CERIAS Security Seminar *Jan. 17, 2001*

Distributed Denial-of-Service Attack Prevention using Route-Based Distributed Packet Filtering

Heejo Lee
heejo@cerias.purdue.edu

Network Systems Lab and CERIAS

This is joint work with Prof. Kihong Park.

Outline

- Introduction to Denial-of-Service (DoS) attacks
- Related works and research motivation
- Route-based distributed packet filtering
- Effectiveness for DDoS attack prevention
- Concluding remarks

Denial-of-Service Attack

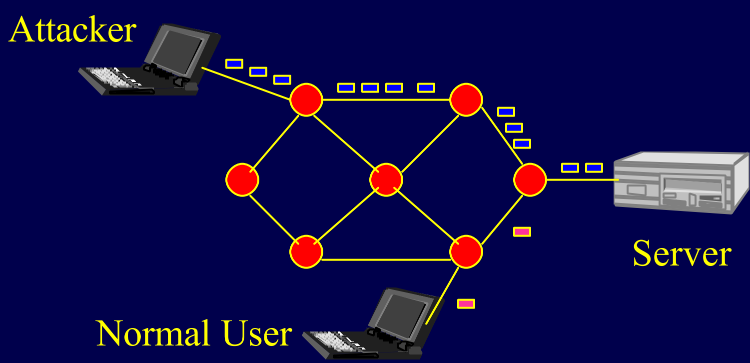

Overwhelming of fake requests consumes all resources on a server or network!

DoS Attack

- DoS Attack Style
  - Demanding more resources than the target system can supply
  - Network-based DoS attacks with IP spoofing
  - Launching a distributed DoS (DDoS) attack
- DoS Attack Impact
  - Complete shutdown a web site.
    - Yahoo, CNN, Amazon, eBay (Feb. 2000)
  - The greatest threat in e-commerce.

DoS Attack Reports

- 2000 Information Security Industry Survey, Sep. 2000
  - 51% companies experienced DoS attacks.
- Top 10 Security Stories of 2000, ZDNet News, Dec. 2000
  - No.1 and No.2 stories are related to DoS.
- New Year's DDoS Advisory, NIPC, Dec. 2000
  - More effective DDoS exploits have been developed.
  - Trin00,Tribal Flood Net, TFN2K,MStream, Stacheldraht, Trinity V3, Shaft, Godswrath…

Intrinsic Problems in DoS Attack

- Vulnerability
  - Any system is susceptible to DoS attacks.
- Traceback Problem
  - IP spoofing enables an attacker to hide his identity.

Easy to attack, hard to protect!

Related Works

- Resource management
  - Mitigating the impact on a victim [Schuba97, Banga99].
  - Does not eliminate the problem.
- Edge filtering
  - Ingress filtering in border gateways [Ferguson00].
  - Requires prolonged period for broad deployment.
- IP traceback
  - Trace back to the origin of the attacking source.
  - Recently a few approaches have been proposed:

    Traffic analysis,ICMP trace messages, packet marking.

IP Traceback Mechanisms

- Traffic analysis [Sager98]
  - Trace via traffic logs at routers
  - High storage and processing overhead
- ICMP traceback messages [Bellovin00]
  - IETF itrace working group
  - Extra traffic and authentication problem
- Probabilistic packet marking [Savage00]
  - Probabilistically inscribe trace information on a packet
  - Efficient and implementable

Probabilistic Packet Marking

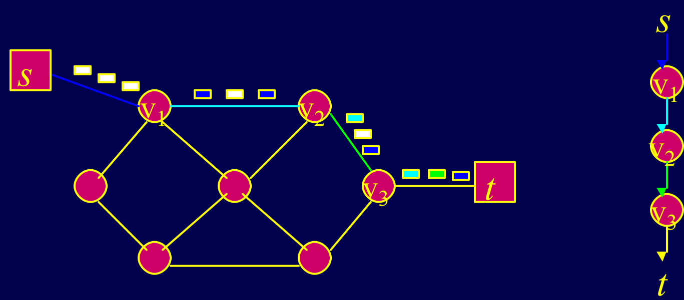

Router v_i inscribes (v_{i-1}, v_i) onto a packet with probability p.

Attack path reconstruction

Probabilistic Packet Marking

- Probabilistic packet marking (PPM)
  - Probabilistically inscribe its local path information
  - Use constant space in the packet header
  - Reconstruct the attack path with high probability
- Merits
  - Efficiency and implementability
- Weaknesses
  - Marking field spoofing problem

Marking Field Spoofing on PPM

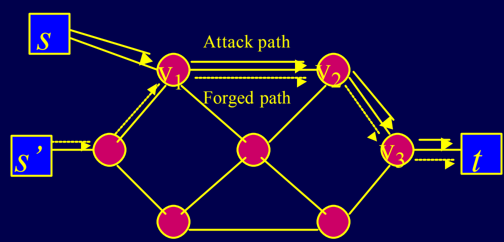

An attacker can use fake marking to forge a path that is equally likely as the true attack path.

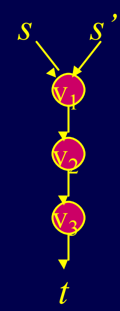

Reconstructed attack path

Effectiveness of PPM

Analysis under marking field spoofing:

- Single source attacks
  - Effective localization to within 2~5 sites.
- Distributed attacks
  - Uncertainty amplification on DDoS.
- Further information
  - Park and Lee, Tech. Rep. CSD-00-013, Purdue University, which will be presented at IEEE INFOCOM 2001.
    http://www.cs.purdue.edu/nsl/ppm-tech.ps

Summary of DoS Attack Study

	Resource Manage	Ingress Filtering	Traffic Analysis	ICMP Messages	PPM	DPF
Cost	X	O	X	Δ	Δ	Δ
Deployment	O	Δ	Δ	Δ	O	O
Traceback	X	X	O	O	O	O
Protection	Δ	Δ	X	X	X	O
Scalability	X	X	X	X	X	O

X: poor, Δ: good, O: excellent

Research Motivation

- Weaknesses of IP Traceback Mechanisms
  - Post-mortem: debilitating effect before corrective actions
  - Bad scalability: susceptible to DDoS
- Demand for DDoS protection
  - Find a protective and incrementally deployable approach

Distributed Packet Filtering (DPF)

- Packet filtering using routing information
  - Filter spoofed packets traveling unexpected routes from their specified addresses.
- Distributed filtering
  - Collective filtering on autonomous systems (AS).

Route-Based Detection of Spoofed Packets

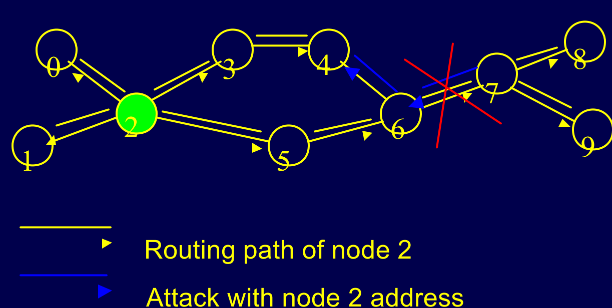

System Model for DPF

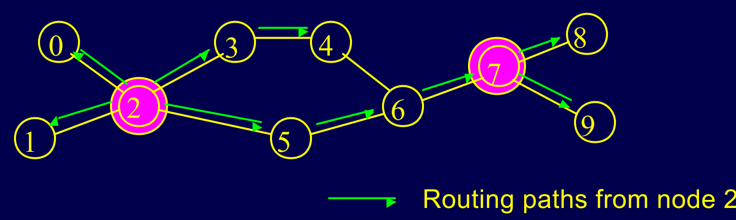

- G: network topology
- T: filtering nodes
- R: routing policies
- F: filtering function

Network (G) and Filtering Nodes (T)

- AS Connectivity Graph *G(V,E)*
  - *V*: a set of nodes, where a node is an AS. $|V|=n$.
  - *E*: a set of links in *G*.
- Node Type
  - *T*-node: a set of filtering nodes.
    - Filter internal traffic as well as incoming traffic
  - *U*-node: a set of nodes without filtering.
    - $V = T \cup U$

Routing Policies (R)

- Routing (R)
  - $R(u,v) \subseteq \mathcal{L}(u,v)$
    where $\mathcal{L}(u,v)$ is set of all loop-free paths from u to v.
- Routing Policies
  - Tight: single shortest-path routing, $|R(u,v)| = 1$.
  - Multipath: multiple routing paths, $1 < |R(u,v)| < |\mathcal{L}(u,v)|$.
  - Loose: any loop-free path routing, $R(u,v) = \mathcal{L}(u,v)$.

Filter (F)

- Filter for a link *e*
  - A function of a source and a destination

$$F_e : V^2 \rightarrow \{0,1\}$$

- Route-based filters
  - Maximal filter
  - Semi-maximal filter

Route-Based Filters

- Maximal filter
  - Use of all (src/dst) pairs of routing paths.
  - Huge filtering table *O(n^2)*, e.g., 4GB for 16bit AS's.

$$F_e(s,t) = \begin{cases} 0, & \text{if } e \in R(s,t); \\ 1, & \text{otherwise} . \end{cases}$$

- Semi-maximal filter
  - Use of only source addresses coming via the link.
  - *O(n),* e.g., 8KB for all AS's.

$$F_e'(s,t) = \begin{cases} 0, & \text{if } e \in R(s,v) \text{ for some } v \in V; \\ 1, & \text{otherwise.} \end{cases}$$

Semi-Maximal Filter Updates

BGP (Border Gateway Protocol) Routing Updates
- Initiated by node 2
- Shortest path routing

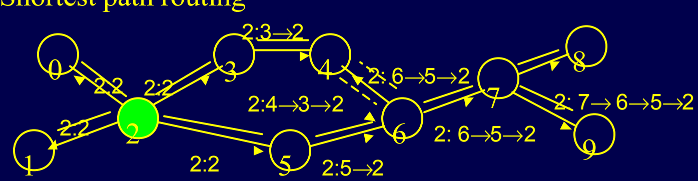

After Completing BGP Updates from Every Nodes

Routing Table of Node 2	Filtering Tables of Node 2	
0: 0	(0,2): 0111111111	0: allow
1:1	(1,2): 1011111111	1: deny
3:3	(3,2): 1110011111	
4: 3 →4	(5,2): 1111100000	
5: 5		
6: 5 →6		
...		

Filtering Effect

- Attack *a:(s,t)*
  - Attacker at node *a* sends *(s,t)* packets to node *t*.
- Spoofing range $S_{a,t}$ – attacker's point of view
  - a set of nodes with which node *a* can send spoofed packets to node *t*.
- Candidate range $C_{s,t}$ – victim's point of view
  - a set of nodes which can send *(s,t)* packets.

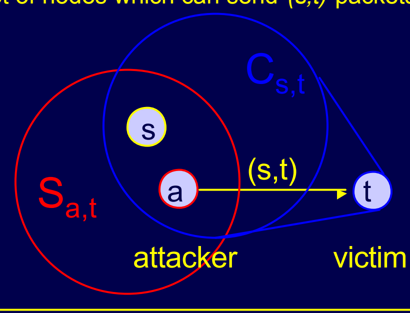

Distributed Filtering Effect

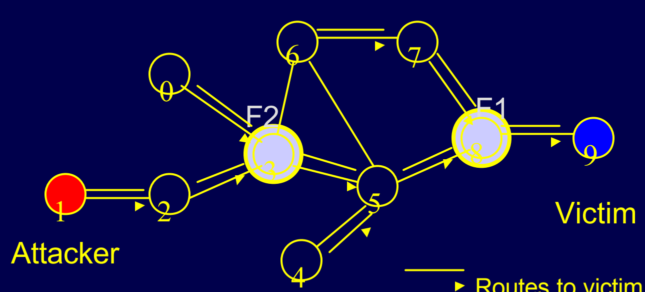

No filtering: $S_{1,9}$={0,1,2,3,4,5,6,7,8}

Filtering at F1: $S_{1,9}$={0,1,2,3,4,5}

Filtering at F1 and F2: $S_{1,9}$={1,2}

Experimental Environments

- Topology *G*
  - Internet AS connectivities from 1997~1999.
  - Random topologies.
- Routing *R*
  - Tight, multi-path routing policies.
- T-nodes *T*
  - R30: 30 percent of nodes chosen randomly.
  - R50: 50 percent of nodes chosen randomly.
  - *VC*: a vertex cover of *G(V,E)*.

Vertex Cover (VC)

- VC of G(V,E)
  - $\forall (u,v) \in E,\ u \in$ VC or $v \in$ VC
- T=VC
  - Any node in U has only T nodes as its neighbors.
- Finding a minimal VC
  - NP-complete problem
  - Two well-known algorithms used for finding a VC

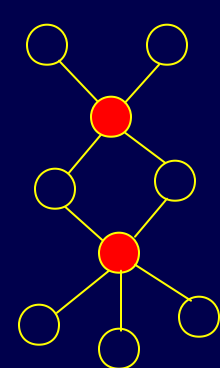

Metrics for Proactive Effect

- Perfect proactivity

$$\Phi_1(t) = \frac{\left|\{t : \forall a \in V, |S_{a,t}| \le t\right|}{n}$$

  - $\Phi_1(1)$: fraction of AS's safe from spoofing attack
- DDoS prevention

$$\Phi_2(t) = \frac{\left|\{a : \forall t \in V, |S_{a,t}| \le t\right|}{n}$$

  - $\Phi_2(1)$: fraction of AS's from which no spoofed packets coming
- Attack volume reduction

$$\Theta = \frac{\left|\{(a,s,t) : s \in S_{a,t}\right|}{n(n-1)^2} = \frac{\left|\{(a,s,t) : a \in C_{s,t}\right|}{n(n-1)^2}$$

  - Θ: penetrating ratio of spoofed packets

How DPF Works on Internet?

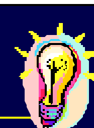

- Impractical perfect proactivity
  - $\Phi_1(1) \approx 1$ is hard to be achieved.
- Effective DDoS attack prevention
  - $\Phi_2(1) \approx 0.88$ renders most attack sites impotent.
- Significant attack volume reduction
  - $\Theta \approx 0$ for random source addresses.

Impractical Perfect Proactivity

- G: 1997 Internet connectivity (n=3015,|E|=5230)
- T: VC→ n
- R: Tight
- F: Semi-maximal

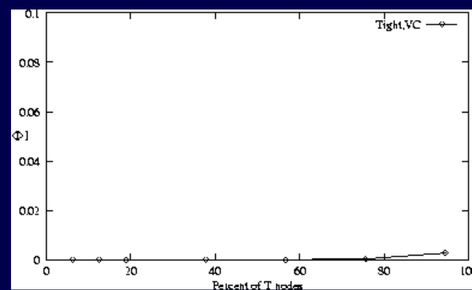

$\Phi_1(1)\approx 1$ is hard to achieve!

Perfect proactivity is practically useless objective.

Distributed Filtering for DDoS Attack Prevention 29/40

DDoS Attack Prevention

- G: 1997~1999 Internet connectivity
- T: VC
- R: Tight
- F: Semi-maximal

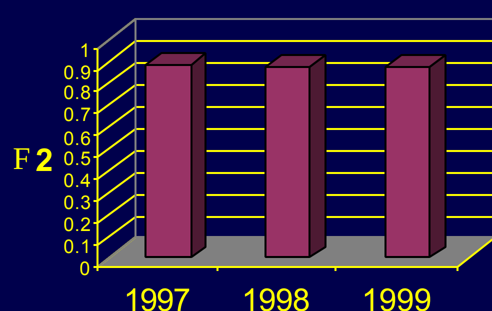

DPF renders 88% of possible attack sites impotent: effectively curtail the ability to mount DDoS attacks.

Distributed Filtering for DDoS Attack Prevention 30/40

Attack Volume Reduction

- Θ =0.0004 when T=VC
- 99.96% attack volume reduction

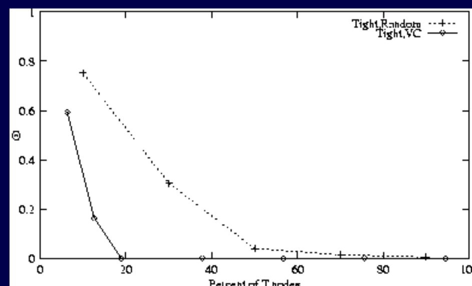

Randomly generated spoofed address has almost zero chance to reach its target!

Reactive Filtering Effect: Traceback

- IP Traceback Capability

$$\Psi_1(t) = \frac{\left|\{t : \forall s \in V, \left|C_{s,t}\right| \leq t\right|}{n}$$

- Localization: meaningful for τ greater than 1.
- $\Psi_1(5)$: fraction of AS's which can resolve the attack location to within 5 possible sites.

IP Traceback Effect

- Traceback capability
  - $\Psi_1(5)=1$ for 1997~1999 AS connectivities
  - Localization to within 5 possible sites

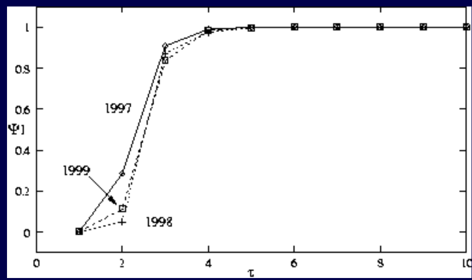

Filtering out many spoofed flows allows source identification of an attack location.

Efficient Semi-Maximal Filter

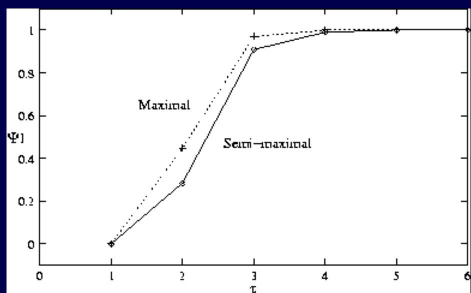

Maximal filter requires quadratic space, but results in marginal enhancement of traceback capability.

Effectiveness on Multi-path Routing

- Gradual reduction of traceback capability
  - $\Psi_1(\tau) \approx 1$ for =5~10 when the number of routing paths are 2~3.

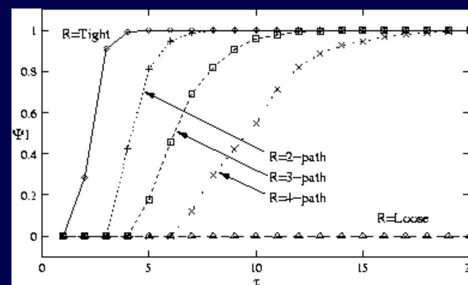

DPF is still effective on multi-path routing policies!

Impact of Network Topology

- Benchmarking network topologies
  - Internet AS connectivities from 1997-1999
  - Random graphs with link probability p
  - Power-law connectivity by Inet generator
- Topological impacts
  - Intimate relation to VC size and filtering performance
  - Internet has good characteristics for DPF
    - small VC and good performance

Internet AS Connectivity

- Small VC on Internet
  - Vertex covering with 18% nodes
  - Incremental deployment feasible

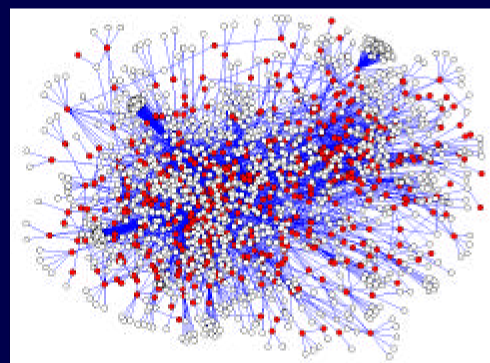

1997 Internet Connectivity
- Red nodes are in VC

Random Graph

- Random graph generation
  - Connecting any two nodes with a link probability p.
  - VC on random graphs requires 55% nodes.
  - Lower performance with more T nodes.

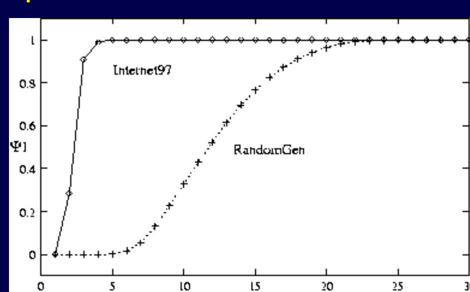

Inet Topology Generator

- Inet Generator (http://topology.eecs.umich.edu/)
  - Generate a graph with power-law connectivity.
  - VC on Inet graphs requires 32% nodes.
  - Small VC has more effectiveness.

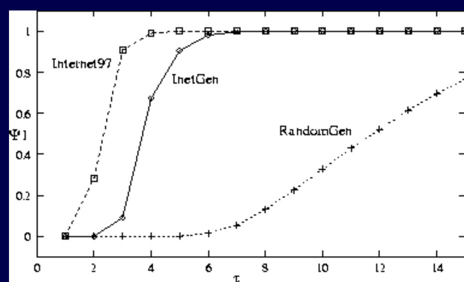

Summary of Dynamic Packet Filtering

- Distributed packet filtering
  - Packet filtering mechanism using routing information
- Practicality
  - Implementable with BGP
  - Incrementally deployable
- Effectiveness
  - Protection from DoS attacks
  - Prevention from DDoS attacks
  - Traceback capability