



Defending against Flooding-Based Distributed Denial-of-Service Attacks: A Tutorial

Rocky K. C. Chang

The Hong Kong Polytechnic University

Presented by
Scott McLaren



Overview

- DDoS overview
- Types of attacks
- Solutions to DDoS attacks
- Internet Firewall
- Comparisons
- Conclusions

DDoS Attacks

- Do not rely on particular network protocols or system weaknesses
- Exploit huge resources of the Internet
 - Many attackers, one victim
- Traffic jams or crashes the victim, or its Internet connection
- Yahoo!, eBay, Amazon, were attacked by DDoS attacks in February 2000

DDoS Attacks

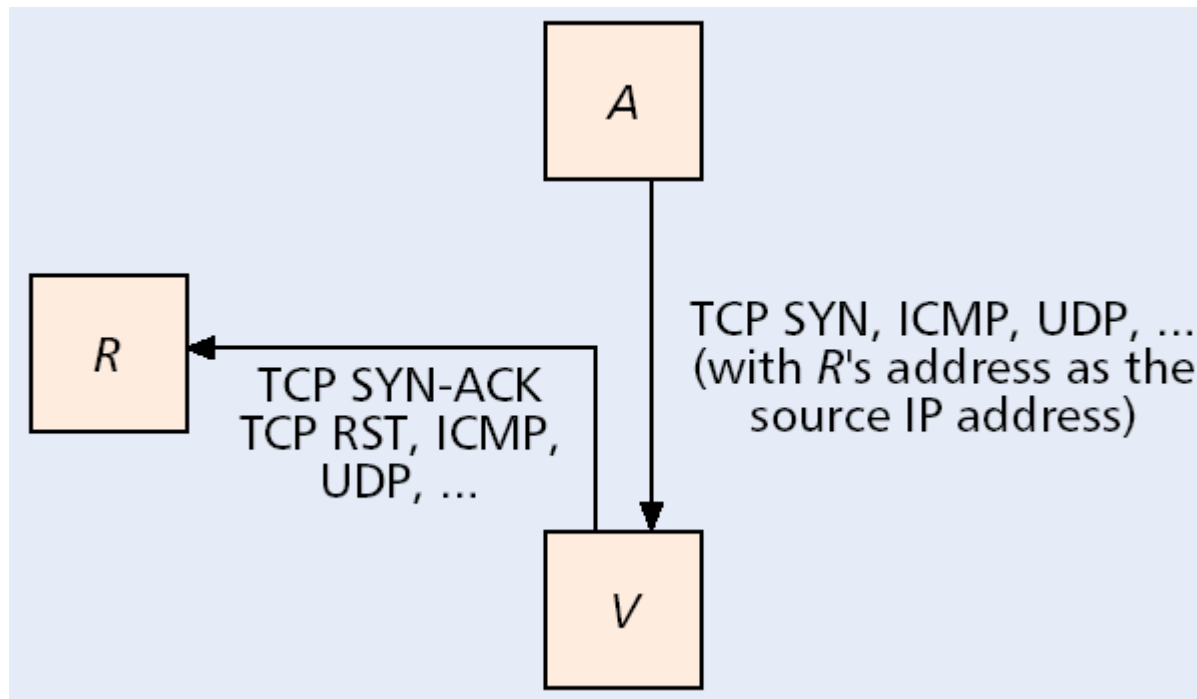
- Are most common form of attacks on the Internet today
- Most go unreported
- A recent study observed more than more than 12,000 DoS (DDos?) attacks during a three-week period
 - Actual number is probably much higher

DDoS Attacks

- Already a major problem
- Attacks are made easy by user friendly tools
- Still a lack of effective defense
 - Aborting attack in progress
 - Tracing back to attack sources
- Expected to become more severe and serious
- Cyber Warfare
 - Disable strategic business, government, public utility and military sites
 - Blackmail
- Companies have appeared in the last 2 years to offer solutions

Direct Attacks

- An attacker sends a large number of attack packets directly to a victim
- Spoofed addresses in packets, so responses go unACKed to R until timeout



SYN flooding

- If port is listening, victim responds with SYN-ACK packets
- Source addresses are spoofed, responses go to other hosts
- Victim retransmits SYN-ACK packet several times
- Half-open connections consume all the resources for pending connections, prevents new requests

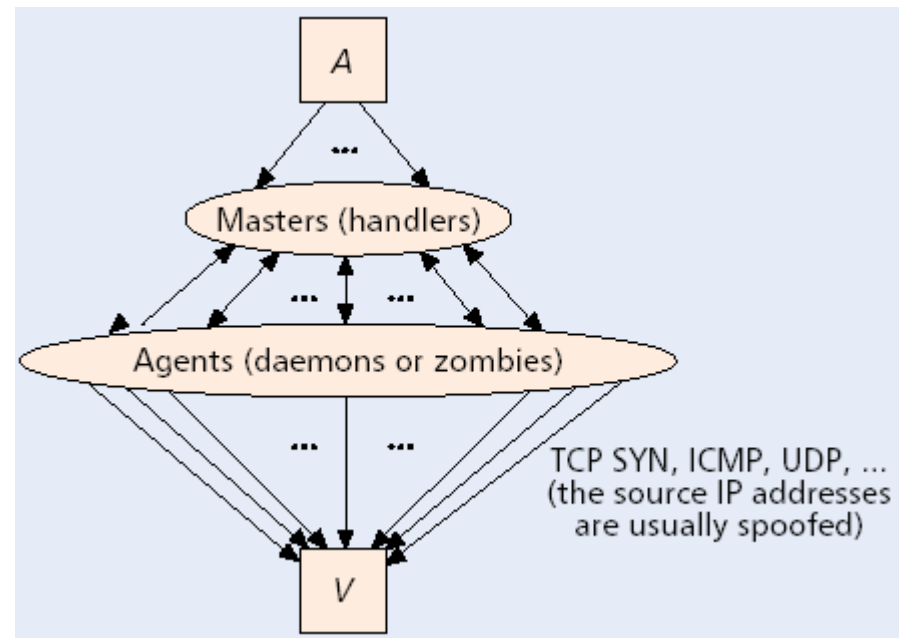
Attacks by protocol

Protocol	Percentage
TCP	94%
UDP	2%
ICMP	2%

- TCP attacks are mainly SYN-ACK based, RST packets, or ICMP error messages

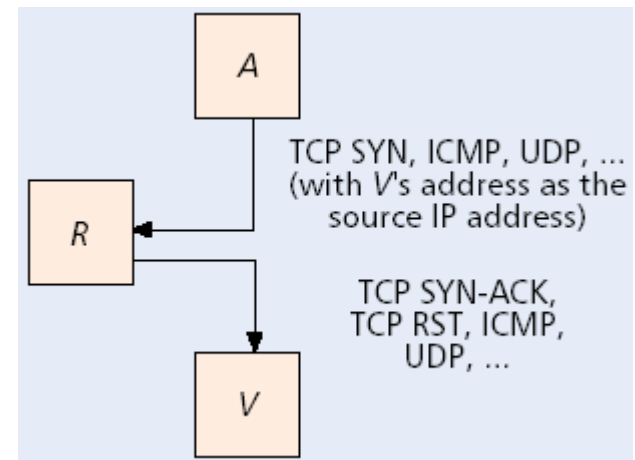
Attack Process

- Attacker sets up attack network
- Attacking host is compromised by attacker
- Attacking host implanted with master and agent programs
 - Trinoo, Tribe Flood Network 2000, Stacheldraht



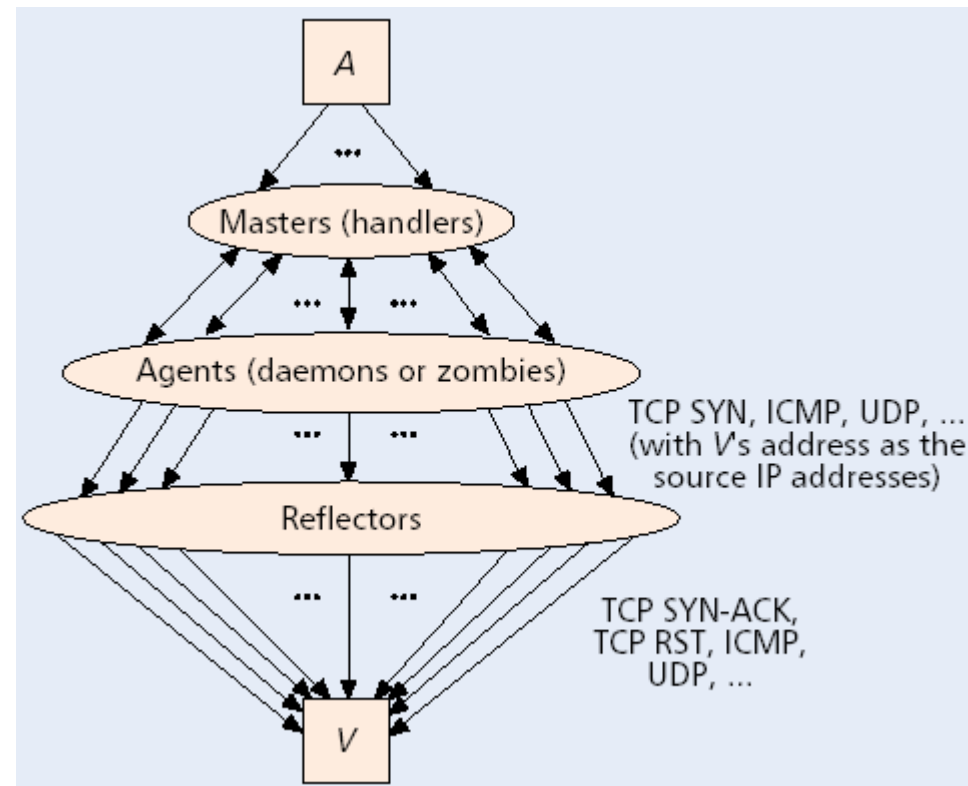
Reflector Attacks

- Intermediary nodes (routers & servers) are used to launch attack
- Attacker sends packets with source address set to victim's
- Reflectors send response to victim



Attack Process

- Based on reflector generating messages in response to other messages
- Any protocol that supports “automatic message generation” can be used
- SYN ACK or RST packets
- When SYN ACK used, reflector behaves like victim of SYN flooding due to $\frac{1}{2}$ open connections
- Clog network link



Types of Reflector Attacks

- Packets with inactive destination ports result in ICMP port unreachable messages
- Packets with small TTL result in ICMP time exceeded messages
- Bandwidth amplification
 - Attack packet results in reflected packet much larger in size (DNS replies)

Analyzing Reflector Attacks


- Cannot be observed by backscatter analysis, because victims do not send back any packets
- Number of reflector attacks unknown
- Reflected packets are normal packets, so they cannot be filtered based on address spoofing or route-based mechanism

Attack Packets Required

- Modeled as a $G/D/\infty/N$ queue
 - G – general arrival process
 - D – lifetime for each $\frac{1}{2}$ open connection
 - N – $\frac{1}{2}$ open connections allowed by victim
 - Infinite server queuing model yields the minimal rate of SYN packets required to exhaust server's resources

Server Comparison

- BSD – retransmission timeout at 6, 24, 48s, gives up after total of 75s
- Linux – 3, 6, 12s, etc. Up to 7 retransmissions, gives up after 309s
- Windows 2000 Advanced Server – retransmits SYN packets at most twice, gives up after 9s



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

Server Comparison

- If SYN packet is 84 bytes long, a 56 kb/s connection will stall Linux and BSD, $N \leq 6,000$
- A 1 Mb/s connection will stall all three with $N \leq 10,000$
- Direct ICMP ping flooding attack requires 5,000 agents for a T1 link
 - Reflector attack requires 5,000 reflectors, but agents are much fewer if each agents sends requests to multiple reflectors



Solutions to DDoS Problems

- Attack prevention and preemption
 - Before the attack
- Attack detection and filtering
 - During the attack
- Attack source traceback and identification
 - During and after the attack

Attack Prevention and Preemption

- Signatures and scanning procedures exist to detect agent implants
- Monitor network traffic for known attack messages between attackers and masters
- Cyber-informants and cyber-spies
- Some users just don't care
- No incentive for ISPs or enterprise networks do not have incentive to monitor for attack packets

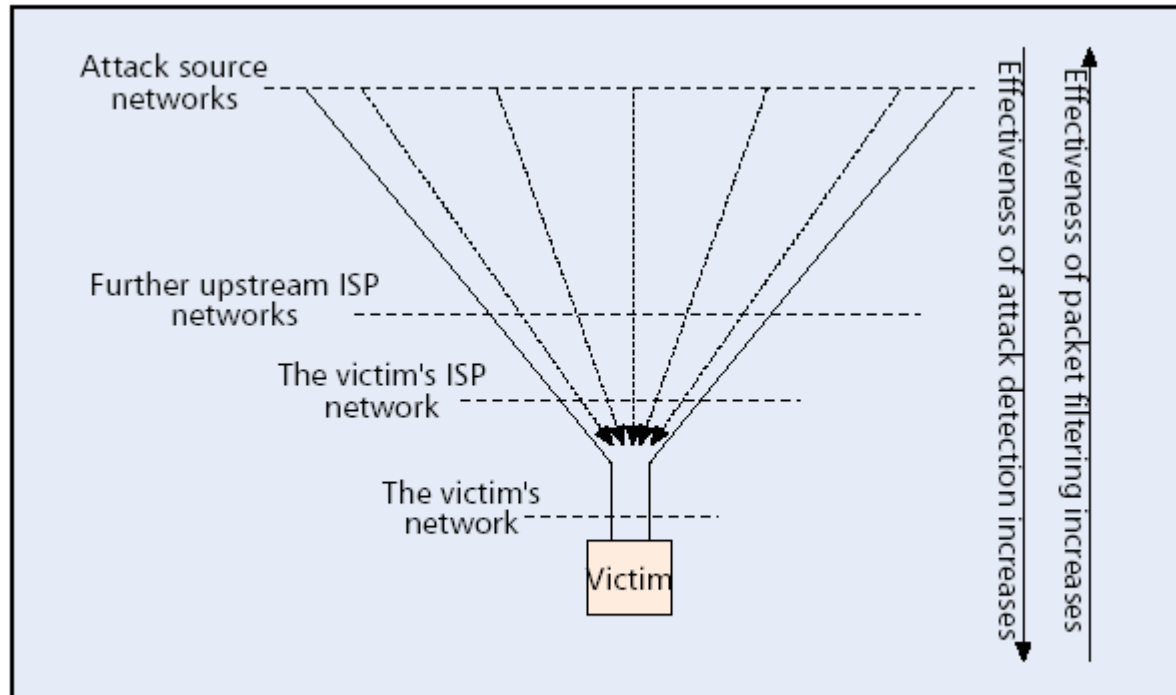
Attack Source Traceback and Identification

- Trackback – identifying the actual source of packets, without relying on header information
- Two approaches
 - Router records information about packets
 - Router sends additional information to destinations, via the packets or ICMP messages
- Cannot be used to stop an ongoing attack
- Packet's origin cannot always be traced (firewalls and NAT)
- Ineffective in reflector attacks – Packets come from legitimate sources in
- Used to collect evidence for post-attack law enforcement

Attack Detection and Filtering

- False positive ratio (FPR)
 - Packets classified as attack packets that are actually normal, divided by total normal packets
- False negative ratio (FNR)
 - Packets classified as normal that are actually attack packets, divided by total attack packets
- Packet filtering drops attack and normal packets
 - Effectiveness measured by normal packet survival ratio (NPSR)

Attack Detection and Filtering



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

Attack Detection and Filtering

- Source Networks – can filter packets
- Victim's Networks – can detect attack
- Victim's Upstream ISP
 - Requested to filter attack packets (by phone)
 - Ideally an intrusion alert protocol would be used
- Further Upstream ISP
 - Networks would have to cooperate and install packet filters when intrusion alerts are received

Internet Firewall

- Detect DDoS attack in the Internet core
- Could maintain a victim's normal service during an attack

Route-based Packet Filtering

- Extends ingress packet filtering to core
- Checks if packet comes from correct link, according to inscribed source and destination
 - If packet is from unexpected source it is dropped
 - Route changes can cause false positives
- Packet filters in 18% of ASs in Internet can significantly reduce spoofed packets
- BGP messages would require source addresses, increasing message size and time
- Currently there are > 10,000 ASs, so 1800 filters would have to be in place

Distributed Attack Detection Approach

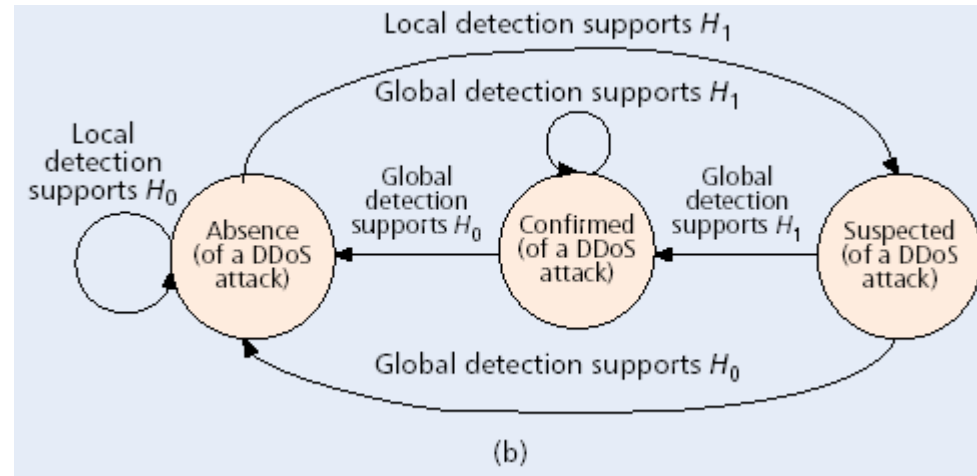
- Extends intrusion detection system to core
- Detects based on network anomalies and misuses observed by detection systems (DSs)
- Anomaly detection determines normal and deviant traffic patterns
- Misuse detection identifies attack signatures

Detection Systems

- Placed in strategic locations
- Nonintrusively monitor traffic
- Exchange attack information from local observations
- Stateful to presence or absence of DDoS attacks
- Need a separate channel to communicate
- Number of DSs is much smaller than RPF, DSs does not rely on routing information
- More DSs would result in a larger delay response

Detection System Design

- Process packets at very high speeds
 - Need a high-speed packet classifier
- Local and global detection
 - H_1 – presence of a DDoS attack
 - H_0 – a null hypothesis
- When H_1 occurs, alerts sent to other DSs
 - Each DS analyzes its results and other DSs results to make a global detection decision
 - Attack confidence level
 - If DS is confirmed, filters are installed, optionally notifies upstream routers



Detection System Design

- Install filters only on suspected switch interfaces
- DSs must always be connected, physically and have usable paths
- Questions remain – best topology, how to reconnect DSs, how does DSs send alerts when it is under attack
- Communication Protocols
 - Intrusion Detection Exchange Protocol
 - Intrusion Detection Message Exchange Format

Quickest Detection

- Studied in signal processing, quality control, and wireless channel monitoring
- DS periodically computes instantaneous traffic intensity
- Objective is to minimize the expected delay in detection, based on thresholds

Limitations and problems

- Need to determine thresholds for local and global thresholds and traffic modeling
- There is a delay to reach global detection, DS network does not detect short attacks
 - DS network should be designed for attacks > 5 min (75% of all attacks in a recent study)
- Flash crowds result in false alarms
 - Unpredictable – major news stories
 - Predictable but nonrepetitive – sports
 - Predictable and repetitive – opening of stock market
 - Use a different traffic model when flash crowd occurs
- Degradation of Service Attacks (DeS)
 - Short bursts of attack packets

Comparison

	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.

Conclusion

- Current defense is inadequate
- Still many insecure areas on the Internet
- More effective detect-and-filter approaches must be developed

What's the big deal?

- Argues for the use of an Internet Firewall
- Compares and contrasts route-based packet filtering and distributed attack detection



Questions