UNIVERSITY OF TWENTE.

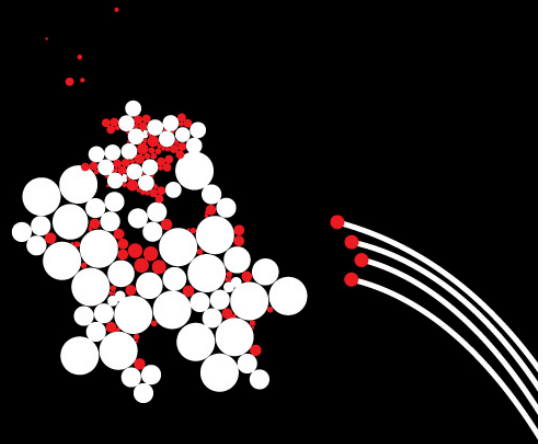

Network Security

Attack and Defense Techniques 2

Anna Sperotto, Ramin Sadre

Design and Analysis of Communication Networks (DACS)

University of Twente

The Netherlands

UNIVERSITY OF TWENTE.

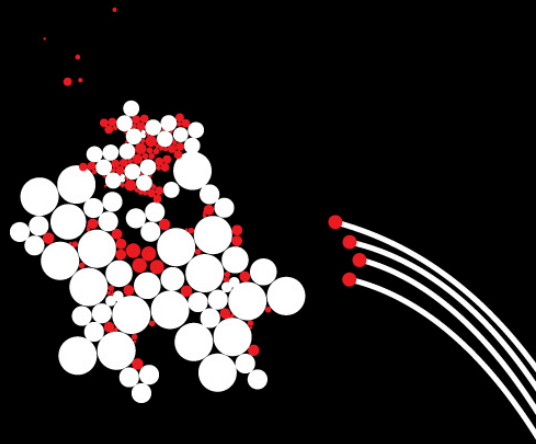

Firewalls

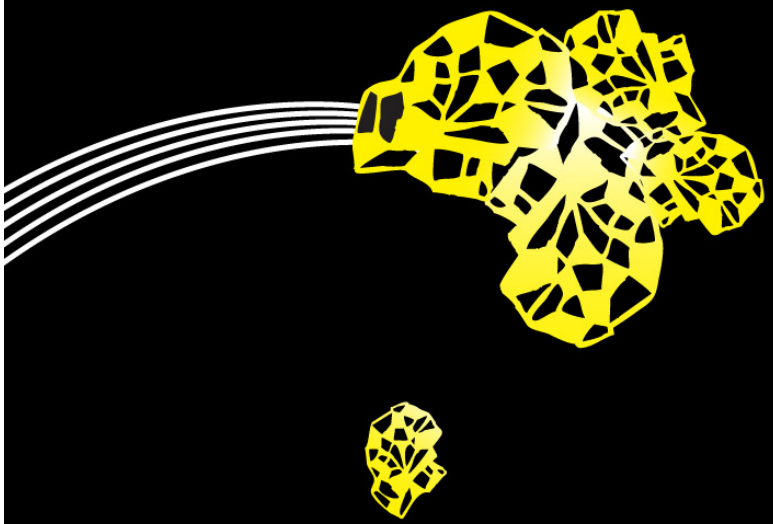

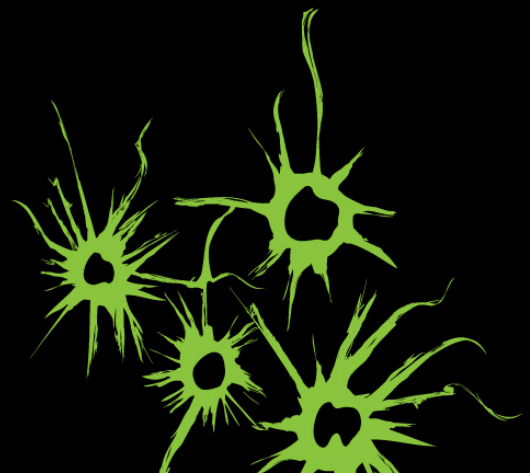

Network firewall

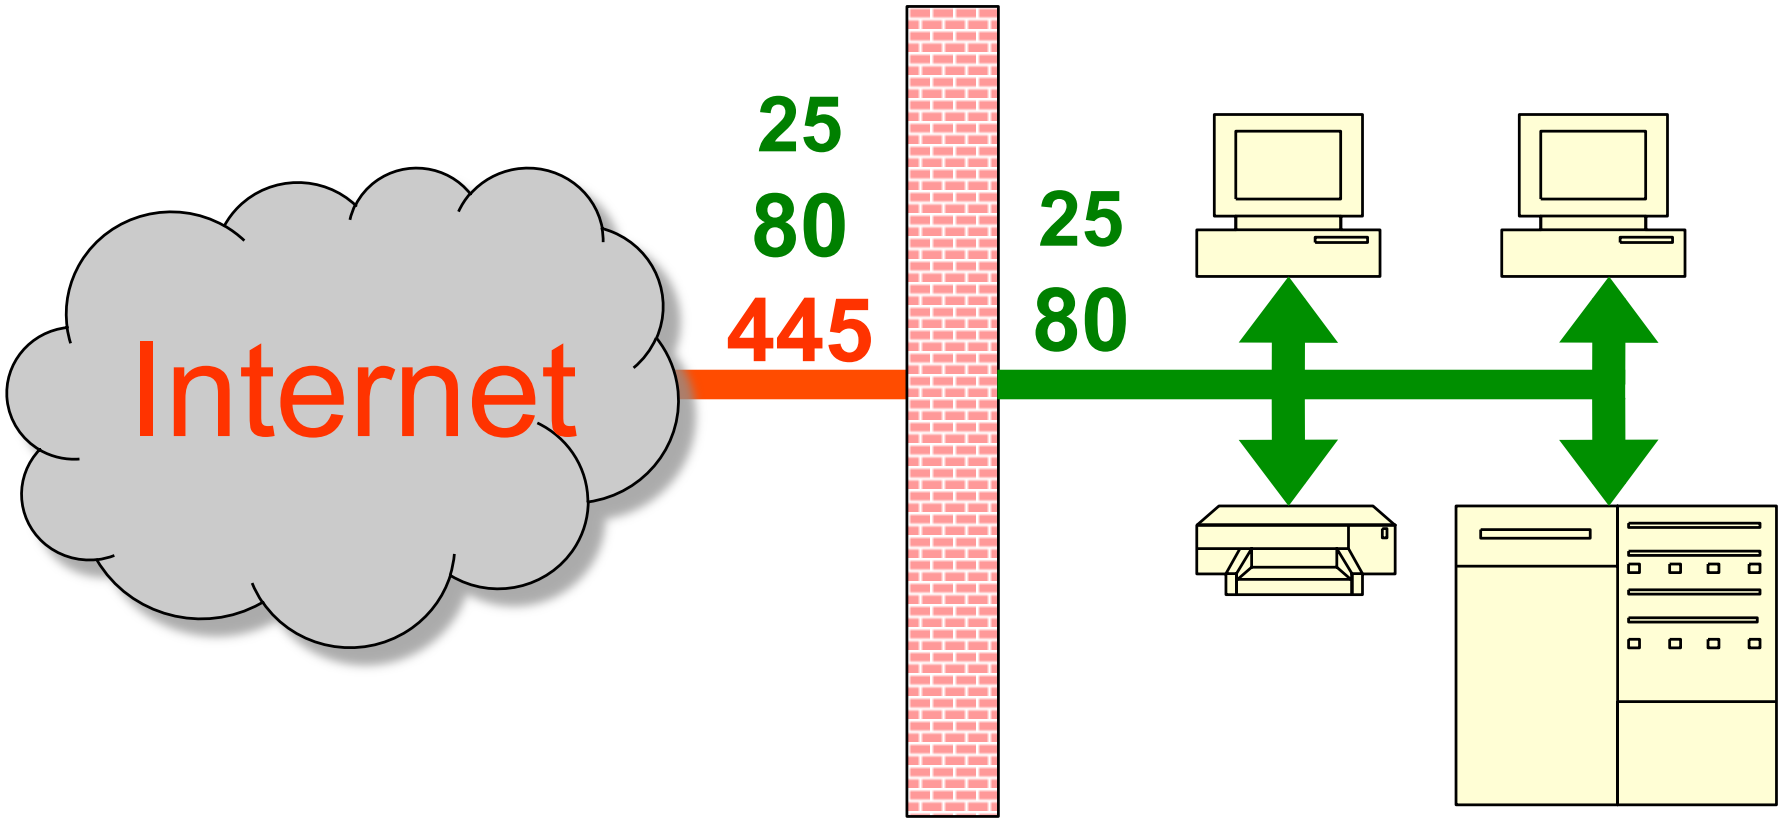

Personal firewall

- Runs on the computer of the user
- Same filtering capabilities as network firewall
- Filter may also distinguish between computer programs

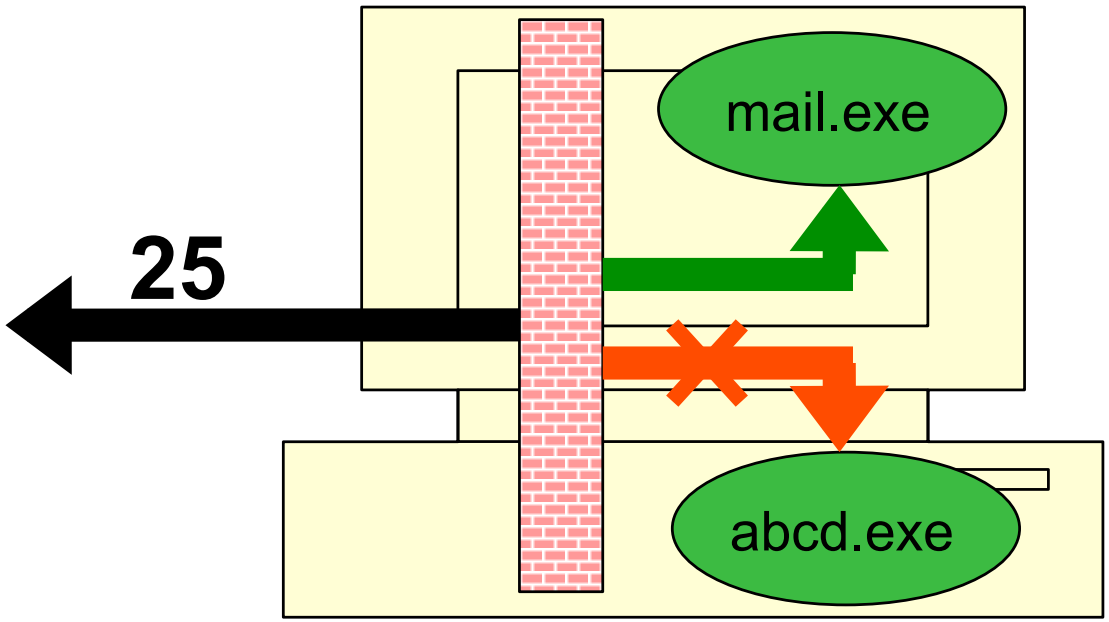

Network-level firewall

- Filters on IP header fields, such as:
  - Source/Destination IP address
  - Type of Transport protocol

- Default policies:
  - *Discard:* what is not explicitly permitted is discarded
  - *Forward:* what is not explicitly prohibited is allowed

- Stallings calls this “Packet-filtering firewall”

Transport-level firewall

- Filters additionally on TCP header fields, such as:
  - Source Port
  - Destination Port
  - Flags (SYN, ACK)
- Stallings calls this “Circuit-level Gateway”

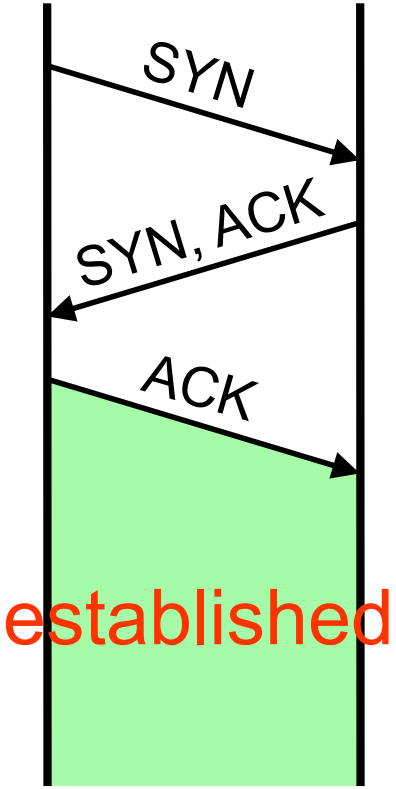

Application-level firewall

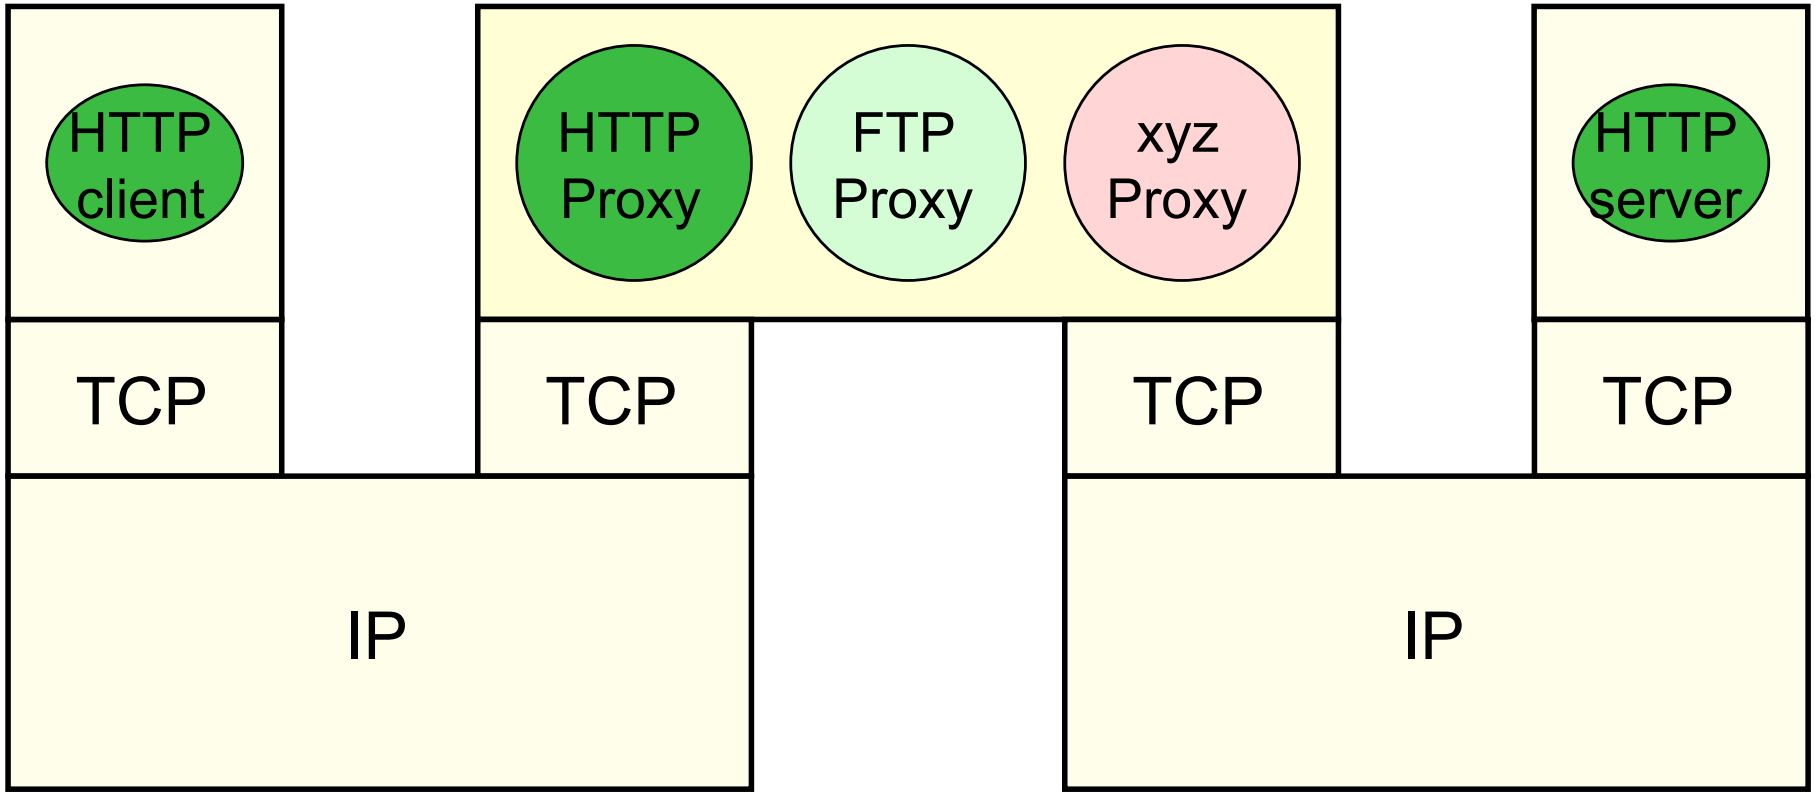

Application-level firewall

- Inspects the contents of packets
- May filter certain websites, mail-viruses etc.
- Firewall may accept only trusted connections
- Logging of accepted connections is easy
- Performance may be problematic
- Since this type of firewall is quite complex, it may become a security risk itself

Stateless firewall

- Treats each packet in isolation
- Has no memory of previous packets
- For each packet checks firewall rules again
- Easy to implement / very efficient
- Can not easily handle protocols that use random ports, such as FTP

Stateless firewall - Example

action	src	port	dest	port	flags
allow	*	*	*	**80**	*
allow	{our hosts}	*	*	*	*
block	*	*	*	*	*

Stateful firewall

- `IF (packet belongs to an existing “association”)`
- `  THEN {accept packet}`
- `  ELSE {checks firewalls rules;`
- `    IF (packet may pass)`
- `      THEN {store “association” in state table}`
- `      ELSE {discard packet}}`

- Time-out inactive connections
- Connections may send “keep alive”
- SYN attack can overflow State table

Stateful firewall

- Associations may be:
- TCP connections
- UDP flows
- ICMP request/response pairs

- Stateful firewalls can, for example, be configured to:
- Allow “associations” initiated by internal systems
- Deny “associations” initiated by external systems

- Stateful firewalls can easily deal with protocols such as FTP

Where to put the firewall

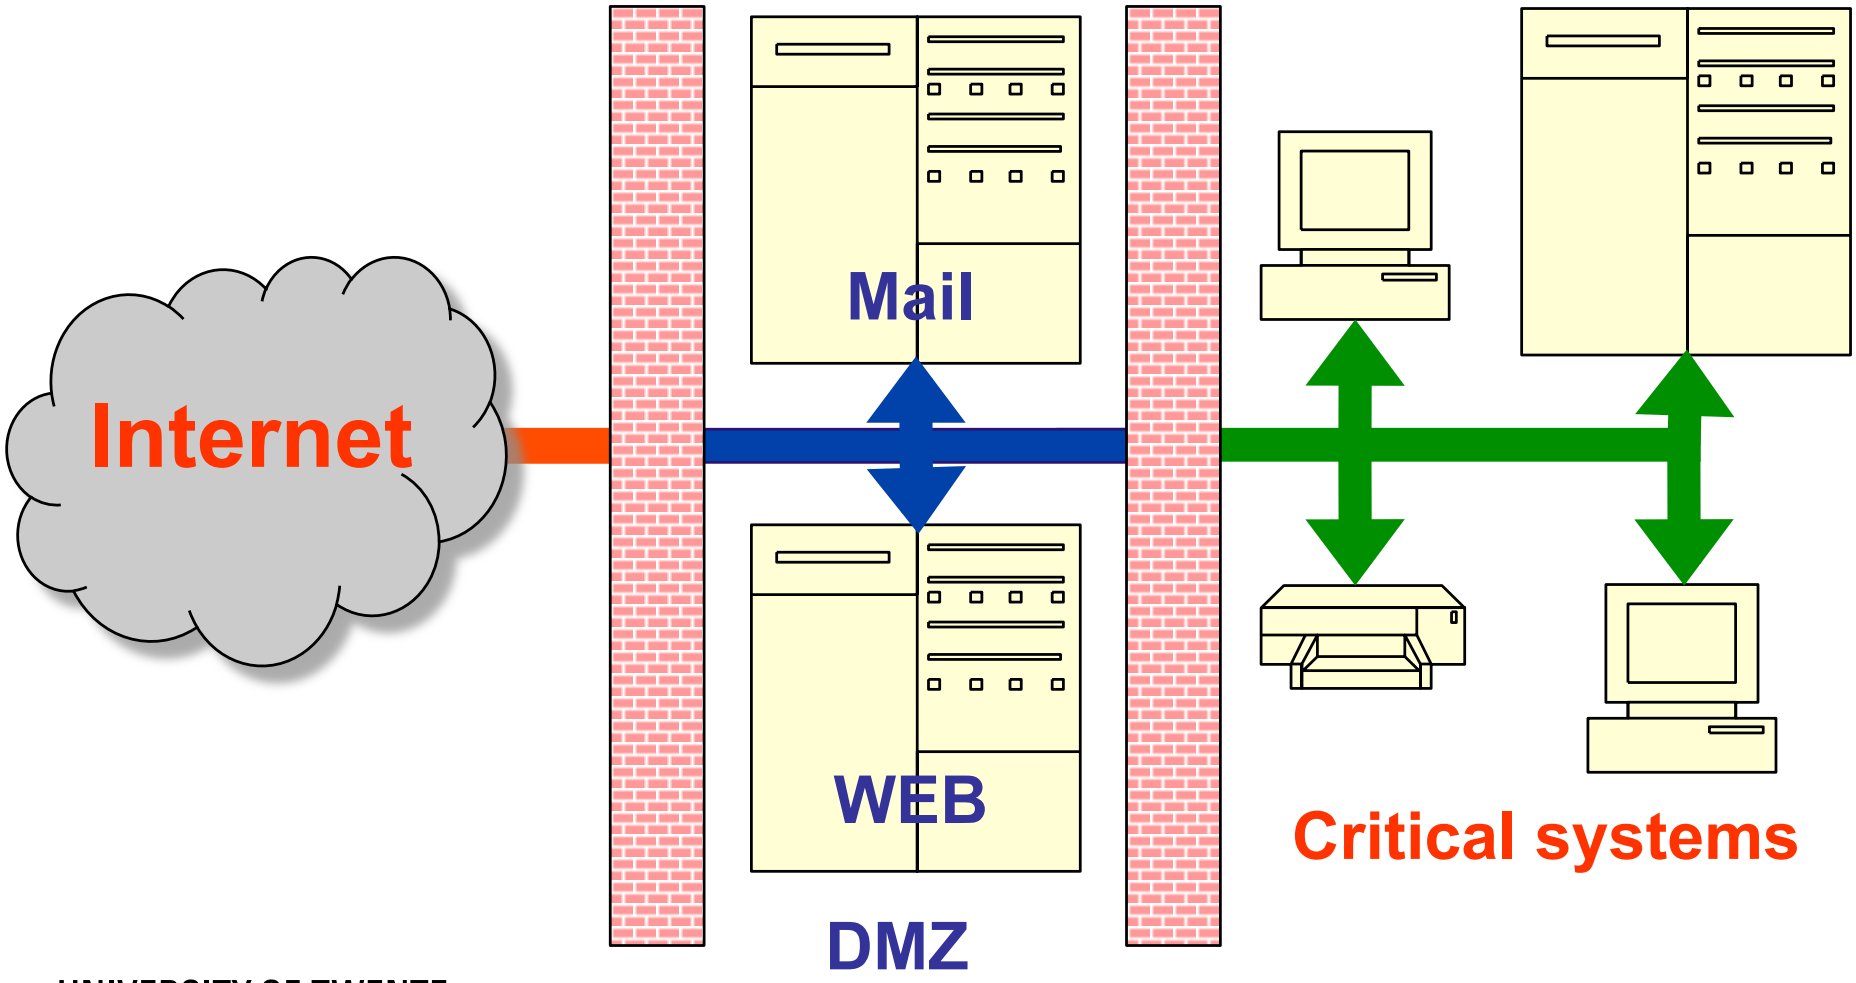

Firewalls versus Network Address Translators

- Origin of NATs is different from that of firewalls
- Like Application-level firewalls, NATs modify IP addresses and Port numbers
- In general, NATs do not inspect application data
- NATs can be compared to transport-level firewalls
- Like certain firewall configurations, certain type of NATs accept incoming data only after an external “connection” has been established
- If both sides have firewalls / NATs, communication may be difficult / impossible

UNIVERSITY OF TWENTE.

DNS tunneling

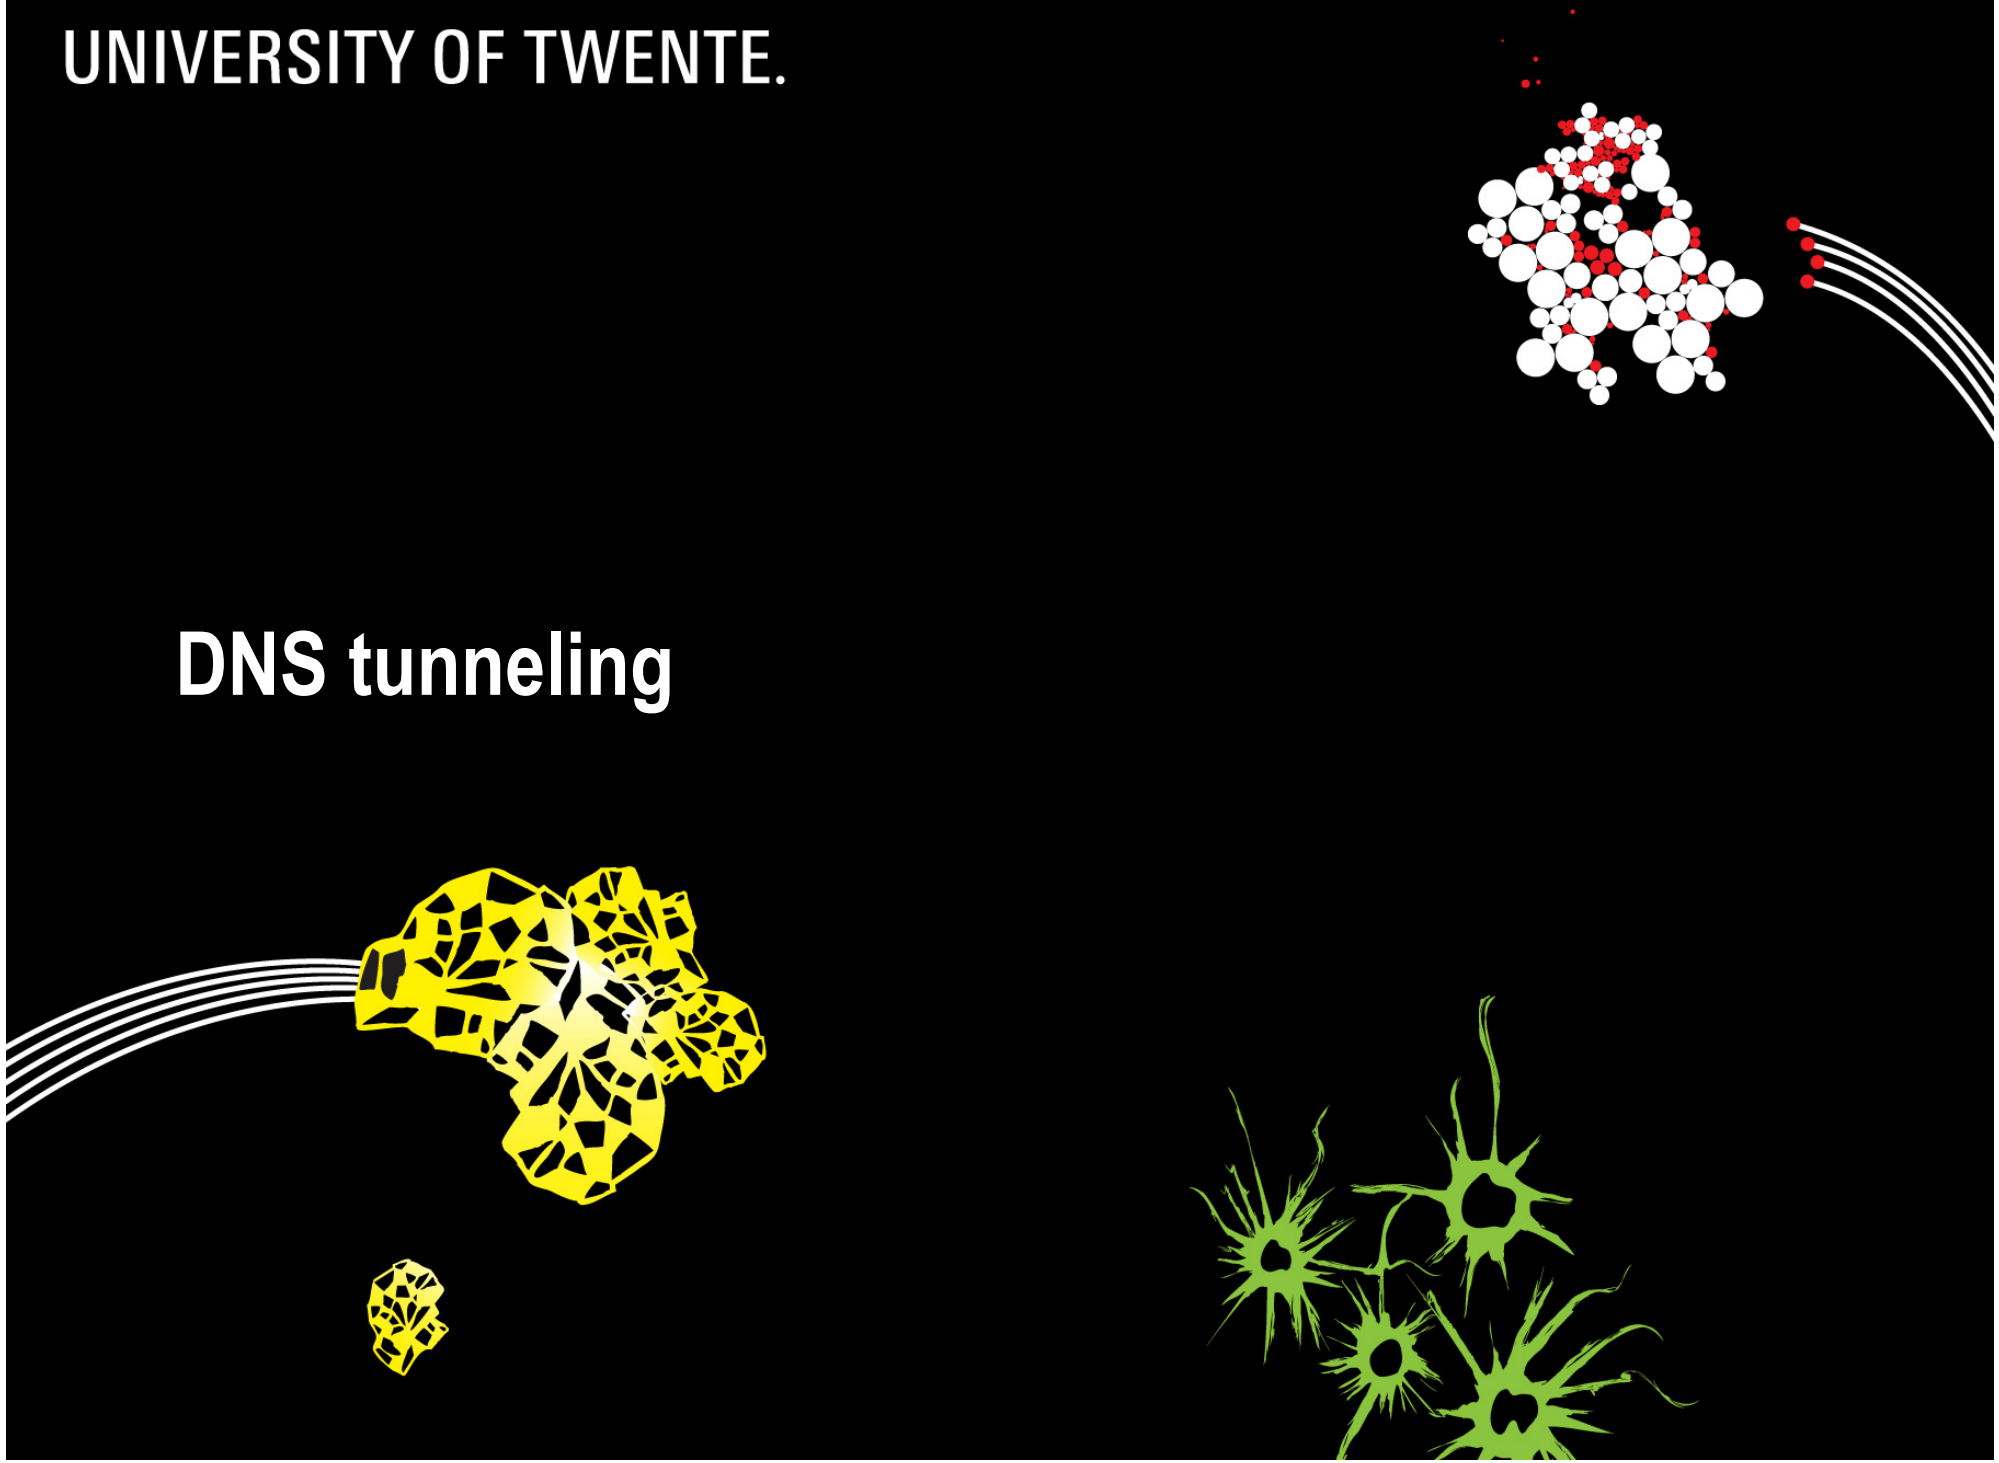

DNS tunneling

- You sit at the airport
- WLAN provided, but any access to a Web server, FTP, P2P,… is chargeable
- Is there a way to avoid the fee?
- Would it be an *attack*?
  - You are bypassing the billing/security policy of your ISP
  - Data exfiltration for cyber-espionage

DNS tunneling: upstream

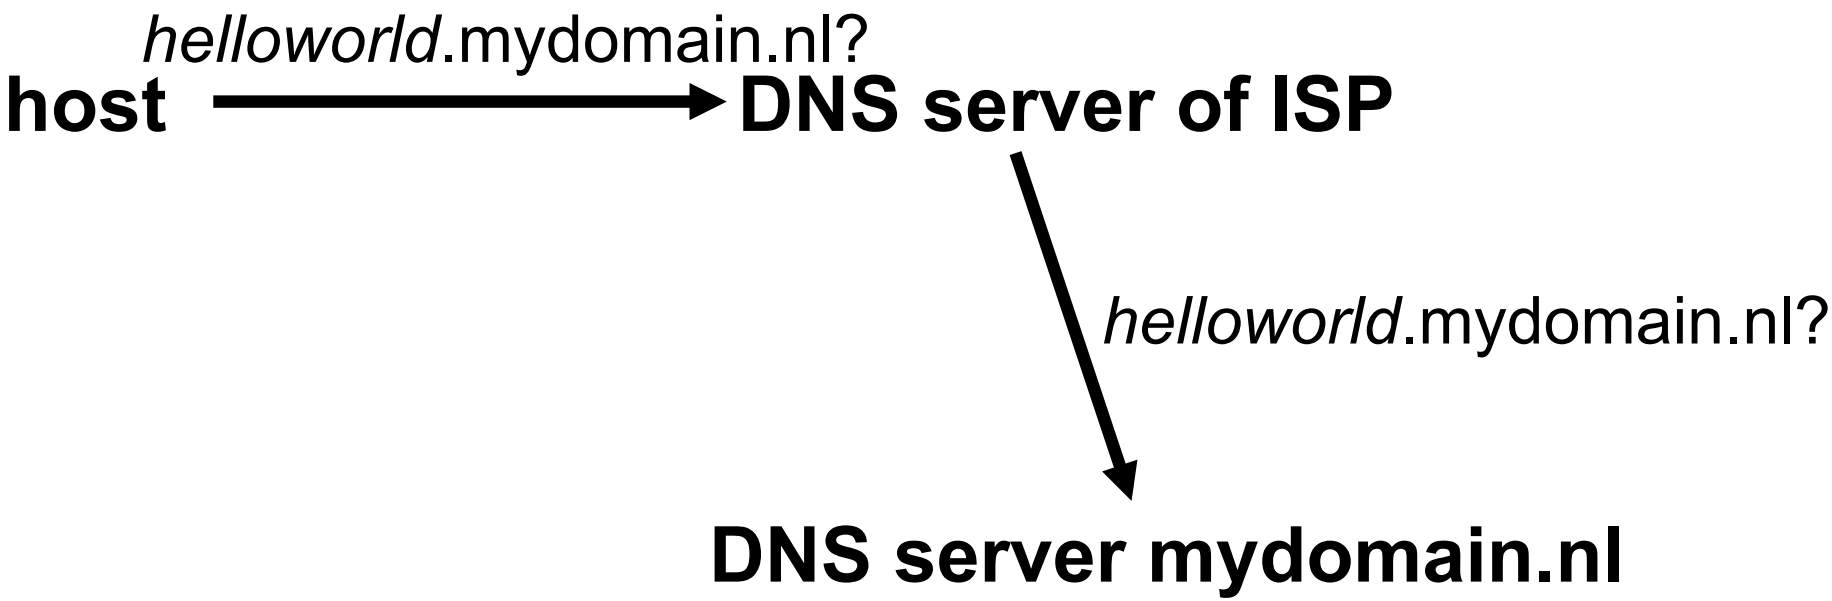

DNS tunneling: upstream

- Query can contain up to 252 characters
- Character set restricted: not case-sensitive,...
- ~5 bit/character, ~110 bytes

DNS tunneling: downstream

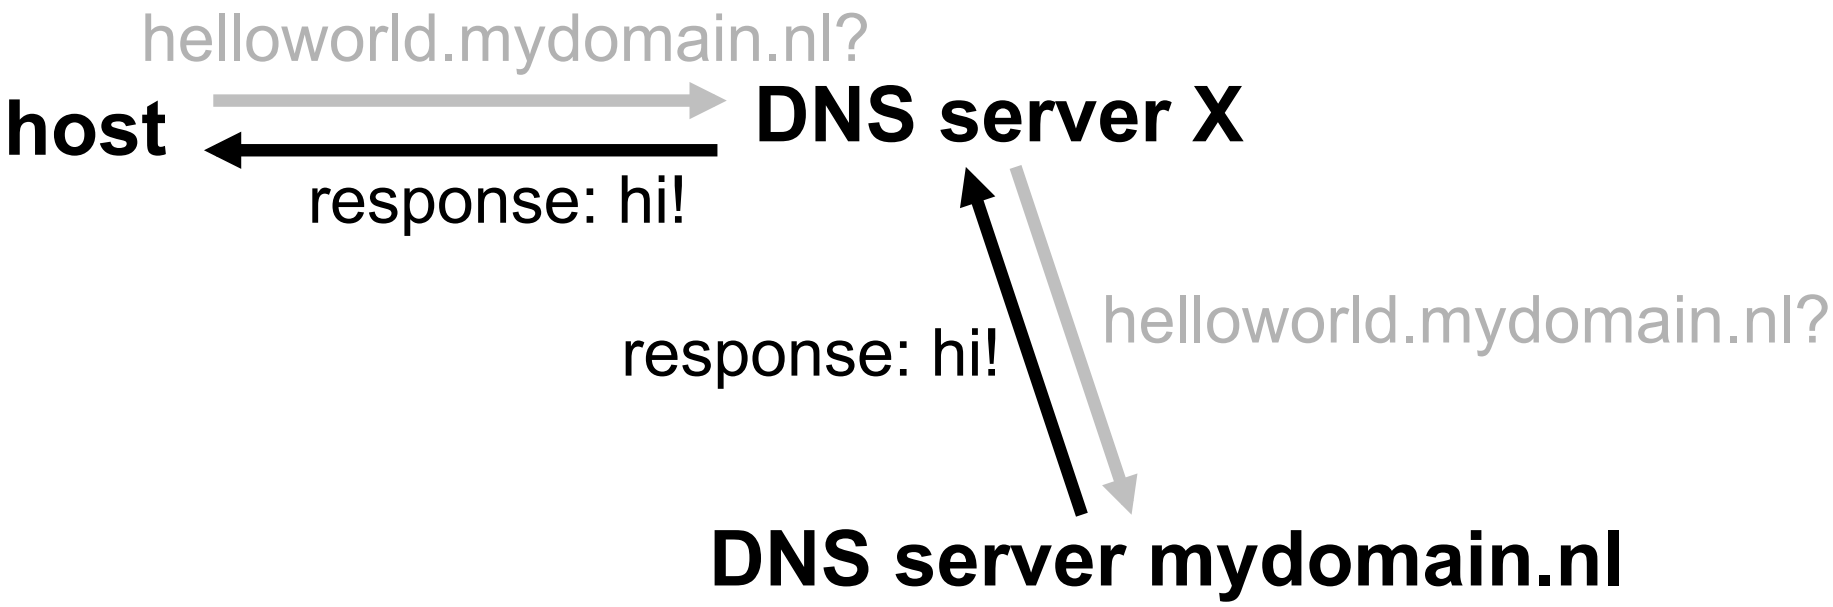

DNS tunneling: downstream

Main limitation:

- Response < 512 bytes to prevent fragmentation

Server responds with TXT-record:

- Character set restricted: 7 bit ASCII
- ~6 bit/character, ~220 bytes

As for amplification, EDNS (RFC 2671) allows larger answers

Using MX-records and A-records is possible, too, but more complicated (data may be reordered)

Example of DNS Tunneling (Iodine)

1329812676.512747 IP 192.168.1.2.60531 > 192.168.1.1.3128: Flags [S]

1329812676.515310 IP 1.1.1.1.51823 > 2.2.2.2.53: 22911+ [1au] NULL? 0eaba82M-J2hbM->M-nYM-VwjM-GM-MRbM-^M-^PM-\M-UM-HcvM-DtimM-eM-`M-KyM-aM-VM-IM-yM-yM-BM-jdilmnuM-iM-bM-ktaM-^XyUwtf.M-BM-^M-o8M-]M-=M-xM-=M-FouZzM-JwaeM-NaM-u….test.domain.nl

1329812676.524541 IP 192.168.1.1.3128 > 192.168.1.2.60531: Flags [S.]

1329812676.525743 IP 2.2.2.2.53 > 1.1.1.1.51823: 15184 1/0/1 NULL (140)M-N.test.domain.nl. (130)

1329812676.524573 IP 192.168.1.2.60531 > 192.168.1.1.3128: Flags [.]

1329812676.526742 IP 1.1.1.1.51823 > 2.2.2.2.53: 30638+ [1au] NULL? 0ibbb82M-J2hbM->M-nYM-VgjM-GM-MBbM-^M-^PM-\M-TM-XcvM-DtimM-eM-`M-KyM-aM-VM-IM-yM-yM-CDYM-eM-X3qWgM-JM-SM-qSM-?M->M-bYyCU.xpM-_M-VM-`M-HEM-LJM->M-nf6upM-{M->.test.domain.nl. (126)

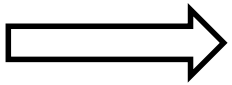

1329812676.557242 IP 2.2.2.2.53 > 1.1.1.1.51823: 22911 1/0/1 NULL (144)

1329812676.525189 IP 192.168.1.2.60531 > 192.168.1.1.3128: Flags [P.], (request web page)

1329812676.558096 IP 1.1.1.1.51823 > 2.2.2.2.53: 38365+ [1au] NULL? 0mbbc82M-J2hbM->M-nYM-VhdM-yEM-rdM-?M->M-q5MM-tcvM-DtimM-eM-`M-KyM-aM-VM-IM-yM-yM-CDYM-eM-X3qWMM-JM-SM-CM-CM-DdbM->M-bM-p4.CM-=wM-icOM-x4oM-YM-kM-gM-SiHM-OM-guM-JcPM-<M-=rM-K0M-rf8M-cM-=M-XPgM-@M-HM-RM-\5FM-SM-uM-yM-CM-PM->GM-]M-hiM-?M-wQM-KFM-HM.0M-wM-zM-_UM-ZM-MwM-RM-C6M-?M-PpWM-tRPM-RM-fWyuM-\qM-FGtM-NBM-sgM-<huuTNl6NQ1FM-KvSkWM-H9ESaIM-AX.M-OHM-OM-bYM-wM-PM-C3MM-MM-dM-HAM-\3rM-bM-LMM-QfM-^ALM-UM-g18UhM-]CQaM-K6M-lM-mlM-IM-`M-naIDM-NM-cM->.test.domain.nl. (274)

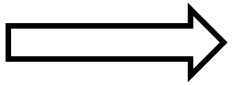

How to detect DNS tunneling?

- Tunnel packets have different characteristics compared to regular DNS packets.

- http://armatum.com/blog/2009/a-study-of-dns/

UNIVERSITY OF TWENTE.

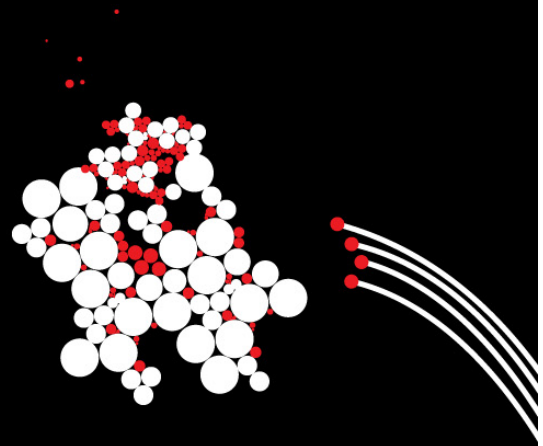

Intrusion Detection Systems

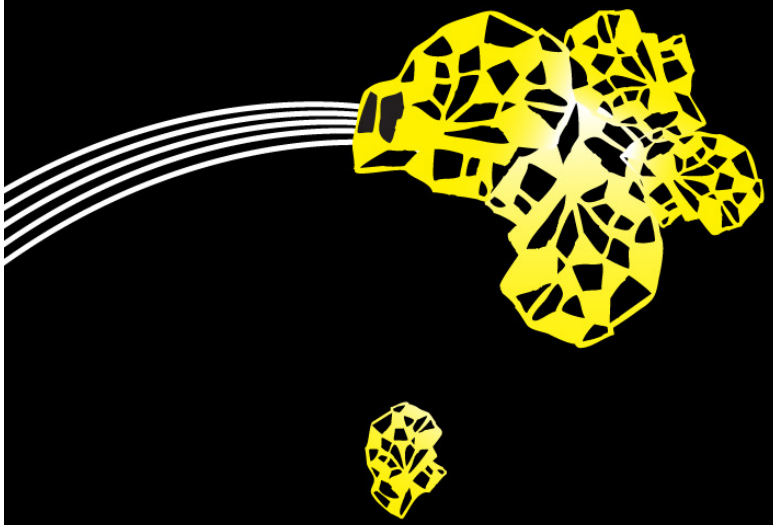

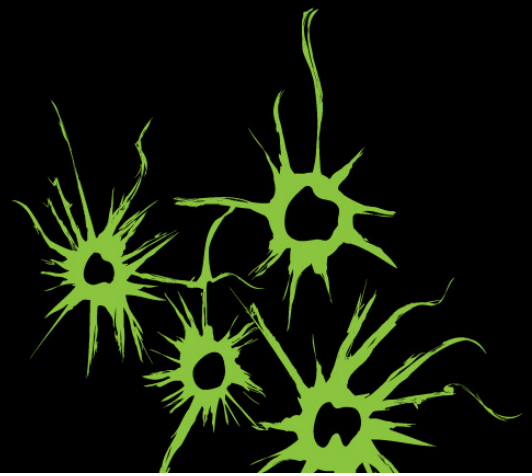

Intrusion Detection System

- Intrusion detection is the process of identifying (and responding to) malicious activities targeted at computing and network resources
- Goal: identify intrusions and report them

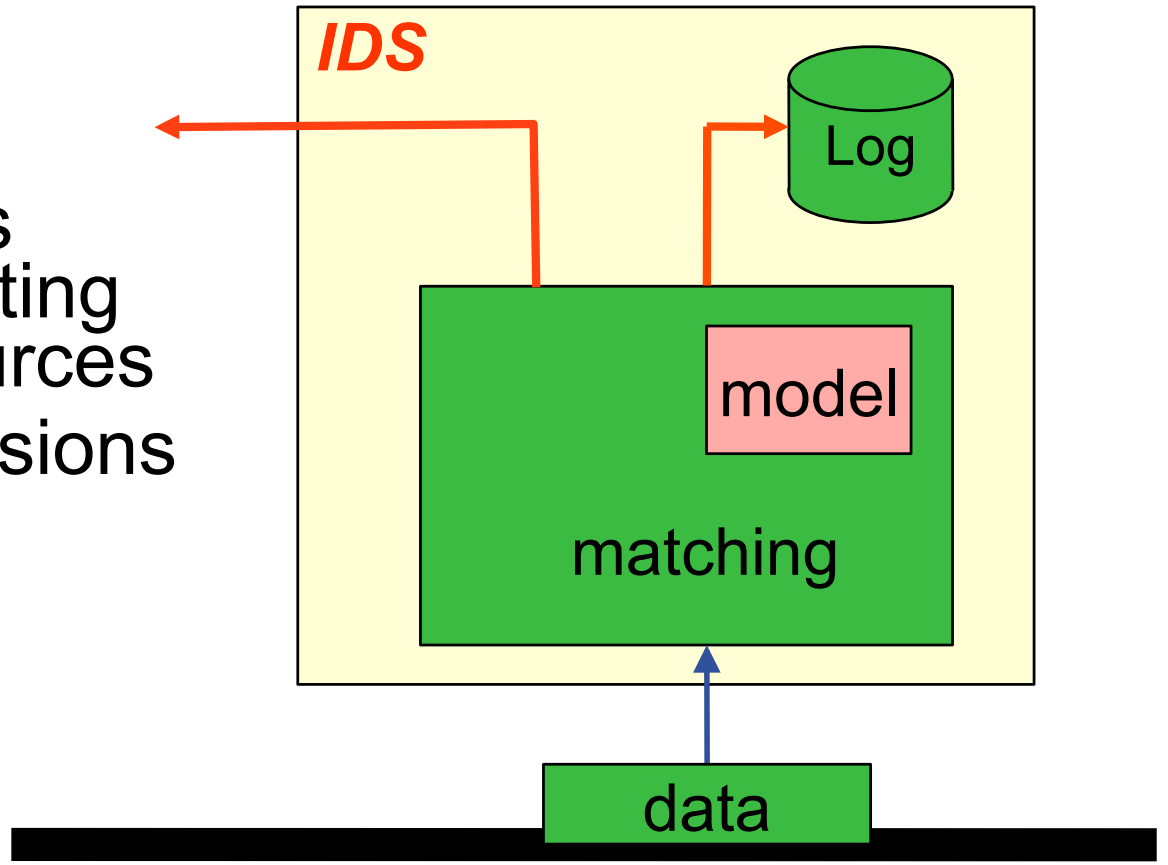

IDS Taxonomy

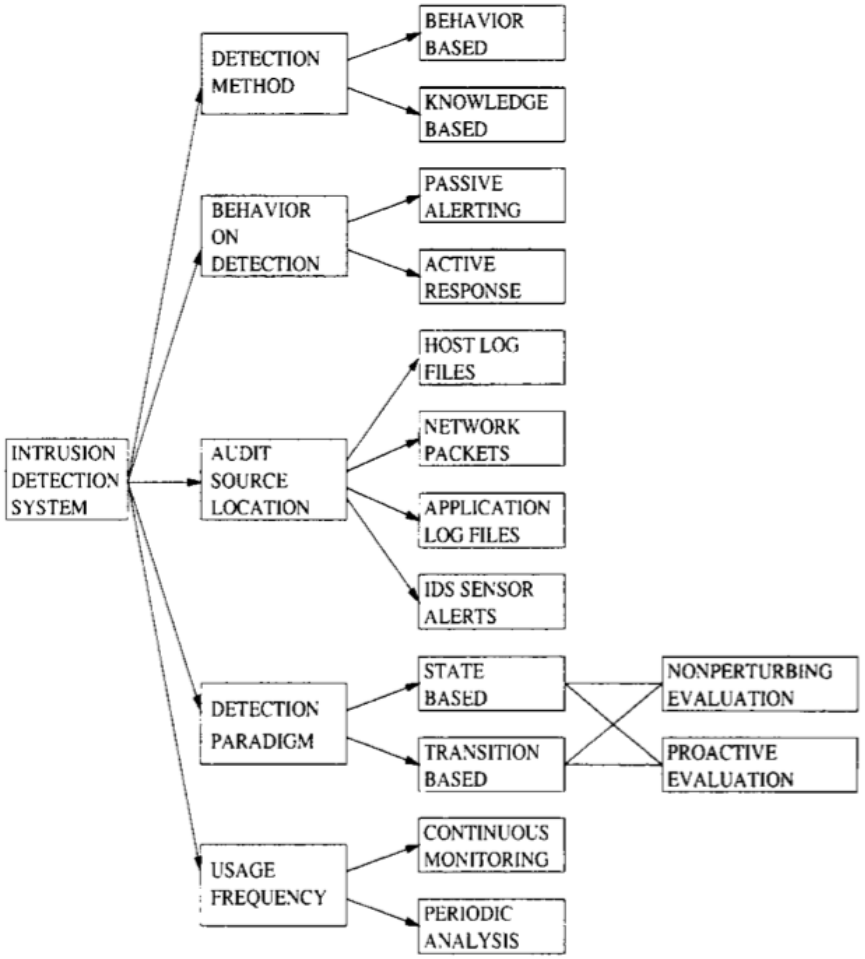

For more info: "A revised taxonomy for intrusion-detection systems", Debar, Dacier, Wespi, 2000
http://www.springerlink.com/content/4xq65ng0l0801626/

Knowledge-based vs Behavior-based

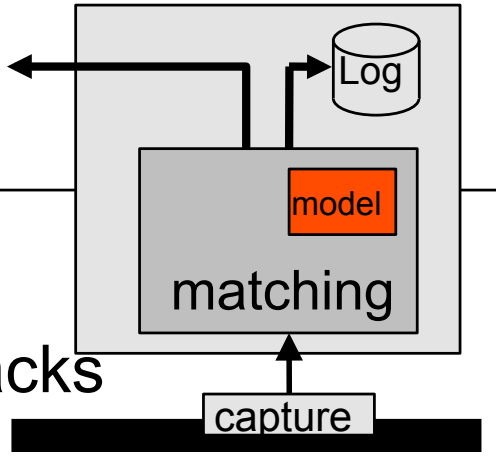

- Also signature-based vs anomaly-based
- Signature-based IDS: Model/definition of attacks
  - Use exploits or attack signatures
  - Can only detect known attacks
    - Example: SNORT

- Anomaly-based: Model of normal behavior
  - Detect deviations
  - Can detect unknown attacks
  - It often needs tuning

IDS - Performance

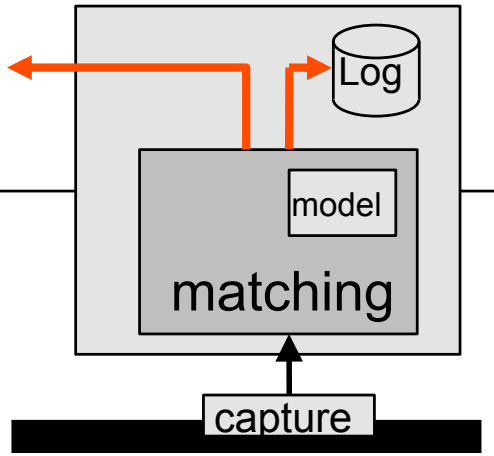

IDS Outputs

Ground Truth	Positive	Negative
Malicious	TRUE POSITIVE This is really an attack	FALSE NEGATIVE IDS fails to report an activity as malicious
Benign	FALSE POSITIVE IDS reports malicious activity, though activity is OK	TRUE NEGATIVE This is really safe traffic

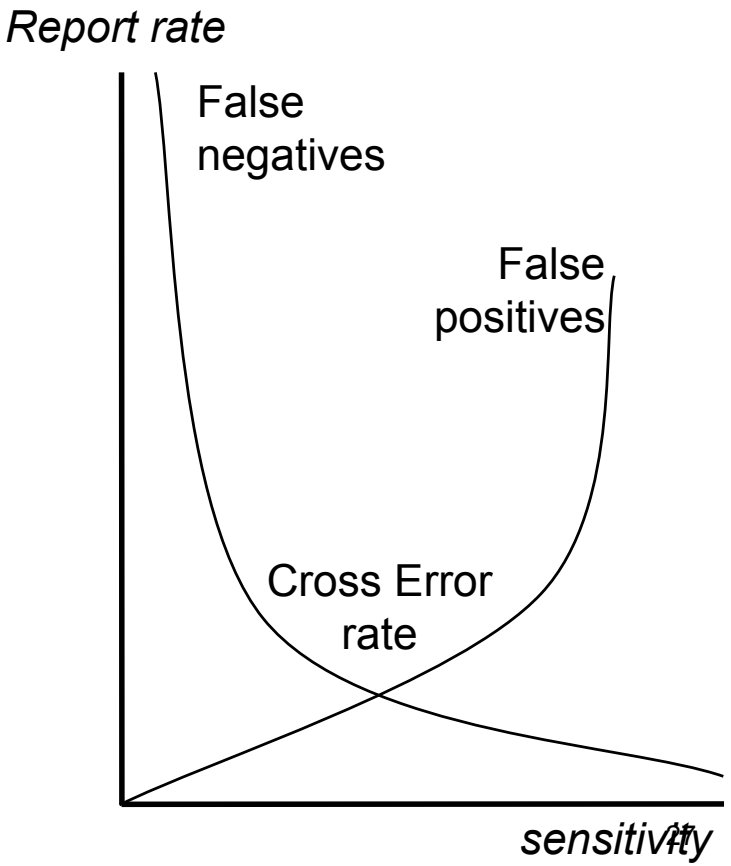

Performance tuning example

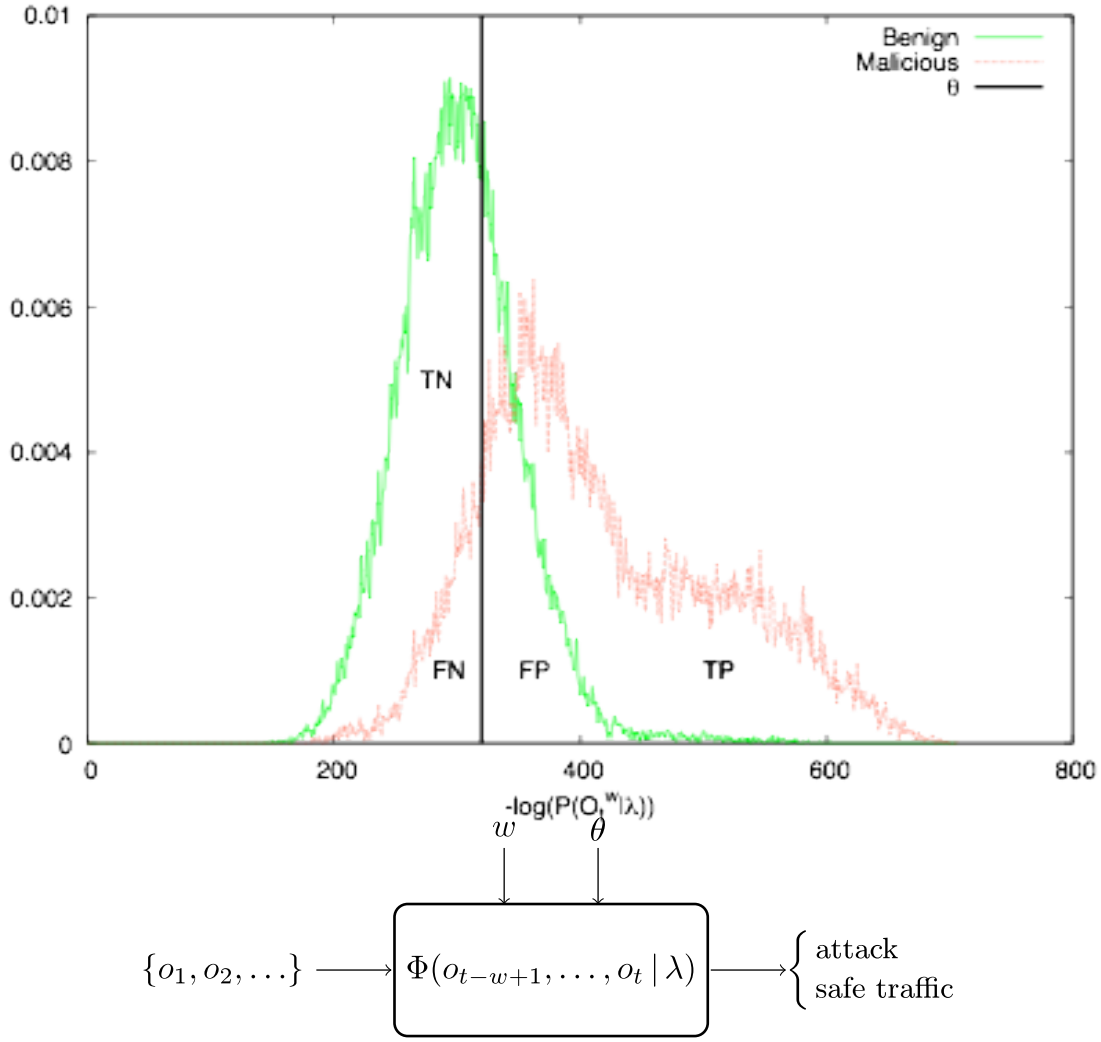

- Tuning can be modeled as an optimization problem

$$\max_{w\theta} \alpha \cdot TN + \beta \cdot TP$$

What is missing in the taxonomy?

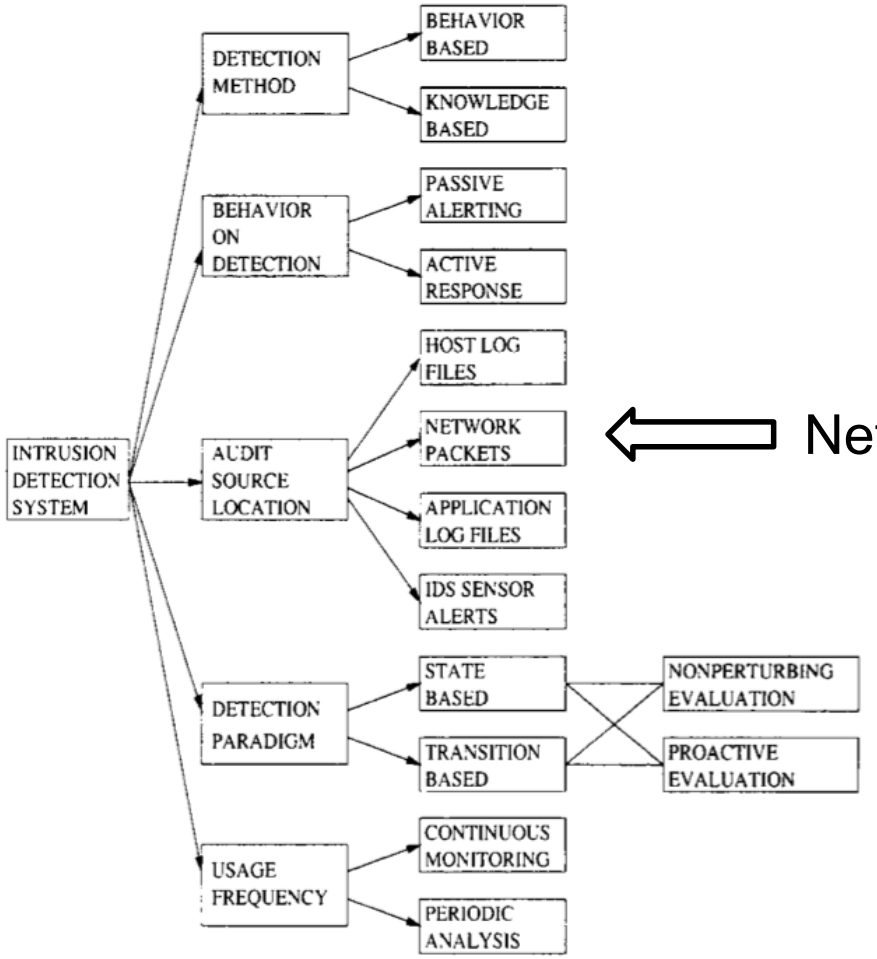

⟸ Network flows

Flow-based intrusion detection

- Can you perform intrusion detection in a backbone network?
  - Several Gbps (link UT-SURFnet 10Gbps; ESnet going towards 100Gbps)
- Deep packet inspection is typically not scalable to these rates

- Aggregation: data reduction!
  - First look only to packet headers
  - Not enough: **network flows**
- Applicable also in presence of data encryption

Network flows

- As defined by the IETF IPFIX working group:

  A set of IP packets passing an observation point in the network during a certain time interval and sharing a set of common properties (RFC 3917).

  - Basic flow definition:

```
(src IP, src port, dst IP,
  dst port, IP protocol,
  number of packets,
  number of bytes)
```

  - Data reduction: 30x (almost as the heights of a Xperia mini vs a phone booth)

Flow-based Intrusion Detection

- The cost of data reduction: no payload available

- What can be detected:
  - attacks that create variations in volume/number of flows
    - Scans
    - DDoS
    - Spam campaigns

Can we detect worm spreading using flows?

- Example: graph-based detection of hit-list worms
- Hit-list:
  - bootstrapping the spreading phase
  - It would change the connection patterns in a network
    - Number of hosts
    - Connected components

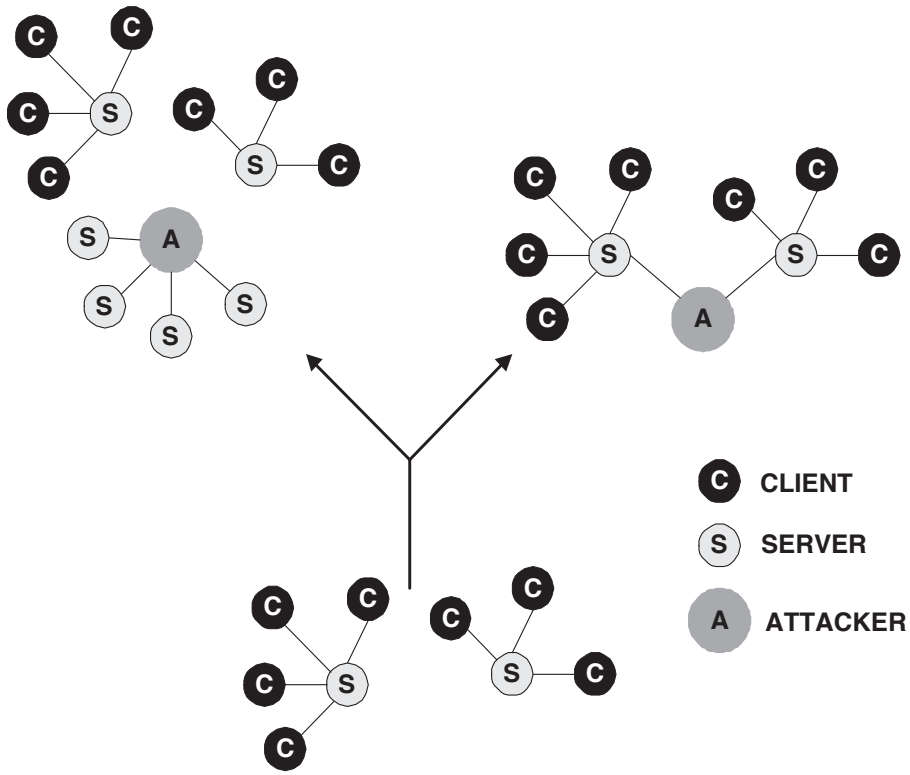

UNIVERSITY OF TWENTE.

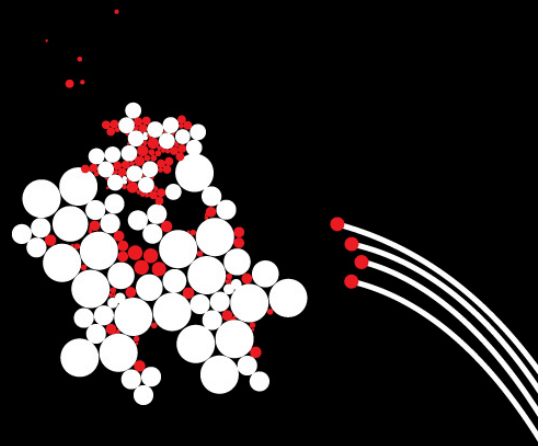

Other defense techiques

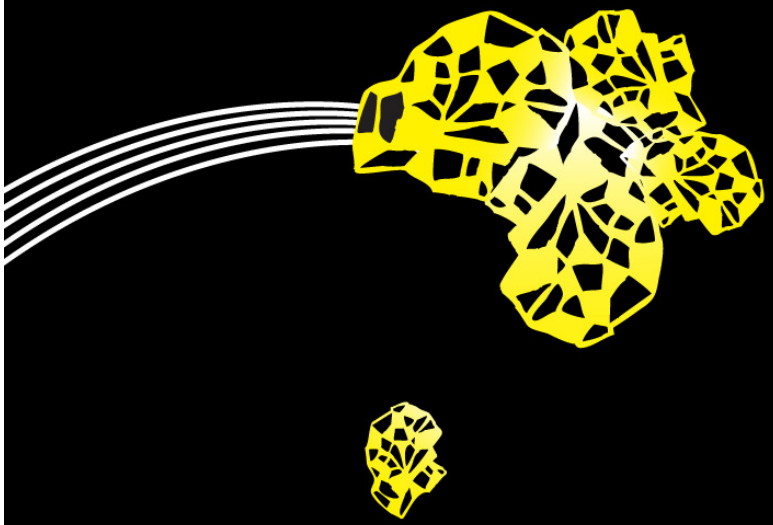

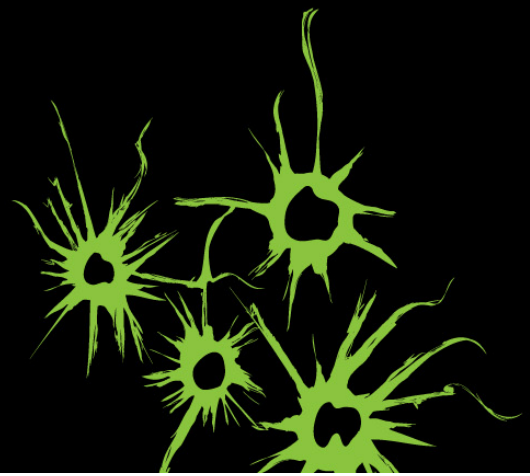

It all starts with monitoring...

- From last week lecture: the network telescope
  - Identify intensity and frequency of attacks (DoS)
  - Misconfigurations
  - Worm spreading (Code Red, Sapphire)
  - Botnet behaviors
  - See http://www.caida.org/publications/papers/

Spam: some numbers

- In May 2009, 58% of all spam e-mails were delivered by botnets
- USA, September, 2009:
  - Zeus: 3.6 million zombies
  - Koobface: 2.9 million zombies

Spam-Campaign on Storm Botnet

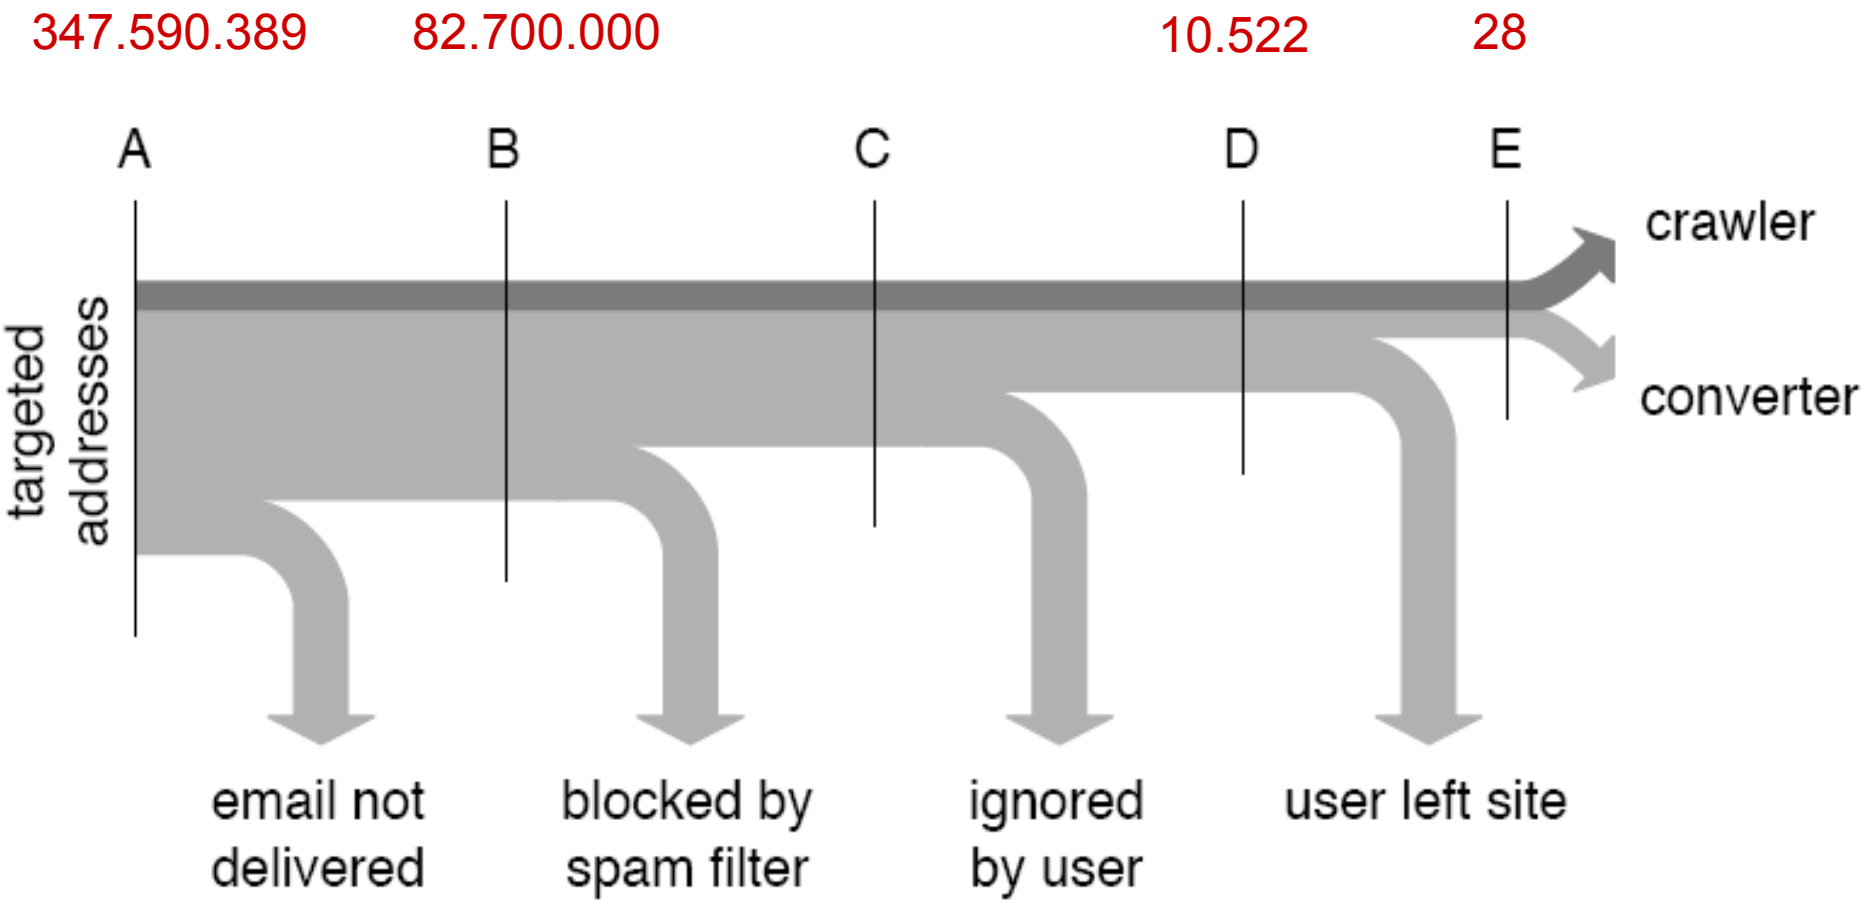

(Source: Spamalytics: An Empirical Analysis of Spam Marketing Conversion, Kanich *et al.*, 2008)

DNS Blacklists

- How do we know if a host has sent SPAM?
  - SPAM filter on the local Mail server
  - SPAM traps: hosts that receive and collect information about SPAM messages
  - DNS-Blacklist: list of IP addresses that sent mail to SPAM traps
    - Periodically updated
    - Many of them are publicly availble (CBL, PSBL etc..)
    - They use DNS as query protocol for retrieving data

Bad Neighborhoods

- Suppose you do not want/ cannot access the body of a mail. Can you say if it is SPAM or not?
- There is a correlation between the source IP address of a message and the amount of malicious activities from the same subnetwork