cs570

The Internet Protocol (IP)

What problem are we trying to solve?

cs570

Since there are numerous DL technologies and protocols, an internetwork is going to need to pass data between subnetworks with different:

- protocols
- addressing schemes
- speeds
- …

How can we manage these problems efficiently in large internets?

The University Of Alabama in Huntsville Computer Science

What is IP?

cs570

- Most widely applied internetworking protocol
- The L3 protocol of the Internet
- Packet-oriented communication
- Best-effort (“unreliable”)

- Two versions we care about:
  - IPv4 -- the version currently in use (mostly)
  - IPv6 -- the next version

The University Of Alabama in Huntsville Computer Science

cs570

IPv4

IPv4 packet format

<table>
<tr><td rowspan="5">Basic Header (20 bytes)</td><td>IP version Number</td><td>H'der Len (in 4B wds)</td><td>Type of Service</td><td colspan="2">Total length of packet in bytes</td></tr>
<tr><td colspan="3">Message ID</td><td>Flags</td><td>Fragment offset</td></tr>
<tr><td colspan="2">Time to live (TTL)</td><td>Higher-level protocol ID</td><td colspan="2">Header Checksum</td></tr>
<tr><td colspan="5">Source Address</td></tr>
<tr><td colspan="5">Destination Address</td></tr>
<tr><td>Optional</td><td colspan="5">Optional fields</td></tr>
<tr><td>Data</td><td colspan="5">Data (up to 64K-20 bytes)</td></tr>
</table>

Notes on some IPv4 header fields

Header Length:	Measured in 32-bit wds. Minimum is 5.
Type of service:	Options for how IP will treat the packet (will discuss when we get to QoS)
Message ID:	Identifies this packet with a particular message between the source and destination. The combination of Source_address, Dest_address, Message_ID, Protocol, and Fragment_number identify this packet uniquely.
Flags:	Only 2 of 3 bits defined. Used to support fragmentation (later chart).
TTL:	Used to ensure that packets will eventually die if not delivered. Originally intended to measure life in seconds; is processed as a hop count (every router decrements TTL until it reaches 0).
Protocol:	Identifies the Transport-level protocol (usually TCP or UDP).
Options:	Used by the sender to request network services (padded to be a multiple of 32 bits)
Data:	The total packet length including header and options can be 64KB.

IP in the protocol stack

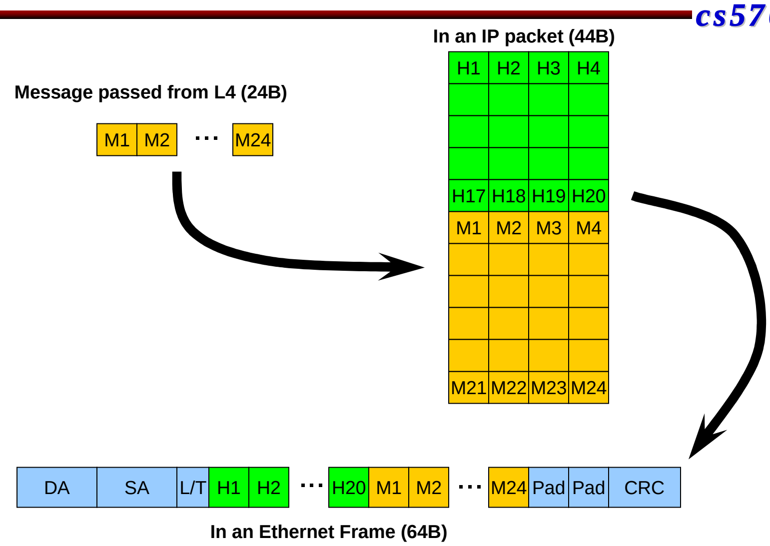

IPv4 addresses

General format:

32 bits

Class ID	Network number	Host number

How they are usually written and talked about:

Dotted decimal notation: Express each byte as its equivalent in decimal.

Example:

11000000	00101001	00000110	00010100
192 .	41 .	6 .	20

IPv4 address formats ("classful" addressing)

		Network number	Host number
A	0	Network number (7 bits)	Host number (24 bits)
B	1 0	Network number (14 bits)	Host number (16 bits)
C	1 1 0	Network number (21 bits)	Host number (8 bits)
D	1 1 1 0	Multicast Address (28 bits)	

Note: Class E("11110") is reserved for future use.

IPv4 addresses

Class	Format (when reading in dotted decimal)	Range of Unreserved Addresses	Approximate number of networks/hosts
A	N.H.H.H	1.0.0.0 to 126.255.255.255	126 / 16M
B	N.N.H.H	128.0.0.0 to 191.255.255.255	16K / 64K
C	N.N.N.H	192.0.0.0 to 223.255.255.255	2M / 256

Some special reserved addresses:

All zeroes:	This host
Network=0 w/ host #:	The indicated host on this network
All ones:	Broadcast on this network.
Network # w/ host=all ones:	Broadcast on the indicated network
Network=127	Loopback

Hierarchical addressing in IP networks

cs570

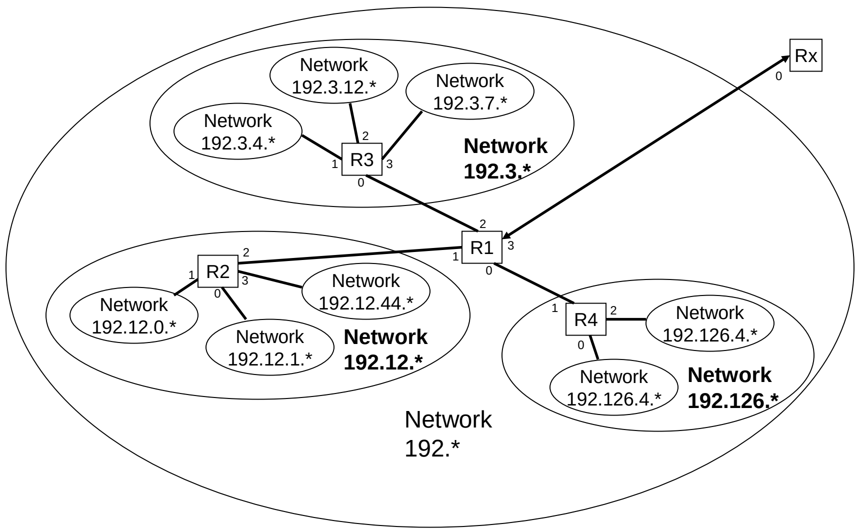

How hierarchical addressing streamlines routing (1)

cs570

Traffic for any address starting with 192 will be sent to R1.

This means that routing tables in other parts of the network can handle the entire 192.* network with one entry.

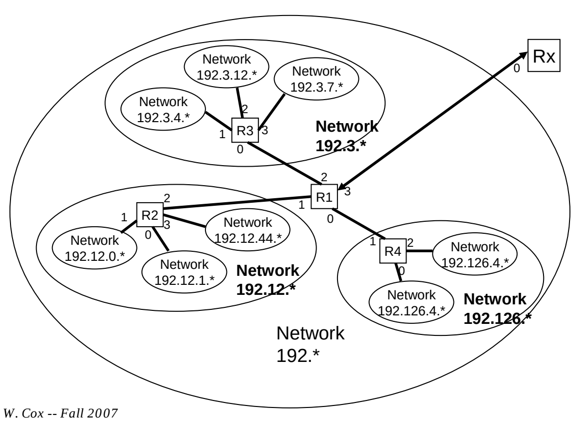

If this was not the case, every router in the Internet would need a routing table with millions of entries.

The University Of Alabama in Huntsville Computer Science

How hierarchical addressing streamlines routing (2)

cs570

R1's routing table

192.3.*	2
192.12.*	1
192.126.*	0
default	3

R2's routing table

192.12.0.*	1
192.12.1.*	0
192.12.44.*	3
default	2

Rx's routing table

192.*	0
•	
•	
•	

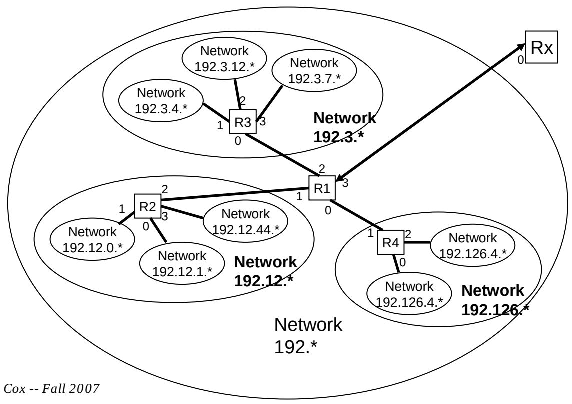

The default route is used when no other route applies, meaning the packet is addressed to some other network or subnet.

The University Of Alabama in Huntsville Computer Science

cs570

Mapping IP addresses to L2 devices

IP addressing over MAC addresses

- IP addresses are "virtual" addresses assigned to a device. They do not relate to the device's "real" address (its MAC address).
- When an IP packet arrives at its destination subnetwork, it needs to be delivered to the connected host having the specified IP address. But in most multidrop subnetworks (e.g, Ethernet), we need to know the MAC address -- the IP address does no good.
- This means that the subnetwork needs a system for translating IP addresses into MAC addresses.

The Address Resolution Protocol (ARP)

- Each host on the multidrop subnetwork maintains a table of the IP address and MAC address of each node on the subnetwork
- When a host receives an IP packet, it:
  - Checks to see if the MAC address is in the table. If it is, it forwards the IP packet to the indicated MAC address.
  - If not, it broadcasts an ARP query on the subnetwork asking which node has the IP address in question
  - The node with the IP address broadcasts its MAC address, allowing all nodes to update their tables.

The University Of Alabama in Huntsville Computer Science

cs570

IP Fragmentation

G. W. Cox -- Fall 2007

IP part A - 17

The University Of Alabama in Huntsville Computer Science

IP Packet Fragmentation

cs570

- Assume we send an IP packet through a subnetwork in which the frame payload size is smaller than the packet size
- We could design to do either:
  (1) L2 Fragmentation: Divide the IP packet among frames when it enters the subnetwork, then recombine them when it leaves the subnetwork
  Problems:
    - May introduce high delay by repeatedly fragmenting and re-assembling the same packet in different subnetworks
    - Have to wait for all frames at the exit of each subnetwork

  (2) L3 Fragmentation: When entering the subnetwork, divide the packet into smaller IP-formatted packets. Re-assembly is doen at the receiver.

G. W. Cox -- Fall 2007

IP part A - 18

Header fields supporting fragmentation

- Source Address
- Destination Address
- Message ID

Uniquely identify the message that the fragment is part of

- Flags:
  - DF (if set, Don't Fragment this packet)
  - MF (if set, More Fragments follow this fragment)

- Fragment Offset -- The offset (in bytes) of the data in this fragment packet referenced to the start of the data in the original packet

The University Of Alabama in Huntsville Computer Science

Fragmentation Example

cs570

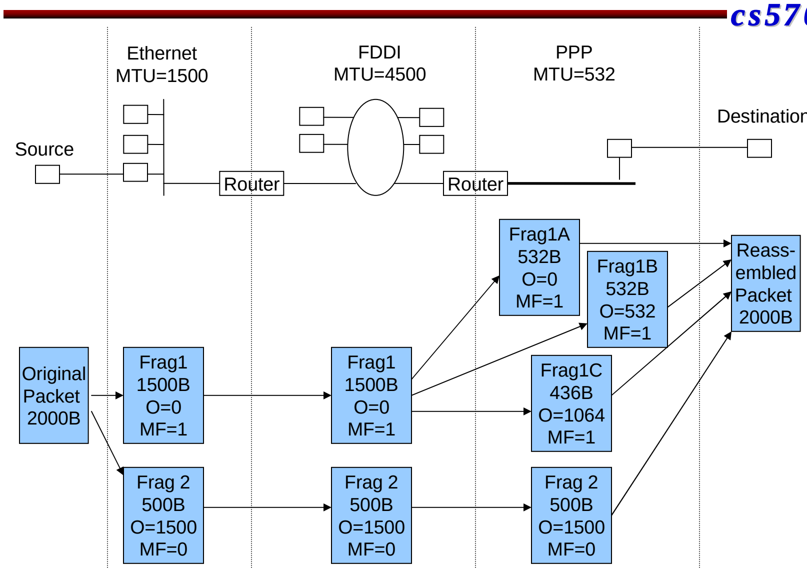

The University Of Alabama in Huntsville Computer Science

cs570

IP Tunneling

G. W. Cox -- Fall 2007

IP part A - 21

The University Of Alabama in Huntsville Computer Science

IP tunneling

cs570

Sometimes, we want to set up a virtual point-to-point link across an IP internet

- Make a virtual “Direct Connection”
- Redirect traffic to other addresses
- Use non-IP protocols
- Security

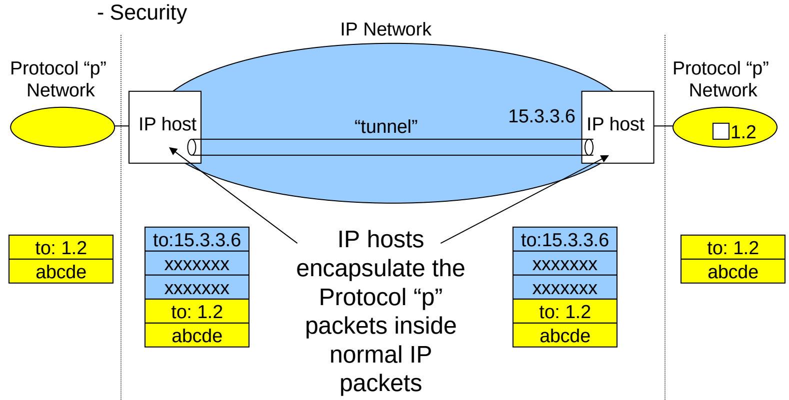

G. W. Cox -- Fall 2007

IP part A - 22