Malice Aforethought

[D]DoS on Today's Internet

Henry Duwe and Sam Mussmann

http://bit.ly/cs538-ddos

What is DoS?

"A denial of service (DoS) attack aims to deny access by legitimate users to shared services or resources." [Gligor 1984]

Two main flavors:

- Exploit vulnerabilities (ping-of-death)

- Use massive amounts of traffic to occupy resources

A Short History of DoS

Observed on the Internet as early as 1990.

Tier-1 attacks involve only the attacker(s) and victim.

- Bandwidth flood
- UDP flood
- SYN flood

Tier-2 attacks begin in mid-1990s.

- Smurf Attack
- Reflection Attack

http://www.uniforum.chi.il.us/slides/ddos/index.html http://goo.gl/5VBCq

A Short History of DoS

Tier-3 attacks use slave computers, called DDoS

- First well-documented case in August 1999 on a University of Minnesota system
- First well-publicized case occurred in February of 2000

Attack techniques are constantly evolving

- Malicious TCP window sizes (see Sockstress on wikipedia)
- Banana Attack

http://www.garykessler.net/library/ddos.html http://goo.gl/JgPjS

Why is the Internet susceptible to DoS?

Packet Switching

End-to-end principle

Multipath routing

Decentralized Internet management

Everything that makes the Internet awesome

Outline

1. Introduction to DoS
2. Scalable Attack Detection in the Network (Primary Paper)
3. Kill-Bots (Detection on the victim)
4. A Novel Approach to Early Attack Detection (Detection from backscatter)
5. Fixing the Incentive Chain (Economic meta-solution)

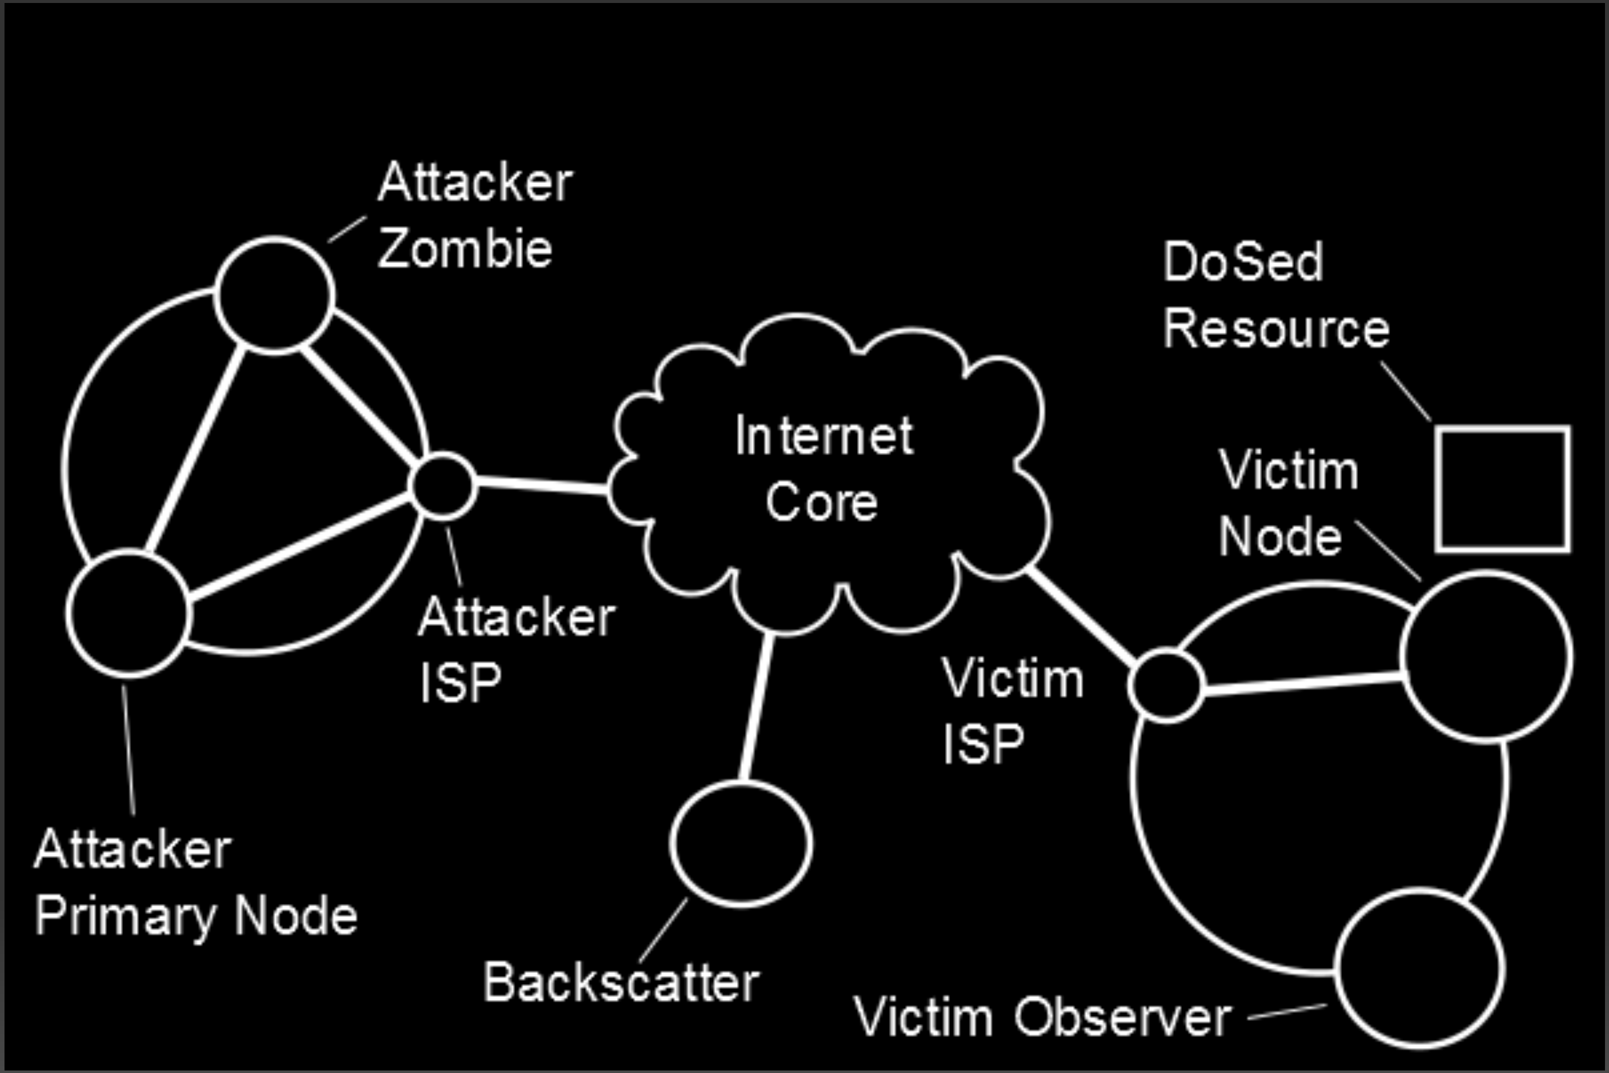

Where can we detect and defend against attacks?

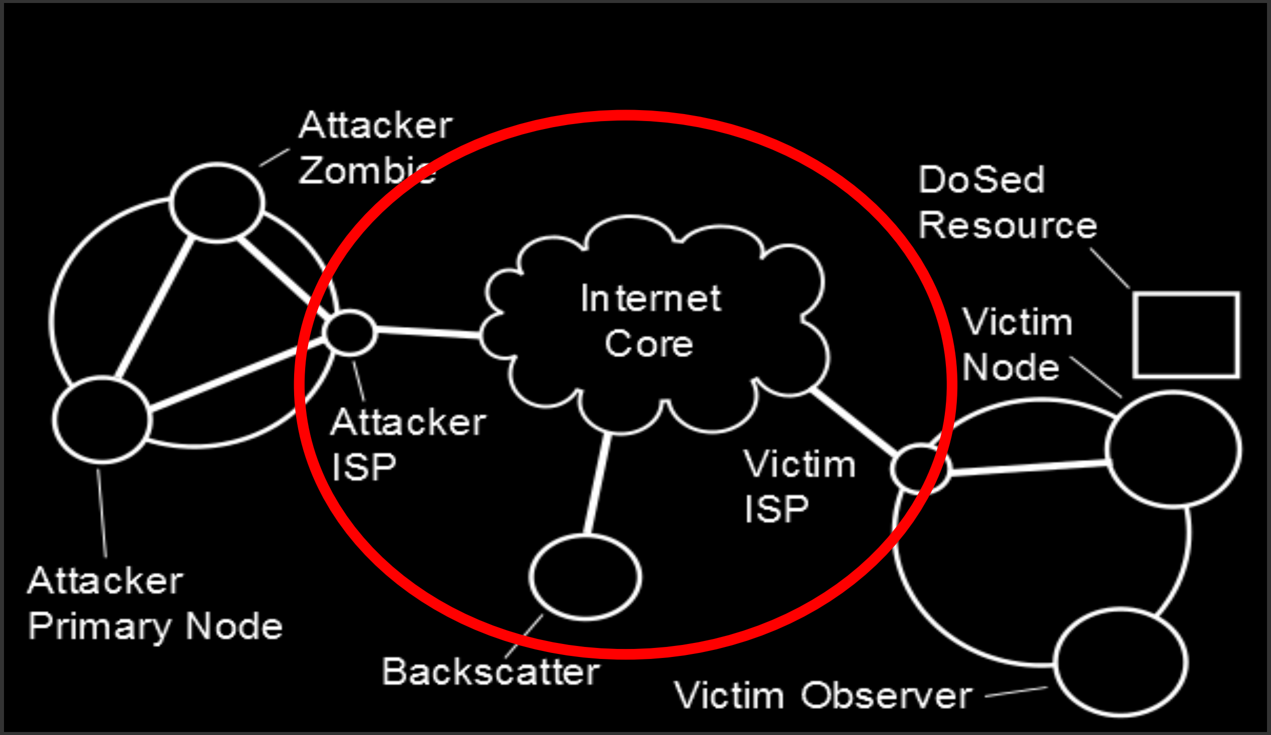

On Scalable Attack Detection in the Network.

Ramana Rao Kompella, Sumeet Singh,
and George Varghese. 2004.
4th ACM SIGCOMM conference on
Internet measurement

TCP Scan detection

Consider detecting TCP scans in the case where there's only one flow.

Most connections during a TCP scan will not receive a SYN/ACK packet, and thus will never complete (FIN packet).

Increment a counter when we see a SYN packet.

Decrement the counter when we see a FIN packet.

If counter grows beyond some bound, we can recognise malicious behavior (a reasonable value for this bound is three times the standard deviation of a non-malicious flow).

TCP Scan Detection

Consider the case where we have two flows.

Imagine one is scanning a number of hosts in your network and the other is connecting to a number of HTTP servers in your network.

One counter for each source IP address.

This is the current state-of-the-art. (Snort, Bro)

But what when we have thousands or millions of flows?

Key {Problem,Question,Contribution}

Current solutions require per-flow state to detect behavioral attacks.

Can we use aggregation to reduce the state required to detect behavioral attacks?

Yes.

Two new problems:

- Behavioral Aliasing -- False Positive
- Behavioral Spoofing -- False Negative

What kind of attacks will this detect?

- Partial Completion (or Claim-and-Hold) attacks
  - e.g. TCP SYN flood
- Scanning Attacks
  - TCP port scan
- Bandwidth Attacks
  - These are already easily covered by current techniques

TCP Scan Detection with Partial Completion Filters

Instead of one entry per source IP, hash the source IP and put all the flows with the same hash in the same bucket.

If an attacker is scanning, a lot of good traffic will be mixed in that hash bucket.

Hash the source IP with 3 different functions, and then increment/decrement the counter for each of those buckets.

We'll keep the buckets for $H_1()$, $H_2()$, $H_3()$ separate and call these separations "stages".

We can detect that a flow is attacking us if all of its counters (one for each stage) are above a bound.

Partial Completion Filters

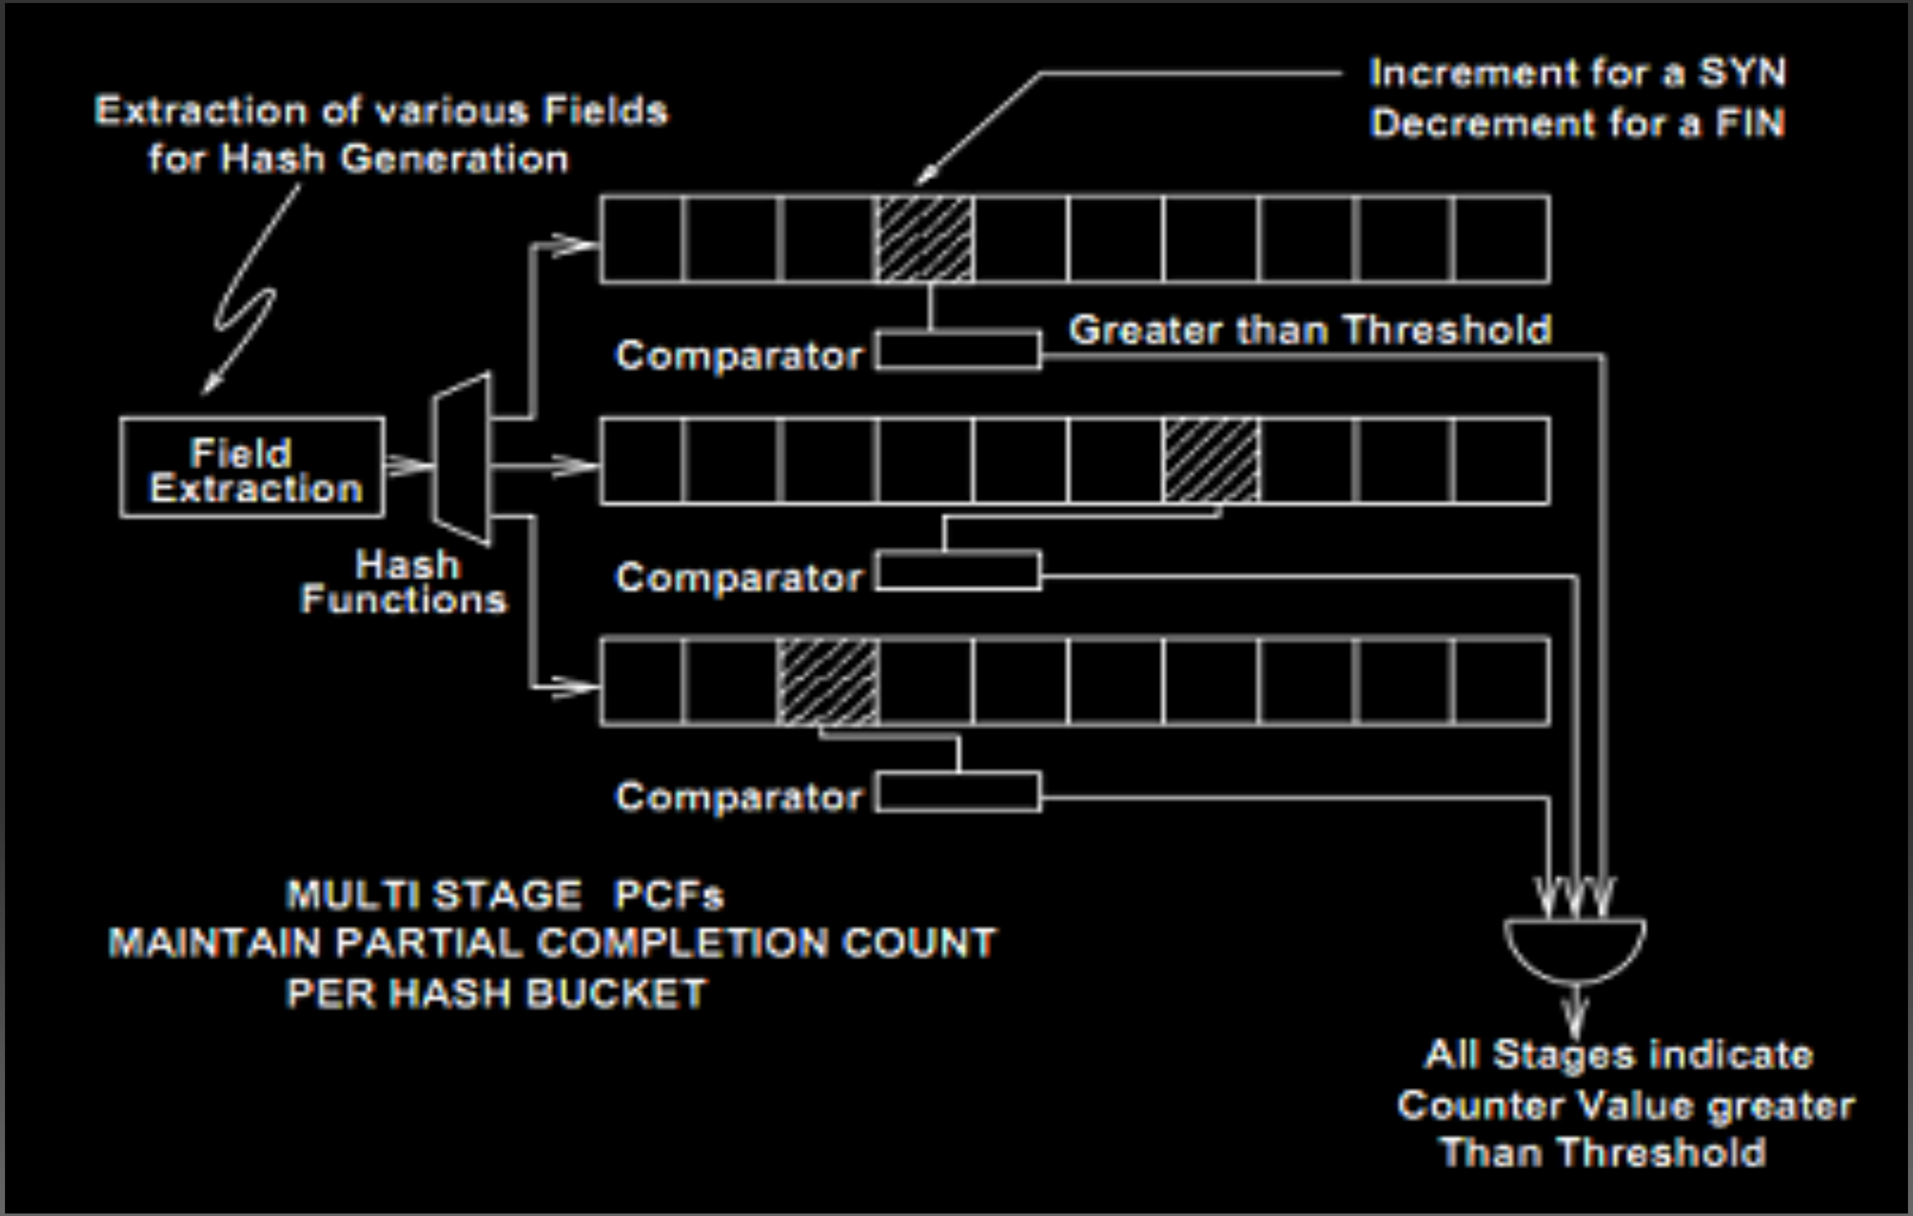

How do we use them?

- TCP SYN flood
  - increment on SYN
  - decrement on FIN
  - hash Destination IP/Port
- Spoof-resistant TCP SYN flood
  - hash Source IP/Port on reverse path
- TCP Scan
  - increment on SYN
  - decrement on FIN
  - hash Source IP

Where can we put them?

In theory, anywhere in the network!

In practice, it's easier if we put them towards the edge of the network.

Economically, we're more motivated to put them near the victim.

We can do both detection and defense if we're on the attacker's path.

Outline

1. Introduction to DoS
2. Scalable Attack Detection in the Network (Primary Paper)
3. Kill-Bots (Detection on the victim)
4. A Novel Approach to Early Attack Detection (Detection from backscatter)
5. Fixing the Incentive Chain (Economic meta-solution)

Discussion

1. What is legitimate traffic?
    - What sort of traffic should be prioritized, etc.
    - How is legitimate traffic identified?
2. Who is responsible for defeating DDoS?
    - Who is in the best tactical position?
    - Who is hurt the most?
3. How should vulnerable entities be best protected?

Botz-4-Sale: Surviving DDoS Attacks That Mimic Flash Crowds [Kandula et al. '05]

- Goal: Maximize **Goodput** (legitimate client throughput)

- Idea:
  - Distinguish human (legitimate) traffic from zombie traffic
    - Reverse Turing tests

  - Don't allocate connection space until authenticated

  - Drop zombie traffic

Kill-Bots Implementation

- Kill-Bots runs at the kernel level of a protected server
- Modified Bloom filter for IP addresses
- Two stage authentication:
  1. Issue reverse Turing test (no access to TCBs or socket buffers), building Bloom filter
  2. Once Bloom filter stabilizes, no Turing tests

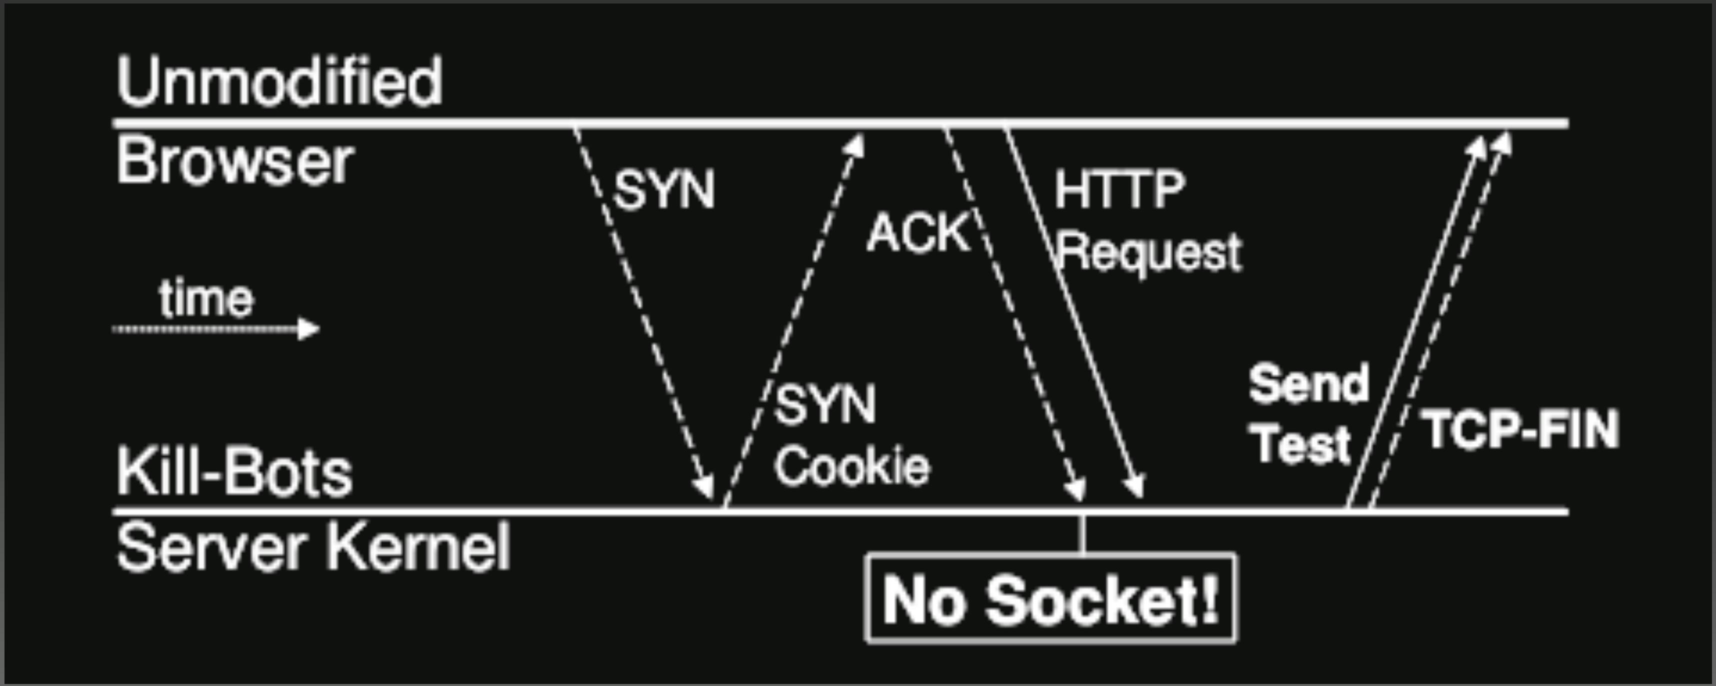

Kill-Bots

- Admission Control
  - Don't allow authentication to consume all resources
  - Authenticate only as many clients as can be served
    - Dynamically determine authentication probability based on server idle time
    - Beneficial even in non DDoS situations

- Results

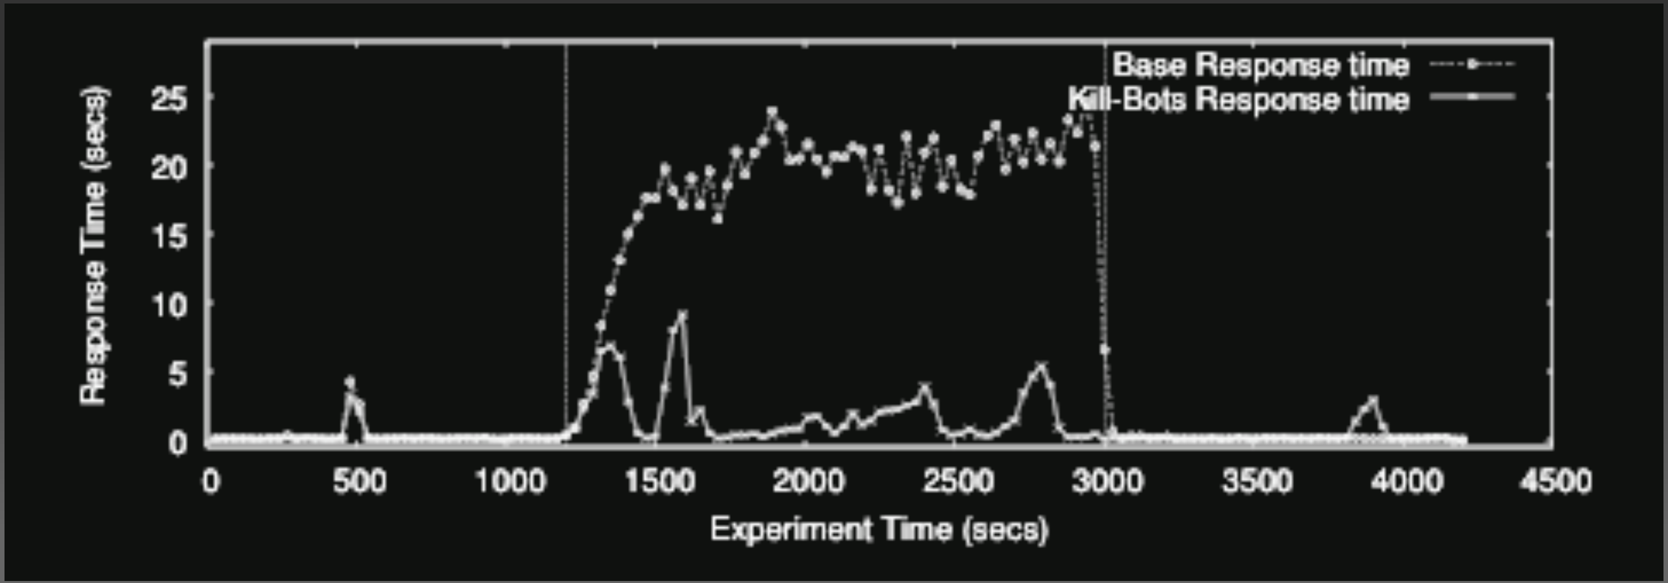

Kill-Bots Issues

- Not all legitimate traffic originates from a human sitting at a browser
  - Sort of solved...
- Potential to be really, really annoying
- IP aliasing
  - Sort of solved...
- Bloom filter false positives (human classified as zombie)
  - Are false negatives an issue?
- Others?

Outline

1. Introduction to DoS
2. Scalable Attack Detection in the Network (Primary Paper)
3. Kill-Bots (Detection on the victim)
4. A Novel Approach to Early Attack Detection (Detection from backscatter)
5. Fixing the Incentive Chain (Economic meta-solution)

A novel approach to detecting DDoS Attacks at an Early Stage [Xiao et al. '06]

- Assume a normal SYN flood attack

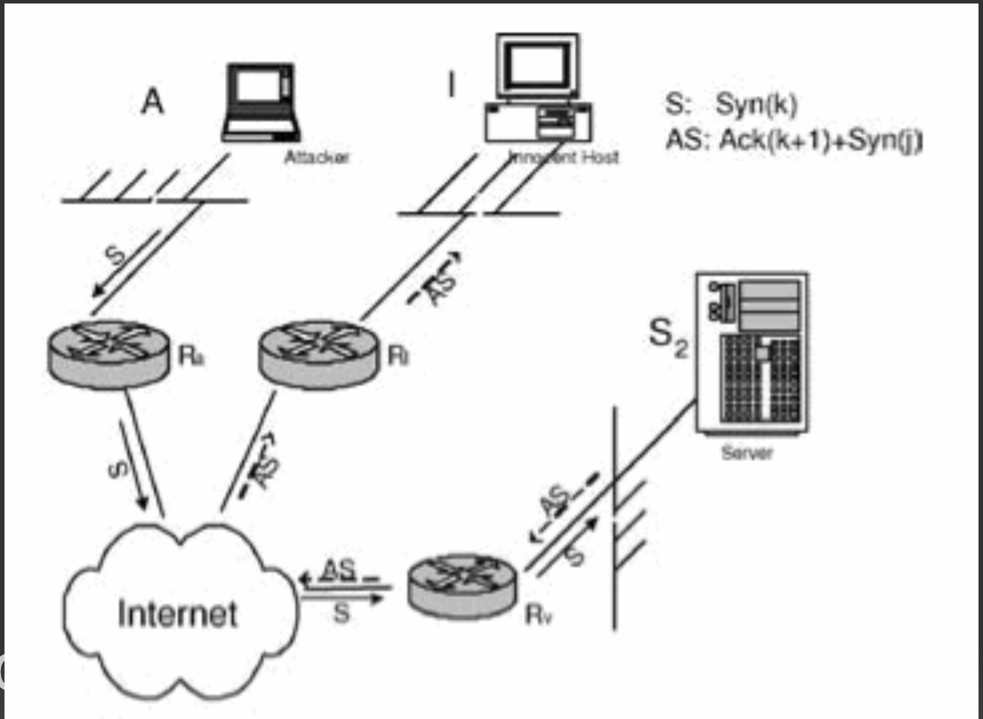

- What abo
  - Backscatter
  - An innocent client can detect these early in an attack

Implementation

- Edge routers contain modified Bloom filter
  - Outgoing SYN --> increment counter
  - Incoming ACK/SYN --> decrement counter
  - Decrement a 0 --> Suspicious Alarm (SA)
- Multiple SAs close together indicates attack, notify victim server

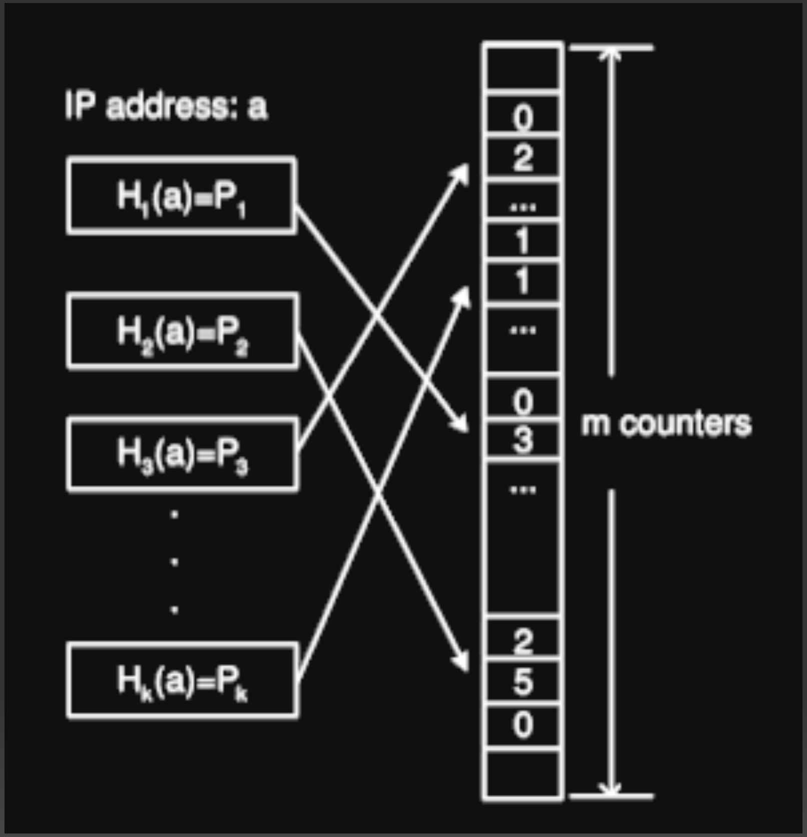

Issues Raised

- Attack specific
- Who is responsible for DDoS protection?
- Edge routers could help stop IP spoofing already, but don't
- Economic incentives

Outline

1. Introduction to DoS
2. Scalable Attack Detection in the Network (Primary Paper)
3. Kill-Bots (Detection on the victim)
4. A Novel Approach to Early Attack Detection (Detection from backscatter)
5. Fixing the Incentive Chain (Economic meta-solution)

Defeating DDoS by Fixing the Incentive Chain

[Huang et al. '07]

- Cooperative filtering and cooperative caching greatly reduce effect of DDoS attacks
  - Rarely used
  - Require ISPs to use precious resources (and potentially reduce performance)
  - Require unscalable numbers of bilateral agreements

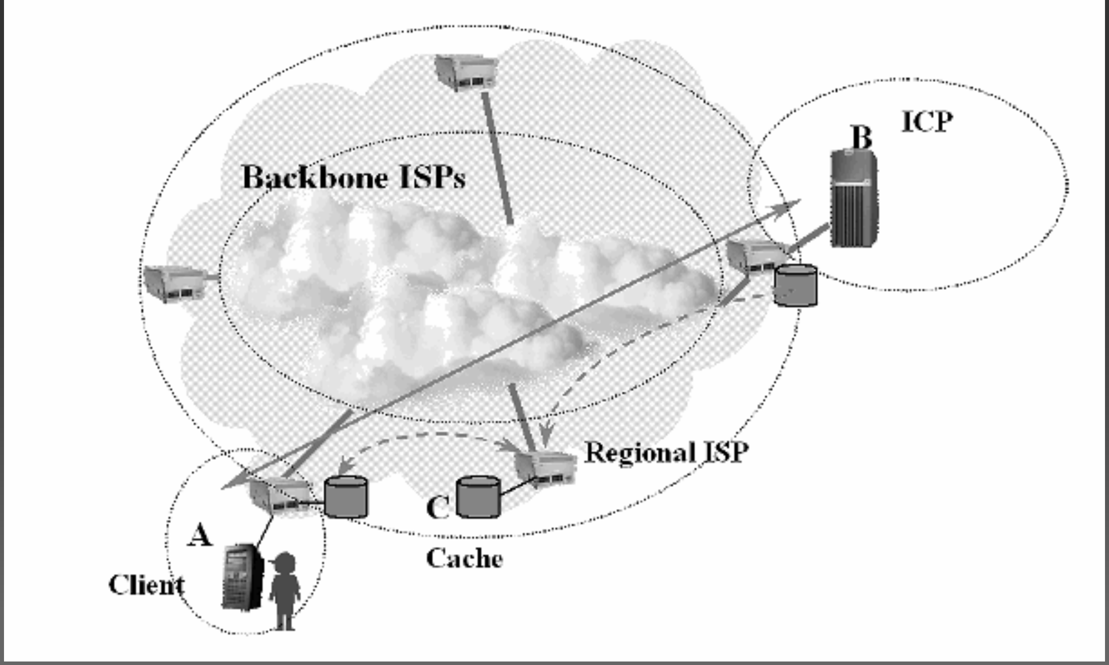

Fixing Incentive Chain

- Usage-based price schemes
  - Users pay flat fee and have volatile usage
  - ISPs overprovision bandwidth --> who cares if DDoS attacks use extra bandwidth?
  - E.g., by how much congestion a user causes :)

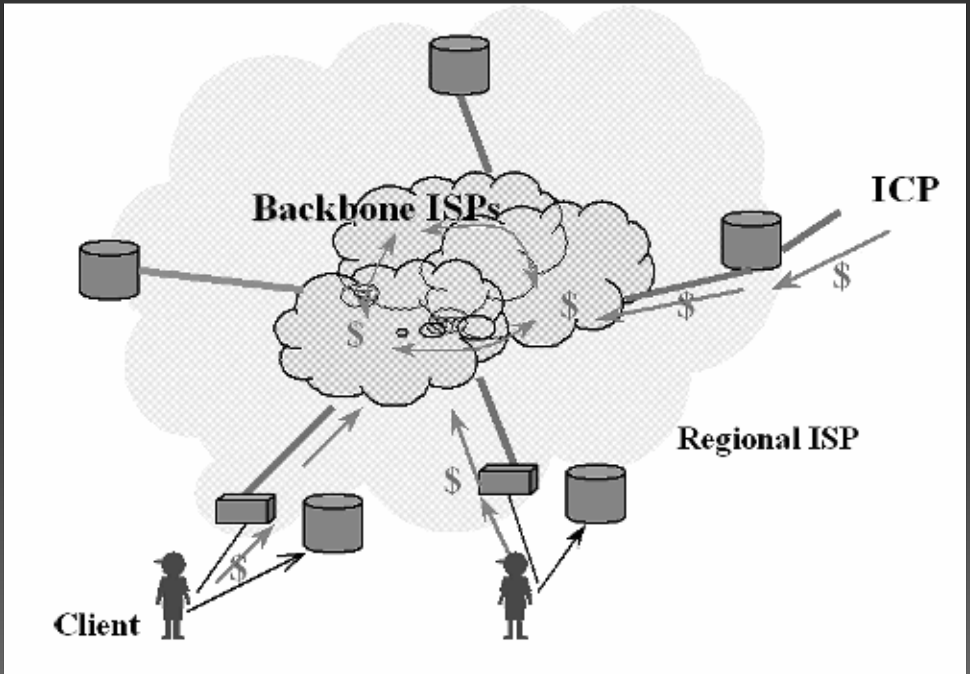

Fixing Incentive Chain

- Capacity Provision Networks (CPNs)
  - CPN owner deals with large numbers of ISPs
  - ICPs contract with CPN owner for cooperative caching during DDoS attack

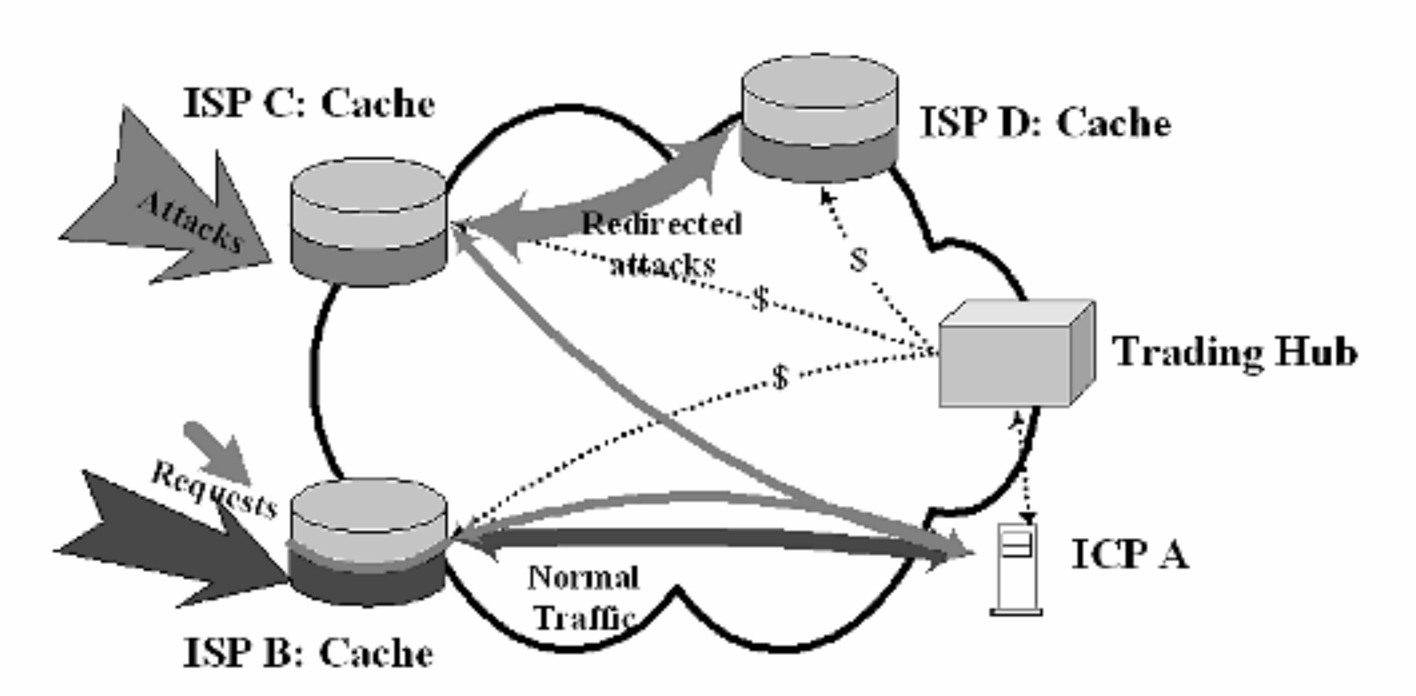

Discussion

1. What is legitimate traffic?
   - What sort of traffic should be prioritized, etc.
   - How is legitimate traffic identified?
2. Who is responsible for defeating DDoS?
   - Who is in the best tactical position?
   - Who is hurt the most?
3. How should vulnerable entities be best protected?

Sources

Yun Huang, Xianjun Geng, and Andrew B. Whinston. 2007. Defeating DDoS attacks by fixing the incentive chain. *ACM Trans. Internet Technol.* 7, 1, Article 5 (February 2007)

Srikanth Kandula, Dina Katabi, Matthias Jacob, and Arthur Berger. 2005. Botz-4-sale: surviving organized DDoS attacks that mimic flash crowds. In *Proceedings of the 2nd conference on Symposium on Networked Systems Design \& Implementation - Volume 2* (NSDI'05), Vol. 2. USENIX Association, Berkeley, CA, USA, 287-300.

Bin Xiao, Wei Chen, and Yanxiang He. 2006. A novel approach to detecting DDoS Attacks at an Early Stage. *J. Supercomput.* 36, 3 (June 2006), 235-248.