

# Internet Protocol (IP)

Raj Jain

**Raj Jain is now at  
Washington University in Saint Louis  
Jain@cse.wustl.edu  
<http://www.cse.wustl.edu/~jain/>**

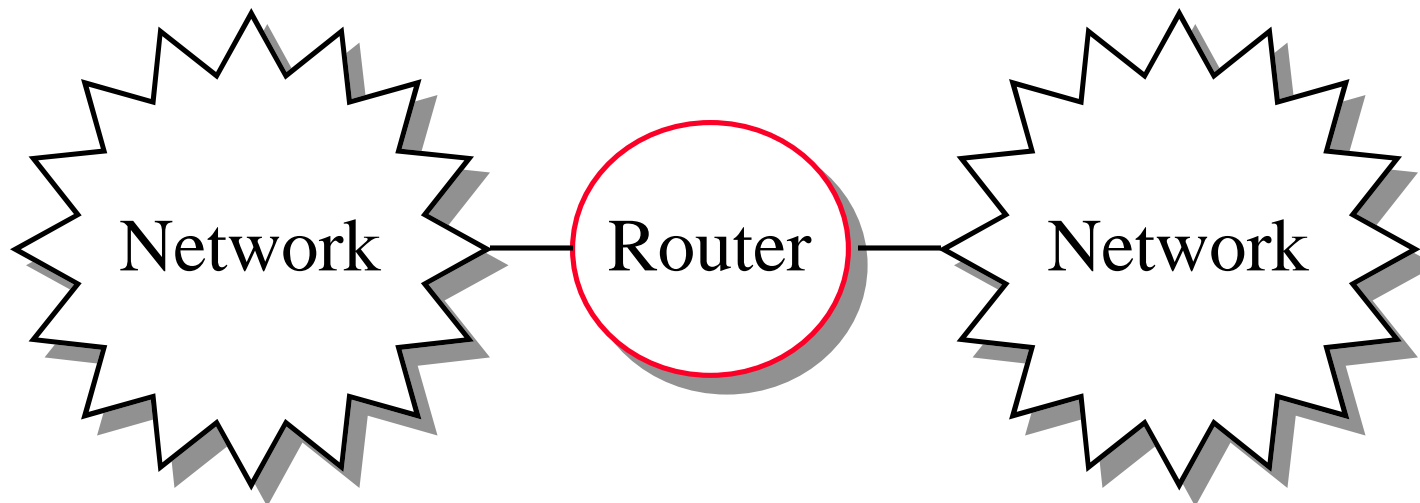


- ❑ Internetworking
- ❑ IP Address format
- ❑ IP data forwarding
- ❑ Fragmentation and reassembly

Ref: Chapters 13, 14, 16, and 17 of Comer's Computer Networks and Internets

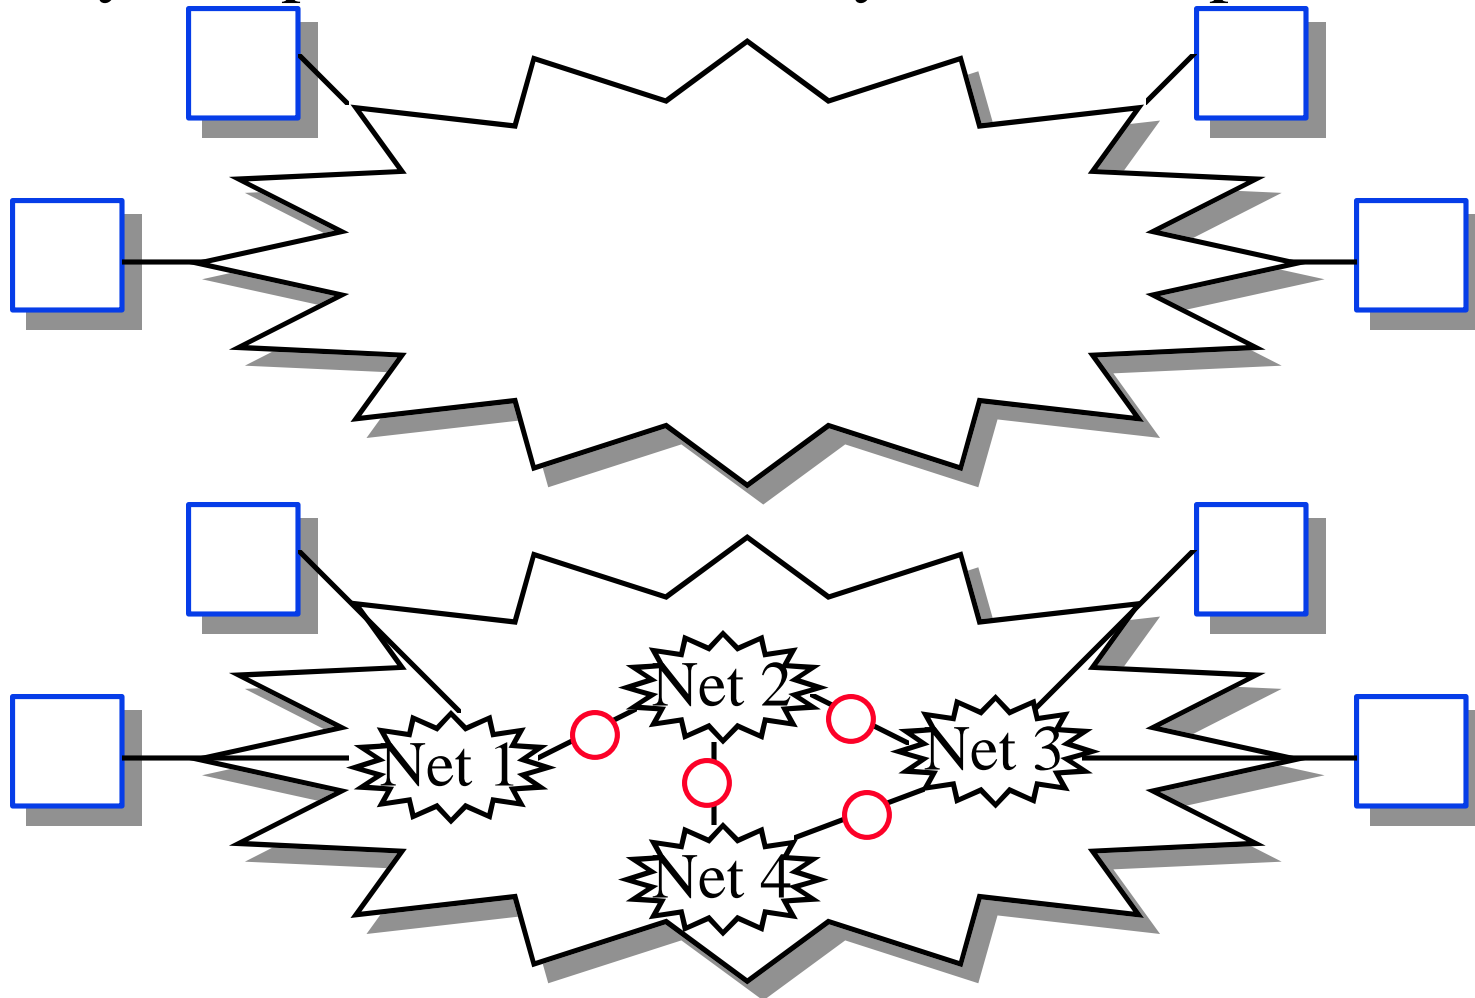
# Internetworking

- ❑ Internetwork = Collection of networks  
Connected via routers




# Internet = Virtual Network

- Any computer can talk to any other computer




# IP Address


❑ Class A:



❑ Class B:




❑ Class C:



❑ Class D:



❑ Local = Subnet + Host (Variable length)



# Computing The Class of an Address

| First 4 bits | Index | Class |
|--------------|-------|-------|
| 0000         | 0     | A     |
| 0001         | 1     | A     |
| 0010         | 2     | A     |
| 0011         | 3     | A     |
| 0100         | 4     | A     |
| 0101         | 5     | A     |
| 0110         | 6     | A     |
| 0111         | 7     | A     |
| 1000         | 8     | B     |
| 1001         | 9     | B     |
| 1010         | 10    | B     |
| 1011         | 11    | B     |
| 1100         | 12    | C     |
| 1101         | 13    | C     |
| 1110         | 14    | D     |
| 1111         | 15    | E     |

# Classes and Dotted Decimal Notation

- Binary: 11000000 00000101 00110000 00000011  
Hex Colon: C0:05:30:03  
Dotted Decimal: 192.5.48.3

| Class | Range           |
|-------|-----------------|
| A     | 0 through 127   |
| B     | 128 through 191 |
| C     | 192 through 223 |
| D     | 224 through 239 |
| E     | 240 through 255 |


## Division of the Address Space

| Class | Bits in Prefix | Max # of Nets | Bits in Suffix | Max # of Hosts per Net |
|-------|----------------|---------------|----------------|------------------------|
| A     | 7              | 128           | 24             | 16,777,216             |
| B     | 14             | 16,384        | 16             | 65,536                 |
| C     | 21             | 2,097,152     | 8              | 256                    |

- ❑ Not all possible addresses can be used.




# An Addressing Example



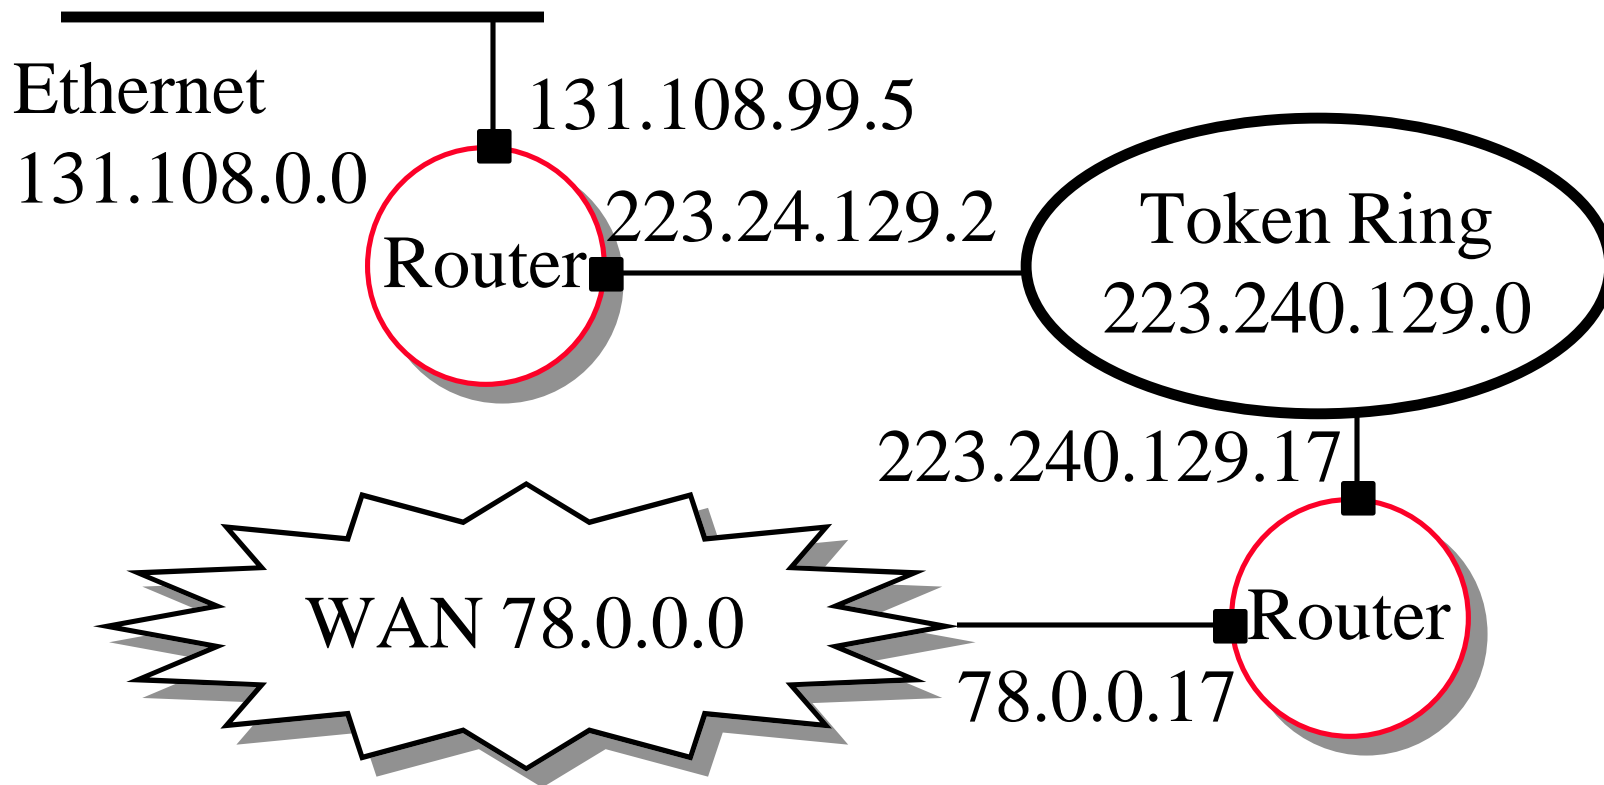
- All hosts on a network have the same network prefix

# Multi-Homed Hosts



- ❑ Each interface has an address.  
Two or more interfaces  $\Rightarrow$  Multi-homed hosts
- ❑ Multihoming is for reliability or performance

# Routers and the IP Addressing Principle



- ❑ Routers have two or more addresses.  
One for each interface.

# Special IP Addresses

- ❑ All-0 host suffix  $\Rightarrow$  Network Address
- ❑ All-0s  $\Rightarrow$  This computer
- ❑ All-1 host suffix  $\Rightarrow$  All hosts on the destination net (directed broadcast),  
All-0 host suffix  $\Rightarrow$  Berkeley directed broadcast address
- ❑ All-1s  $\Rightarrow$  All hosts on this net (limited broadcast)
- ❑ 127.\*.\*  $\Rightarrow$  Looback through IP layer

# Homework

- ❑ Read chapters 13, 14
- ❑ Submit answer to exercise 14.8  
(requires 14.6 and 14.7 also)
- ❑ Use your program to printout characteristics of the following addresses using a subnet mask of 255.255.255.0: 164.107.61.200, 164.107.61.0, 164.107.61.255, 127.107.61.200, 0.0.0.0

# IP Features

- ❑ Connectionless service
- ❑ Variable size datagrams
- ❑ Best-effort delivery: Delay, out-of-order, corruption, and loss possible. Higher layers should handle these.
- ❑ Handles only data forwarding  
Uses routing tables prepared by other protocols, e.g.,  
Open Shortest Path First (OSPF),  
Routing Information Protocol (RIP)
- ❑ Provides only “Send” and “Delivery” services  
Error and control messages generated by  
Internet Control Message Protocol (ICMP)

# Forwarding an IP Datagram

- ❑ Delivers datagrams to destination network (subnet)
- ❑ Routers maintain a “routing table” of “next hops”
- ❑ Next Hop field does not appear in the datagram




Table at R2:


Destination      Next Hop

|       |                |
|-------|----------------|
| Net 1 | Forward to R1  |
| Net 2 | Deliver Direct |
| Net 3 | Deliver Direct |
| Net 4 | Forward to R3  |

Fig 16.2

# IP Addresses and Routing Table Entries

- IF ((Mask[i] & Destination Addr) == Destination[i])  
*Forward to NextHop[i]*



| Destination | Mask          | Next Hop       |
|-------------|---------------|----------------|
| 30.0.0.0    | 255.0.0.0     | 40.0.0.7       |
| 40.0.0.0    | 255.0.0.0     | Deliver direct |
| 128.1.0.0   | 255.255.0.0   | Deliver direct |
| 192.4.10.0  | 255.255.255.0 | 128.1.0.9      |



# IP Datagram Format

|                             |        |              |                 |  |
|-----------------------------|--------|--------------|-----------------|--|
| Vers                        | H. Len | Service Type | Total Length    |  |
| Identification              |        | Flags        | Fragment Offset |  |
| Time to live                | Type   |              | Header Checksum |  |
| Source IP Address           |        |              |                 |  |
| Destination IP Address      |        |              |                 |  |
| IP Options (May be omitted) |        |              | Padding         |  |
| Data                        |        |              |                 |  |

# IP Format

- ❑ Version (4 bits)
- ❑ Internet header length (4 bits): in 32-bit words.  
Min header is 5 words or 20 bytes.
- ❑ Type of service (8 bits): Reliability, precedence, delay, and throughput
- ❑ Total length (16 bits): header + data in bytes  
Total must be less than 64 kB.
- ❑ Identifier (16 bits): Helps uniquely identify the datagram during its life for a given source, destination address

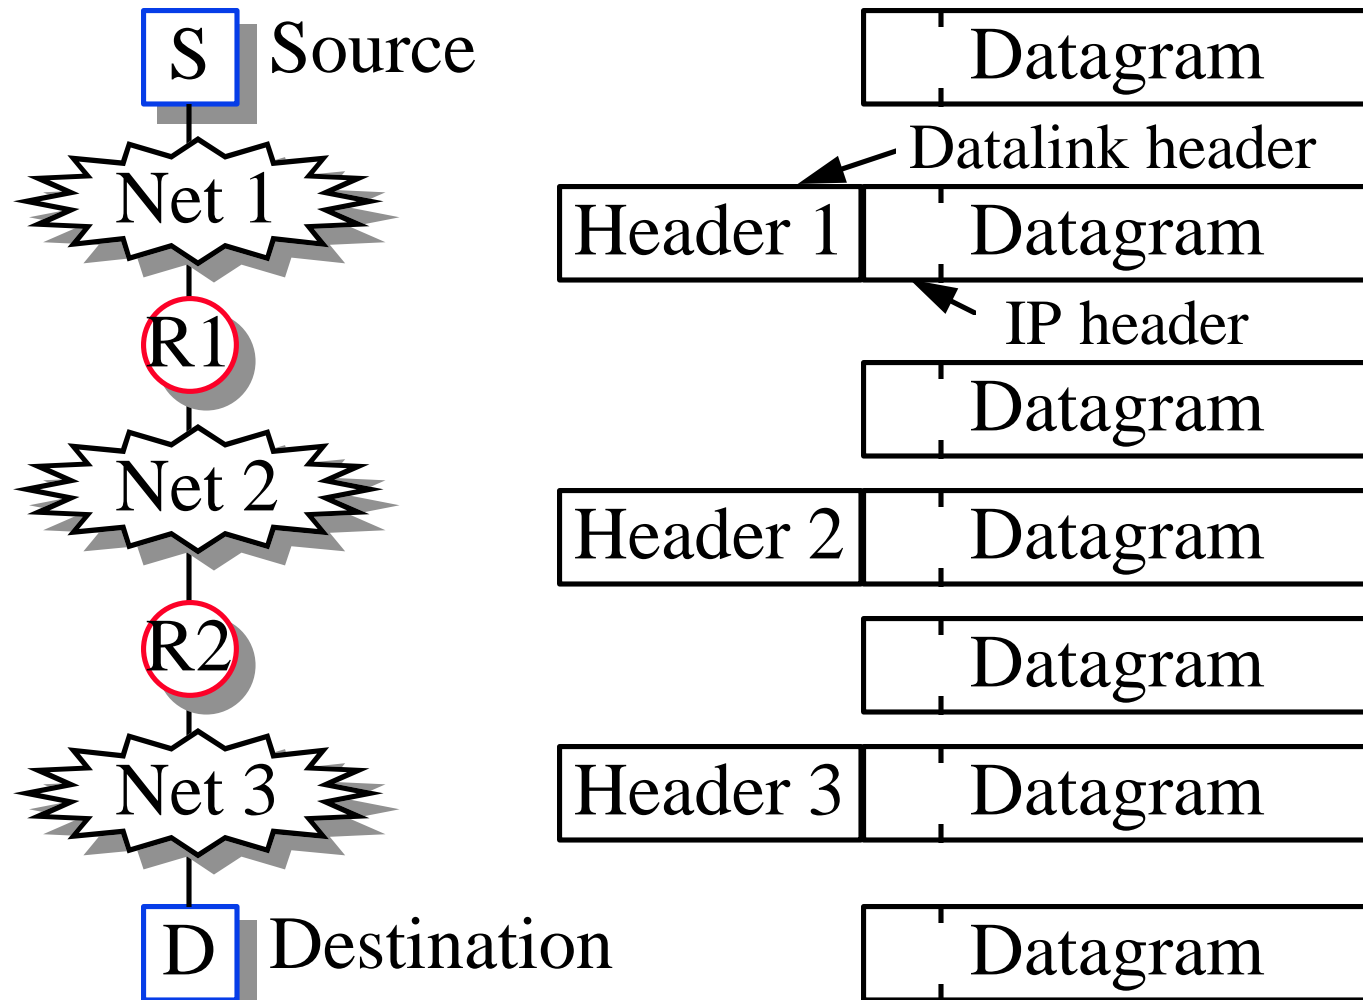
## IP Header (Cont)

- ❑ Flags (3 bits):
  - More flag - used for fragmentation
  - No-fragmentation
  - Reserved
- ❑ Fragment offset (13 bits): In units of 8 bytes
- ❑ Time to live (8 bits): Specified in router hops
- ❑ Protocol (8 bits): Next level protocol to receive the data
- ❑ Header checksum (16 bits): 1's complement sum of all 16-bit words in the header

## IP Header (Cont)

- ❑ Source Address (32 bits): Original source.  
Does not change along the path.
- ❑ Destination Address (32 bits): Final destination.  
Does not change along the path.
- ❑ Options (variable): Security, source route, record route, stream id (used for voice) for reserved resources, timestamp recording
- ❑ Padding (variable):  
Makes header length a multiple of 4
- ❑ Data (variable): Data + header  $\leq 65,535$  bytes

# Transmission Across An Internet



- Datalink header changes at every hop

# Maximum Transmission Unit

- ❑ Each subnet has a maximum frame size  
Ethernet: 1518 bytes  
FDDI: 4500 bytes  
Token Ring: 2 to 4 kB
- ❑ Transmission Unit = IP datagram (data + header)
- ❑ Each subnet has a maximum IP datagram length: MTU




Fig 17.3

# IP Protocol Numbers

| Decimal | Key word | Protocol                           |
|---------|----------|------------------------------------|
| 0       |          | Reserved                           |
| 1       | ICMP     | Internet Control Message Protocol  |
| 2       | IGMP     | Internet Group Management Protocol |
| 4       | ST       | Stream Protocol                    |
| 5       | TCP      | Transmission Control Protocol      |
| 8       | EGP      | Exterior Gateway Protocol          |
| 9       | IGP      | Interior Gateway Protocol          |
| 17      | UDP      | User Datagram Protocol             |

# IP Options Coding

|      |        |       |
|------|--------|-------|
| Type | Length | Value |
| 1B   | 1B     | $n$ B |

|           |       |        |
|-----------|-------|--------|
| Flag Copy | Class | Number |
| 1b        | 2b    | 5b     |

- ❑ Flag Copy: 0 = Copy the option only into the first fragment of a fragmented datagram  
1 = Copy into all fragments
- ❑ Class: 0=User or control, 1=Reserved, 2=Diagnostics, 3=reserved



# IP Options

| Class | Number | Length | Description           |
|-------|--------|--------|-----------------------|
| 0     | 0      | 0      | End of Options        |
| 0     | 1      | 0      | No Op                 |
| 0     | 2      | 11     | Security              |
| 0     | 3      | Var    | Loose Source Routing  |
| 0     | 7      | Var    | Record Route          |
| 0     | 8      | 4      | Stream ID (obsolete)  |
| 0     | 9      | Var    | Strict Source Routing |
| 2     | 4      | Var    | Internet Time-Stamp   |

# IP Source Routing

| Code | Length | Pointer | Router Data |
|------|--------|---------|-------------|
|------|--------|---------|-------------|


|   |           |           |             |
|---|-----------|-----------|-------------|
| P | 128.2.3.4 | 128.7.8.9 | 128.10.4.12 |
|---|-----------|-----------|-------------|

|   |           |           |             |
|---|-----------|-----------|-------------|
| P | 128.2.3.4 | 128.7.8.9 | 128.10.4.12 |
|---|-----------|-----------|-------------|


# Route Recording

| Code | Length | Pointer | Route Data |
|------|--------|---------|------------|
|------|--------|---------|------------|


|   |           |       |       |
|---|-----------|-------|-------|
| P | 128.2.3.4 | Empty | Empty |
|---|-----------|-------|-------|



|   |           |           |       |
|---|-----------|-----------|-------|
| P | 128.2.3.4 | 128.7.8.9 | Empty |
|---|-----------|-----------|-------|

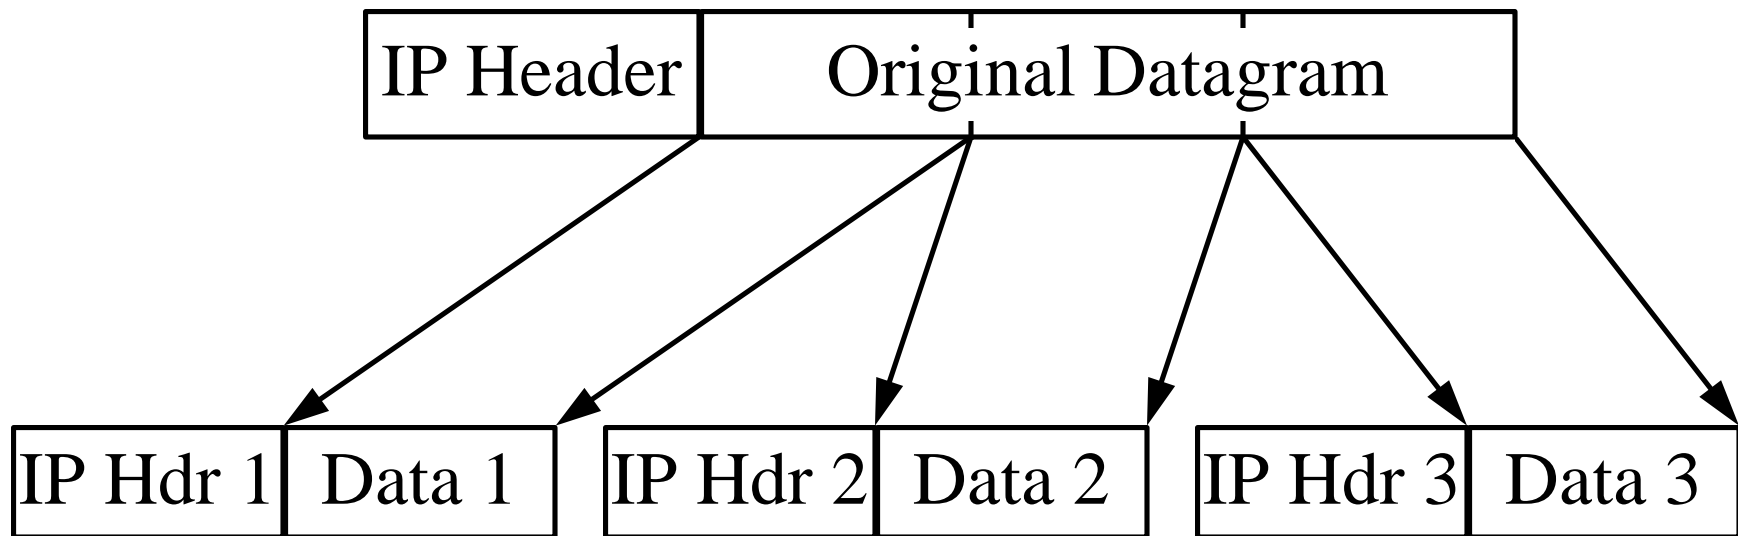


# Timestamp Option




# Fragmentation

- ❑ Datagrams larger than MTU are fragmented
- ❑ Original header is copied to each fragment and then modified (fragment flag, fragment offset, length,...)



# Fragmentation



ID = 12345, More = 1  
Offset = 160W, Len = 1500B

ID = 12345, More = 1  
Offset = 0W, Len = 512B

ID = 12345, More = 1  
Offset = 0W, Len = 256B

ID = 12345, More = 1  
Offset = 32W, Len = 256B

ID = 12345, More = 1  
Offset = 64W, Len = 512B

ID = 12345, More = 1  
Offset = 64W, Len = 256B

ID = 12345, More = 1  
Offset = 96W, Len = 256B


ID = 12345, More = 1  
Offset = 128W, Len = 476B

ID = 12345, More = 1  
Offset = 128W, Len = 256B


ID = 12345, More = 1  
Offset = 160W, Len = 220B

# Reassembly

- ❑ Reassembly only at the final destination
- ❑ Partial datagrams are discarded after a timeout
- ❑ Fragments can be further fragmented along the path. Subfragments have a format similar to fragments. It is not possible to tell how many times fragmented.
- ❑ Minimum MTU along a path  $\Rightarrow$  Path MTU



# Summary



- ❑ IPv4 uses 32-bit addresses organized as network prefix and host suffix.
- ❑ Four classes of networks: A, B, C, D
- ❑ Routers determine next hop using routing tables
- ❑ IP provides connectionless unreliable service



# Homework

- ❑ Read Chapters 16, 17
- ❑ Submit answers to exercises 16.7, 17.4