detection rates, computation and communication overhead and deployment complexity.

#### 3.1 Logging Approach

This general approach proposes to log packets at various points throughout the network and then use some extraction (“data mining”) techniques to find the path packets traversed, see [11]. Snoeren et al [1] proposed a modification to this approach called the Source Path Isolation Engine (SPIE). Their idea is to hash and store only the first 28 bytes of a packet. This avoids 99.9% of all collisions (false positives due to many-to-one mapping of the hash), while saving tremendous amounts of storage space. This approach has greater deployment complexity and (memory) resource requirements (especially in the interior of the network) as compared to BRPM.

#### 3.2 Ingress Filtering

Ferguson and Senie [7] suggested that border routers check the source IP addresses of all IP datagrams to see whether they fall within the address space of the network from which the datagrams originated, i.e., “ingress filtering.” Even if the source IP addresses are spoofed from within the range allowed by their provider of Internet connectivity, it is still possible to determine the subnetwork from which the illegal packets are emanating. Park and Lee [2] proposed Route-Based Distribution Packet Filtering (RBDPF) that uses ingress filtering on some gateways for coverage of backbone routers. It functions by analyzing the routes packets usually traverse and determining whether they are among those routes typically used. They then can begin to determine which TCP requests are fake. This scheme is not, however, 100% effective. While many routers employ the ingress filtering technique, the implementation of ingress filtering is opposed by some ISPs, especially the larger high-volume providers. Packet filtering increases CPU utilization and measurably lowers throughput, leading to potential performance degradation, because routers need to check each outgoing IP datagram for invalid IP source address. Moreover, “spoofed source addresses are legitimately used by network address translators (NATs), Mobile IP, and various unidirectional link technologies such as hybrid satellite architectures” [1].

#### 3.3 ICMP Traceback Messages

Bellovin [9] proposed a scheme known as the ICMP Traceback Messages that was later extended by Wu et al [10]. In this scheme, routers, with low probability, generate a Traceback message that is carried in an ICMP packet and is sent along the path of the packet. With sufficient numbers of Traceback messages from enough routers on the path, it is easy to reconstruct the path of the illegal packets. The Traceback messages can help to identify the message generator, the link that the traced packet arrived from, or the link it was forwarded on. ICMP Traceback Messages can be easily used in conjunction with any packet-marking scheme to deliver even a more powerful traceback mechanism. Clearly, this technique may significant communication overhead within the network (unlike BRPM) and, if packet volume constituting illegal content distribution is moderate to low, ICMP Traceback Messages will suffer from high missed detection rates (because of low packet sampling rates).

### 3.4 Probabilistic Packet Marking (PPM)

The PPM scheme [4, 5, 8] requires that a router, with specified probability, inscribes its local path information into the packet header. The path of the packets is reconstructed starting from the packets received from the closest routers moving up to the ISPs' border routers. Two prominent varieties have been proposed: Fragment Marking Scheme (FMS) by Savage et al [4] and Advanced Marking Scheme (AMS) by Song and Perrig [5].

Under FMS, each router's IP address is bit interleaved with its hash value. The resulting 64-bit quantity is split into eight fragments. Each router probabilistically marks an IP packet it forwards with one of the eight fragments.

Under AMS, each router's IP address is hashed into an 11-bit or 8-bit value (according to whether AMS version I or II is used) and probabilistically inscribed in forwarded IP packets. However, unlike FMS, AMS requires the knowledge of a topological map of the Internet a priori to be able to reconstruct a 32-bit router IP address from the 11-bit or 8-bit hash values.

Performance and complexity disadvantages of PPM were discussed in the introductory section. BRPM clearly has lower associated deployment and computation complexity and is less vulnerable to malicious false positives (from fake marks).

## 4. Summary Discussion

The proposed Border Router Packet Marking scheme is a simple and effective packet marking strategy for traceback to the sources of unlawful distributed copyrighted material. Ingress border routers mark all packets that forward into the Internet. The small packet marks can be stored in feasibly-sized logs near the destination end-systems. Traceback (address reconstruction) could be initiated upon the detection of an unlawful transmission or after a hoard of unlawfully obtained copyrighted material is discovered. In general, the BRPM scheme offers significant enhancement to Internet security by mitigating the anonymity of malicious activity.

## References

- [1] A.C. Snoeren, C. Partridge, L.A. Sanchez, C.E. Jones, F. Tchakountio, S.T. Kent and W. T. Strayer. Hash-Based IP Traceback. Proc. ACM SIGCOMM, 2001.
- [2] K. Park and H. Lee. On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack Prevention in Power-Law Internets. Proc. ACM SIGCOMM, 2001.
- [3] K. Park and H. Lee. On the Effectiveness of Probabilistic Packet Marking for IP Traceback under Denial of Service Attack. Proc. IEEE INFOCOM, 2001.
- [4] S. Savage, D. Wetherall, A. Karlin and T. Anderson. Practical Network Support for IP Traceback. Proc. ACM SIGCOMM, Aug. 2000.
- [5] D.X. Song and A. Perrig. Advanced and Authenticated Marking Schemes for IP Traceback. Proc. IEEE INFOCOM, Apr. 2001.
- [6] I. Stoica and H. Zhang. Providing Guaranteed Services Without Per Flow Management. Proc. ACM SIGCOMM, Aug. 1999.
- [7] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing. IETF RFC 2267, Jan. 1998.
- [8] T.W. Doepfner, P.N. Klein and A. Koyfman. Using router stamping to identify the source of IP packets. Proc. ACM Conference on Computer and Communications Security, 2000.
- [9] S.M. Bellovin. ICMP Traceback Messages. IETF Internet Draft: draft-bellovin-itrace-00.txt, 2000.
- [10] S.F. Wu, L. Zhang, D. Massey, and A. Mankin. Intention Driven ICMP Traceback. IETF Internet Draft: draft-wu-itrace-intention-00.txt, Feb. 2001.
- [11] G. Sager. Security fun with OCxmon and cflowd. Internet 2 Working Group Meeting, Nov. 1998. <http://www.caida.org/projects/ngi/content/security/1198/>.

- [12] S. McCreary and K. Claffy. Trends in wide area IP traffic patterns: a view from Ames Internet exchange. ITC Specialist Seminar on IP Traffic Modeling, Measurement and Management, 2000.