Distributed Denial of Service (DDoS)

Defending against Flooding-Based DDoS Attacks: A Tutorial
Rocky K. C. Chang

Presented by
Adwait Belsare (adwait@wpi.edu)
Suvesh Pratapa (suveshp@wpi.edu)

Modified by Bob Kinicki
18 April 2012

Outline

- Introduction
- The DDoS Problems
- Solutions to the DDoS Problems
- An Internet Firewall?
- A Comparison of Four Detect and Filter Approaches
- Conclusions

Introduction

- A typical **DDoS attack** consists of amassing a large number of compromised hosts to send useless packets to jam a victim or its Internet connection or both.

- Can be done in following ways:
  - To exploit system design weaknesses such as ping to death .
  - Impose computationally intensive tasks on the victim such as encryption and decryption.
  - **Flooding-based DDoS Attack**.

DDoS Attacks

- Do not rely on particular network protocols or system design weaknesses.
- Consist of sufficient number of compromised hosts amassed to send useless packets toward a victim around the same time.
- Have become a major threat due to availability of a number of user-friendly attack tools on one hand and lack of effective solutions to defend against them on the other.

Attacks Reported

- *May/June, 1998*

  First primitive DDoS tools developed in the underground - Small networks, only mildly worse than coordinated point-to-point DoS attacks.
- *August 17, 1999*

  Attack on the University of Minnesota reported to UW network operations and security teams.
- *February 2000*

  Attack on Yahoo, eBay, Amazon.com and other popular websites.
- **One study observed more than 12,000 attacks during a three week period.**

Reference: http://staff.washington.edu/dittrich/misc/ddos/timeline.html

The DDoS Problems

The attacks can be classified into:

Direct Attacks.

Reflector Attacks.

Direct Attacks

- Consists of sending a large number of attack packets directly towards a victim.
- Source addresses are usually spoofed so the response goes elsewhere.
- Examples:
  - **TCP-SYN Flooding**: The last message of TCP's 3 way handshake never arrives from source.
  - Congesting a victim's incoming link using ICMP messages, RST packets or UDP packets.
  - Attacks use TCP packets (94%), UDP packets (2%) and ICMP packets(2%).

Direct Attack

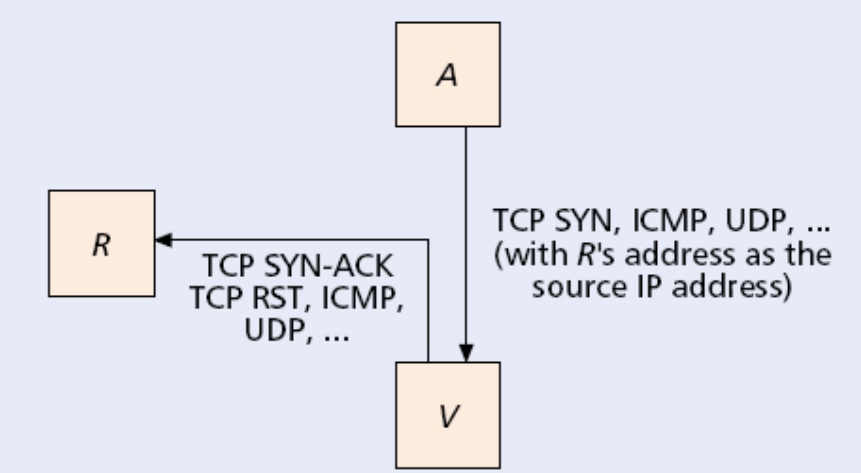

Figure 1. (a)

Agent Programs: Trinoo, Tribe Flood Network 2000, and Stacheldraht

Reflector Attacks

- Uses innocent intermediary nodes (routers and servers) known as reflectors.
- An attacker sends packets that require responses to the reflectors with the packets' inscribed source address set to victim's address.
- Can be done using TCP, UDP, ICMP as well as RST packets.
- Examples:
  - **Smurf Attacks**: Attacker sends ICMP echo request to a subnet directed broadcast address with the victim's address as the source address.
  - SYN-ACK flooding: Reflectors respond with SYN-ACK packets to victim's address.

Reflector Attack

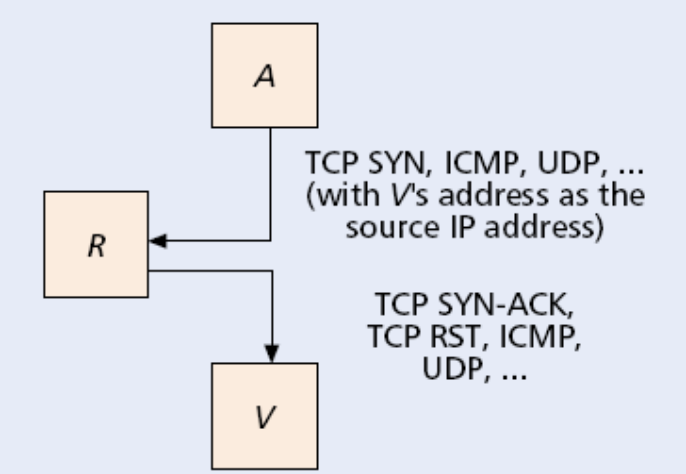

Figure 1.(b)

- Cannot be observed by backscatter analysis, because victims do not send back any packets.
- Packets cannot be filtered as they are legitimate packets.

DDoS Attack Architectures

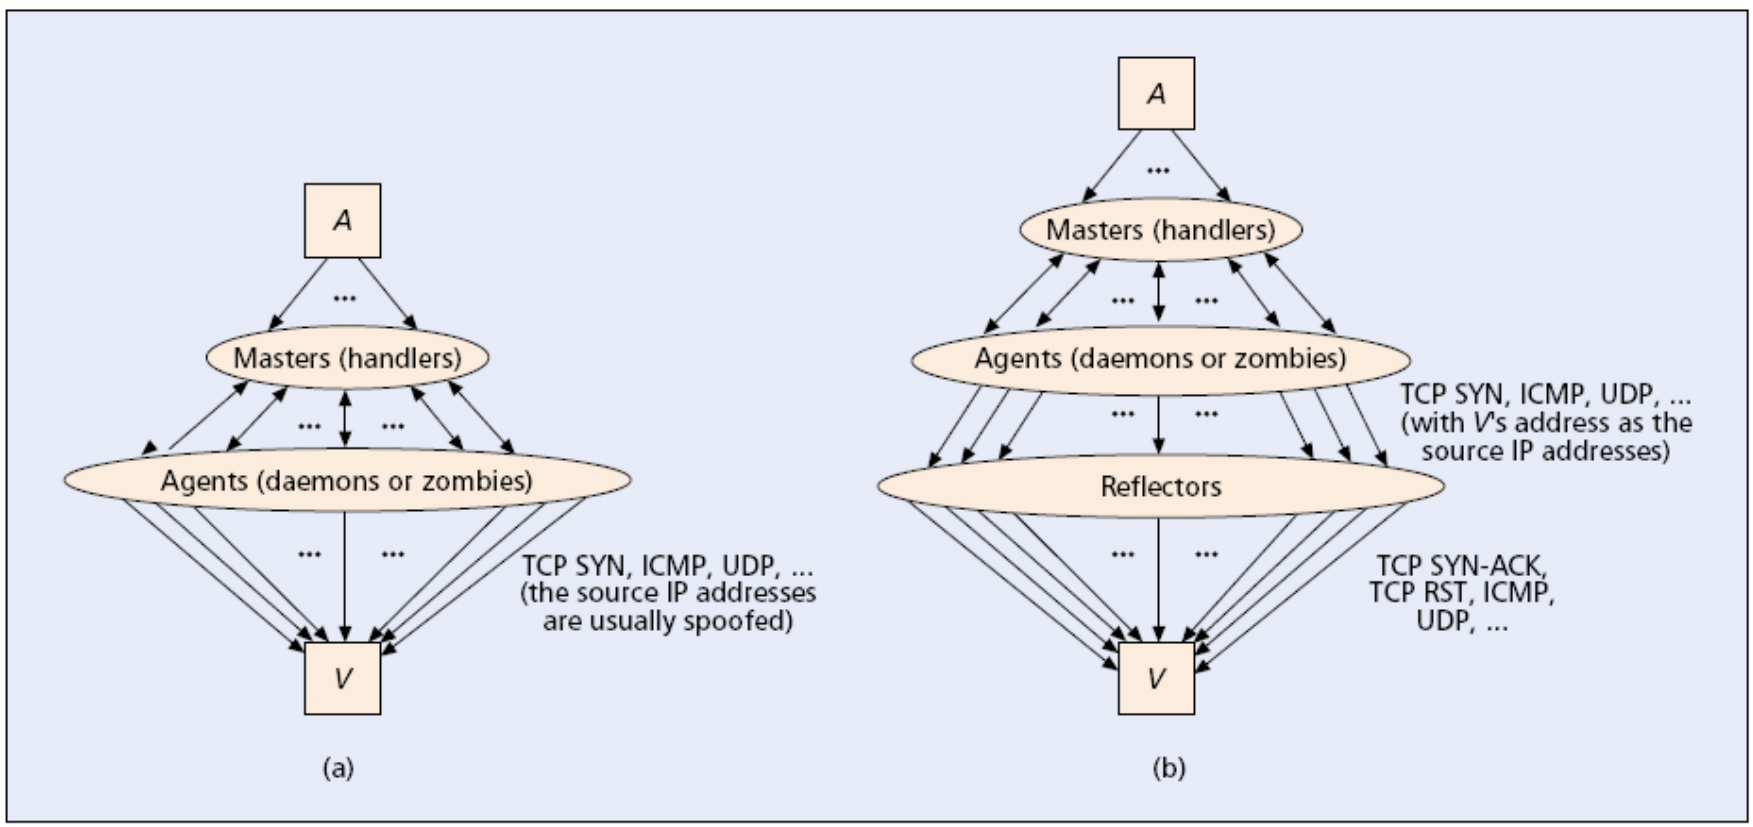

Figure 2. *DDoS attack architectures for a) direct and b) reflector attacks.*

Some Reflector Attack Methods

	Packets sent by an attacker to a reflector (with a victim's address as the source address)
Smurf	ICMP echo queries to a subnet-directed broadcast address
SYN flooding	TCP SYN packets to public TCP servers (e.g., Web servers)
RST flooding	TCP packets to nonlistening TCP ports
ICMP flooding	• ICMP queries (usually echo queries) • UDP packets to nonlistening UDP ports • IP packets with low TTL values
DNS reply flooding	DNS (recursive) queries to DNS servers

	Packets sent by the reflector to the victim in response
Smurf	ICMP echo replies
SYN flooding	TCP SYN-ACK packets
RST flooding	TCP RST packets
ICMP flooding	• ICMP replies (usually echo replies) • ICMP port unreachable messages • ICMP time exceeded messages
DNS reply flooding	DNS replies (usually much larger than DNS queries)

How many attack packets are needed?

If a victim has resources to admit N half open connections, its capacity of processing incoming SYN packets can be modeled as a G/D/INFINITY/N queue where :

G = General arrival process for the SYN packets.

D = Deterministic lifetime of each half-open connection if not receiving the third handshaking message.

Minimal rates of SYN packets to stall TCP servers in SYN flooding attacks

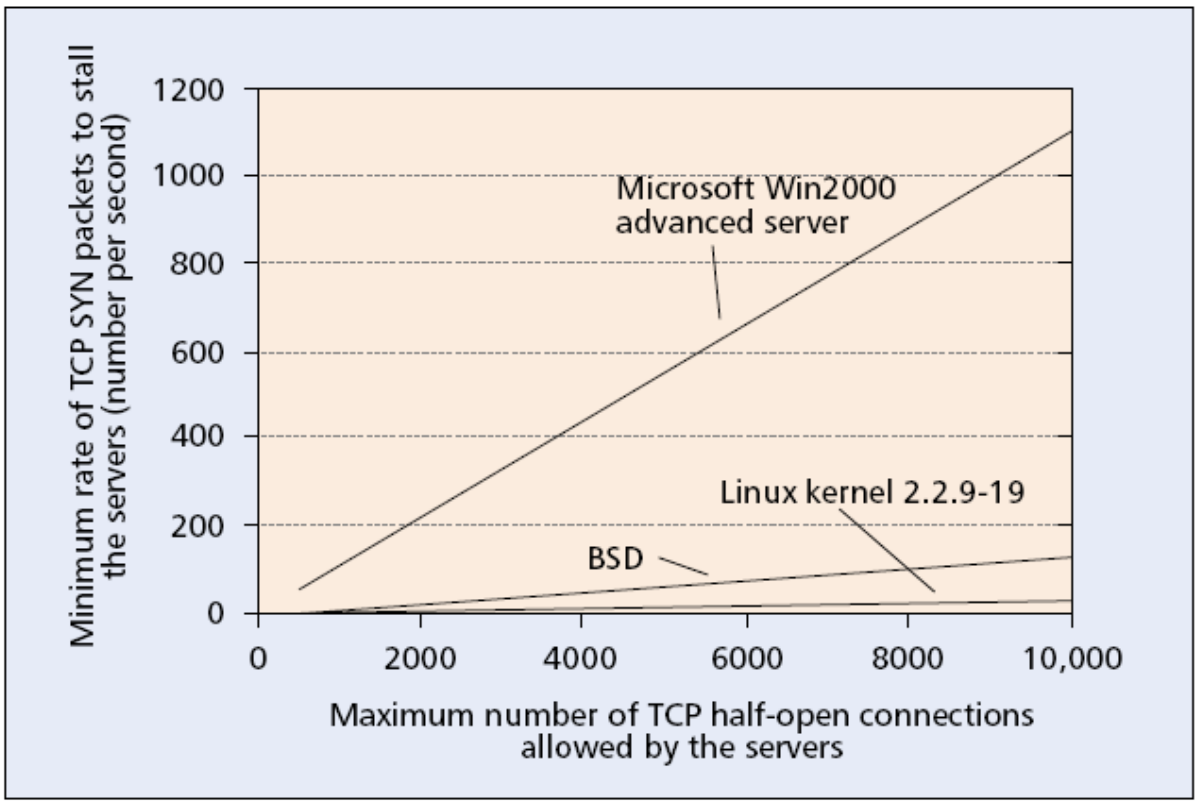

Figure 3. *Minimal rates of SYN packets to stall TCP servers in SYN flooding attacks.*

WIN system offers better protection against SYN flooding based on maximum lifetimes of half-open connections.

1Mb/s connection is sufficient to stall all three servers with N<= 10,000.

Solutions to the DDoS Problems

- There are three lines of defense against the attack:
  - Attack Prevention and Preemption *(before the attack)*
  - Attack Detection and Filtering *(during the attack)*
  - Attack Source Traceback and Identification *(during and after the attack)*

- **A comprehensive solution should include all three lines of defense.**

Attack Prevention and Preemption

- On the passive side, protect hosts from master and agent implants by using signatures and scanning procedures to detect them **{essentially an IDS strategy}.**
- Monitor network traffic for known attack messages sent between attackers and masters.
- On the active side, employ cyber-informants and cyber-spies to intercept attack plans (e.g., a group of cooperating agents).
- **This line of defense alone is inadequate**.

Attack Source Traceback and Identification

- An after-the-fact response.
- **IP Traceback**: Identifying actual source of packet without relying on source information.
  - Routers can record information they have seen.
  - Routers can send additional information about seen packets to their destinations.
- **Infeasible to use IP Traceback during ongoing attack**. Why?
  - Cannot always trace packets' origins. (NATs and Firewalls!)
  - IP Traceback also ineffective in reflector attacks.
- Nevertheless, it is at least a good idea and is helpful for post-attack law enforcement.

Attack Detection and Filtering

- Two phases:
  - **DDoS Attack Detection**: Identifying DDoS attack packets.
  - **Attack Packet Filtering**: Classifying those packets and dropping them.

    (Overall performance depends on effectiveness of both phases.)

- Effectiveness of Detection
  - FPR (False Positive Ratio):

    No. of *false positives*/Total number of confirmed normal packets
  - FNR (False Negative Ratio):

    No. of *false negatives*/Total number of confirmed attack packets

Both metrics should be low!

Attack Detection and Filtering

- Effectiveness of Filtering
  - *Effective attack detection ≠ Effective packet filtering

    Detection phase uses victim identities (Address or Port No.), so even normal packets with same signatures can be dropped.
  - NPSR (Normal Packet Survival Ratio):

    Percentage of normal packets that can *survive* in the midst of an attack

NPSR should be high!

Attack Detection and Filtering

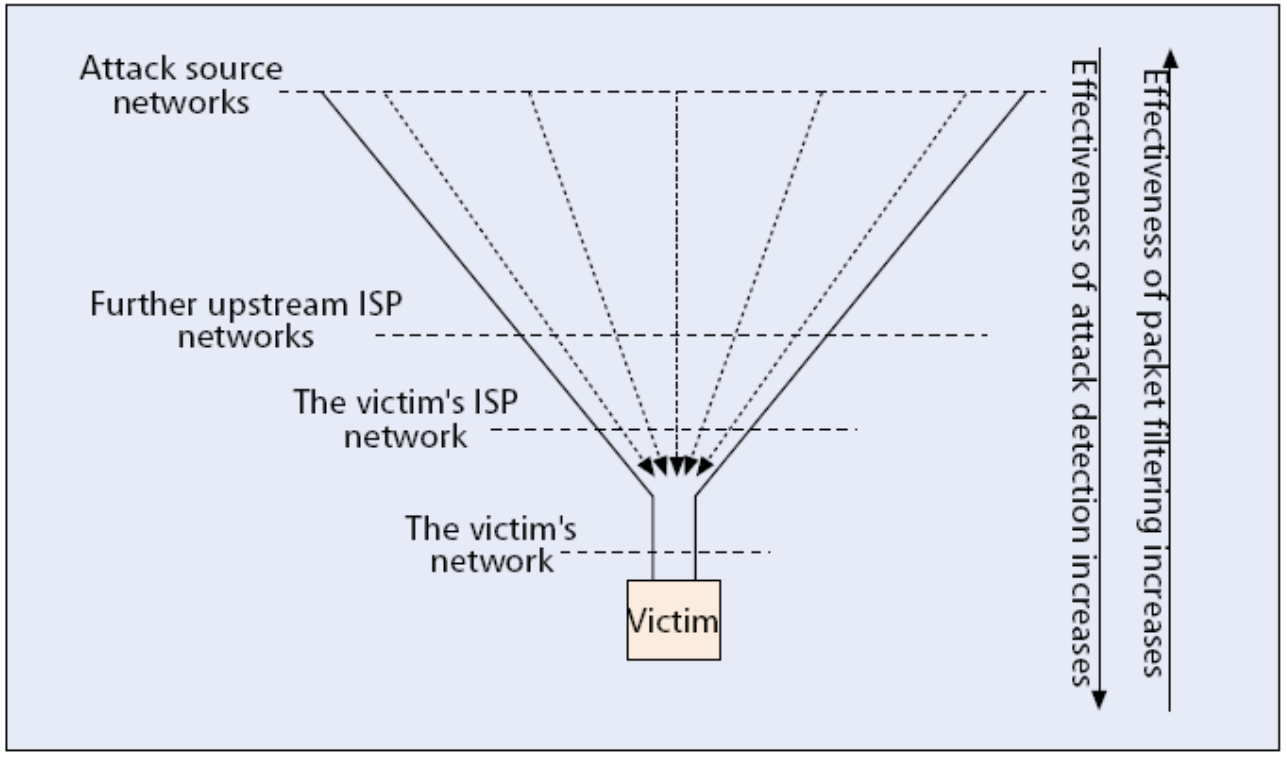

Figure 4. *Possible locations for performing DDoS attack detection and filtering.*

Attack Detection and Filtering

- At Source Networks:
  - Can filter packets based on address spoofing.
  - Direct attacks can be traced easily, difficult for reflector attacks.
  - Need to ensure all ISPs have **ingress packet filtering**. Very difficult (Impossible?)

- At the Victim's Network:
  - DDoS victim can detect attack based on volume of incoming traffic or degraded performance. Commercial solutions available.
  - Other mechanisms: *IP Hopping* (Host frequently changes it's IP address when attack is detected. DNS tracing can still help the attackers)
  - Last Straw: If incoming link is jammed, victim has to shut down and ask the upstream ISP to filter the packets.

Attack Detection and Filtering

- At a Victim's Upstream ISP Network:
  - Victim requests frequently to filter packets.
  - Can be automated by designing intrusion alert systems, which should be designed carefully.
  - Not a good idea though. Normal packets can still be dropped, and this upstream ISP network can still be jammed under large-scale attacks.
- At further Upstream ISP Networks:
  - The above approach can be further extended to other upstream networks.
  - Effective only if ISP networks are willing to co-operate and install packet filters.

An Internet Firewall

- A bipolar defense scheme cannot achieve both effective packet detection and packet filtering.

  Hence a proposal to **deploy a global defense infrastructure**.

  The plan is to detect attacks right at the Internet core!

- Two methods, which employ a set of distributed nodes in the Internet to perform attack detection and packet filtering.
  - **Route-based Packet Filtering Approach (RPF)**
  - **Distributed Attack Detection Approach (DAD)**

Route-Based Packet Filtering (RPF)

- Extends the ingress packet filtering approach to the Internet.
  - Distributed packet filters examine the packets based on source addresses and BGP routing information.
  - A packet is considered an attack packet if it comes from an unexpected link. {Attack packets are then dropped!!}
- Major Drawbacks
  - Requiring BGP messages to carry the needed source addresses - Overhead inside BGP messages!
  - Deployment is still tough! – Filters need to be placed in almost 1800 AS (when there were 10,000 AS's) and the number of AS is continuously increasing.
  - Cannot filter reflected packets because source address is legitimate.

Distributed Attack Detection (DAD)

- Deploys a set of distributed Detection Systems (DSs) to observe network anomalies and misuses.
- ***Anomaly detection*:** Observing and detecting traffic patterns that significantly deviate from normal (e.g., unusual traffic intensity for specific packet types).
- ***Misuse detection*:** Identifying traffic that matches a known attack signature.
- DSs rely mainly on anomaly detection. Various DSs **exchange attack information** from local observations. This is stateful in respect to the DDoS attacks.
- Designing an effective and deployable architecture for the DAD approach is a challenging task.

DS Design Considerations

- Packet processing needs to be high speed.
- Attack detection process has two levels: local detection and global detection.
- Binary hypothesis - H1 and H0 in Figure 5 is tested on a set of packet flows with the same destination IP address.
- If local detection supports H1, DS floods attack alert to all other DS's. DS's then consolidate and analyze local detection result with attack alerts to make a global detection decision.

Distributed Attack Detection

DS Design Considerations

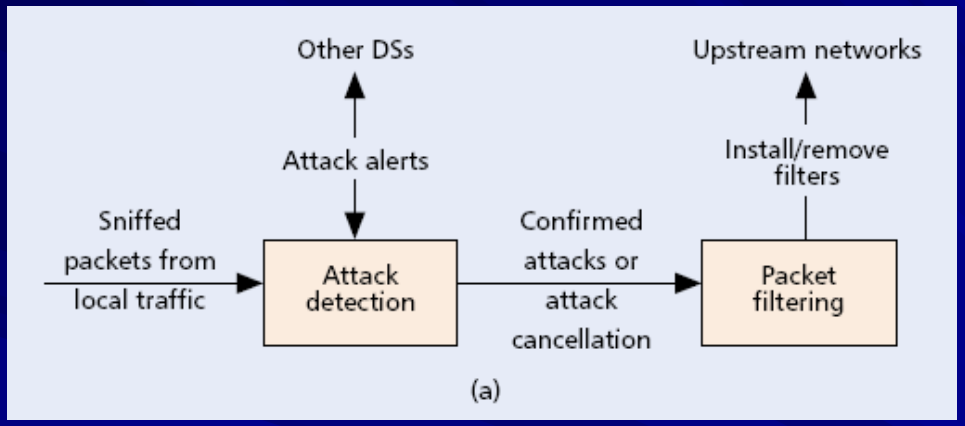

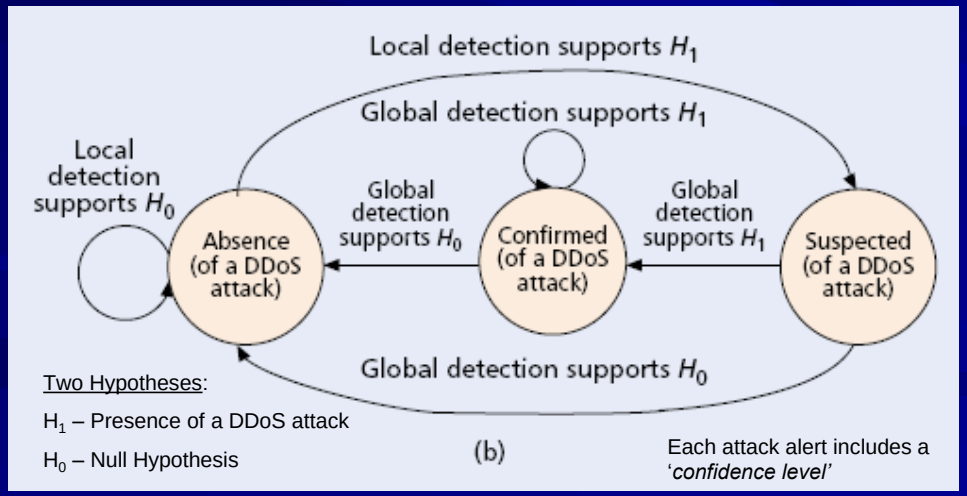

Other considerations:

- Filters should be installed only on attack interfaces when in 'CONFIRMED' state
- All DSs should be connected '*always*'
- Works in Progress:
  Intrusion Detection Exchange Protocol
  Intrusion Detection Message Exchange Format

Figure 5. *a) High-level DS architecture and b) a state diagram of two-level attack detection in the distributed detection approach.*

Distributed Attack Detection

Quickest Detection Problem Formulation

*Let i*th Sample of instantaneous traffic intensity be A_i

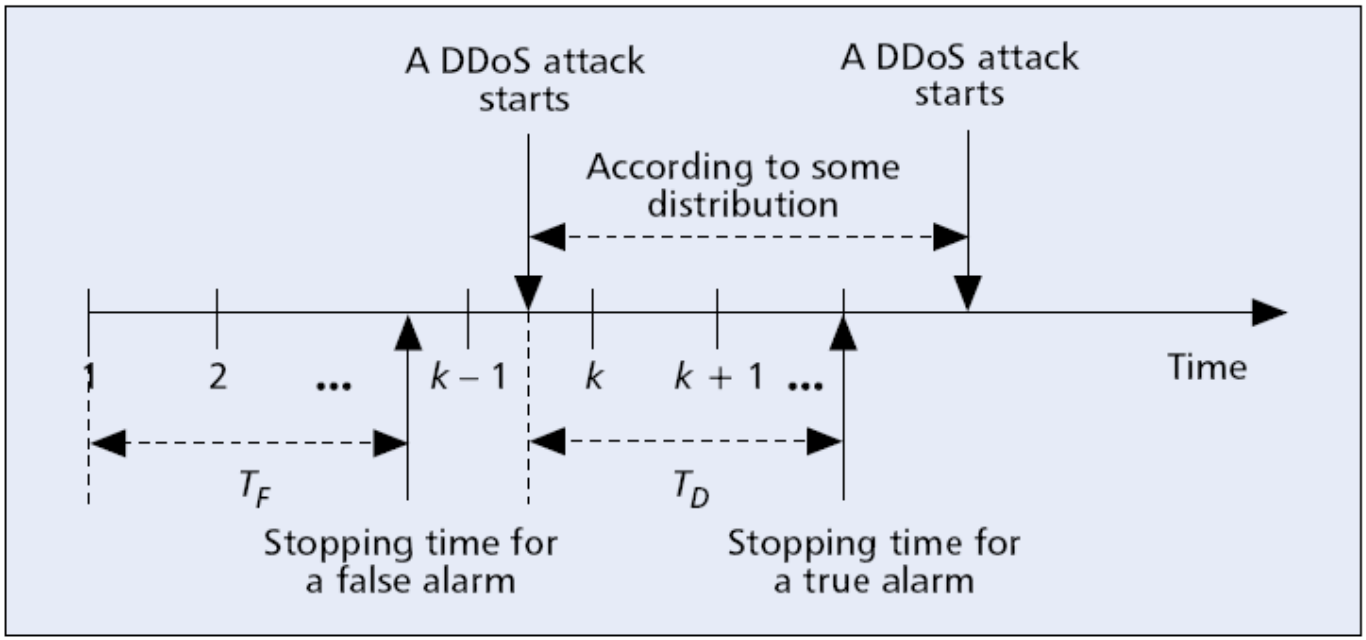

Figure 6. *Relevant events in DDoS attack detection.*

Goal – detect abrupt change in distribution $P_0 \rightarrow P_1$ as soon as possible.

Disorder event (arrival of attack packets) triggers distribution change.

Limitations and Open Problems

- Limitations of Mathematical Nature:

  Choices of global / local thresholds, traffic modeling, etc.

- Performance Aspects:
  - Two-level detection not useful for DDoS attacks of short durations.
  - **Flash crowds** can trigger false alarms. Algorithm should adapt to this new ‘normality’.

- Other attack patterns:
  - DeS (degradation of service attacks) use ‘pulsing agents’ with short bursts.
  - Using different sets of attack agents each time.

Comparison of Four Detect-And-Filter Approaches

	Ubiquitous ingress packet filtering (UIPF)	Route-based packet filtering (RPF)	Local attack detection (LAD)	Distributed attack detection (DAD)
1. Detection locations	All ISP networks that are connected to leaf networks in the Internet	A set of packet filters distributed in the Internet	Potential victims' networks and/or their upstream ISP networks	A set of detection systems distributed in the Internet
2. Filtering locations	Same as the detection locations	Same as the detection locations	Same as the detection locations and further upstream ISP networks if backpressure is used	Same as the detection locations and other upstream networks
3. Attack signatures	Spoofed source IP addresses	Spoofed source IP addresses according to the BGP routing information	Traffic anomalies and misuses detected by local intrusion detection systems	Mainly traffic anomalies observed from the set of distributed detection systems
4. False positive ratio (FPR)	= 0	= 0 if the BGP routes are correct	≥ 0 (= 1 in a sufficiently large-scale DDoS attack)	≥ 0 (high if the detection algorithms are overly sensitive)
5. False negative ratio (FNR)	≥ 0 (= 0 if all attack packets use spoofed addresses)	≥ 0 (small if most attack packets use spoofed addresses)	≥ 0 (= 0 in a sufficiently large-scale DDoS attack)	≥ 0 (high if the detection algorithms are not sensitive enough)
6. Normal packet survival ratio (NPSR)	≥ 0 (= 1 if all attack packets use spoofed addresses)	≥ 0 (large if most attack packets use spoofed addresses and the number of the AS nodes involved in the packet filtering is sufficiently large)	≥ 0 (= 0 in a sufficiently large-scale DDoS attack)	≥ 0 (high if both the false negative and positive ratios are low, and the set of detection systems are placed optimally in the Internet)
7. New communication protocols	Not required	Modifications to BGP protocols	Attack alert protocols between victims and their upstream ISP networks if backpressure is used	Protocols between detection systems
8. Computation requirement	Low	Moderate	Low	High
9. Deployment difficulty	Very high	High	Moderate without backpressure mechanisms	High
10. Technical complexity	Low	High	Moderate without backpressure mechanisms	High

Table 2. *A comparison of four approaches to detecting and filtering DDoS attack packets.*

Conclusions

- Current defense mechanisms are far from adequate.
- One promising direction is to develop a global infrastructure, an Internet Firewall.
- Deployment and design considerations should be worked upon.

We see that DDoS Defense is possible through careful planning, and this tutorial covered defense mechanisms which try to discover and slow down bad clients.