ALBERT-LUDWIGS-
UNIVERSITÄT FREIBURG

Communication Systems

Firewalls

University of Freiburg
Computer Science
Computer Networks and Telematics
Prof. Christian Schindelhauer

CoNe Freiburg

IIF
INSTITUT FÜR
INFORMATIK
FREIBURG

Organization

- I. Data and voice communication in IP networks
- **II. Security issues in networking**
- III. Digital telephony networks and voice over IP

Network Security – “the magic device”: Firewall

- Take a completely new track now ...
- Firewalls are traffic / packet filters that operate on different layers of our OSI protocol stack
- Try for a definition: “A Firewall is a network security device designed to restrict access to resources (information or services) according to a security policy”
- Important remark is to be made here:
  - Firewalls are not a “magic solution” to network security problems, nor are they a complete solution for remote attacks or unauthorized access to data!!
  - Firewalls could be circumvented in several ways and may increase the complexity of network and this way decrease the level of security!

Network Security – Firewalls

- A Firewall is a often a network security device, but can be or simply is implemented directly into the end systems
- It serves to connect two parts of a network a control the traffic (data) which is allowed to flow between them
- Often installed between an entire organization's network and the Internet
- A Firewall is always the single path of communication between protected and unprotected networks
  - Of course there are special cases of multiple Firewalls, redundant connections, fault-tolerant failover etc.
  - A Firewall can only filter traffic which passes through it
  - If traffic can get to a network by other means, the Firewall cannot block it

Network Security – Firewalls

- Types of firewalling concepts:
  - (MAC / Ethernet frame filter)
  - Packet filter
  - Circuit-level proxy
  - Stateful packet filter
  - Application-level proxy
- Filtering on data link layer
  - Ethernet packets contain source and destination addresses: MAC
  - Allow only frames to be delivered from known sources, block frames with unknown MACs

Network Security – Firewalls

- Filtering on network layer
  - Source & destination IP addresses
    - Source address
    - Destination address
      - Both are numerical – it is not easy for a Firewall to deal with machine or domain names
      - e.g. www.hotmail.com
    - Request: client = source, server = destination
    - Response: server = source, client = destination

Network Security – Firewalls

- Filtering on transport level
  - This is where we deal with (mostly) TCP and UDP port numbers
    - e.g.: 25 SMTP – sending email (TCP)
    - 110 POP3 – collecting email (TCP)
    - 143 IMAP – collecting email (TCP)
    - 389 LDAP – directory service (TCP)
    - 636 LDAPS – TLS secured directory service (TCP)
    - 80 HTTP – web pages (TCP)
    - 443 HTTPS – secure web pages (TCP)
    - 53 DNS – name lookups (UDP)
    - 68, 69 DHCP – dynamic end system IP config (UDP)

Network Security – Firewalls

- Most Firewalls and their administrators assume that the port number defines the service – not necessarily
  - who could stop me from sending or receiving mail over the HTTP port
  - who could stop users from tunneling all their IP traffic over an open port (demonstration of tunnels in Christmas lecture)
- Here we get major problem: If users are blocked from using a service and try to avoid the blocking firewall they might find a way through – the admin still thinks all is fine with the network, but the situation might be even worse than without firewall at all ...

Network Security – Firewalls

- Layer 7 – Application
  - There is where we find all the 'interesting' stuff ...
    - Web requests
    - Images
    - Executable files
    - Viruses
    - Email addresses
    - Email contents
    - Usernames
    - Passwords

Network Security – Firewalls

- Packet filter – a special router that have the ability to throw packets away independently of network congestion
- Examines TCP/IP headers of every packet going through the Firewall, in either direction
- Choice of whether to allow or block packet based on:
  - (MAC source & destination)
  - IP source & destination addresses (layer 3)
  - TCP / UDP source & destination ports (layer 4)
- Stateful filter
  - Same as a packet filter, except initial packets in one direction are remembered, and replies are automatically allowed fo
  - Simpler rules than simple port based packet filter

Network Security – Filtering of Packets

- Packet filter use rules specify which packets are allowed through the Firewall, and which are dropped
  - Rules must allow for packets in both directions
  - Rules may specify source / destination IP addresses, and source / destination TCP / UDP port numbers
  - Certain (common) protocols are very difficult to support securely (e.g. FTP, IRC, SIP, ...)
  - Low level of security
- Stateful packet filter
  - Packet filter which understands requests and replies (e.g.: for TCP: SYN, SYN-ACK, ACK)

Network Security – Packet Filters

- Stateful packet filter
  - Rules need only specify packets in one direction (from client to server – the direction of the first packet in a connection)
  - Replies and further packets in the communication are automatically processed
  - Supports wider range of protocols than simple packet filter (eg: FTP, IRC, H323)
  - Medium-high level of security
- But how to handle the packets traveling through the network stack?
- Packet Classification Problem
  - Individual entries for classifying a packet are called **rules**

Network Security – Filter Rules

- Rules
  - Each rule is a combination of K values (one for each header field in the packet), a priority and an action Ai.
  - For each entry in a rule different kind of matches are allowed:
    - exact match (e.g. protocol or packet flags/options)
    - prefix match (e.g. blocking subnetwork)
    - range match (e.g. port number ranges)
  - The classifier or rules database consists of a finite set of rules (R_1, ... R_n) ordered by descending priority
  - A packet P matches R_i if all the header fields F_j, (j = 1...K) match the corresponding fields in R_i
  - The Packet Classification Problem is to determine the matching rule with highest priority for each incoming packet

Network Security – Packet Classification

- Packet Classification Problem
  - Todays network links easily reach 1 GBit/s
  - Fiber optic links can operate at over 40 GBit/s
  - A huge amount of current Internet traffic is TCP which transmits ACK packets (40 Bytes each) or VoIP RTP/ UDP packets of ~80 Byte
  - Therefore a worst-case scenario could be a constant stream of ACK packets:
    - E.g. a saturated 10 GBit/s link carries more then 31/15 million packets per second

Network Security – Classification Algorithms

- Packet Classification Problem: Algorithm design
  - Exhaustive Search
    - Test all rules (e.g. linear search)
  - Decision Tree
    - Construct a decision tree from filter rules. Use packet fields to traverse the tree
  - Decomposition
    - Decompose multiple field search in instances of single field searches and perform independent searches. Finally combine results
  - Tuple Space
    - Partition filter set according to the number of specified bits in filter. Probe partitions or subsets with simple exact match searches

Network Security – Classification Algorithms

- Packet Classification Problem: Algorithm design (examples)

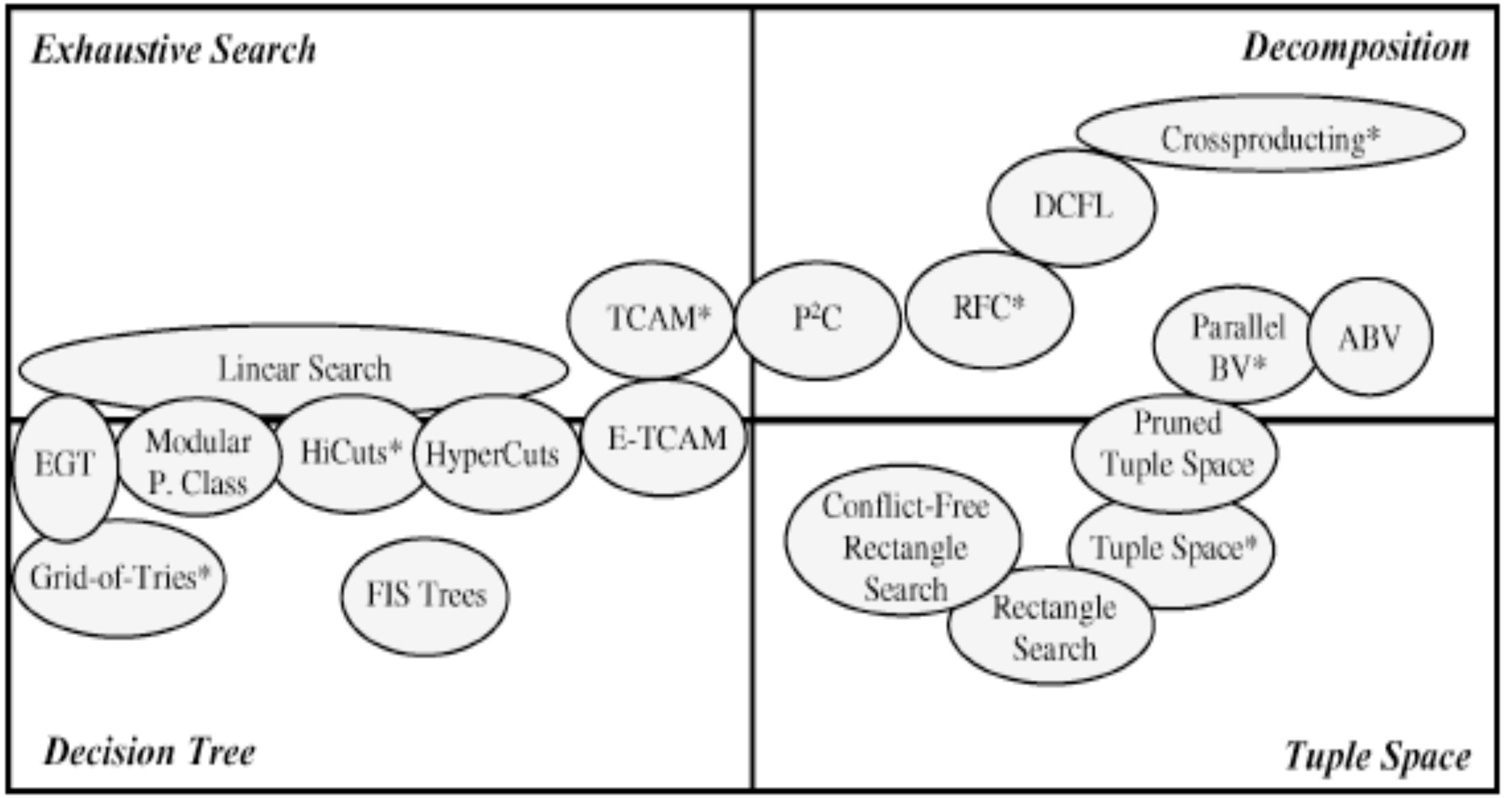

Network Security – Exhaustive Search

- Packet Classification Problem: Exhaustive Search
  - Perform linear search through a list of filter
    - has O(N) storage requirements
    - has O(N) memory accesses per packet
  - where N is the number of filter rules
  - In general a slow solution even for modest-size filter sets
  - But a popular solution in combination with other techniques for the final stage of a lookup, when the set of possible matching filters is reduced to a bounded constant
- Using Decision Trees discussed in next lecture
- This practical course will apply different rule sets for packet classification in Linux Netfilter for firewalling, next practical we will use them for QoS/traffic shaping

Network Security – Decision Trees

- Packet Classification Problem: Decision Trees
  - Idea:
    - Leaves contain filter (or subsets of filters)
    - Construct a search key from the packet fields
    - Traverse the tree by using individual bits or subset of bits from the search key to make branching decisions at each inner node of the tree
    - If we reach a leaf node the best matching filter (or subset of filters) was found
  - Since different search types are possible (e.g. prefix, range match) construction of the tree is difficult

Network Security – Decision Trees

- Packet Classification Problem: Decision Trees
  - Constructing a decision tree
    - Convert all match condition into bit vectors of the following with values 1, 0, * (don't care)
  - Example:
    - Filter table with 3-bit address prefix, arbitrary range of port numbers (3-bit) and an exact 2-bit value:

10*	[0:2]	01
0*	[3:7]	01
111	[0:7]	*
11*	[0:2]	10
*	[0:7]	10

Network Security – Decision Trees

- Packet Classification Problem: Decision Trees
  - Constructing a decision tree
    - Convert the five filters to bit vectors with arbitrary bit masks

10*	[0:2]	01
0*	[3:7]	01
111	[0:7]	*
11*	[0:2]	10
*	[0:7]	10

a	10*	00*	01
b	10*	010	01
c	0**	011	01
d	0**	1**	01
e	111	***	**
f	11*	00*	10
g	11*	010	10
h	11*	010	10

    - Construct the tree by expand the tree path until the node covers only one filter or the bit vector is exhausted

10*	[0:2]	01
0*	[3:7]	01
111	[0:7]	*
11*	[0:2]	10
*	[0:7]	10

convert to arbitrary mask bit vector

a	10*	00*	01
b	10*	010	01
c	0**	011	01
d	0**	1**	01
e	111	***	**
f	11*	00*	10
g	11*	010	10
h	11*	010	10

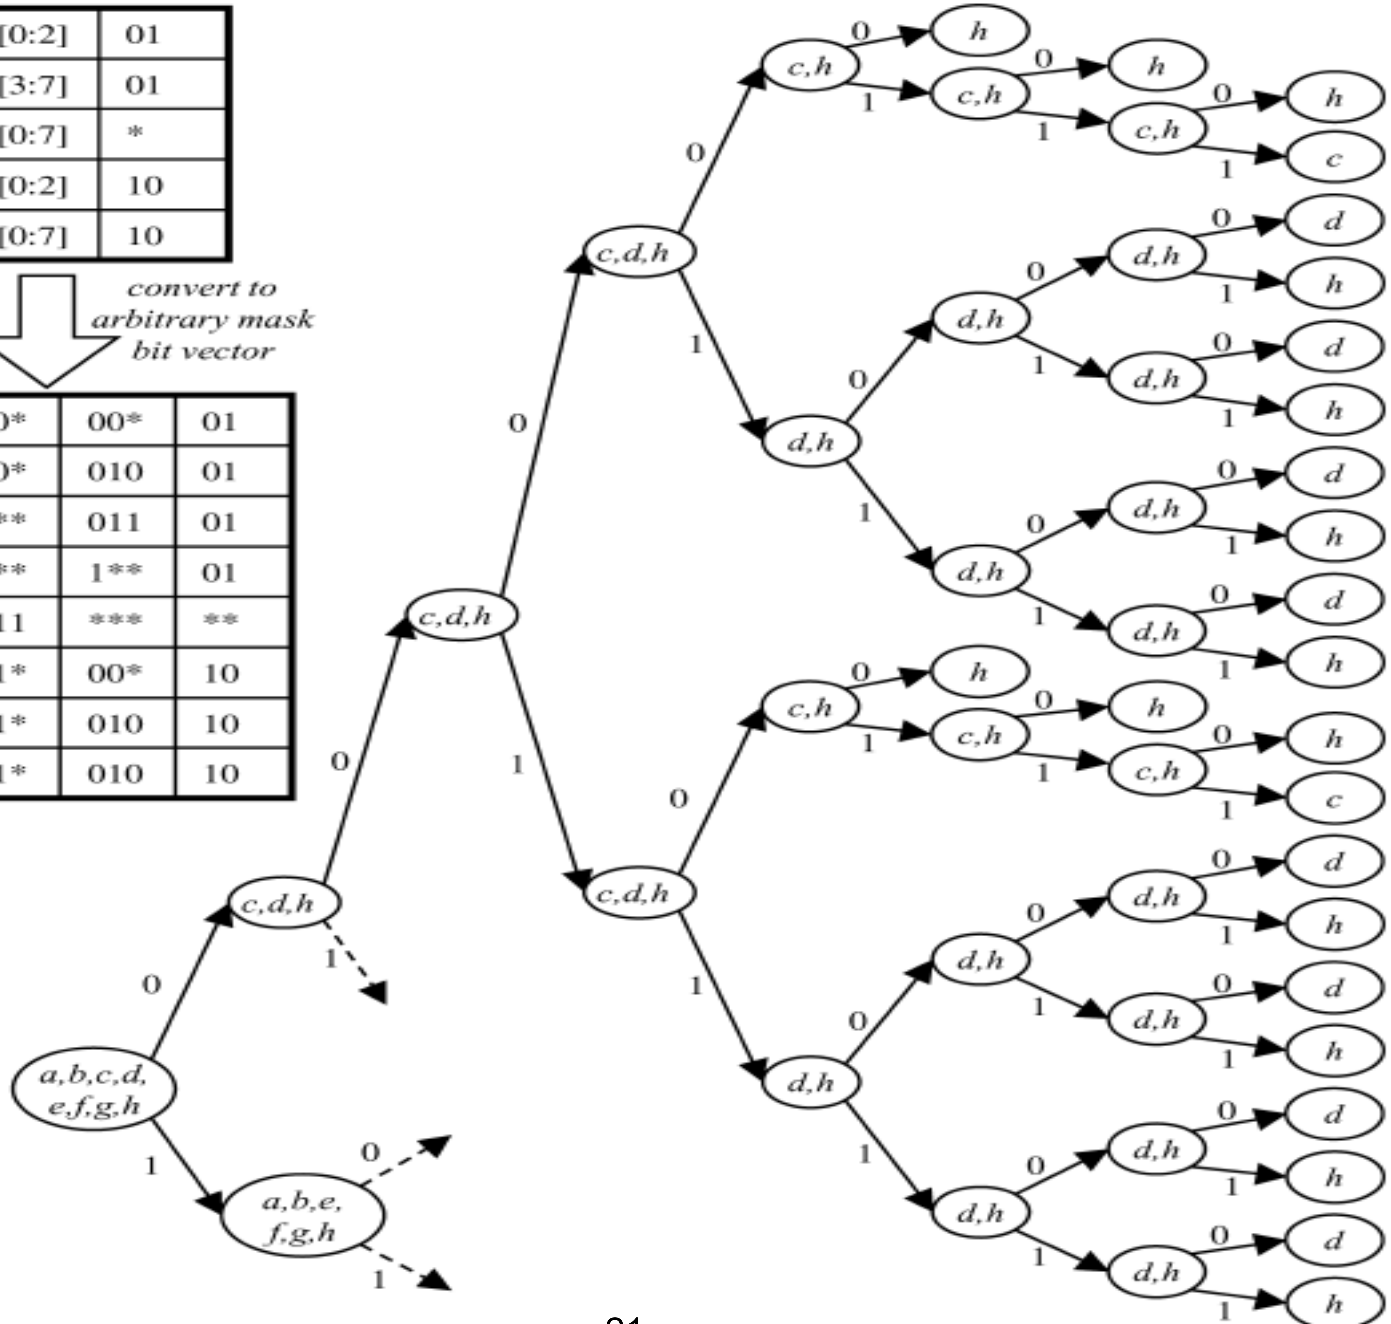

Network Security – Decision Trees

- Packet Classification Problem: Decision Trees
  - Complexity of Decision Trees
    - Lookup time O(W)
    - Memory requirement of the naive approach $O(2^{W+1})$
  - where W is the number of bits used to specify the filter
  - Several improved and optimized decision tree construction algorithms were developed
  - Example: Hierarchical Intelligent Cuttings (Gupta/McKeown 1999)
    - Geometric approach: each filter defines a d-dimensional rectangle in d-space, where d is the number of fields in the filter
    - Specialized heuristics on the filter set are used to minimize tree-depth and memory resource requirements

Network Security – Classification Problem: hi cuts

- Decision Trees Example: Hierarchical Intelligent Cuttings (Gupta/McKeown 1999)
  - Filter set and geometric representation:

Filter	Address	Port
a	1010	2 : 2
b	1100	5 : 5
c	0101	8 : 8
d	*	6 : 6
e	111*	0 : 15
f	001*	9 : 15
g	00*	0 : 4
h	0*	0 : 3
i	0110	0 : 15
j	1*	7 : 15
k	0*	11 : 11

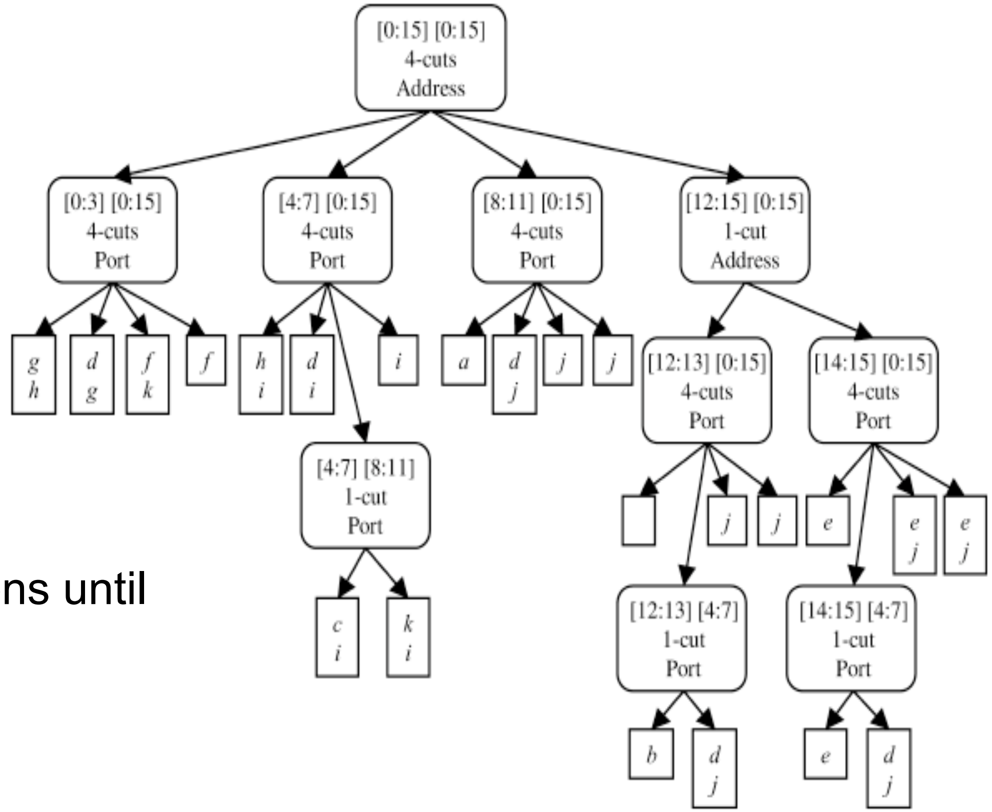

Example:
* Cut every dimension into 4 partitions until leave-node contains at most 2 filter.
Here:
* Start with address dimension.

Network Security – Decision Trees, hi cuts

- Decision Trees as hi cuts
  - Example: Hierarchical Intelligent Cuttings (Gupta/ McKeown 1999)
    - Geometric represent

Filter	Address	Port
a	1010	2 : 2
b	1100	5 : 5
c	0101	8 : 8
d	$*$	6 : 6
e	111$*$	0 : 15
f	001$*$	9 : 15
g	00$*$	0 : 4
h	0$*$	0 : 3
i	0110	0 : 15
j	1$*$	7 : 15
k	0$*$	11 : 11

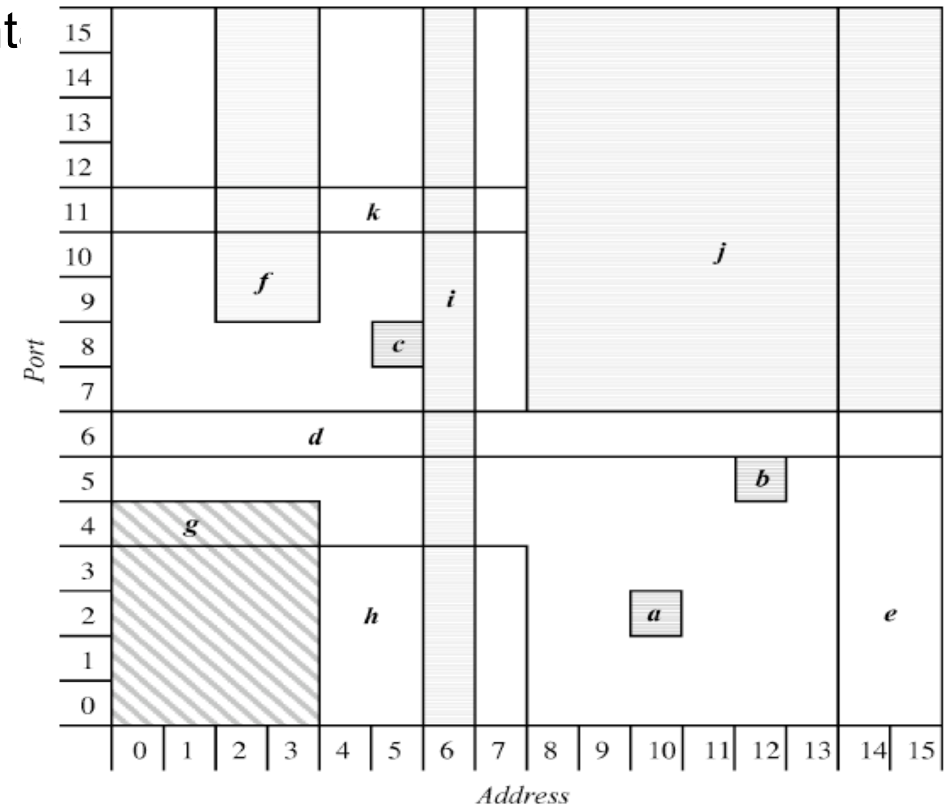

Network Security – Decomposition

- Packet Classification Problem: Decomposition
  - Single field searches can be issued independently. Therefore opportunity to leverage parallelism available on modern hardware
  - Idea:
    - Decompose multiple field search in instances of single field searches and perform independent searches. Finally combine results.
  - Challenge:
    - Efficiently combine results from single field searches

Network Security – Firewalls

- Packet Classification Problem: Decomposition
  - Example: Crossproducting (Srinivasan et. al. 1999)
  - Observation: the number of unique field specifications is significantly less than the number of filters in a filter set
  - Crossproduction starts by constructing d sets of unique match conditions:
    - e.g. all destination address prefixes from all filters.
  - Each such set returns a single best match entry for a given packet field
  - In order to resolve the best matching filter from the set of best matching entries, a table of crossproducts is necessary:
    - For every possible combination of results the best matching filter is precomputed

Network Security – Decomposition

Filter Set

Filter	*Address*	*Port*	*Protocol*
a	000*	[0:1]	TCP
b	001*	[0:1]	TCP
c	1101	[1:1]	UDP
d	10*	[5:15]	UDP
e	001*	[5:15]	UDP
f	111*	[0:15]	UDP
g	000*	[5:15]	UDP
h	10*	[0:1]	TCP
i	001*	[1:1]	TCP
j	*	[0:15]	UDP
k	*	[0:15]	*

Field Sets

Address
000*
001*
1101
10*
111*
*

Port
[0:1]
[1:1]
[5:15]
[0:15]

Protocol
TCP
UDP
*

Table of Crossproducts

Address	*Port*	*Protocol*	*Best Match*
000*	[0:1]	TCP	*a*
000*	[0:1]	UDP	*j*
000*	[0:1]	*	*k*
000*	[1:1]	TCP	*a*
000*	[1:1]	UDP	*j*
000*	[1:1]	*	*k*
...	...	...	...
111*	[5:15]	TCP	*k*
111*	[5:15]	UDP	*f*
111*	[5:15]	*	*k*
...	...	...	...
*	[0:15]	*	*k*

- Packet Classification Problem: Decomposition
  - Example: Crossproducting (Srinivasan et. al. 1999)

Network Security – Decomposition

- Packet Classification Problem: Decomposition
  - Example: Crossproducting (Srinivasan et. al. 1999)
  - Performance:
    - A parallel implementation can provide high throughput
    - But at the cost of exponential memory requirements:
    - For N filters and d fields, the crossproduct table can grow up to $O(N^d)$

Network Security – Tuple Space

- Packet Classification Problem: Tuple Space
  - Similar to the Decomposition approach the observation that the number of unique field specifications is significantly less than the number of filters in a filter set
  - A Tuple defined as the number of specified bits in each filter.
  - Example:
    - For address prefix fields, the number of specified bits is the number of non-wildcard bits of the address
    - For protocols, the value is 1 if e protocol is specified, 0 if not.
    - etc.

Network Security – Tuple Space

- Packet Classification Problem: Tuple Space
  - Example:

Filter	SA	DA	SP	DP	Prot	Tuple
a	0*	001*	2 : 2	0 : 15	TCP	[1, 3, 2, 0, 1]
b	01*	0*	0 : 15	0 : 4	UDP	[2, 1, 0, 1, 1]
c	0110	0011	0 : 4	5 : 15	TCP	[4, 4, 1, 1, 1]
d	1100	*	5 : 15	2 : 2	UDP	[4, 0, 1, 2, 1]
e	1*	110*	2 : 2	0 : 15	UDP	[1, 3, 2, 0, 1]
f	10*	1*	0 : 15	0 : 4	TCP	[2, 1, 0, 1, 1]
g	1001	1100	0 : 4	5 : 15	UDP	[4, 4, 1, 1, 1]
h	0011	*	5 : 15	2 : 2	TCP	[4, 0, 1, 2, 1]
i	0*	110*	2 : 2	0 : 15	UDP	[1, 3, 2, 0, 1]
j	10*	0*	2 : 2	2 : 2	TCP	[2, 1, 2, 2, 1]
k	0110	1100	0 : 15	0 : 15	ICMP	[4, 4, 0, 0, 1]
l	1110	*	2 : 2	0 : 15	*	[4, 0, 2, 0, 0]

The computed Tuples can now be used for a fast, hash-like, exact match. From the Tuple [1, 3, 2, 0, 1] a search key is build by concatenating the first bit of the packet's source address, the first 3 bits of the packet's destination address, etc.

Network Security – Tuple Space

- Packet Classification Problem: Tuple Space
  - Probes to separate tuples can be performed in parallel
  - But in general the tuple space is not predictable, therefore performance can vary widely
  - Tuple Space implementation can be very memory efficient, due to efficient encoding and storage of filter rules. In general memory requirement is O(N)

Network Security – Filtering, Conclusion

- Packet Classification Problem: Conclusion
  - No optimal solution for the general case yet
  - Still an active and vivid research topic.
  - Due to steady hardware improvements and cheap multi-core systems algorithms with parallelized lookups seems very promising
- Other concepts – no direct client connection to the outside service
- Layer-7 proxy server – application level proxy
  - Client and server in one box
  - For every supported application protocol
    - SMTP, POP3, HTTP, SSH, FTP, NNTP, Q3A, ...
  - Packets are received and processed by server
  - New packets generated by client

Network Security – Protocol Proxies

- Prevents the need for direct network connection of clients - advantage
  - no client packet is directly routed into the Internet
  - no packet from Internet is directly handed to the client
- Disadvantage
  - Re-implement every protocol tp be proxied
  - Difficult for proprietary/closed and encrypted protocols
- More general approach: Use a special proxy protocol supported by many applications which offers authentication: socks5

Network Security – Proxies

- Complete server & client implementation in one box for every protocol which can be expected through it
  - Client connects to Firewall
  - Firewall validates request
  - Firewall connects to server
- Response comes back through Firewall and is also processed through client/server
- Large amount of processing per connection
- High level of security but lots of juridical implications – doing stuff with the traffic of other people
- E.g.: lot of funny modifications could be tried with filtering (SPAM, Ads, porno sites, ...) - would you like it if your Computer Center, Provider, Minister of the Interior is deciding for you?

Network Security – Firewall Taxonomy

- Packet filters, classification, circuit-level proxies and stateful packet filters are like telephone call-barring by number
  - block or allow mobile calls
  - block or allow international calls
  - block or allow 0190/0900 calls
  - from different internal extensions
- Application level proxy is like telephone call monitoring by listening to the conversations
  - conversations may still be encoded, or in a foreign language !!

Firewalls - Conclusion

- Firewalls control network traffic to and from the protected network
- Can allow / block access to services (both internal and external)
- Can enforce authentication before allowing access to services
- Can monitor traffic in/out of network
- Can classify and re-route traffic based on these classifications
- Firewalls typically defend a protected network against an attacker, who tries to access vulnerable services which should not be available from outside the network

Firewalls - Conclusion

- Firewalls are also used to restrict internal access to external services, for many different reasons:
  - security (don't want people downloading and installing unknown applications)
  - productivity (don't want people wasting time on non-work related websites etc)
  - cost (many Internet connections, e.g.: Dial-Up are charged by data transferred – ensure this is all necessary)
- But firewalls could mislead to total control and monitoring
  - or distract admins from more important security issues ...

Literature

- Lots of online resources, like www.netfilter.org
- Overview article
  - Taylor, D. E.
    - Survey and taxonomy of packet classification techniques.
    - ACM Computing Surv. 37, 3 (Sep. 2005), 238-275.
  - Singh, Baboescu, Varghese, Wang
    - Packet Classification Using Multidimensional Cutting
    - SIGCOMM 2003, Karlsruhe, Germany

ALBERT-LUDWIGS-UNIVERSITÄT FREIBURG

Communication Systems

University of Freiburg
Computer Science
Computer Networks and Telematics
Prof. Christian Schindelhauer

CoNe Freiburg

IIF INSTITUT FÜR INFORMATIK FREIBURG