

Coloring the Internet: IP Traceback

M. Muthuprasanna, G. Manimaran
Iowa State University,
Ames, IA, USA - 50014
{muthu, gmani}@iastate.edu

Mansoor Alicherry, Vijay Kumar
Bell Labs, Lucent Technologies,
Murray Hill, NJ, USA - 07974
{mansoor, vijay}@bell-labs.com

Abstract

Several IP Traceback schemes employing packet marking have been proposed to trace DoS/DDoS attacks that use source address spoofing. The major challenges in the design of an efficient traceback technique are to minimize the number of packets required for successful traceback, and also to reduce the number of bits marked per packet by any router along the attack path. We propose a graph-coloring approach here that specifically addresses these issues. We propose to view the deployment of the traceback-enabled routers as an Internet Traceback Overlay Network, which not only provides easy scalability and incremental deployment, but also allows for the spatial reuse of the router labels used for packet marking, directly resulting in a reduced bit-space, and hence in fewer packets required for successful traceback. We additionally propose an enhanced (logical) partitioned coloring technique to achieve an order of magnitude improvement over the best known schemes today. We also propose a 2-tier architecture that provides greater incentives for deployment to different ISP networks worldwide. We analyze the proposed techniques using real Internet AS-level topologies obtained from various sources.

1. Introduction

Denial-of-Service (DoS) attacks pose an increasing threat to today's Internet. DDoS attacks on several sites including Yahoo and eBay, and against root DNS servers had virtually paralyzed the Internet then. Recent attacks motivated by political and economic reasons on SCO, RIAA, 2Checkout etc. have established a disturbing trend, where the victims face considerable downtime and huge financial losses. Unless this issue is properly addressed, these attacks may not stop or even scale down. The stateless nature and destination-oriented routing of the Internet makes tracking of attackers, employing simple source address spoofing, a difficult problem referred to as *IP Traceback* [1] [2] [3]. Finally, a traceback scheme should not only trace attackers, but also aid in effective mitigation of the ongoing attack [4].

To be realistically applicable in an Internet environment, a traceback mechanism must be incrementally deployable, scalable, require minimal changes to existing hardware, maintain high accuracy both in terms of failing to identify true attack sources and also incorrectly implicating non-attacking hosts, require very few packets to complete traceback, and resist tampering due to spoofed information injected by the attackers or compromised routers, amongst other requirements. The traceback schemes in literature can be classified as: *Packet Marking* techniques where the packets are marked with partial path information deterministically/probabilistically by the intermediate routers, and *Packet Logging* techniques where the routers store packet digests in the form of Bloom filters.

Our Contribution: In this paper, we propose a graph-coloring approach to the traceback problem employing packet marking, wherein we associate a color to every traceback-enabled router subject to certain criteria. The traceback operation then identifies the attack path as a sequence of colors (of the routers) from the source to the destination. The use of colors allows for intentional spatial reuse of the conventional bit-space, thus requiring fewer bits to be marked on a packet by a certain router, and also requiring fewer packets to achieve successful traceback. We also propose an enhanced (logical) partitioning coloring technique, to achieve an order of magnitude reduction in the bit-space representing the router labels. We also propose a 2-tier traceback architecture to provide greater incentives for deployment to the different ISP networks worldwide.

The rest of the paper is organized as follows. Section 2 reviews the different traceback schemes known. Section 3 presents the basic graph coloring problem, while Section 4 provides the details of the proposed coloring-based traceback technique. Section 5 presents the theoretical analysis and experimental results, while Section 6 concludes.

2. Related Work

The earliest work in literature can be traced to the concept of network traceback [5] by Burch and Cheswick.

Bellovin et.al. proposed ICMP-based out-of-band messaging in iTrace [2]. Snoeren et.al. proposed SPIE [3] employing packet logging, which was subsequently improved by Li et.al. in [6]. Belenky and Ansari proposed a deterministic packet marking scheme in [7], while Savage et.al. proposed a probabilistic packet marking (PPM) technique in [1]. Subsequently, various improvements to PPM schemes have been proposed in [8] [9] [10]. IP address fragmentation for efficient packet marking has been studied in [11]. The vulnerability of PPM schemes to attacker induced noise (GOSSIB & rumors) has been studied in [12].

Recently, various encoding techniques have been used to progressively improve the performance of PPM schemes, as in Tabu marking [13], Local Topology marking [14], Space-Time encoding [15], and use of Huffman Codes [16], Algebraic Geometric Codes [17] etc. Additionally various architectures for traceback have been explored, such as inter-domain traceback [18] and hybrid traceback [19], [20].

3. Graph Coloring Problem

The graph coloring problem¹ forms the crux of our proposed traceback technique. In this section, we define the relevant graph coloring problems, and also provide the theoretical bounds and different algorithms for these problems.

Proper Vertex Coloring: *It is a graph coloring of G, where the colors are assigned to the vertices, such that no two vertices are assigned the same color.*

Distance-k Coloring: *It is a proper vertex coloring of G such that no two vertices have the same color if the shortest path between them is less than or equal to k-hops.*

Star Coloring: *It is a proper vertex coloring of G such that no path of length three in G is bi-colored. It is also called the Distance-2 coloring problem.*




Figure 1. Lower & Upper Coloring Bounds

In the star coloring problem, not only do adjacent vertices in the graph have different colors, but also vertices that are adjacent to a (common) third vertex. Stated differently, for every vertex in the graph, a unique color is assigned to itself and to all its 1-hop neighbors. We use this alternate definition in proposing a solution to the traceback problem.

¹For clarity, the figures represent the different colors as unique numbers respectively (say Red, Green, Yellow etc. as 1, 2, 3 etc. respectively).

Theoretical Lower Bound: Consider a graph G having a maximum node degree Δ (each vertex V_i has a node degree Δ_i), and requiring $\chi(G)$ colors for a star coloring. For any vertex V_i in G, as a unique color is assigned to itself and all its 1-hop neighbors, a minimum of $\Delta_i + 1$ colors are required. Thus a minimum of $\Delta + 1$ colors are required for a star coloring of the graph G, as in Fig. 1(a).

Theoretical Upper Bound: For any vertex V_i in G, a path of length three consists of one of its 1-hop neighbors and one of its corresponding 2-hop neighbors. Thus, in the worst case, to color the neighborhood of a vertex, we would need colors one more than the number of its 1-hop and 2-hop neighbors. As the 2-hop neighbors are bounded by $(\sum_{\forall j \in \Delta_i} \Delta_j - \Delta_i)$, the number of colors is bounded by $(\sum_{\forall j \in \Delta_i} \Delta_j + 1)$. Thus star coloring of G would require $(\Delta^2 + 1)$ colors in the worst-case, as shown in Fig. 1(b). The readers are referred to [21], [22], [23] for a formal derivation and proof of these bounds. Thus, the number of colors required for star coloring is as shown in Eqn. 1, Eqn. 2.

$$(\Delta + 1) \leq \chi(G) \leq (\Delta^2 + 1) \quad (1)$$

$$\Omega(\Delta) \leq \chi(G) \leq O(\Delta^2) \quad (2)$$

The star coloring problem has been shown to be NP-complete, for general graphs and also when restricted to planar graphs [21], [23]. We present a greedy algorithm in Algorithm 1. A formal analysis of this algorithm is presented in [21]. Various other approximation and parallel distributed algorithms have been analyzed in [24], [25].

Algorithm 1 Star Coloring Greedy Algorithm

```

input: A graph  $G = (V, E)$ 
output: A coloring  $c : V \rightarrow 1, 2, 3, \dots$ 

for  $i = |V|; i > 0; i--$  do
   $v \leftarrow$  vertex with least degree in subgraph of unlabeled vertices
  assign label  $i$  to  $v$ 
  set  $c(v) \leftarrow 0$ 
end for

for  $j = largest-label; j > 0; j--$  do
   $\mu \leftarrow \phi$ 
  for all  $v$  such that  $(u, v) \in E$  do
     $\mu \leftarrow \mu \cup c(v)$ 
    for all  $w$  such that  $(w, v) \in E$  do
       $\mu \leftarrow \mu \cup c(w)$ 
    end for
  end for
   $c(u) \leftarrow$  least color  $\notin \mu$ 
end for

```

4. Proposed Traceback Mechanism

We motivate the need for a traceback overlay network and star coloring here, while also describing the 2-tier traceback architecture and (logical) partitioned coloring technique. We finally build a single comprehensive traceback protocol that exploits each of these individual features.




Figure 2. Traceback Overlay




Figure 3. Why Star Coloring?




Figure 4. 2-tier Architecture

4.1. Traceback Overlay Network

Consider the network as in Fig. 2. The physical topology consists of 9 nodes labeled (A, B, \dots, I), while the logical topology (overlay network) consists of 6 of these nodes labeled (A, B, \dots, F), satisfying certain special properties. A unique overlay network can be constructed by restricting only non-overlapping shortest paths between the overlay nodes in the physical topology, to be included as links in the overlay topology. Hence in Fig. 2, link AE does not exist, while link DE exists in the overlay topology.

The nodes in a traceback overlay network may thus correspond to traceback-enabled routers (ASes) forming a router-level (AS-level) Internet overlay topology. Thus partial and incremental deployment of traceback schemes can be supported on this overlay infrastructure, implicitly allowing co-existence of both traceback-enabled and non-traceback-enabled routers. The overlay can be managed in a distributed manner by appropriate protocol messaging, thus lending well for a scalable and incrementally deployable Internet-wide design of a traceback mechanism.

4.2. Star Coloring

In the original PPM scheme, a router is identified by multiple hash fragments of its IP address. The number of unique labels and hence the bit-space needed to represent them is thus bounded by $O(n)$, where n is number of overlay nodes. If we now associate a label or color with each node in the traceback overlay such that any attack path can be represented as a sequence of labels in the overlay, then star coloring the overlay topology bounds the total number of colors to $[\Omega(d), O(d^2)]$, where d is the maximum overlay node degree, and $d \ll n$. Thus the spatial reuse of the bit-space reduces both the number of bits marked per packet and the number of unique labels, and hence we require far fewer packets for successful traceback.

We now show that star coloring (and not any naive vertex coloring) is necessary to guarantee correctness. Due to spatial reuse of colors and the greedy star coloring algorithm, there exist multiple paths having the same sequence of col-

ors in the overlay topology. In Fig. 3, the victim can identify nodes 1 and 2, but given a label 3, it cannot identify which of the two nodes actually sent that label. As it is always possible for another path to satisfy the same sequence of colors, we might induce false positives by implicating a legitimate user as an attacker. Star coloring ensures that no two neighbors of a given node have the same color, and hence using the pseudo-uniqueness introduced, the false positives and false negatives can be eliminated.

4.3. 2-tier Traceback Architecture

The convergence time of any traceback technique is determined by the number of routers that mark the packet along the attack path, and also the number of fragments per marking router. Using star coloring, we reduced the number of hash fragments (labels) from the customary k to 1. Additionally, we propose a 2-tier traceback architecture to reduce the number of marking routers, within permissible limits of traceback inaccuracies.

The maximum diameter of a router-level (AS-level) Internet overlay graph is approximately 30 (8-10), as reported in [14], [26] etc. ([27], [28], [29] etc.). Thus, an AS-level traceback overlay would have far fewer packet marking entities for traceback. However, inaccuracies are now introduced in the traceback process due to lack of knowledge of the intra-AS path, and the reality that ASes can peer at multiple points in the Internet. Hence, we propose a 2-tier traceback architecture as in Fig. 4, where the first tier identifies the AS-level path, while the second tier identifies a corresponding intra-AS path by pinning down the ingress and egress points for the AS. Every AS edge (border) router is assigned a tier 1 AS color satisfying star coloring properties, and also an unique tier 2 color. In Fig. 4, the attack path is represented in tier 1 as $A-B-D$, and in tier 2 as $1-2-1-2-1-2-1-2-2$. For any packet, the AS border routers mark either of the two tier colors probabilistically.

The AS-level overlay essentially requires only the edge (border) routers to be traceback-capable, while the other routers can still remain legacy routers, thus reducing the deployment cost for ISP networks today. Additionally, us-

ing the AS-level traceback overlay does not reveal the internal topology of any ISP network, which is best kept secret by network providers for security and competitive reasons. Thus, the proposed 2-tier architecture provides greater incentives for practical deployment than most other traceback techniques in literature. It is important to note here that the use of the current 16-bit Autonomous System Numbers (ASNs) for router labels as opposed to star coloring is not advisable, as it limits scalability in the wake of the proposed use of 32-bit ASNs [30] in the future.

4.4. (Logical) Partitioned Coloring

Although star coloring provides a significant improvement over the naive hash fragments, the size of the palette (set of all colors) is highly skewed by the high degree nodes in the network. We propose an efficient graph coloring technique called (logical) partitioned coloring that eliminates the effect of these high degree nodes.




Figure 5. Normal and Logical Coloring

Consider an overlay node having 8 neighbors as in Fig. 5, requiring a minimum of 9 distinct colors. However, replacing this node with a logical set of 4 mesh-connected nodes, reduces the palette size to 6 colors. As the logical coloring increases every 3-hop path to a 4-hop path in the vicinity of that node, greater spatial reuse of colors is possible, which in turn reduces the total number of colors as the original neighbors of the node can now be identically colored.

$$\Omega(t) \leq \chi(G) \leq O(t^2) \quad (3)$$

$$\Omega(p + \lceil \frac{t}{p} \rceil) \leq \chi(G) \leq O(t^2) \quad (4)$$

We now generalize this concept of (logical) partitioned coloring. Let a node with (threshold) degree t be logically partitioned into p mesh-connected nodes. Uniform distribution of its neighbors among the different logical nodes reduces the effective node degree to $\lceil \frac{t}{p} \rceil$. The number of colors required for simple star coloring and (p, t) (logical) partitioned coloring are given in Eqn. 3 and Eqn. 4 respectively. We thus see a greater reduction in the palette size even for highly skewed topologies such as the Internet.

$$\frac{\text{star coloring}}{\text{logical coloring}} = \frac{(k) \text{ color path}}{(k+m) \text{ color path}} \quad (5)$$

Although (logical) partitioned coloring reduces the number of bits marked by each router along the attack path, the number of packets required for traceback is increased slightly, as at least two logical nodes replace each physical node. Let there be k routers on the attack path, m of which are logically partitioned. The sequence of colors now representing the attack path is as shown in Eqn. 5. For a router-level overlay topology, the router interfaces can be logically grouped into distinct colors, and the router can probabilistically mark either the ingress or egress logical color. For an AS-level overlay topology, the border routers can be logically grouped into distinct tier 1 colors, and the edge routers can then probabilistically mark their associated logical tier 1 color or the unique tier 2 color on the packet.

4.5. Protocol Specifics

We overload the 16-bit IP Identification field in the IP headers for packet marking as shown in Fig. 6. The use of this field has been extensively discussed in literature, and we refer the reader to previous work [1] for detailed discussions. We set aside a 1-bit field to denote either tier 1 or tier 2. A k -bit color field is used for marking the node color for that tier. The size of this field is left variable, but it is advisable to use a sufficiently large number of bits to account for scalability. Additionally, a 1-bit field is used for a distance metric. We assume the use of the novel TTL-based distance measurement technique as discussed in FIT [26]. This bit basically represents the number of routing hops between the marking router and the victim as an offset from the TTL field in the IP headers, the exact details of which are not restated here. The remaining bits of the IP ID field can be used for representing a random number generated by the marking node on the overlay. As there might exist multiple nodes having the same color at a particular distance from the victim, this random number can be used to distinguish between these different nodes. Additionally, this random number and distance metric can be used to pair-wise associate the different tier colors of a particular marking router that arrive in different packets.




Figure 6. IP ID Field Format for Traceback

Once an ongoing attack has been detected, the victim requires to start the attack path reconstruction process, using the information received from the traceback-capable routers in the Internet. We assume the existence of an (up-stream) overlay network map at the victim for successful traceback. It leverages the fact that an endhost can group together packets that traverse the same path during a TCP connection. As explained in the FIT map reconstruction

algorithm [26], when receiving packet markings from the same distance and TCP connection, an endhost can assume that the markings come from the same router in the Internet. However, as opposed to the FIT algorithm, we do not try to associate an IP address with every router identified, but store the path information (sequence of colors) as an attack tree, rooted at the endhost. This map can then be used for attack path reconstruction during traceback, and a simple pushback mechanism can be used to throttle traffic at upstream routers, to mitigate the effects of the attack.

5. Protocol Evaluation and Analysis

We provide a detailed theoretical analysis of the proposed traceback scheme here and also evaluate its performance on real Internet AS-level topologies. The performance metrics studied include the number of packets required for successful traceback, and the number of bits to be marked on a packet by any router along the attack path.

5.1. Theoretical Analysis

We consider 2 types of traceback possible with the proposed 2-tier architecture, namely *minimal traceback* and *complete traceback*. The *minimal* traceback corresponds to identification of the AS-level attack path from the tier 1 labels alone. The intra-AS path and the inter-AS path (in case of multiple peering points) are not known here. The *complete* traceback corresponds to identification of the entire attack path from both the tier 1 and tier 2 labels. It can be trivially inferred that the *minimal* traceback would almost always finish faster than the *complete* traceback as fewer packets are required for successful traceback. As the *minimal* traceback acts as a preview to the *complete* traceback information obtained apriori, mitigation strategies and the resulting negotiations with other collaborating Internet filters can thus be initiated as soon as the *minimal* traceback completes. However, the actual deployment of the mitigation strategy can be deferred till the *complete* traceback is successful, as it provides the precise information required to appropriately start throttling of attack packets.

Let us assume that there are d routers along the attack path from the attacker to the victim, where each router has a marking probability of p . Let the probability of marking a tier 1 label be q , and consequently the probability of marking a tier 2 label be $(1 - q)$. Let us denote by $P_i(M)$, $P_i(M_1)$ and $P_i(M_2)$ the probability that the packet received at the victim is marked, marked with a tier 1 label, and marked with a tier 2 label respectively, by a router i hops from the victim, as in Eqns. 6(a), 6(b) and 7 respectively.

$$P_i(M) = p(1 - p)^{i-1} \quad P_i(M_1) = pq(1 - p)^{i-1} \quad (6)$$

$$P_i(M_2) = p(1 - q)(1 - p)^{i-1} \quad (7)$$

Let us now conservatively assume that (marked) packets from all the routers appear with the same likelihood as the furthest router. Since these probabilities are disjoint, the probability that a given packet delivers a mark from some router is $P(M)$ (Eqn. 8(a)). Similarly the probability that a given packet delivers a tier 1 (tier 2) mark from some router is $P(M_1)$ ($P(M_2)$), as in Eqn. 8(b) (Eqn. 9).

$$P(M) \geq dp(1 - p)^{d-1} \quad P(M_1) \geq dpq(1 - p)^{d-1} \quad (8)$$

$$P(M_2) \geq dp(1 - q)(1 - p)^{d-1} \quad (9)$$

Let us now consider the Generalized Coupon Collector Problem [31] [32] in Probability Theory, which examines the problem where we wish to sample with replacement from a population of k different items until, and only until, all k items are represented in the sample. If every item has an associated probability p_i with which it will be selected on any draw, then it gives us a bound on the expected size of the sample. For our 2-tier traceback technique, the sample space consists of the d tier 1 labels and the d tier 2 labels. The *minimal* traceback consists of obtaining the d tier 1 labels, while the *complete* traceback consists of obtaining the entire set of $2d$ labels.

For *minimal* traceback, let X_m^1 denote the number of trials required to select one of each of the d equi-probable labels (Eqn. 10). Therefore, the number of packets, X_m , to reconstruct an AS-path of length d is given in Eqn. 11.

$$E[X_m^1] = d \sum_{i=1}^d \frac{1}{i} \approx d \ln(d) + O(1) \quad (10)$$

$$E[X_m] < \frac{\ln(d)}{pq(1 - p)^{d-1}} \quad (11)$$

For *complete* traceback, let X_c^1 denote the number of trials required to select one of each of the $2d$ equi-probable labels (Eqn. 12 to the second approximation).

$$E[X_c^1] \approx \sum_{1 \leq i_1 \leq 2d} \frac{1}{p_{i_1}} - \sum_{1 \leq i_1 \leq i_2 \leq 2d} \frac{1}{p_{i_1} + p_{i_2}}$$

$$\text{where } p_i = q \quad (\forall 1 \leq i \leq d),$$

$$\text{and } p_i = (1 - q) \quad (\forall (d + 1) \leq i \leq 2d) \quad (12)$$

However, both the tier 1 and tier 2 labels are equally important in obtaining the complete information about the attack path, and hence we now make the simplifying assumption that $q = \frac{1}{2}$. If $q > \frac{1}{2}$, *minimal* traceback would finish earlier, but *complete* traceback process would take much longer. Similarly, if $q < \frac{1}{2}$, *minimal* traceback would take longer to finish. As we would ideally like to minimize finish time for *minimal* traceback, and then subsequently minimize time taken for *complete* traceback, we choose the

	BGP Tables		BGP Updates		Skitter		WHOIS	
Number of nodes	17446		17417		9204		7485	
Number of edges	40805		42484		28959		56949	
Coloring	Normal	Logical	Normal	Logical	Normal	Logical	Normal	Logical
Avg. 1-Hop degree	4.68	4.68	4.88	4.83	6.29	6.04	15.21	14.66
Max. 1-Hop degree	2498	171	2627	146	2070	132	1079	101
Avg. 2-Hop degree	1074.32	157.91	1147.83	167.18	1247.44	216.56	613.80	503.76
Max. 2-Hop degree	12066	7492	12095	8467	8714	7291	5443	9258
Number of colors	2499	191	2628	196	2071	164	1081	112
(Color) Bit-Space	12	8	12	8	12	8	11	7

Table 1. Normal and (Logical) Partitioned Coloring of AS Topology Data Sets

optimal value of q to be $\frac{1}{2}$. Consequently, the number of packets, X_c , required to reconstruct the entire attack path of length d is as shown in Eqns. 13, 14. The value of p , that minimizes the number of attack packets is given by Eqn. 15.

$$E[X_c^1] = 2d \sum_{i=1}^{2d} \frac{1}{i} \approx 2d \ln(2d) + O(1) \quad (13)$$

$$E[X_c] < \frac{2 \ln(2d)}{p(1-p)^{d-1}} \quad (14)$$

$$p = \frac{1}{d} \quad (\text{for minimal \& complete traceback}) \quad (15)$$

AS-path Length	3	4	5	6	7	8	9
Overlay Hops (d)	6	8	10	12	14	16	18
Minimal Traceback	54	85	119	156	194	234	275
Complete Traceback	75	113	155	199	245	292	341

Table 2. Number of Packets for Traceback

Table 2 shows the number of packets required for *minimal* traceback, and also for *complete* traceback. The 2-tier architecture along with logical star coloring of the traceback overlay, provides us significant improvement over the traditional PPM schemes in literature. The traditional traceback schemes in literature require a few thousand packets, while even the best known schemes require hundreds of packets. The proposed scheme requires just over 150 packets (AS-path length ≈ 5) from each attack source to trace the attack source and also the attack path.

5.2. Experimental Evaluation

We evaluate the feasibility of the proposed techniques here by accurately quantifying the size of the palette required to star color the Internet AS-level topology.

There are various ways in which one can infer the AS-level topology in today's Internet. Traceroute-based techniques have been used in the *skitter* tool developed by CAIDA, to make Internet topology measurements [33]. The RouteViews Project collects and archives both the static snapshots of the BGP routing tables and the dynamic data

in the form of BGP message dumps, and can be used to infer the AS-level topology [34]. WHOIS is a collection of databases containing a wide range of information useful to network operators, and RIPE's WHOIS database can also be used to construct an AS-level topology [35]. The various AS-level topologies obtained by these different techniques have been extensively studied in [27], [36]. We thus use four different AS-level topology data sets in our experimental evaluation, namely static BGP Tables, dynamic BGP Updates, Skitter and WHOIS.

Let us now define a few notations for simplicity sake - the number of 1-hop neighbors of a node as its *1-degree* and the number of distinct nodes that are at most 2-hops away from a node as its *2-degree*. On similar lines, let us denote by *1-color* and *2-color*, the number of *1-degree* and *2-degree* nodes having distinct colors respectively. We study both normal star coloring and (p,t) logical coloring techniques here. Every node having degree greater than t is logically partitioned into p partitions, as explained above.

Internet AS-level Topology Analysis: As the size of the palette required for star coloring is bounded on the lower side by a node's *1-degree*, and on the higher side by its *2-degree*, we evaluate these metrics for all the four data sets. The *1-degree* and *2-degree* distributions are presented in Fig. 7 and Fig. 9 respectively, and also tabulated in Table 1.

The *1-degree* distribution shows that very few nodes have a high node degree. As the palette size for star coloring is at least as high as the maximum node degree in the network, we would ideally like to eliminate these nodes from the network. Thus employing logical partitioning of the nodes would result in multiple nodes each having smaller node degree, and hence limiting the growth of the palette size. For a highly connected network, it is evident that the palette size would be governed more by a node's *2-degree*, as illustrated in Section 3. Although the AS-level topology is hardly a complete graph, it would still be preferable to reduce the *2-degree* in the network to limit the growth of the palette size. As employing logical partitioning not only reduces a node's *1-degree*, but also its *2-degree*, it is bene-




Figure 7. Normal Coloring: AS 1-Hop Degree Distribution




Figure 8. (15,50) Coloring: AS 1-Hop Degree Distribution




Figure 9. Normal Coloring: AS 2-Hop Degree Distribution




Figure 10. (15,50) Coloring: AS 2-Hop Degree Distribution




Figure 11. Tables




Figure 12. Updates




Figure 13. Skitter




Figure 14. WHOIS

ficial to apply logical partitioning for the high degree nodes in the different AS-level topologies.

The *1-degree* and *2-degree* distributions for the (15,50) (logical) partitioned coloring are presented in Fig. 8 and Fig. 10 respectively. The *1-degree* distribution is clipped above the specified threshold and the *2-degree* distribution also shows a quantum shift of the cloud towards lesser *2-degree* values. It is important to note that the threshold scheme is run for only the original physical topology nodes and not recursively for the (logical) partitioned nodes, i.e. we limit our algorithm to a single level of partitioning, as higher levels of partitioning implies multiple colors to be marked by each router and thereby increasing the number of packets required for successful traceback.

Internet AS-level Topology Coloring Analysis: The palette size for a general (p, t) threshold scheme is presented in Fig. 11, Fig. 12, Fig. 13 and Fig. 14, for the different data sets respectively. These graphs indicate that increasing number of partitions or decreasing the threshold beyond a certain limit, only provides incremental benefit at the cost of increased partitioning and longer traceback delay.

Fig. 15 and Fig. 18 present the total number of ASes having any particular color in the palette, for the normal and (logical) partitioned coloring respectively. The AS color distribution graph indicates that the number of nodes having higher color indices is very small, and is as expected due to the greedy nature of the algorithm (see Section 3) used for star coloring. The AS color distribution also shows a significant reduction in the palette size, achieved due to (logical) partitioned coloring, as against normal star coloring.

The *1-color* and *2-color* distributions for normal color-

ing are presented in Fig. 16 and Fig. 17 respectively, while that for logical coloring are presented in Fig. 19 and Fig. 20 respectively. The *1-color* and *2-color* distributions show similar behavior as the *1-degree* and *2-degree* distributions respectively, and arguing along similar lines, we conclude that logical partitioning of AS-level topologies is necessary to avoid explosion of the palette size.

6. Conclusions

Although IP Traceback has emerged as a promising means of defense against large scale DDoS attacks in the Internet, many traceback techniques proposed in literature suffer from problems such as scalability in size as well as number of attackers, expected duration to complete traceback successfully, and the expected size of the encoded traceback information. In this paper, we propose a graph-coloring approach to address these issues specifically, wherein the attack path is identified as a sequence of colors (of the routers). The intentional spatial reuse of the bit-space in representing the colors thus requires fewer bits to be marked on a packet, while also requiring fewer packets to achieve traceback. The enhanced (logical) partitioned coloring technique and the 2-tier traceback architecture provide greater flexibility and better performance as compared to other known techniques. We have additionally validated our proposed claims by analyzing the proposed techniques using real AS-level topologies of the Internet.

References

- [1] S. Savage et. al., "Practical Network Support for IP Traceback," in *Proc. ACM SIGCOMM*, pp. 295-306, Aug. 2000.
- [2] S. M. Bellovin, "ICMP traceback messages," *Internet Draft: draft-bellovin-itrace-00.txt*, Mar. 2000.




Figure 15. Normal Coloring, AS Color Distribution




Figure 16. Normal Coloring, AS 1-Hop Color Distribution




Figure 17. Normal Coloring, AS 2-Hop Color Distribution




Figure 18. (15,50) Coloring, AS Color Distribution




Figure 19. (15,50) Coloring, AS 1-Hop Color Distribution




Figure 20. (15,50) Coloring, AS 2-Hop Color Distribution

- [3] A. C. Snoeren et al., "Hash-Based IP Traceback," in *Proc. ACM SIGCOMM*, Aug. 2001.
- [4] M. Sung, J. Xu, "Intelligent Packet Filtering: A Novel Technique for defending against DDoS Attacks", in *IEEE TPDS*, 2003.
- [5] H. Burch, B. Cheswick, "Tracing Anonymous Packets to their approximate source", in *Proc. USENIX LISA*, Dec. 2000.
- [6] Li et al., "Large-Scale IP Traceback in High-Speed Internet: Practical Techniques and Theoretical Foundation", in *Proc. IEEE Symposium on Security and Privacy*, 2004.
- [7] A. Belenky, N. Ansari, "IP Traceback with Deterministic Packet Marking", in *IEEE Communication Letters*, Apr 2003.
- [8] D. Song, A. Perrig, "Advanced and Authenticated Marking Schemes for IP traceback", in *Proc. IEEE INFOCOM*, 2001.
- [9] M. Adler, "Tradeoffs in Probabilistic packet marking for IP traceback", in *Proc. STOC*, pp. 407-418, 2002.
- [10] T. Peng, C. Leckie, K. Ramamohanarao, "Adjusted Probabilistic Packet Marking for IP Traceback", in *Proc. Networking*, 2002.
- [11] I. Hamadeh, G. Kesidis, "Performance of IP Address Fragmentation Strategies for DDoS traceback", in *Proc. IEEE IPOM*, 2003.
- [12] M. Waldvogel, "GOSSIB vs Traceback Rumors", in *ACSAC*, 2002.
- [13] M. Ma, "Tabu Marking Scheme for Traceback", in *IPDPS*, 2005.
- [14] B. Al-Duwairi, T. Daniels, "Topology Based Packet Marking," in *Proc. IEEE ICCCN*, 2004.
- [15] M. Muthuprasanna, G. Manimaran, "Space-Time Encoding for DDoS Attack Traceback", in *Proc. IEEE GLOBECOM*, 2005.
- [16] K. Choi, H. Dai, "A Marking Scheme using Huffman Codes for IP Traceback", in *Proc. ISPAN*, 2004.
- [17] C. Bai, G. Feng, G. Wang, "Algebraic Geometric Code Based IP Traceback", in *Proc. IEEE IPCCC*, 2004.
- [18] Y. Sawai, M. Oe, K. Iida, Y. Kadobayashi, "Performance Evaluation of Inter-Domain IP Traceback", in *Proc. IEEE ICT*, 2003.
- [19] B. Al-Duwairi, G. Manimaran, "Novel Hybrid Schemes employing Packet Marking & Logging for Traceback", in *IEEE TPDS*, 2005.
- [20] C. Gong, K. Sarac, "IP Traceback based on Packet Marking and Logging", in *Proc. ICC*, 2005.
- [21] E. Lloyd, S. Ramanathan, "On the complexity of distance-2 coloring", in *Proc. IEEE ICCL*, 1992.
- [22] Sajal Das, Irene Finocchi, "Star-coloring of graphs for conflict-free access to parallel memory systems", in *Proc. IEEE IPDPS*, 2004.
- [23] S. T. McCormick, "Optimal approximation of sparse Hessians and its equivalence to a graph coloring problem", in *Math. Prog.*, 1983.
- [24] A. Gebremedhin, F. Manne, A. Pothen, "Parallel distance-k coloring algorithms for numerical optimization", in *Proc. Euro-Par*, 2002.
- [25] D. Bozdag et al., "A parallel distance-2 graph coloring algorithm for distributed memory computers", in *Proc. HPCC*, 2005.
- [26] A. Yaar, A. Perrig, D. Song, "FIT: Fast Internet Traceback", in *Proc. IEEE INFOCOM*, 2005.
- [27] P. Mahadevan et al., "Lessons from three views of the Internet topology", CAIDA Technical Report TR-2005-02.
- [28] D. Magoni, J. Pansiot, "Analysis of the Autonomous System Network Topology", in *ACM CCR*, July 2001.
- [29] M. Fayed et al., "On the size distribution of Autonomous Systems", Technical Report, Boston University, Jan 2003.
- [30] Q. Vohra, E. Chen, "BGP support for Four-octet AS Number Space", *Internet Draft: draft-ietf-idr-as4bytes-12.txt*, Nov. 2005.
- [31] H. von Schelling, "Coupon Collecting for Unequal Probabilities", in *American Mathematical Monthly*, 1954.
- [32] S. Lu, S. Skiena, "Filling a Penny Album", in *CHANCE*, 2000.
- [33] K. C. Claffy, T. E. Monk, D. McRobb, "Internet tomography", Nature, Jan. 1999, <http://www.caida.org/tools/measurement/skitter/>
- [34] "RouteViews Project", <http://www.routeviews.org/>
- [35] "Internet Routing Registries", <http://www.irr.net/>
- [36] "Comparative analysis of the Internet AS-level topologies", <http://www.caida.org/analysis/topology/>