IP – The Internet Protocol

Fields of the IP Header

- **Time To Live (TTL) (1 byte):**
  - Specifies longest paths before datagram is dropped
  - Role of TTL field: Ensure that packet is eventually dropped when a routing loop occurs

  Used as follows:
  - Sender sets the value (e.g., 64)
  - Each router decrements the value by 1
  - When the value reaches 0, the datagram is dropped

Fields of the IP Header

- **Protocol (1 byte):**
  - Specifies the higher-layer protocol.
  - Used for demultiplexing to higher layers.

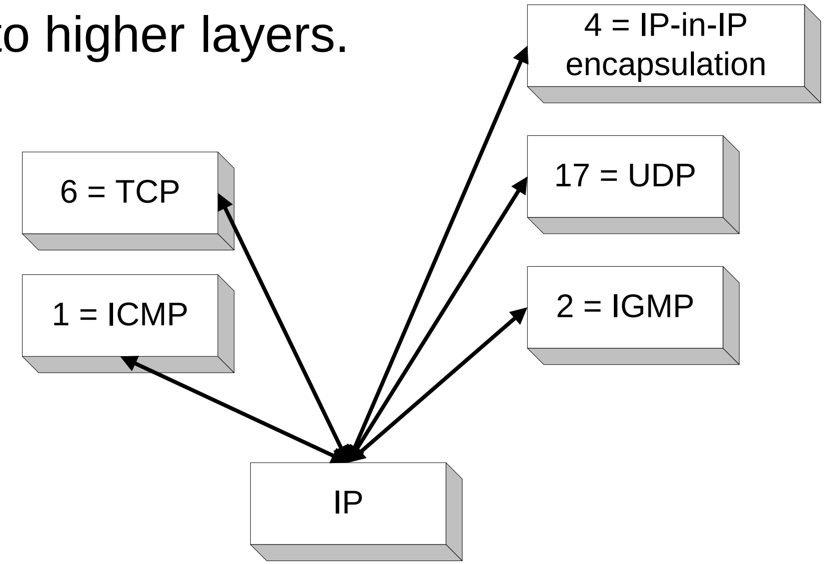

- **Header checksum (2 bytes):** A simple 16-bit long checksum which is computed for the header of the datagram.

Fields of the IP Header

- **Options:**
    - Security restrictions
    - Record Route: each router that processes the packet adds its IP address to the header.
    - Timestamp: each router that processes the packet adds its IP address and time to the header.
    - (loose) Source Routing: specifies a list of routers that must be traversed.
    - (strict) Source Routing: specifies a list of the only routers that can be traversed.
- **Padding:** Padding bytes are added to ensure that header ends on a 4-byte boundary

Maximum Transmission Unit

- Maximum size of IP datagram is 65535, but the data link layer protocol generally imposes a limit that is much smaller

- Example:
  - Ethernet frames have a maximum payload of 1500 bytes

    → IP datagrams encapsulated in Ethernet frame cannot be longer than 1500 bytes

- The limit on the maximum IP datagram size, imposed by the data link protocol is called **maximum transmission unit (MTU)**

- MTUs for various data link protocols:

Ethernet:	1500	FDDI:	4352
802.3:	1492	ATM AAL5:	9180
802.5:	4464	PPP:	negotiated

IP Fragmentation

- What if the size of an IP datagram exceeds the MTU?
  IP datagram is fragmented into smaller units.

- What if the route contains networks with different MTUs?

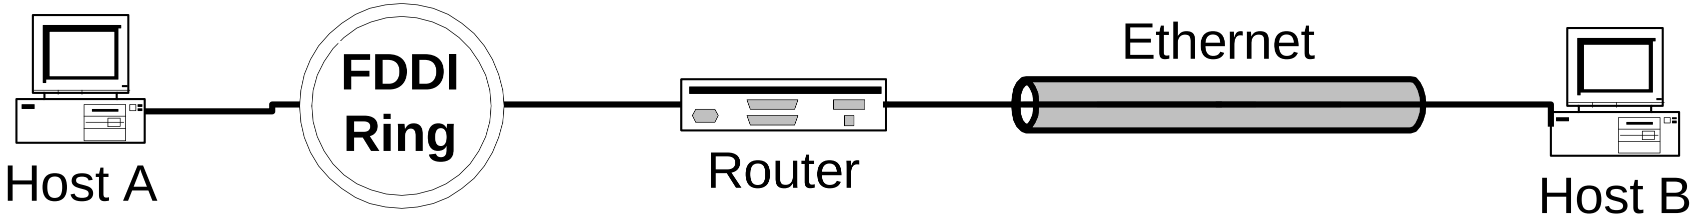

MTUs: FDDI: 4352 Ethernet: 1500

- **Fragmentation**:
  - IP router splits the datagram into several datagram
  - Fragments are reassembled at receiver

Where is Fragmentation done?

- Fragmentation can be done at the sender or at intermediate routers
- The same datagram can be fragmented several times.
- Reassembly of original datagram is only done at destination hosts !!

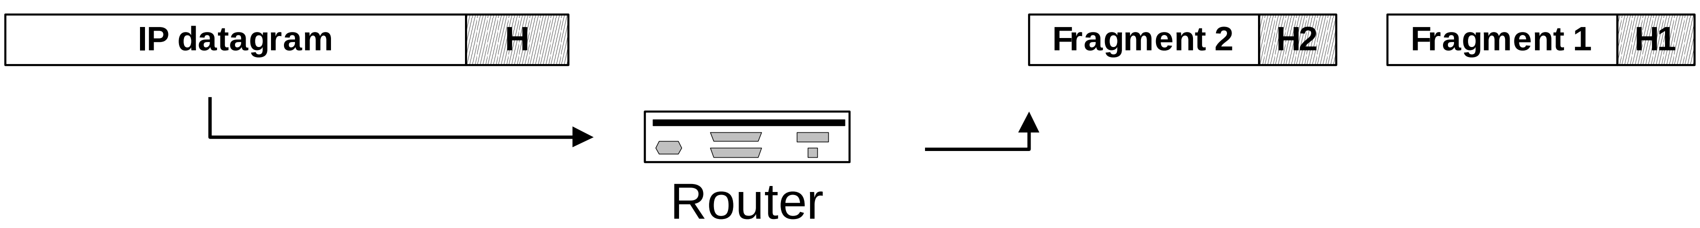

What's involved in Fragmentation?

- The following fields in the IP header are involved:

<table>
<tr><td>version</td><td>header length</td><td>DS</td><td>ECN</td><td colspan="4">total length (in bytes)</td></tr>
<tr><td colspan="4">Identification</td><td>0</td><td>DF</td><td>MF</td><td>Fragment offset</td></tr>
<tr><td colspan="2">time-to-live (TTL)</td><td colspan="2">protocol</td><td colspan="4">header checksum</td></tr>
</table>

Identification	When a datagram is fragmented, the identification is the same in all fragments
Flags	
DF bit is set:	Datagram cannot be fragmented and must be discarded if MTU is too small
MF bit set:	This datagram is part of a fragment and an additional fragment follows this one

What’s involved in Fragmentation?

- The following fields in the IP header are involved:

<table>
<tr><td>version</td><td>header length</td><td>DS</td><td>ECN</td><td colspan="4">total length (in bytes)</td></tr>
<tr><td colspan="4">Identification</td><td>0</td><td>DF</td><td>MF</td><td>Fragment offset</td></tr>
<tr><td colspan="2">time-to-live (TTL)</td><td colspan="2">protocol</td><td colspan="4">header checksum</td></tr>
</table>

Fragment offset	Offset of the payload of the current fragment in the original datagram
Total length	Total length of the current fragment

Example of Fragmentation

- A datagram with size 2400 bytes must be fragmented according to an MTU limit of 1000 bytes

IP datagram	Fragment 3	Fragment 2	Fragment 1
Header length: 20	Header length: 20	Header length: 20	Header length: 20
Total length: 2400	Total length: 448	Total length: 996	Total length: 996
Identification: 0xa428	Identification: 0xa428	Identification: 0xa428	Identification: 0xa428
DF flag: 0	DF flag: 0	DF flag: 0	DF flag: 0
MF flag: 0	MF flag: 0	MF flag: 1	MF flag: 1
Fragment offset: 0	Fragment offset: 244	Fragment offset: 122	fragment offset: 0

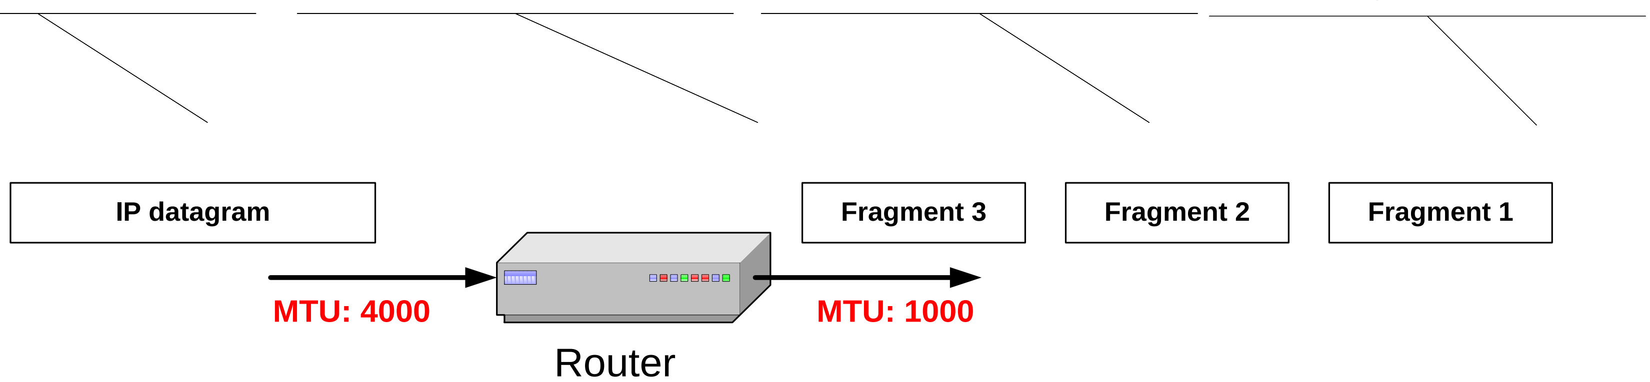

Determining the length of fragments

- To determine the size of the fragments we recall that, since there are only 13 bits available for the fragment offset, the offset is given as a multiple of eight bytes. As a result, the first and second fragment have a size of 996 bytes (and not 1000 bytes). This number is chosen since 976 is the largest number smaller than 1000–20= 980 that is divisible by eight. The payload for the first and second fragments is 976 bytes long, with bytes 0 through 975 of the original IP payload in the first fragment, and bytes 976 through 1951 in the second fragment. The payload of the third fragment has the remaining 428 bytes, from byte 1952 through 2379. With these considerations, we can determine the values of the fragment offset, which are 0, 976 / 8 = 122, and 1952 / 8 = 244, respectively, for the first, second and third

Internet Addressing

- Today's topics:
  - What are Internet addresses?
  - Why do we need them?
  - How are they used?

Universal Communication Service

- A *universal communication service* allows any host to communicate with any other host
- This requires a globally accepted method of identifying each host connected to the internet
- Internetworking strives to provide an universal communication service

How to Uniquely Identify a Host?

- Name – what an object is
  - Well suited for humans
  - Example: www.cs.jmu.edu
- Address – where an object is
  - Well suited for machines
  - Example: 134.126.20.50
- Route – how to get to an object

Internet Addresses

- Also called *IP addresses*
- 32-bit integers
- Global host identifiers
- Chosen to make routing efficient
- IP address = (netid, hostid)

The Original Classful Addressing Scheme

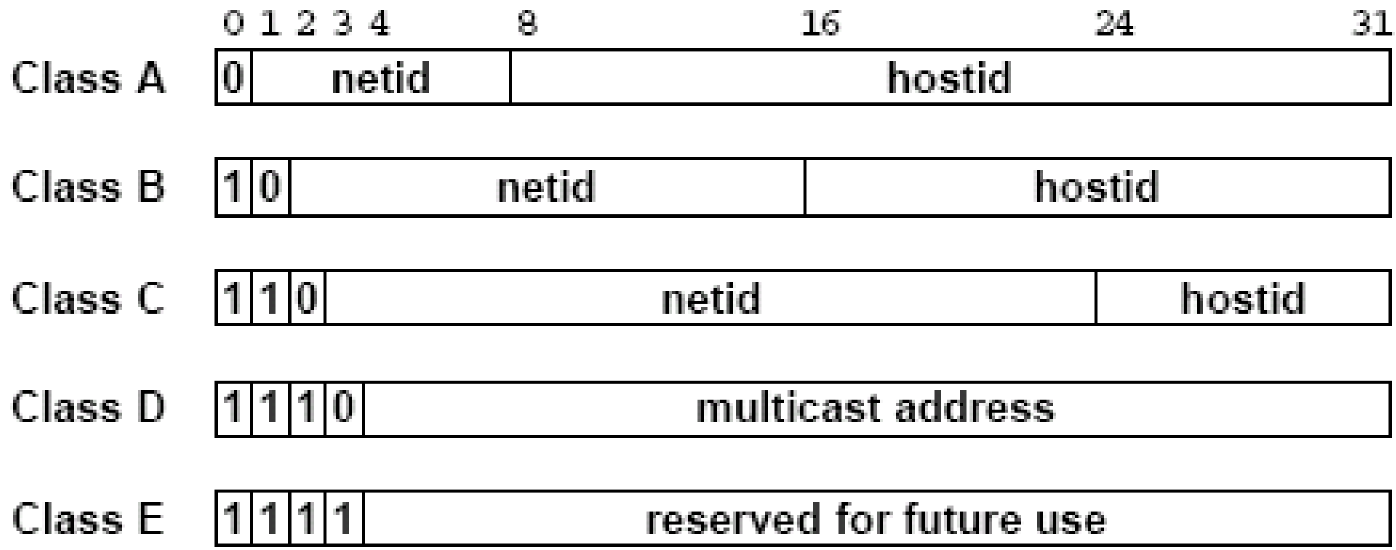

Classes of IP Addresses (cont)

- A small number ($\sim 2^{7}$) of class A networks with a large number ($\sim 2^{24}$) of hosts
- A medium number ($\sim 2^{14}$) of class B networks with a medium number ($\sim 2^{16}$) of hosts
- A large number ($\sim 2^{21}$) of class C networks with a small number ($\sim 2^{8}$) of hosts

Addresses Specify Network Connections

- What is R1's address?

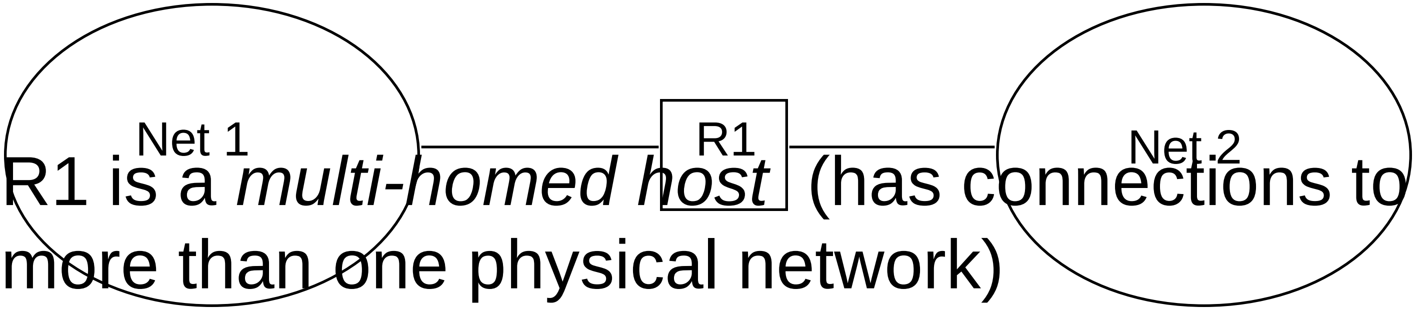

- R1 is a *multi-homed host* (has connections to more than one physical network)
- R1 requires multiple IP addresses

Addresses Specify Network Connections (cont)

- Each address corresponds to **one** of the machine’s connections
- Because IP addresses encode both a network and a host on that network, they do not specify an individual computer, but a connection to a network
- A router connecting *n* networks has *n* distinct IP addresses

IP Addresses Conventions

- The smallest hostid (all 0's) is never assigned to an individual host, instead it is used to refer to the network
- IP addresses can refer to hosts or networks
- Examples:
  - (6,8) = host #8 on network #6
  - (9,0) = network #9

IP Addresses Conventions (cont)

- The largest hostid (all 1's) is never assigned to an individual host, instead it is used to refer to a *directed broadcast*

- Example:
  - (00000101,111111111111111111111111) = all hosts on network #5

IP Addresses Conventions (cont)

- There is also a limited broadcast address (all 1’s for both netid and hostid) that broadcasts on the local network

- Example:
  - (11111111111111111111111111111111) = all hosts on the local network

IP Addresses Conventions (cont)

- Summary:
  - A field of all 1’s can be interpreted to mean “all”
    - (00000100,111111111111111111111111111) = all hosts on network #4
  - A field of all 0’s can be interpreted to mean “this”
    - (00000000,00000000000000000000000011) = host # 3 on this network

Subnet and Supernet Extensions

- Recall: each physical network must have its own unique netid
- Problem: the number of physical networks grew so fast that all netids would be exhausted (especially class B)
- Solutions (to be discussed later):
  - Subnetting – allows multiple physical networks to share the same netid
  - Supernetting – allows more complete utilization of the address space

Weaknesses on Internet Addressing

- If a computer moves from one network to another, its IP address must change

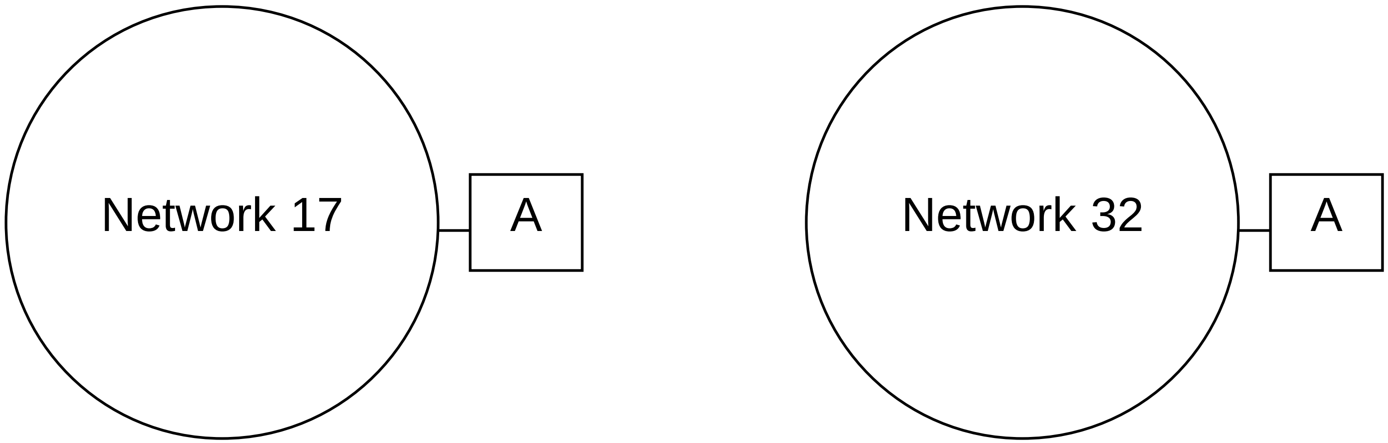

Weaknesses on Internet Addressing (cont)

- If a class C network grows to more than 255 hosts, it must have its address changed to a class B address
- Difficult:
  - Abruptly stop using one network address
  - Change the addresses of all machines
  - Resume operation using the new addresses
  - Debug problems with programs/services still using the old addresses

Weaknesses on Internet Addressing (cont)

- Recall: routers base routing decisions on the netid portion of the destination address
- Consider a host with two internet connections (and therefore two addresses):

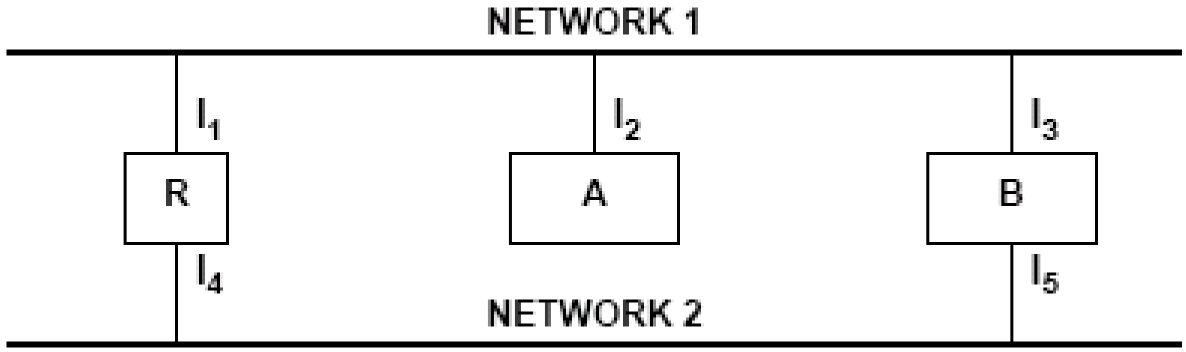

Weaknesses on Internet Addressing (cont)

- Result: the path taken by packets traveling to a host with multiple IP addresses depends on which address is used as the destination

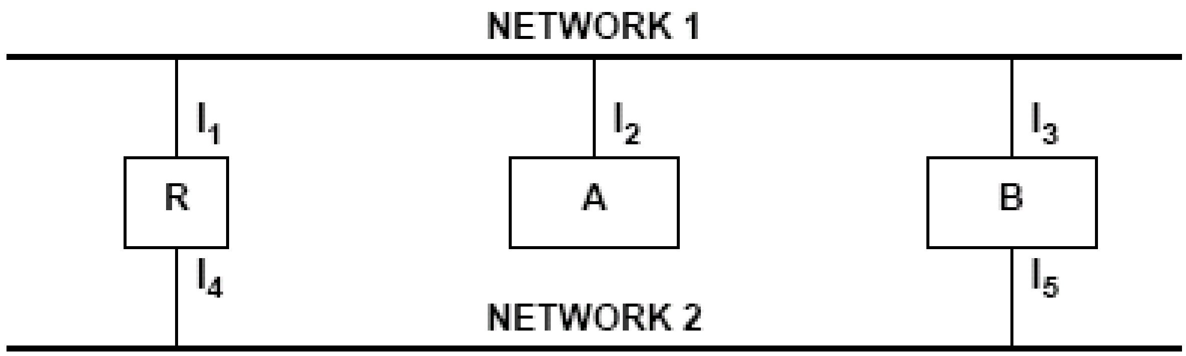

Weaknesses on Internet Addressing (cont)

- Packets sent to the same machine using different addresses may behave differently
- A machine may be reachable by on of its addresses and unreachable by another

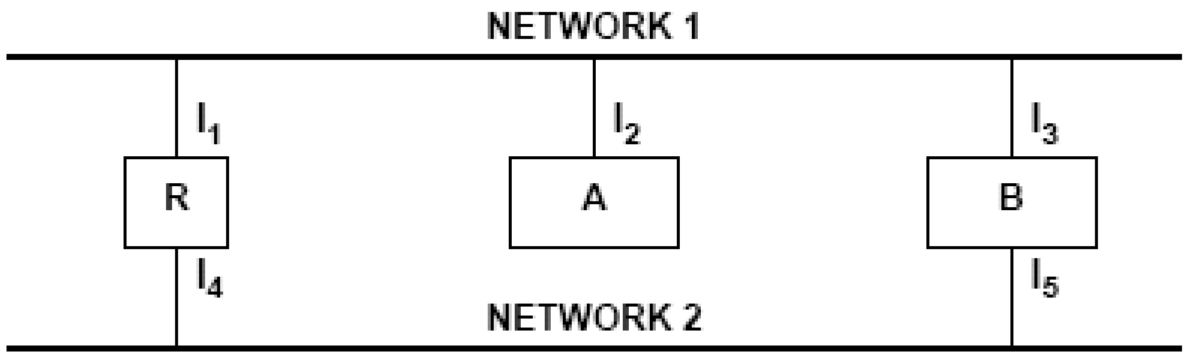

Dotted Decimal Notation

- To make them easier to read (and write) IP addresses are usually written as four decimal integers separated by decimal points
- Each decimal integer gives the value of one octet of the IP address
- Example:
  - 10000000 00000010 00000000 00001111 = 128.2.0.15

The Loopback Address

- Not all IP addresses have been assigned to classes
- In particular, netid 127, a value from the class A range, is reserved for *loopback*
- Used for testing TCP/IP and for inter-process communication