Is It Congestion or Denial of Service?

Nirwan Ansari
(joint work with Amey Shevtekar)
Advanced Networking Laboratory
Department of Electrical and Computer Engineering
New Jersey Institute of Technology
Newark, NJ 07102-1982, USA
Email: Nirwan.Ansari@njit.edu
http://web.njit.edu/~ansari

@ICCCN 2009, San Francisco

10 Most Wanted Botnets[†]

1. **Zeus**: 3.6 mil compromised US computers
2. **Koobface**: 2.9 mil compromised US computers
3. **TidServ**: 1.5 mil compromised US computers
4. **Tojan.Fakeavalert**: 1.4 mil compromised US computers
5. **TR/Dldr.Agent.JKH**: 1.2 mil compromised US computers
6. **Monkif**: 520K compromised US computers
7. **Hamweq**: 480K compromised US computers
8. **Swizzor**: 370K compromised US computers
9. **Gammima**: 230K compromised US computers
10. **Conficker**: 210K compromised US computers

[†]http://www.networkworld.com/news/2009/072209-botnets.html

Outline

- DoS/DDoS Attacks
- IP Traceback & DDoS Defense Schemes
- Low Rate DoS Attacks
- The Perfect Storm
- The Quiet Attack

DDoS Headlines

- 7/27/2009: AT&T blocked 4chan.org
- US & South Korea's government websites under DDoS attack. [July 2009]
- Massive Botnet DDoS Attack on mininova.org [March 2009]
- Estonia's government websites under Botnet DDoS attack. [May 2007]
- Cyber-warfare is here to stay..
- DoS attacks are *consistently placed among top 4 attacks* observed by various US industries since 2005. [CSI/FBI survey 2005-2008]
- Since last year a new entry, *Bots*, is making its way into the list of attacks. [CSI sruvey 2008]

Evolution of DoS Attacks

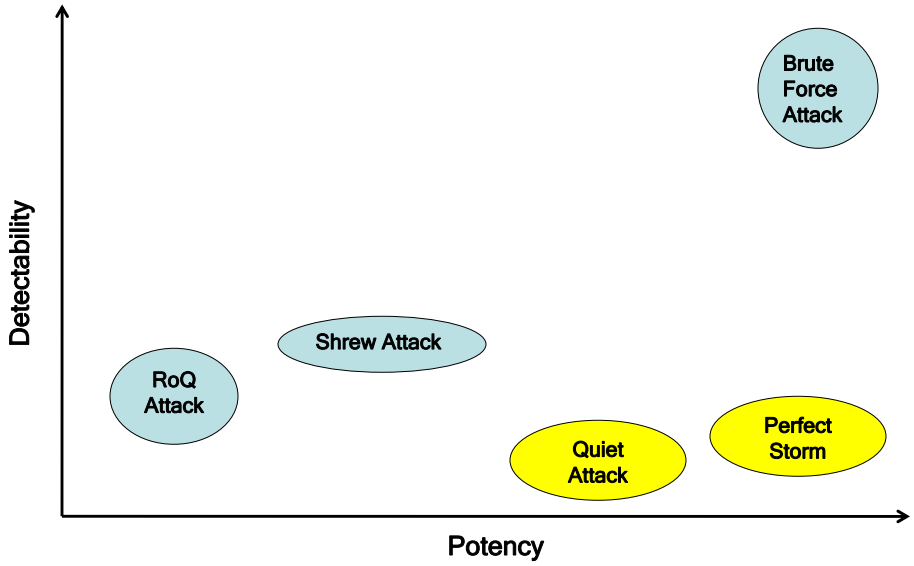

Motivation

One of the best methods of understanding your network security posture is to try to defeat it.

--Gordon “Fyodor” Lyon (Creator of famous Nmap Software)

Snap from Matrix Reloaded

Source: nmap.org

What is a DDoS Attack?

- Not a single attack source, but hundreds of attack sources concerted by the attacker.

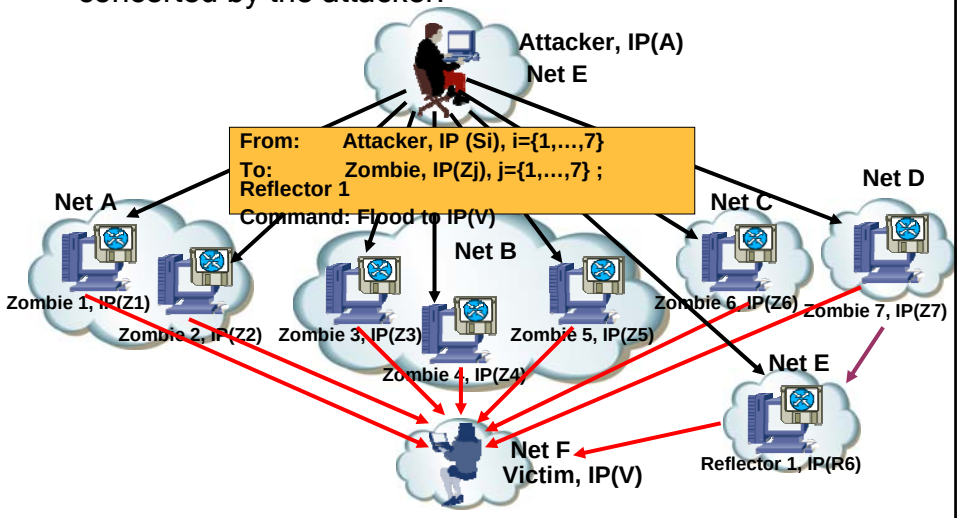

Characteristics of DDoS Attacks

- Differences from other attacks
  - ✓ Don't need to penetrate the target ahead of time
  - ✓ No sensitive information exposed
  - ✓ Usually, a large number of packets are sent to the victim from multiple sources
  - ✓ Source IP addresses are routinely forged
  - ✓ Source IP address spoofing not required with use of botnets.

Outline

- DoS/DDoS Attacks
- IP Traceback & DDoS Defense Schemes
- Low Rate DoS Attacks
- The Perfect Storm
- The Quiet Attack

Some of the Existing DDoS Defense Schemes

- Defending against Flooding Attacks
  - ✓ IP-traceback based filtering
  - ✓ RED-PD
  - ✓ Pushback
- Network Monitoring
  - ✓ PacketScore
  - ✓ SYN detecting
- Honeypot

Probabilistic Packet Marking (PPM)

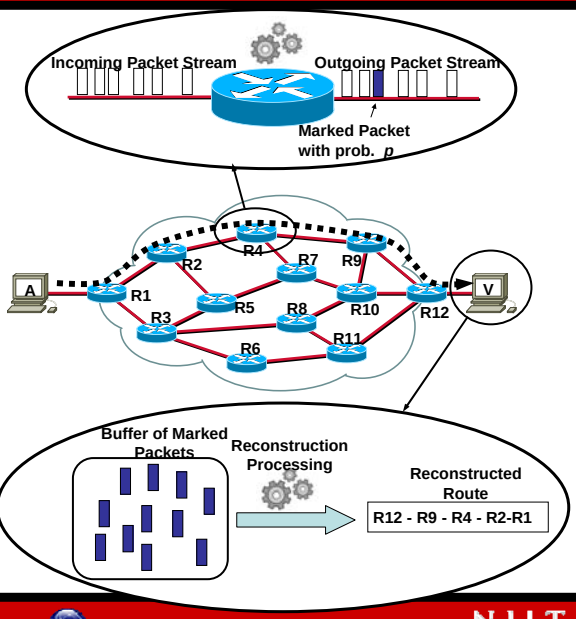

IP Traceback with DPM

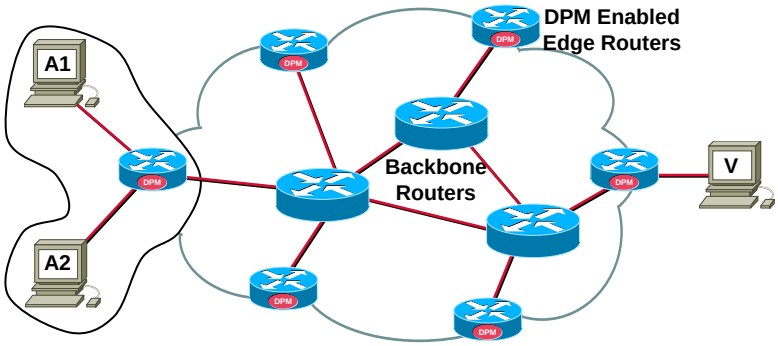

DPM Principles

- ***Interface***, not the Router is a unit of Traceback

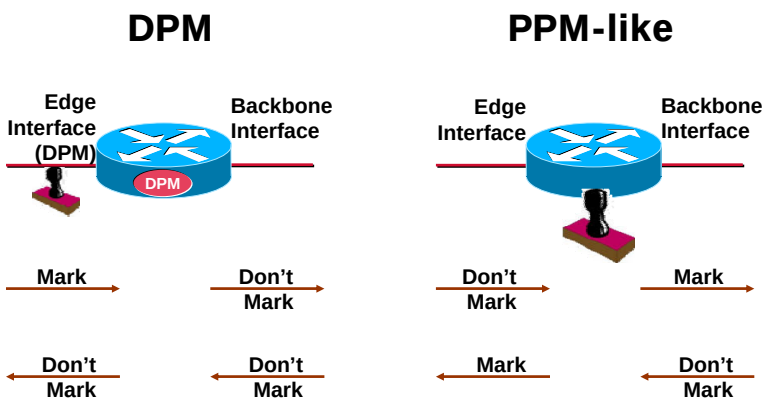

DPM Principles (cont'd)

- Only the ingress DPM-enabled edge interface marks packets
  - ✓ DPM Traceback = Ingress Address
  - ✓ For datagram networks (e.g. Internet), Ingress Address is as good as full-path traceback
- ALL packets are marked by DPM-enabled interface only
  - ✓ Prevents mark spoofing
  - ✓ Decreases traceback time

Defense against Flooding Attacks

- Basic Assumption:
  Sources sending a large number of packets are probably malicious
- Defending Approach:
  QoS-based Filtering/Dropping
  Pushback (to punish attack sources)

RED-PD Router

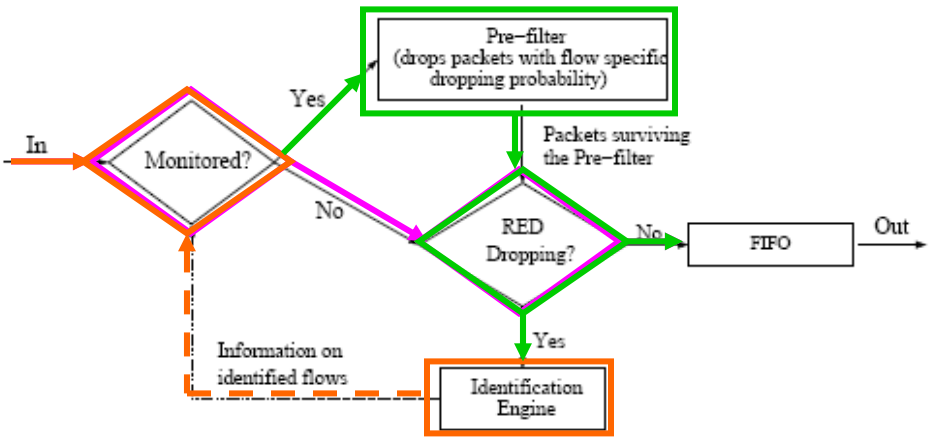

- RED-PD uses the RED drop history (mark history in the presence of ECN) to identify high-bandwidth flows and to confirm that the identified flow has in fact received loss events.
- Drop history represents flows that have been sent congestion signals.

Network Monitoring

- Basic Assumption:
  - The ratio between the number of ingress packets and the number of egress packets should lie in a certain range.
- Defending Approach:
  - Whether to accept the current packet is determined by the current network condition and the score distribution (PacketScore).

PacketScore: Statistics based-Overload Control Against DDoS Attacks

PACKETSCORE consists of 3 phases:

1. Detect the onset of an attack and identify the victim by monitoring four key traffic statistics of each protected target
2. Differentiate between legitimate and attacking packets based on a Bayesian-theoretic metric of each packet. The metric is called conditional legitimate probability (CLP).
3. Discard packets selectively by comparing the CLP of each packet with a dynamic threshold. The threshold is adjusted according to
   - the distribution of CLP of all suspicious packets and
   - the congestion level of the victim.

A Deployment of PacketScore

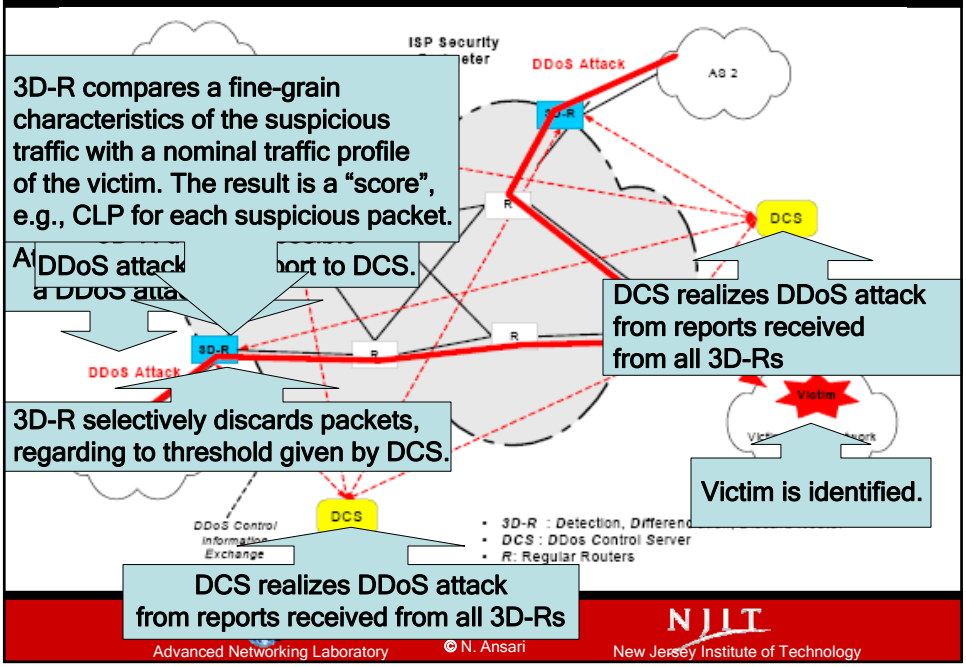

HONEYPOT

- A HONEYPOT is a system that is specifically designed to allow being probed or attacked. It is basically a trap set to detect, deflect, or counteract attempts of any attack from outsiders.
- It may consist of a stand-alone computer, a group of computers, or a network or subnetwork that appears to be part of private networks. However, it is actually isolated and protected.
- Servers and computers in a honeypot may carry "fake" data that seem to be "real" or valuable to attackers.
- Can be used to study the attack techniques and methodologies: how attack begins, what attacking tools are used, how frequent attacks occur, and so on.

HONEYPOT Illustration

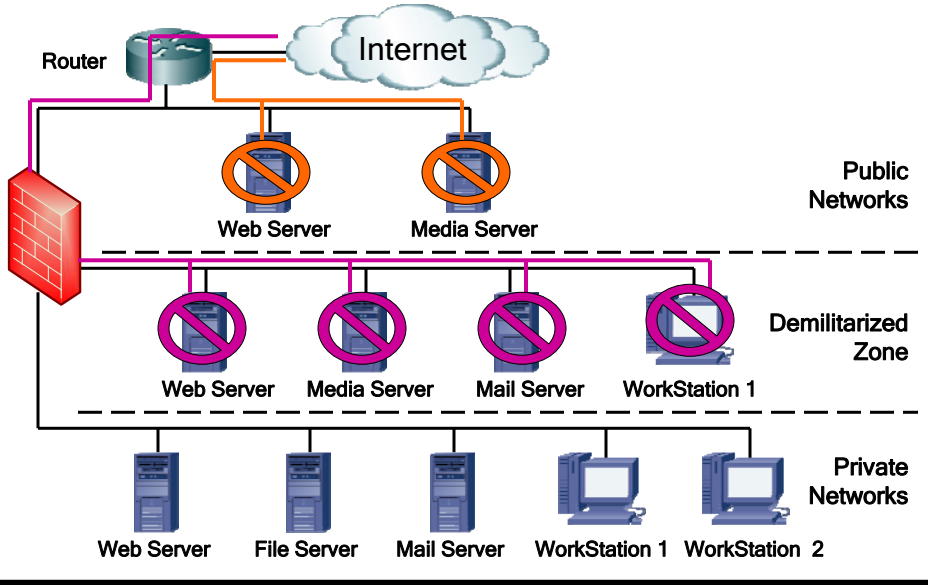

Outline

- DoS/DDoS Attacks
- IP Traceback & DDoS Defense Schemes
- Low Rate DoS Attacks
- The Perfect Storm
- The Quiet Attack

How Does Low Rate DoS Attack Work?

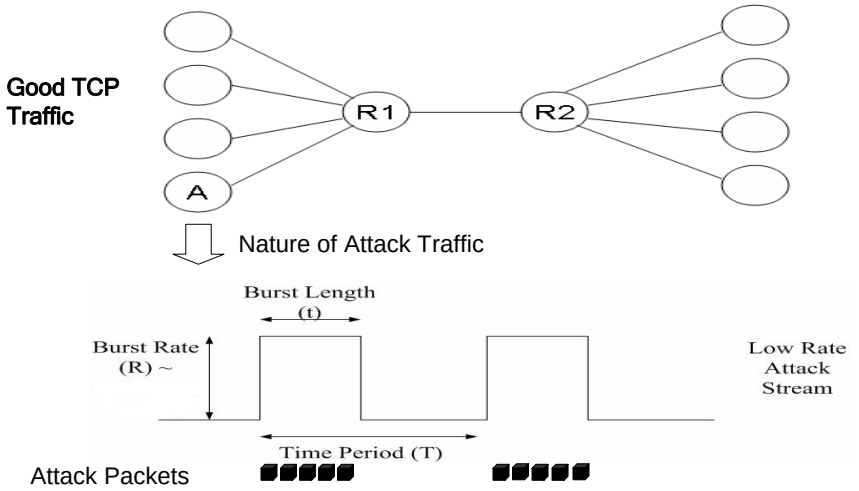

Characteristics of Low Rate DoS Attacks

- Average attack traffic rate remains low
- Factors leading to DoS:
  - ✓ High Rate UDP ~ Bottleneck link capacity or greater
  - ✓ t ~ RTTs of normal connections in the Internet
  - ✓ T ~ 1 second (shrew attack exploits minimum RTO property, Kuzmanovic *et al.*, Sigcomm 2003)
  - ✓ 1 < T < 10 seconds (RoQ attack, Guirguis *et al.*, ICNP 2004)

Time Difference Technique Illustration

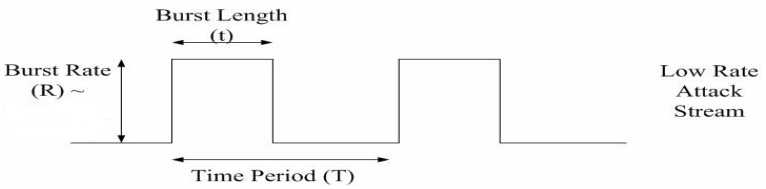

t1 t2 t3 t4 t5 t6 t7 t8 t9 t10

Results

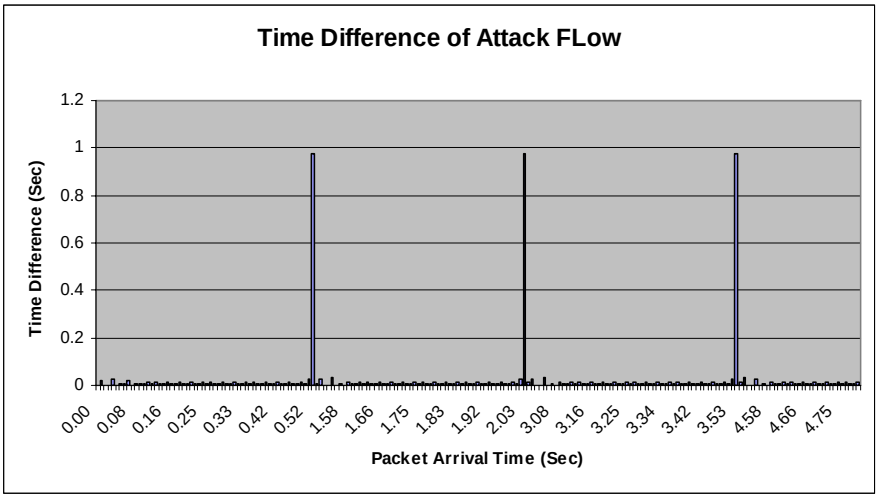

Results

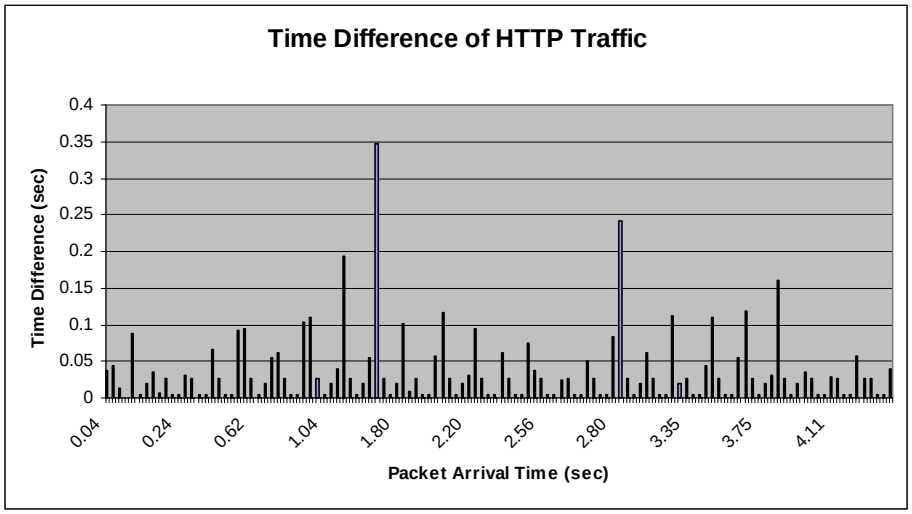

Time Difference Technique

- Per-flow approach.
- Proposed to work mostly at network edges.
- Important shortcoming is inability to find periodicity if attacker uses IP address spoofing.

Outline

- DoS/DDoS Attacks
- IP Traceback & DDoS Defense Schemes
- Low Rate DoS Attacks
- The Perfect Storm
- The Quiet Attack

Perfect Storm

- Short-lived TCP flows are known to be bursty.
- Use TCP along with limited UDP as the attack traffic.
- Incorporate network feedback mechanism in the attack process.
- Distribute the attack sources using botnets to launch the attack.
- TCP traffic coming from multiple sources that adapts to the bottleneck capacity cannot be easily classified as anomalous.

Perfect Storm Model

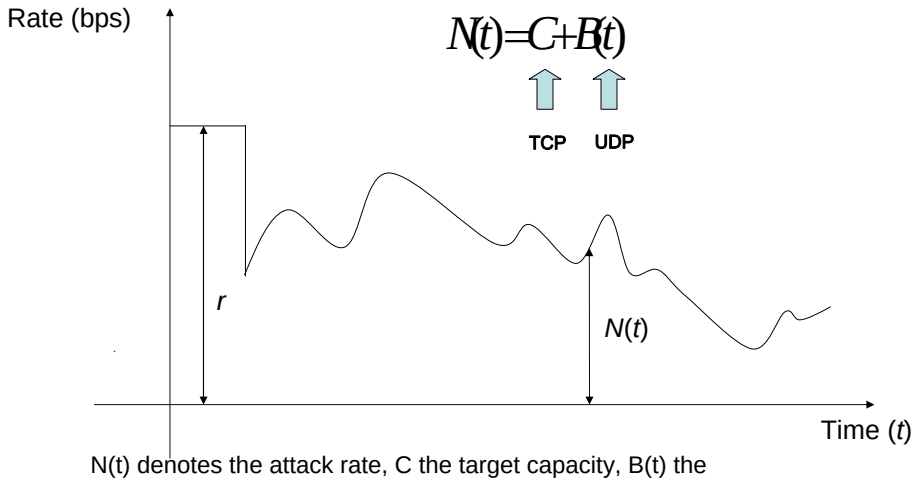

N(t) denotes the attack rate, C the target capacity, B(t) the available bandwidth at the bottleneck, and r is high rate UDP burst sent at the beginning of the attack.

Attack Mechanism

- Initial burst leads to instantaneous high packet loss, that creates more available bandwidth which is filled by attack short lived TCP flows.
- Attack TCP flow's congestion control mechanism will gradually increase the TCP rate.
- Meanwhile attack UDP traffic proportional to the available bandwidth is injected into the network.
- This constant presence of attack flows leads to significant damage to good traffic causing legitimate TCPs to enter into multiple timeouts.

Perfect Attack in the Internet

- Botnet -- Network of attack clients
- Target Router
- Web Server -- Source of data for short-lived TCP flows

Botnet

- Estimated size of a typical botnet: in thousands.
- Software exploits occurring frequently directly fuels the size of botnets.
- Studies have revealed botnets are active per timezone.
- We assume each timezone has a few big ISPs and those bots would be part of an attack.
- Consider EST in which AT&T & Comcast are the big ISPs.
- A perfect attack requires a botnet like clickbot to access webpages. In addition, few bots to send UDP traffic too.

Target Router

- A botmaster can issue bots to run a command tracert www.yahoo.com. Bots are assumed in the same timezone.
- Tracert works on windows boxes (typical bots) even with non-administrative accounts.
- A common IP address in all tracert outputs is selected as a target router.
- An attacker can always rely on network mapping tools like cartoreso (http://cartoreso.campus.ecp.fr/) or nmap.
- These famous free tools can provide precise network topology maps and details like IP addresses of routers.

Tracert Output

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Amey>tracert www.yahoo.com

Tracing route to www.yahoo-ht3.akadns.net [69.147.76.15]
over a maximum of 30 hops:

  1    <1 ms    <1 ms    <1 ms  c-76-124-80-91.hsd1.nj.comcast.net [76.124.80.91
]
  2     *        *        *     Request timed out.
  3    10 ms    20 ms    10 ms  ge-2-4-ur01.narlington.nj.panjde.comcast.net [68
.85.193.29]
  4    10 ms     9 ms    12 ms  po-10-ur02.narlington.nj.panjde.comcast.net [68.
86.158.178]
  5    13 ms    10 ms    12 ms  po-70-ar01.verona.nj.panjde.comcast.net [68.86.2
09.254]
  6    11 ms    11 ms    11 ms  be-80-crs01.plainfield.nj.panjde.comcast.net [68
.86.208.5]
  7    14 ms    13 ms    14 ms  pos-0-5-0-0-cr01.newyork.ny.ibone.comcast.net [6
8.86.90.29]
  8    15 ms    13 ms    13 ms  xe-10-3-0.edge1.NewYork2.Level3.net [4.71.184.1]

  9    12 ms    15 ms    18 ms  vlan69.csw1.NewYork1.Level3.net [4.68.16.62]
 10    22 ms    39 ms    15 ms  ae-63-63.ebr3.NewYork1.Level3.net [4.69.134.97]

 11    32 ms    18 ms    35 ms  ae-3.ebr3.Washington1.Level3.net [4.69.132.89]
 12    19 ms    18 ms    18 ms  ae-63-63.csw1.Washington1.Level3.net [4.69.134.1
62]
 13    20 ms    19 ms    18 ms  ae-11-69.car1.Washington1.Level3.net [4.68.17.3]

 14    21 ms    18 ms    28 ms  4.79.228.2
 15    18 ms    20 ms    22 ms  ae1-p140.msr1.re1.yahoo.com [216.115.108.17]
 16    19 ms    20 ms    19 ms  ge-1-41.bas-a2.re3.yahoo.com [66.196.112.201]
 17    21 ms    19 ms    22 ms  f1.www.vip.re1.yahoo.com [69.147.76.15]

Trace complete.
```

A Sample Cartoreso Map

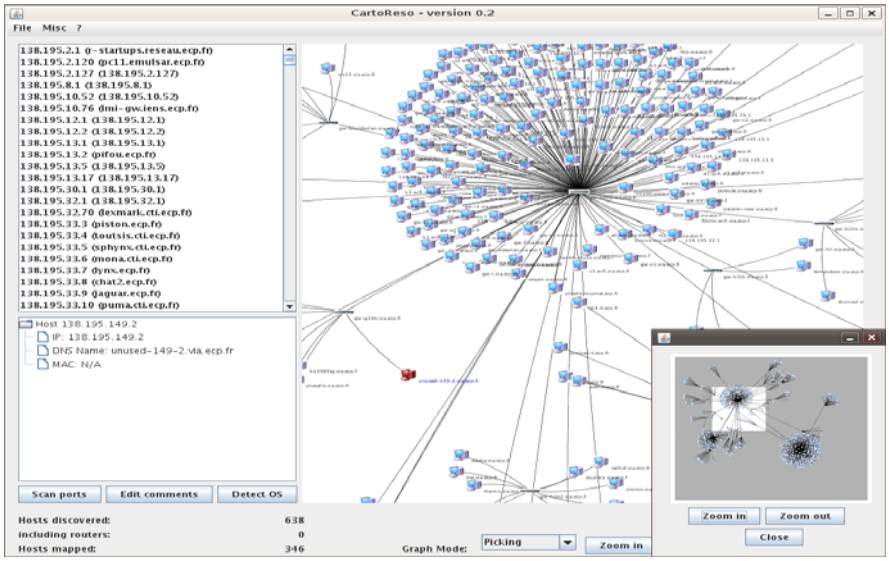

Source: http://cartoreso.campus.ecp.fr/screenshot1.png

Web Servers

- Bots use servers to orchestrate attack short-lived TCP flows by requesting webpages.
- There are free scripts that can estimate web page sizes on web servers.
- In a more complex technique, if a website uses CAPTCHA, bots can employ CAPTCHA cracking tools.
- Apparently, there are several websites which do not use CAPTCHA.
- Assume a 1Gbps link under attack.
- Consider bots using DSL/Cable that get around 500Kbps throughput for TCP.
- Then, the number of flows to fill up 1Gbps is 2000 (i.e., bots).

Webpage access strategy

- Consider 100 pages of 2000 websites are used in attack. Say non-commercial websites that would not use CDN like all universities, all govt organizations,etc.
- An attacker should also avoid websites that use CDN. Akamai is one of the biggest providers of CDN and since they carry 20% of web traffic , i.e., 80% of web traffic does not pass through Akamai, i.e., a lot of websites do not use CDN yet.
- Consider the following webpage(p) access pattern from 2000 websites (W) for each second:

  [W1p1 W2p1.. W2000p1] [W1p2 W2p2…W2000p2] ..
- Thus a single site experiences only one page worth traffic every sec which is clearly non-anomalous.

Web crawlers

```
Crawling site www.njit.edu using wget, with depth=8, limit=1MB, and looking for files
greater than 10KB

Running wget... Please wait
wget finished

No suitable files found (>10KB), for depth=8 and limit=1MB
Please try again with greater depth, greater limit or smaller file size
alizarin-48 abs6>: python crawler.py www.njit.edu 8 1

Usage: python crawler.py host [depth] [size] [limit]

Crawling site www.njit.edu using wget, with depth=8, limit=1MB, and looking for files
greater than 1KB

Running wget... Please wait
wget finished

Found 69 files greater than 1KB (host, filename, size):
www.njit.edu /corporate/uicomponents/mstoner/scripts/prototype.js 99594
www.njit.edu /atoz/index.html 48259
www.njit.edu /news/index.html 35256
www.njit.edu /corporate/uicomponents/mstoner/css/main.css 33839
www.njit.edu /features/student/savio-rozario.php 33771
www.njit.edu /admissions/index.php 31576
www.njit.edu /admissions/index.html 31569
www.njit.edu /alumni/index.html 29762
www.njit.edu /features/sceneandheard/bfa-art.php 29621
www.njit.edu /features/student/nickkintos.php 28735
www.njit.edu /admissions/undergraduate/parents.php 28485
www.njit.edu /academics/index.html 28218
www.njit.edu /admissions/includes/styles.css 27591
www.njit.edu /facultystaff/index.html 26407
www.njit.edu /about/key-contacts.php 26140
www.njit.edu /currentstudents/index.html 25925
www.njit.edu /admissions/openhouse/index.html 24939
www.njit.edu /cds/index.html 24581
www.njit.edu /studentlife/index.html 24552
www.njit.edu /news/2008/2008-375.php 21074
www.njit.edu /research/index.html 20625
www.njit.edu /about/index.html 19428
www.njit.edu /giving/index.html 19025
www.njit.edu /academics/commence/2008/mayne.php 18966
www.njit.edu /news/2008/2008-266.php 18786
www.njit.edu /workingandlearning/corporatepartners.php 18401
www.njit.edu /news/2008/2008-367.php 18387
```

Simulation Results

Attack effects on throughput of long lived TCP flows

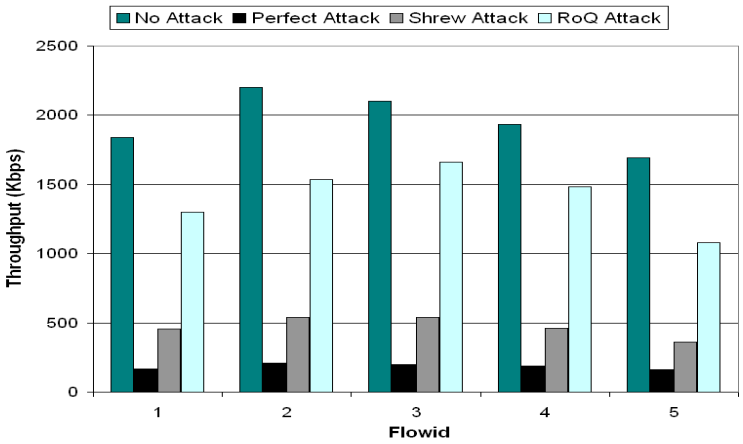

Effects on VoIP Traffic

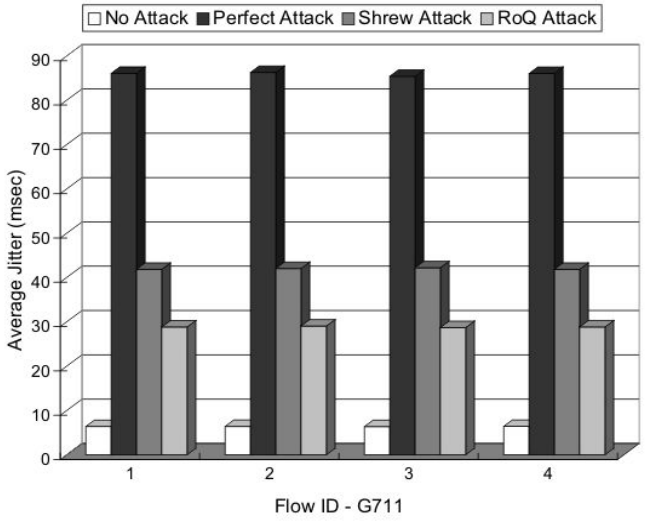

Effects on VoIP Traffic (Cont'd)

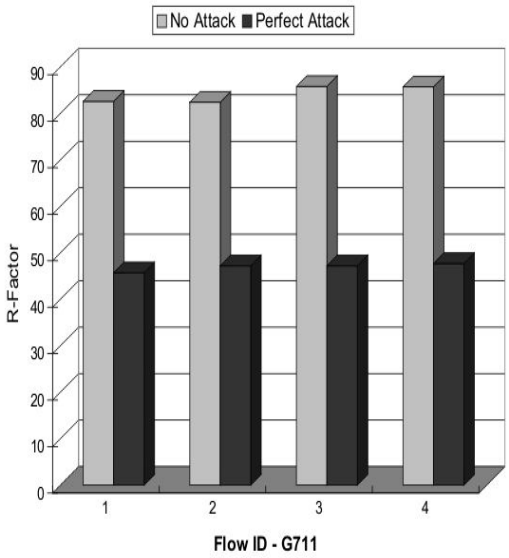

Voice Rating as per Emodel	R-factor
Best	90<R<100
High	80<R<90
Medium	70<R<80
Low	60<R<70
Poor	50<R<60

Impact on Web Traffic

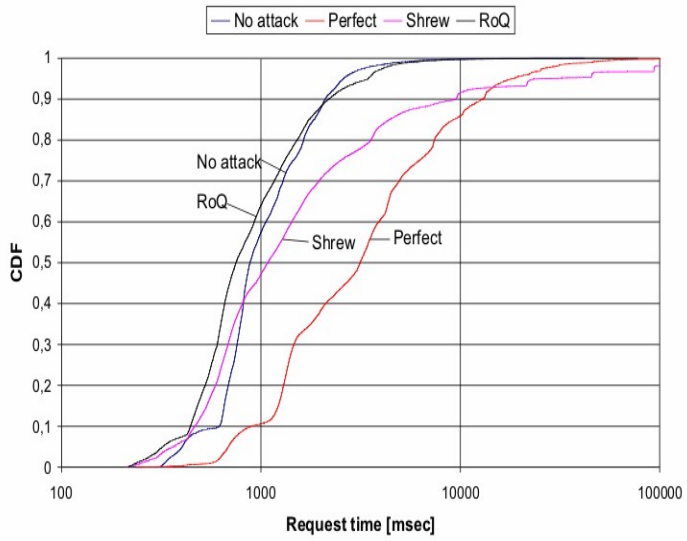

Median Response Times
No attack: 882ms
Shrew:1085ms
RoQ: 750ms
Perfect Attack: 3148ms
In 1000ms, only 10% completed.
User satisfication is high if web downloads are completed within 5000ms.

Existing Defenses and Perfect Attack

- Active Queue Management (AQM) are router based congestion control schemes.
- They find early signs of congestion from queue lengths or virtual queues.
- AVQ uses virtual queue C' < C. If virtual queue overflows, a packet is dropped.
- It is rate based marking which regulates link utilization rather than queue lengths like RED.
- RED has large queue lengths problem over long periods in presence of short-lived flows.

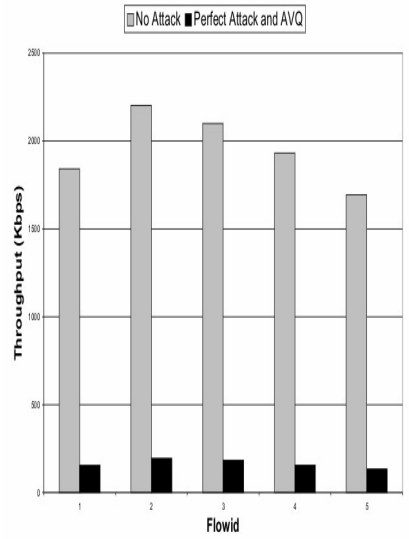

Existing Defenses and Perfect Attack (cont'd)

- Pushback is a router based aggregate congestion control scheme proposed to defend DDoS attacks.
- As the name implies, it finds aggregates which participate in a DDoS attack, and rate limits those aggregates to penalize the attack traffic.
- It can also instruct upstream routers to rate limit aggregates participating in the DDoS attack near the source.

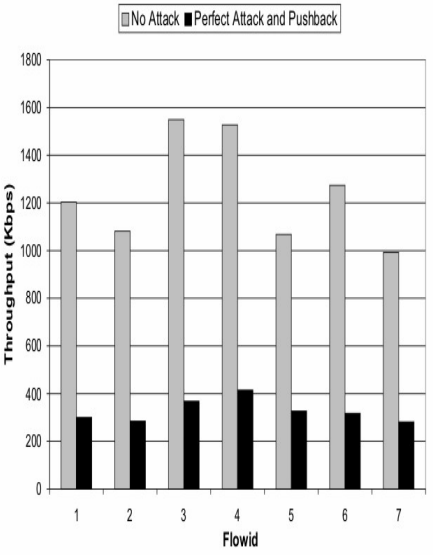

Outline

- DoS/DDoS Attacks
- IP Traceback & DDoS Defense Schemes
- Low Rate DoS Attacks
- The Perfect Storm
- The Quiet Attack

Quiet Attack

- In the Perfect attack, the majority of the attack traffic is TCP traffic, but a small amount of UDP traffic is also used.
- Do we need even a little bit of UDP traffic in the Perfect attack?
- If we eliminate UDP traffic, the attack can become extremely stealthy?

Quiet Attack

- Short-lived TCP flows are known to be bursty.
- Our attack model premise is changing the typical arrival nature of short-lived TCP flows, thus leading to a different traffic mix in the network.
- Such *intentional* use of *just* short-lived TCP flows is shown to be harmful.

Basic Theory

- A set of short-lived TCP flows approximately proportional to link capacity enter the network periodically (T).
- T is random between 0 to 1s – selected to look continuous.
- Steady influx of short-lived flows causes persistent congestion.
- We use active probing tools to monitor available bandwidth & measure link capacity.
- Knowledge of available bandwidth allows adjusting the number of short-lived flows & is particularly useful if ISPs do load sharing.
- Persistent congestion leads to random packet drops at the bottleneck droptail queue.
- The packet drops lead to reduction in the throughput of the legitimate TCP flow.

Recall

- TCP uses end to end window flow control which has some shortcomings.
- If the number of connections increases, then delay T or congestion will roughly increase proportional to the number of connections.
- Basic trade off in end to end window technique is selecting the size of window: too big leading to congestion and too small leading to poor link utilization and less throughput.

Quiet Attack in the Internet

- Quiet attack uses same techniques like those for the Perfect attack:
- Use of Botnets
- Locating a Target Router
- Determining Web Servers

Effects on long-lived TCP Flows like FTP

Average 70% reduction of throughput

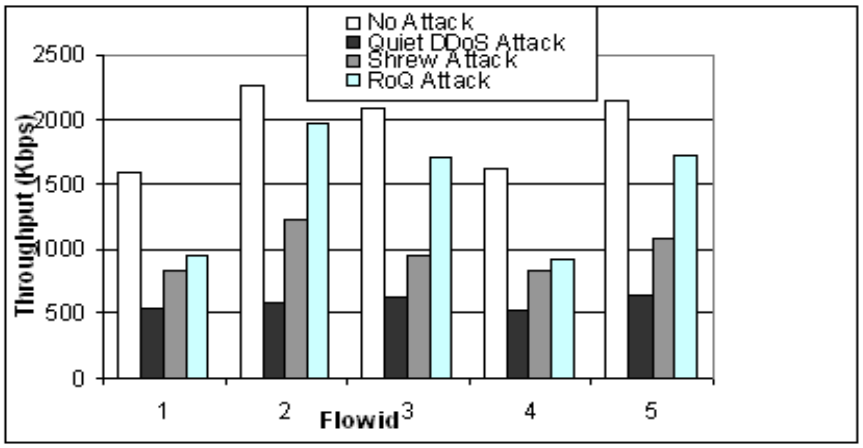

Existing Defenses and Quiet Attack

- Active Queue Management (AQM) are router based congestion control schemes.
- They find early signs of congestion from queue lengths or virtual queues.
- AVQ uses virtual queue C' < C. If virtual queue overflows a packet is dropped.
- It is rate based marking which regulates link utilization than queue lengths like RED.
- RED has large queue lengths problem over long periods in presence of short-lived flows.

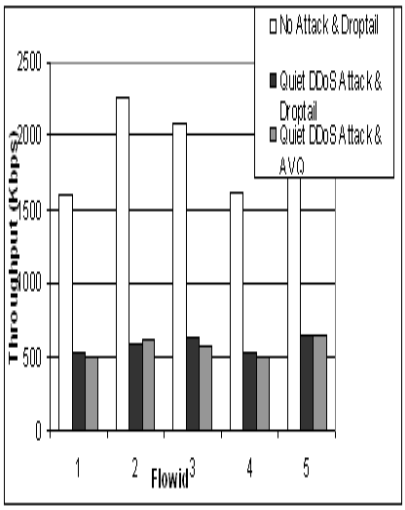

Existing Defenses and Quiet Attack

- Pushback is a router based aggregate congestion control scheme proposed to defend DDoS attacks.
- As the name implies, it finds aggregates which participate in the DDoS attack, and rate limits those aggregates to penalize the attack traffic.
- It can also instruct upstream routers to rate limit aggregates contributing to DDoS attack near the source.
- We did a experiment with a single TCP flow.
- The throughput of TCP flow was 277 Kbps with the quiet attack and pushback scheme, and it was 954Kbps without the attack.

Who can be impacted by these attacks?

- Quiet attack and Perfect attack pose threats to the ISP routers.
- Business model: Rival ISPs can target other ISPs to compete for customers.
- Routers of major websites who routinely get lot of traffic & have big infrastructure can be targeted by using the Perfect Attack & Quiet Attack.
- Could be lucrative…

Future Work

- Better CAPTCHAs to prevent bots accessing network resources..

- Better software would lead to less vulnerabilities and correspondingly less bots…
- Immediate patching of vulnerabilities…
- More awareness among public…
- Network defenses to isolate perfect attack traffic…

Questions ?