Lecture Computer Networks

Internet Protocol Version 4 (IPv4)
Address Resolution Protocol (ARP)
Internet Control Message Protocol (ICMP)

Prof. Dr. H. P. Großmann mit M. Rabel sowie
H. Hutschenreiter und T. Nau | Sommersemester 2012 |
Institut für Organisation und Management von
Informationssystemen

Thomas Nau, kiz

Content

- Specifications
- Internet Protocol Version 4 (IPv4)
  - Motivation
  - Overview
  - Tasks
    - Fragmentation
    - Addressing
    - Routing
- Address Resolution Protocol (ARP)
- Internet Control Message Protocol (ICMP)

Specifications

The formal specification of the Internet protocols is standardized in the “Request for Comments“ (RFCs).

IP	RFC 791
ICMP	RFC 792
ARP	RFC 826
TCP	RFC 793
UDP	RFC 768

OSI and IP Protocol Stacks

OSI Model

7. Application
6. Presentation
5. Session
4. Transport
3. Network
2. Data Link
1. Physical

Internet Protocol Suite

Application (e.g. FTP, TELNET)
TCP or UDP
IP
LAN or WAN Technology (e.g. Ethernet)

Motivation

We have

LANs with their local limitations (e.g. maximum segment length) and own addressing (e.g. Ethernet - 48 bit MAC Addresses)

We want

a connection between the LANs and hosts distributed all over the world

Connected LANs – Motivation (2)

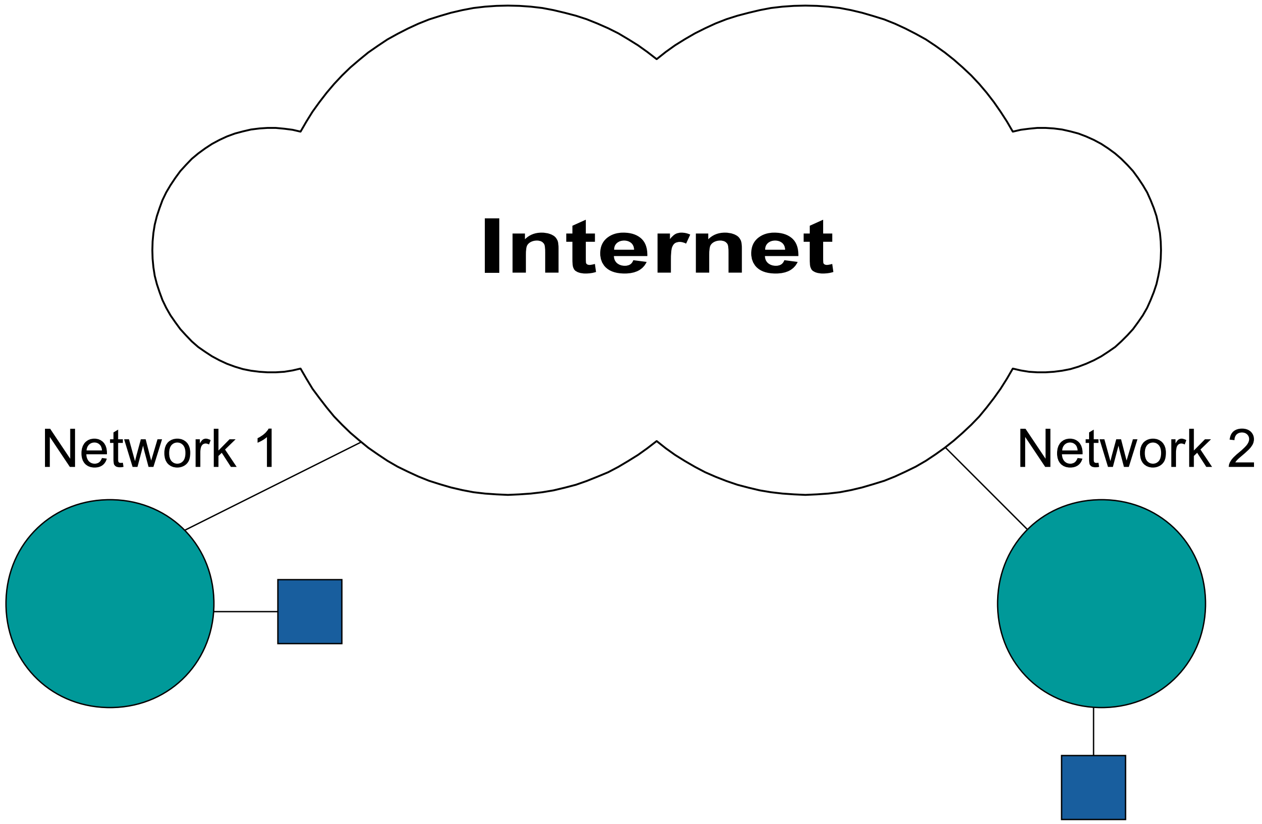

Motivation (3)

What do we need?

- a protocol that supports as many kinds of LANs as possible
- uniform addresses
- routing

Overview

IP (*Internet Protocol*)
is a *connectionless* and *unreliable* packet delivery system

Connectionless Communication Service

(also called datagram transmission)

- the data are divided into segments
- each segment is supplied with a header
  that contains destination and source address
- each datagram is transferred independently through the network

Unreliable Service

- no guarantee that a packet is delivered correctly or doesn’t get lost
- reliability must be provided by the upper layers (e.g. TCP)

Tasks of IP

- Fragmentation
- Addressing and Routing

Fragmentation

In networks which are IP-connected, different frame sizes can be used (e.g. Ethernet and FDDI)

⇨ when the size of a IP datagram exceeds the size of the underlying MTU, fragmentation occurs

IP Addresses

- To every interface there can be assigned an IP address, which is a unique 32 bit address

- IP addresses are usually written as four decimal numbers, separated by points:

  W . X . Y . Z

  W: 1. Byte,
  X: 2. Byte,
  Y: 3. Byte,
  Z: 4. Byte

  e.g. 134.60.40.100

- the IP addresses are assigned by a central authority – the Network Information Center (NIC)

Addressing (at the beginning)

5 Classes of IP Addresses

- class A – C scheme: 2 parts
  first part for the network ID
  second part for the host ID

- class D Multicast Address (assigning a group of hosts)

- class E Addresses: reserved for future use

IP Address Classes

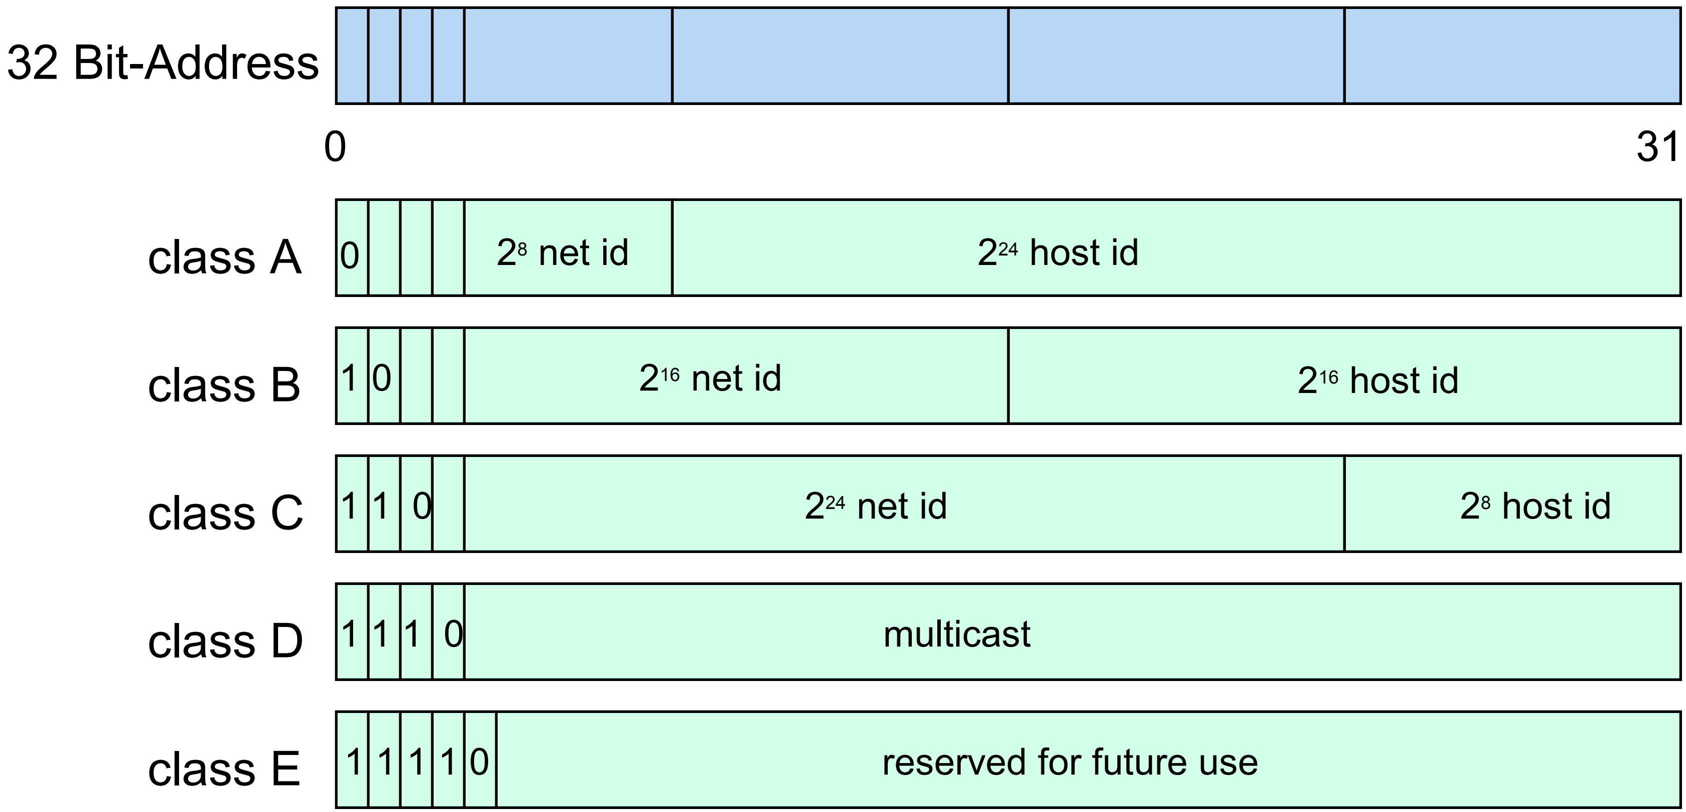

Scaling Problems

- Exhaustion of the class-B network address space.

  (One fundamental cause of this problem is the lack of a network class of a size which is appropriate for mid-sized organization; class-C, with a maximum of 254 host addresses, is too small while class-B, which allows up to 65534 addresses, is too large to be widely allocated.)

- Growth of routing tables in Internet routers is beyond the ability to be managed effectively by current software (and people).

- Eventual exhaustion of the 32-bit IP address space.

Addressing example

Four "class c" addresses for one network:

	decimal	binary
Sub 1	192.168.0.0	11000000 . 10101000 . 00000000 . 00000000
Sub 2	192.168.1.0	11000000 . 10101000 . 00000001 . 00000000
Sub 3	192.168.2.0	11000000 . 10101000 . 00000010 . 00000000
Sub 4	192.168.3.0	11000000 . 10101000 . 00000011 . 00000000

⇨ four networks with the netmask

Mask	255.255.255.0	11111111 . 11111111 . 11111111 . 00000000

Supernetting

Organizations with more than one class C network can merge these networks by moving bits from the network portion of the address into the host portion of the address.

Example:

	decimal	binary
Sub 1	192.168.0.0	11000000 . 10101000 . 00000000 . 00000000
Sub 2	192.168.1.0	11000000 . 10101000 . 00000001 . 00000000
Sub 3	192.168.2.0	11000000 . 10101000 . 00000010 . 00000000
Sub 4	192.168.3.0	11000000 . 10101000 . 00000011 . 00000000

⇨ one network with the netmask

Mask	255.255. 252.0	11111111 . 11111111 . 11111100 . 00000000

Supernetting Chart

equal bits	Supernet Mask	Number of “Class C“s	Number of Hosts
14	255.252.0.0	1024	262144
15	255.254.0.0	512	131072
16	255.255.0.0	256	65536
17	255.255.128.0	128	32768
18	255.255.192.0	64	16384
19	255.255.224.0	32	8192
20	255.255.240.0	16	4096
21	255.255.248.0	8	2048
22	255.255.252.0	4	1024
23	255.255.254.0	2	512

Subnetting

Any organization with a network of any size can subdivide the available host address space according to its network topology

Example: University of Ulm

class B address: 134.60.X.Y

X is used for local subnets

Y is used for the hosts

RFC 950, “Internet Standard Subnetting Procedure”

Classless Inter-Domain Routing (CIDR)

- Bitwise Variable-Length Subnetting
  - a.b.c.d/n   n: from 0 to 32

- RFC1519, an address assignment and aggregation strategy
  - For example: 192.24.0.0/18
    - Mask   255.255.192.0
    - Networks   192.24.0.0 – 192.24.63.0
    - Hosts   16384

- Routing prefix aggregation
  - Two or more contiguous CIDR classes can be aggregated and advertised together

Private Address Space

There are 3 blocks of the IP address space reserved for private Internets:

10.0.0.0 - 10.255.255.255	(class A)
172.16.0.0 - 172.31.255.255	(class B)
192.168.0.0 - 192.168.255.255	(class C)

RFC 1918, „Address Allocation for Private Internets“

Single and Multihomed Host

- a host which is connected to one LAN is called single homed host
- a host which is connected to more than one LAN is called a multihomed host

  the host needs an IP address for each network

Example

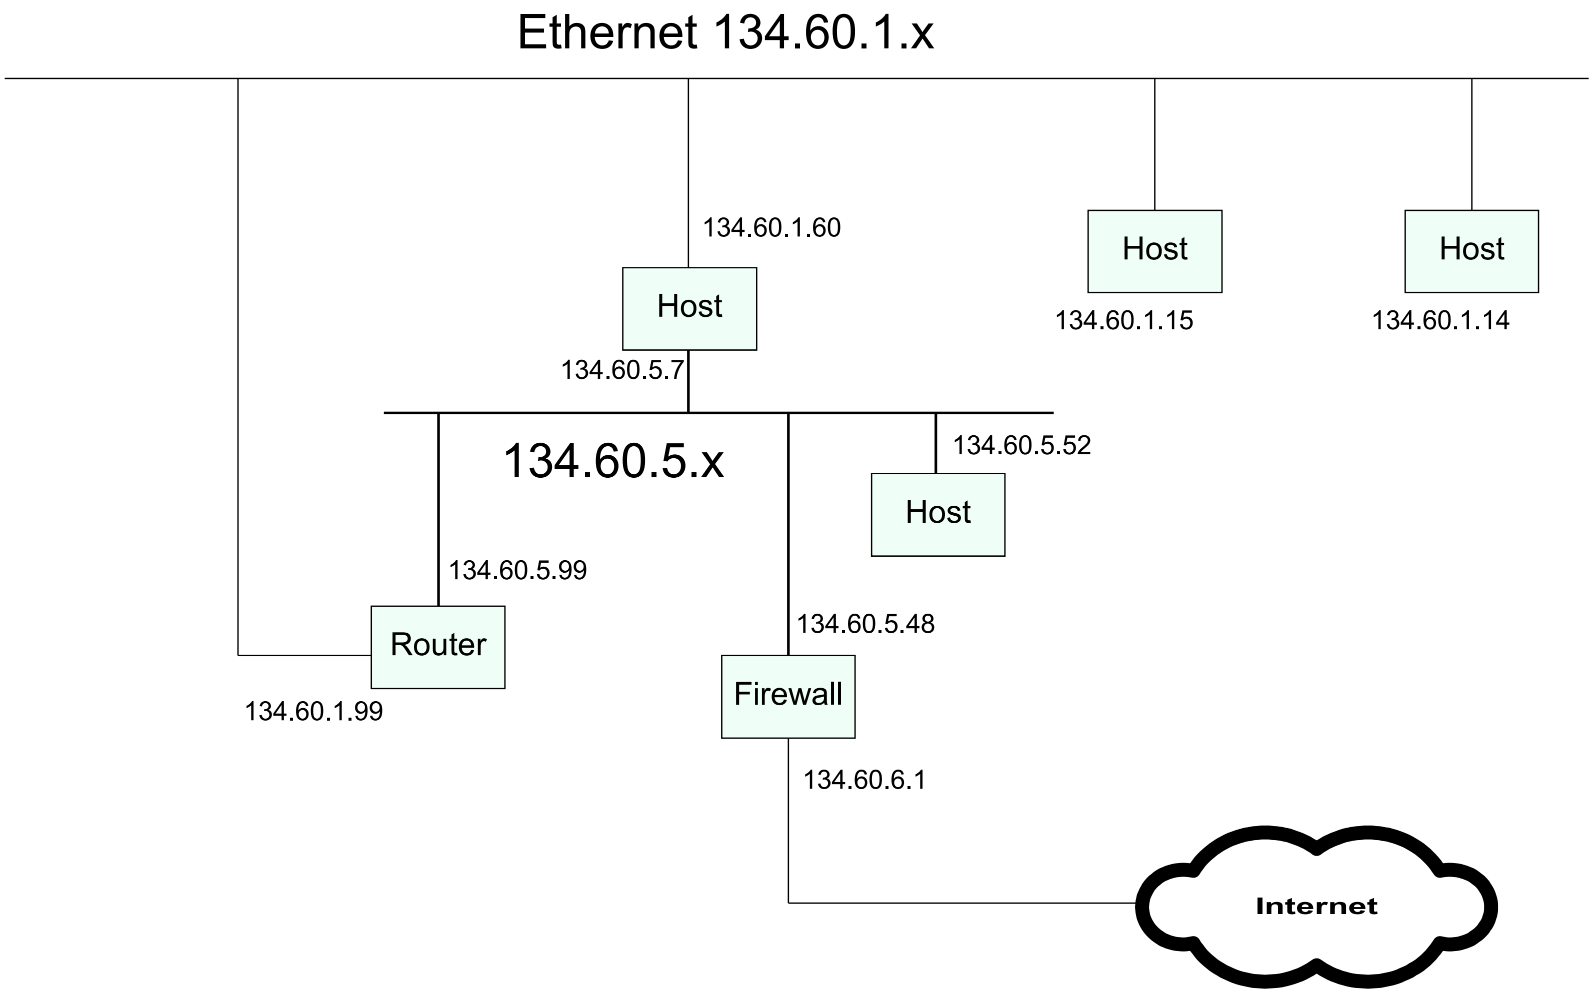
Ethernet 134.60.1.x
134.60.1.60
Host
Host
134.60.1.15
Host
134.60.1.14
134.60.5.7
134.60.5.x
134.60.5.52
Host
134.60.5.99
134.60.5.48
Router
134.60.1.99
Firewall
134.60.6.1
Internet

IP Header

<table>
<tr><th colspan="2">1 Byte</th><th colspan="2">1 Byte</th><th colspan="2">1 Byte</th><th colspan="2">1 Byte</th></tr>
<tr><td>Version</td><td>IHL</td><td colspan="2">Type of Service</td><td colspan="4">Total Length</td></tr>
<tr><td colspan="4">Identification</td><td>Flags</td><td colspan="3">Fragment Offset</td></tr>
<tr><td colspan="2">TTL
(Time to Live)</td><td colspan="2">Protocol</td><td colspan="4">Header Checksum</td></tr>
<tr><td colspan="8">Source Address</td></tr>
<tr><td colspan="8">Destination Address</td></tr>
<tr><td colspan="8">Options</td></tr>
<tr><td colspan="8">Data</td></tr>
</table>

IP Header Explanation (1)

- Version (4 bits): IPv4 (at the moment) or IPv6

- IHL (4 bits): Internet Header Length (minimum 20 Bytes)

- Service Type (8 bits): Not widely supported.
  It specifies the wants of the service datagram.
  It contains
    - 3-bit precedence from 0 (normal) to 7 (network control)
    - 1-bit D: for delay preferred
    - 1-bit T: for throughput preferred
    - 1-bit R: for reliability preferred
    - 2 remaining bits are unused

Type of Service (8 Bits)

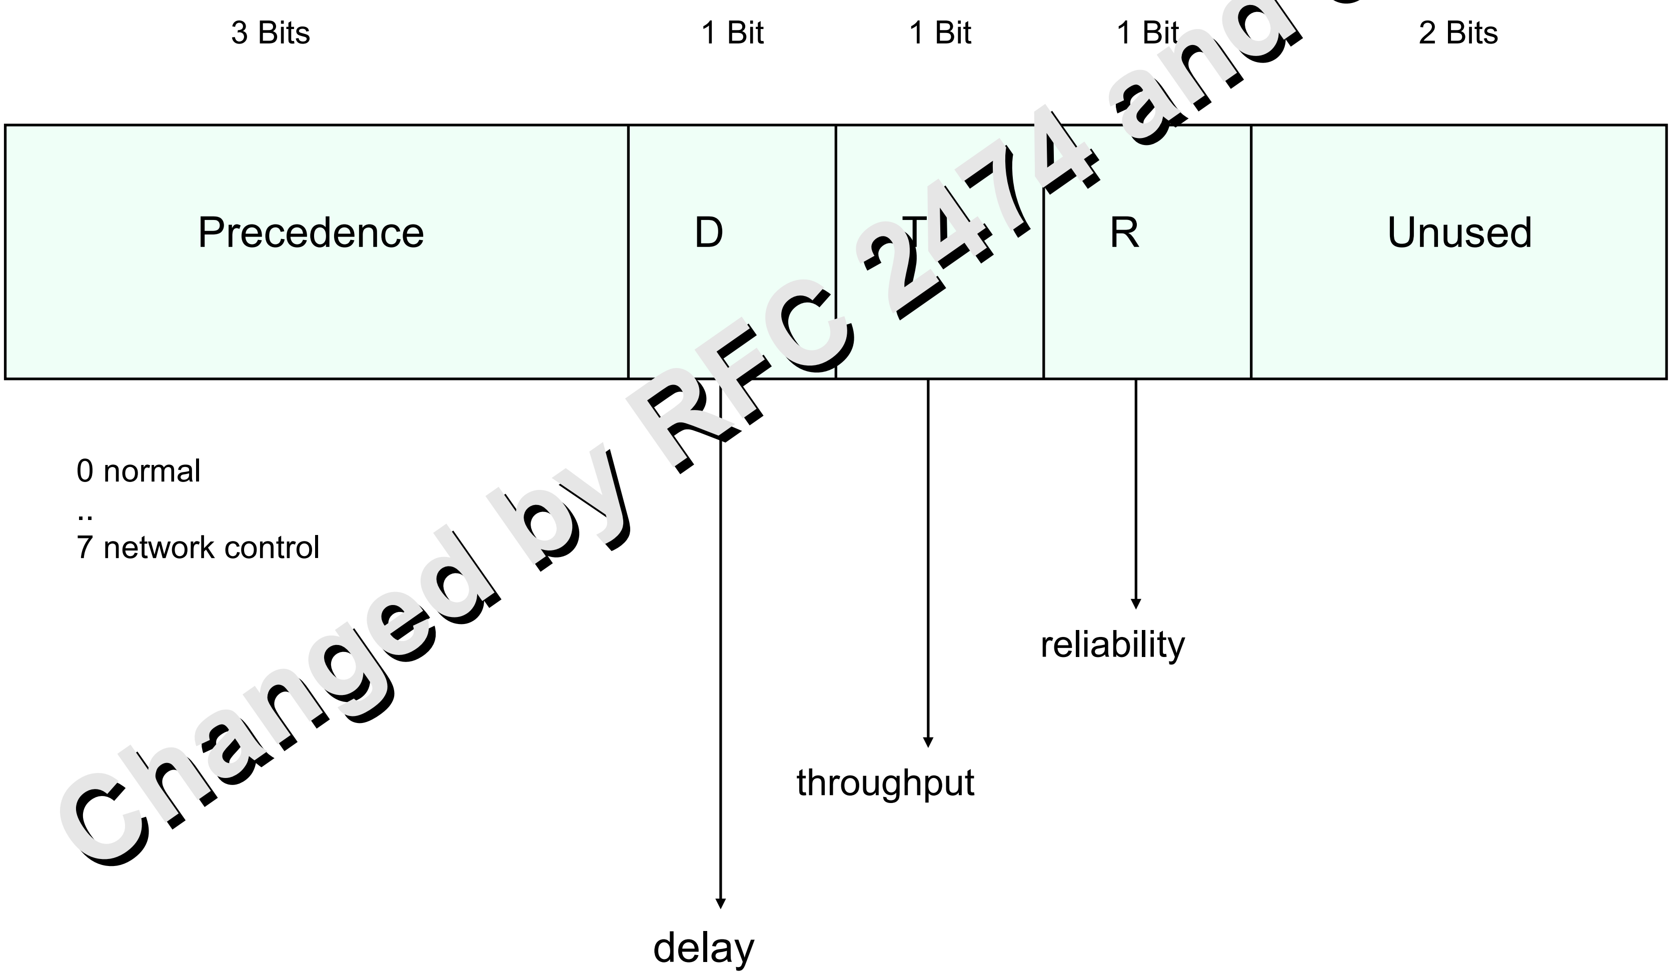

IP Header Explanation (2)

- Total Length (16 bits):
  Length of header and data
  Each host must be able to handle a length of 576 Bytes, today most hosts accept longer packets, too

- Datagram Identification (16 bits):
  Used to allow the destination to re-assemble fragments.
  All fragments of a datagram contain the same identification.

- Fragmentation Flags: MF More Fragmentation
  DF Don’t fragment

- Fragment Offset: tells where this fragment belongs in the current datagram

IP Header Explanation (3)

- Time to Live (TTL) (8bits):
  Maximum number of routers to pass.
  Each router decrements that number; when it hits zero, the datagram is discarded.
  It was intended to record seconds, but now it is used to count hops.

- Protocol: layer 4 protocol (e.g. TCP, UDP)

- Header Checksum:
  If an error is found in the checksum, the datagram is discarded.
  A higher layer protocol has to care for retransmission.
  It has to be calculated by every router, because of the change of the TTL.

Routing

Each host and gateway in the Internet has a routing table

Example for a routing table of a host:

Destination	Gateway	Netmask	Iface
134.90.1.1	134.60.50.1	255.255.255.255	ppp0
134.60.40.0	0.0.0.0	255.255.255.0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	lo
0.0.0.0	134.60.40.99	0.0.0.0	eth0

The entry, for which the following statement results in the longest match of true values, is chosen:
IP address of the datagram to send (binary AND) Netmask = Destination

ARP (Address Resolution Protocol)

Interface Cards (e.g. for Ethernet) only recognize MAC addresses

⇨ they accept only Broadcasts and Frames with their own MAC address

A host or gateway that wants to send an IP datagram to a host (with a known IP address) must find out which MAC address corresponds to that host.

⇨ ARP (Address Resolution Protocol) is used.

ARP (2)

The host or gateway sends a broadcast message on the Ethernet that asks the host with the specified IP address to respond with its Ethernet address. Every host on the Ethernet receives this broadcast packet, but only the specified one will respond.
For the future the result is cached in an ARP table.

ARP (3)

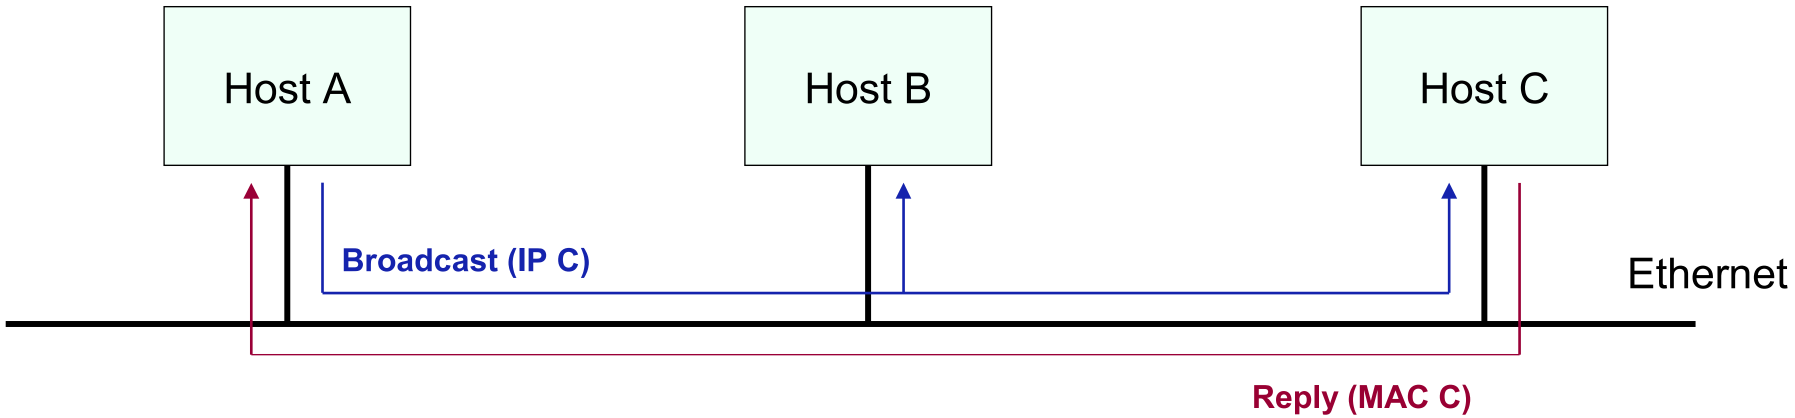

ICMP (Internet Control Message Protocol)

- ICMP messages, delivered in IP packets, are used for out-of-band messages related to network operation or mis-operation.
- Since ICMP uses IP, ICMP packet delivery is unreliable, so hosts cannot count on receiving ICMP packets for any network problem.

Some ICMP Functions (1)

- **Announce network errors**,
  such as a host or entire portion of the network being unreachable, due to some type of failure.

- **Announce network congestion**.
  When a router begins buffering too many packets, due to an inability to transmit them as fast as they are being received, it will generate ICMP Source Quench messages. Directed at the sender, these messages should cause the rate of packet transmission to be slowed. Of course, generating too many Source Quench messages would cause even more network congestion, so they are used sparingly.

Some ICMP Functions (2)

- **Assist Troubleshooting**.

  ICMP supports an Echo function, which just sends a packet on a round-trip between two hosts.

  Ping, a common network management tool, is based on this feature. Ping will transmit a series of packets, measuring average round-trip times and computing loss percentages.

- **Announce Timeouts**.

  If an IP packet's TTL field drops to zero, the router discarding the packet will often generate an ICMP packet announcing this fact.