Motivation for Internetworking

- LAN technologies provide high speed communication across short distances
- WAN technologies serves large areas
- No single networking technology is best for all needs
- Ex: Ethernet might be the best solution for connecting computers in an office
- Ex: Frame relay might be the best solution for interconnecting computers in one city to another

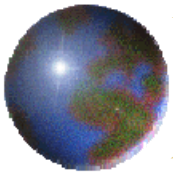

Universal Service

- Allows arbitrary pair of computers to communicate
- Increases individual productivity
- Incompatibilities among network hardware and physical addressing prevents universal service to extend across multiple networks that use multiple technologies
- Solution is Internetworking or Internet
  - Provides universal service among heterogeneous networks
  - Uses both hardware and software
  - Is not restricted in size

Router

- The basic hardware component used to connect heterogeneous networks
- Has a conventional processor and memory as well as separate I/O interface for each network
- Can connect two LANs, a LAN and a WAN or two LANs
- Can interconnect networks that use different technologies, media, physical addressing schemes or frame formats

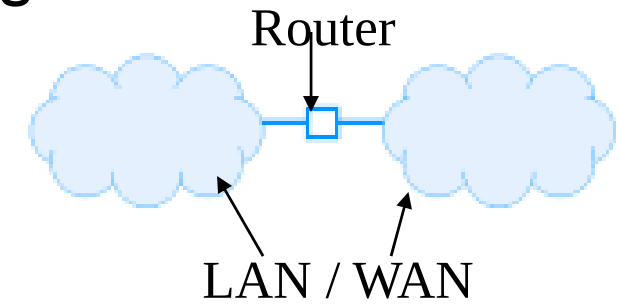

Internet Architecture

- Consists of a set of networks interconnected by networks
- Commercial routers can connect more than two networks
- A single router is seldom used because
  - CPU and memory is insufficient
  - Redundancy improves internet reliability
- The internet scheme allows to choose
  - The number and type of networks
  - The number of routers
  - The exact interconnection topology

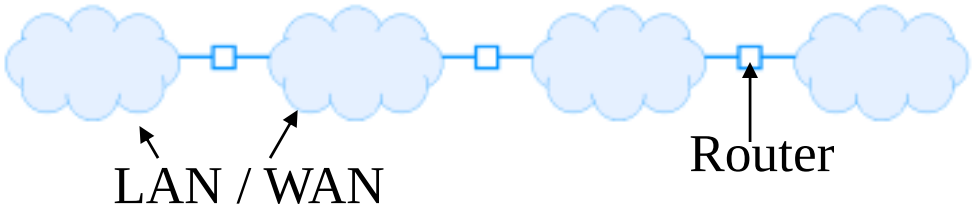

Virtual Network

- Offers universal service
  - Each computer is assigned an address and can communicate with any computer
- Internet is a virtual network system

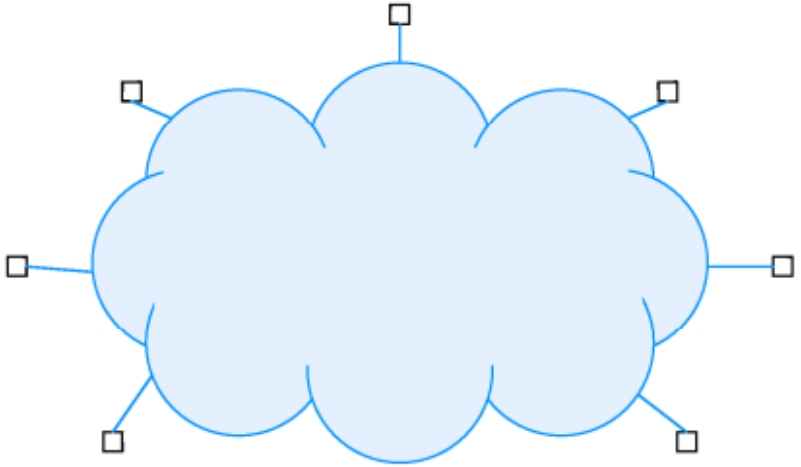

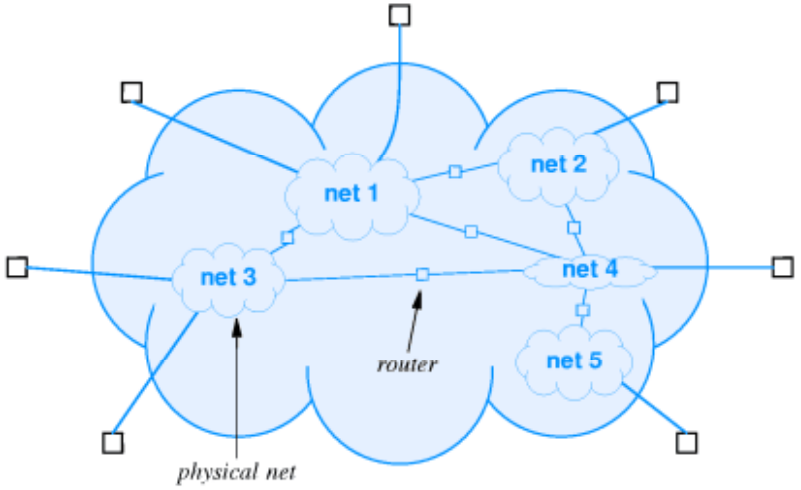

Protocols for Internetworking

- TCP/IP
  - Widely used for internetworking
- The TCP/IP layering model contains five layers
- Also called Internet

  Layering Model or Internet reference Model
- Four layers of TCP/IP reference model correspond to layers of the ISO model
- The ISO model has

  no Internet layer

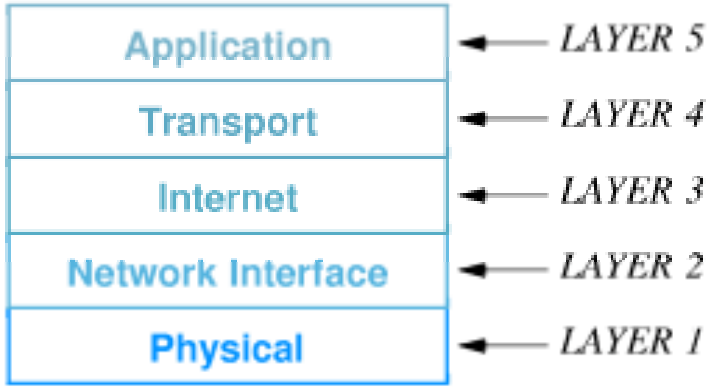

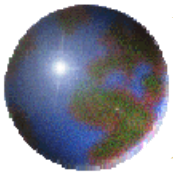

Layers of TCP/IP

- Layer 1: Physical
  - Corresponds to basic network hardware
- Layer 2: Network Interface
  - Specifies how to organize data into frames and how a computer transmits frames over a network
- Layer 3: Internet
  - Specifies the format of packets sent across an internet
  - Mechanisms used to forward packets from a computer through one or more routers to a final destination
- Layer 4: Transport
  - Specifies how to ensure reliable transfer

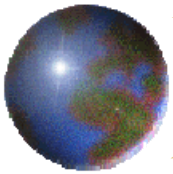

TCP/IP (Cont.)

- Layer 5: Application
  - Specifies how one application uses the internet
- Host computer
  - Any computer system that connects to an internet and runs applications
- Both hosts and routers needs TCP/IP
- A router does not need layer 5 protocols because router do not run applications
  - Ex: a file transfer application

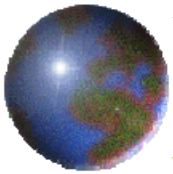

Addresses for the Virtual Network

- Internet is merely an abstraction created entirely by software
- To guarantee uniform addressing, protocol software defines an addressing scheme
- Uniform addressing helps create the illusion of a large, seamless network
- The abstract addressing scheme assigns each host a unique address to communicate
- Users, application programs and higher levels of protocols software use the abstract addresses to communicate

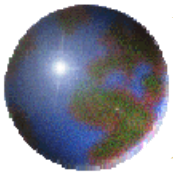

The IP Addressing Scheme

- Addressing is specified by the Internet Protocol (IP)
- Internet protocol address or IP address
  - A unique 32-bit binary number
  - Used for all communication with the host
- Each 32-bit IP address is divided onto two parts
  - A prefix and a suffix
- Network number
  - A unique value assigned to each physical network
- The IP address hierarchy guarantees that
  - Each computer is assigned a unique address
  - Suffixes can be assigned locally without global coordination

Classes of IP Addresses

- The class of an address determines the boundary between the network prefix and host prefix
- IP divides host address into their primary classes A, B and C
- The first four bits of an address determines the class
- To use IP multicasting, a set of hosts must agree to share a multicast address

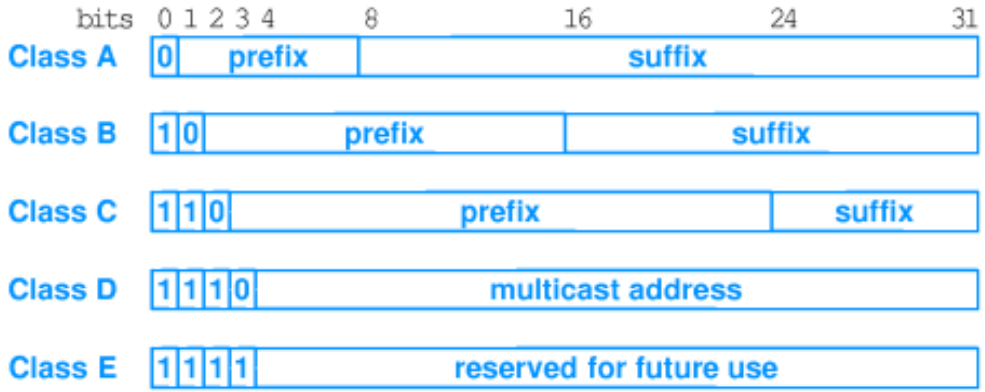

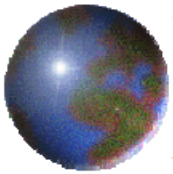

Computing the Class

- IP addresses are self identifying
  - Class of an address can be computed from address itself

First Four Bits Of Address	Table Index (in decimal)	Class of Address
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

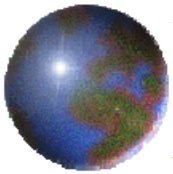

Dotted Decimal Notation

- A syntactic form that IP software uses when interacting with humans
- Expresses each 8-bit section of a 32-bit number as a decimal value
- Uses a dot to separate octets
- The class is recognized by the decimal value of the first octet

32-bit Binary Number	Equivalent Dotted Decimal
10000001 00110100 00000110 00000000	129 . 52 . 6 . 0
11000000 00000101 00110000 00000011	192 . 5 . 48 . 3
00001010 00000010 00000000 00100101	10 . 2 . 0 . 37
10000000 00001010 00000010 00000011	128 . 10 . 2 . 3
10000000 10000000 11111111 00000000	128 . 128 . 255 . 0

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

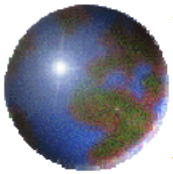

Addresses

- The classes do not contain the same type of network

Address Class	Bits In Prefix	Maximum Number of Networks	Bits In Suffix	Maximum Number Of Hosts Per Network
A	7	128	24	16777216
B	14	16384	16	65536
C	21	2097152	8	256

- Each network prefix must be unique
- An organization obtains network numbers from ISPs
- ISPs coordinate with a central organization, the Internet Assigned Number Authority

An Addressing Example

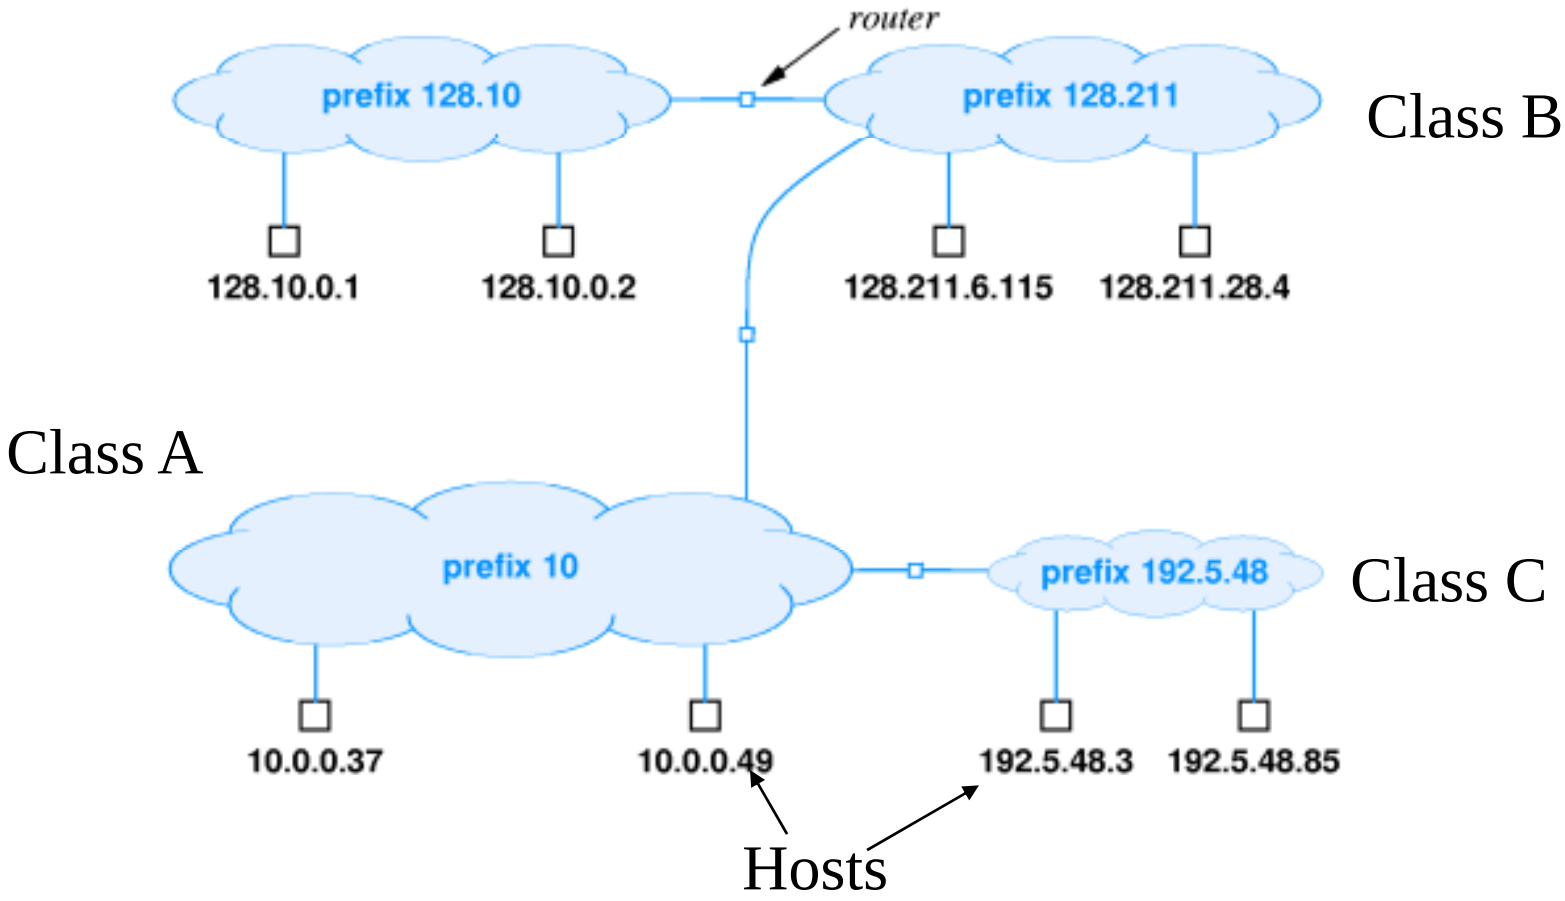

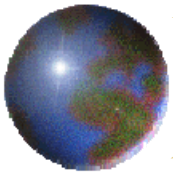

Special IP Addresses

- IP defines a set of special address forms that are reserved

Prefix	Suffix	Type of Address	Purpose
All 0s	All 0s	This computer	Used during bootstrap
Network	All 0s	Network	Identifies a network
Network	All 1s	Direct broadcast	Broadcast on a specified net
All 1s	All 1s	Limited broadcast	Broadcast on a local net
127	any	loop back	testing

Routers and IP Addresses

- Each IP address identifies a connection between a computer and a network
- Each router is assigned two or more IP addresses
- IP does not require that the same suffix be assigned to all interfaces of a router

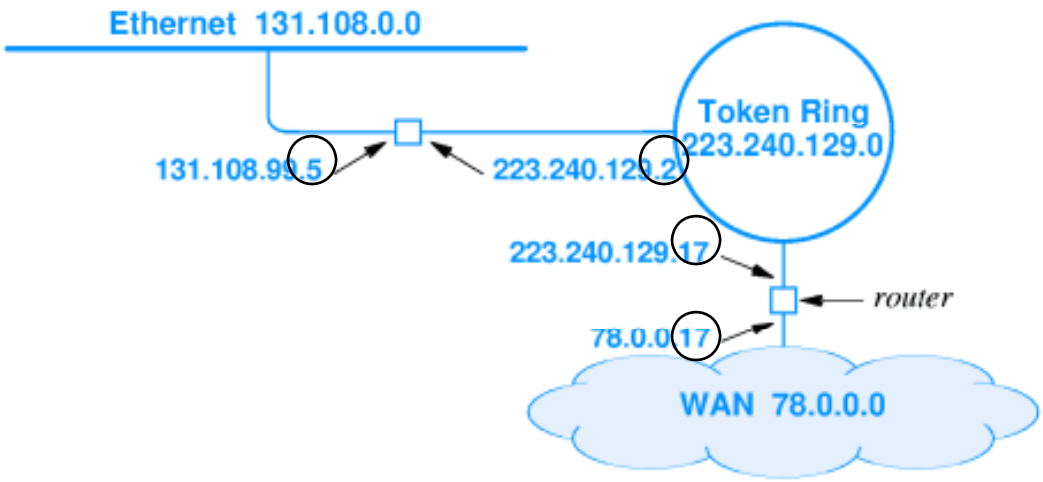

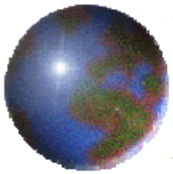

Multi-homed Host

- Multi-Homed
  - A host computer that connects to multiple networks
- Increases reliability
  - If one network fails, the host can still reach the internet through the second connection
- Increases performance
  - Traffic can be directed to avoid congested routers
- Has multiple addresses, one for each network connection

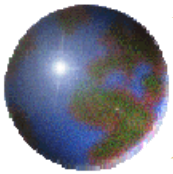

Protocol Addresses

- A frame transmitted across a physical network must contain the hardware address
- The next-hop and the packet's destination address are IP address
- Physical network address does not understand IP addressing
- A frame sent across a given physical network must
  - Use the hardware's frame format
  - Use hardware addresses

Address Resolution

- Mapping between a protocol address and a hardware address
- Address resolution is local to a network
- A computer never resolves the address of a computer that attaches to a remote network.
- Each computer that handles a packet resolves a next-hop address before sending

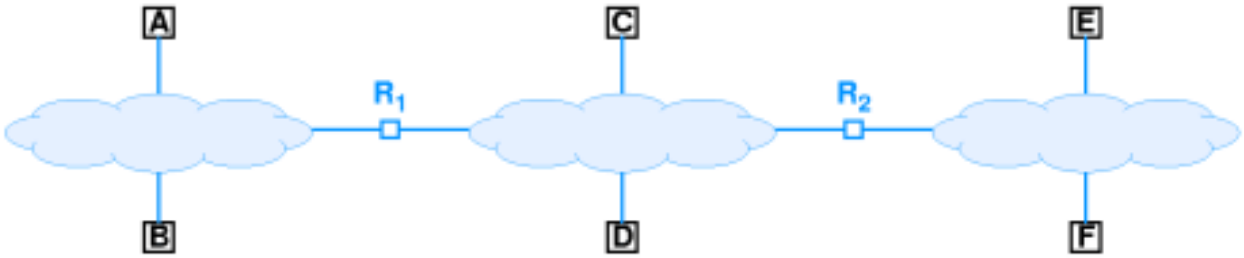

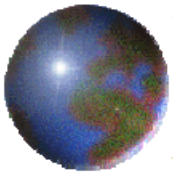

Address Resolution Techniques

- Depends on protocol and hardware addressing scheme
- Table look-up
  - Binding or mappings are stored in a table in memory, which the software searches when it needs to resolve an address
- Closed-form computation
  - Computer's hardware address can be computed from the protocol address using basic Boolean and arithmetic operations
- Message exchange
  - Computers exchange messages across a network to resolve an address

Table Look-up Technique

- The table consists of an array containing a pair of protocol and equivalent hardware addresses
- A separate binding table is used for each physical network
- For small networks sequential search is used
- For large networks
  - Hashing or Direct indexing

IP Address	Hardware Address
197.15.3.2	0A:07:4B:12:82:36
197.15.3.3	0A:9C:28:71:32:8D
197.15.3.4	0A:11:C3:68:01:99
197.15.3.5	0A:74:59:32:CC:1F
197.15.3.6	0A:04:BC:00:03:28
197.15.3.7	0A:77:81:0E:52:FA

Direct look-up for C class

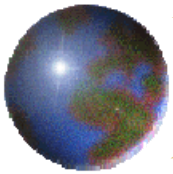

Closed-form Technique

- Some technologies use configurable addressing
- The local network administrator chooses both the hardware an IP address
- Closed form method computes a mathematical function that maps an IP address to a hardware address
- Values are chosen to optimize the translation
- Ex: host portion of a computer's IP address can be chosen to be identical to the computer's hardware address
  - hardware_address = ip_address & 0xFF

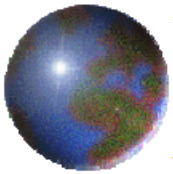

Message Exchange Technique

- To resolve an address, send a message across a network and receive a reply
  - Message carries protocol address
  - Reply carries hardware address
- An address resolution request is sent to
  - One or more resolution serves, or
  - Each computer on a network

Feature	Type Of Resolution
Useful with any hardware	T
Address change affects all hosts	T
Protocol address independent of hardware address	T, D
Hardware address must be smaller than protocol address	C
Protocol address determined by hardware address	C
Requires hardware broadcast	D
Adds traffic to a network	D
Produces resolution with minimum delay	T, C
Implementation is more difficult	D

Address Resolution Protocol

- The ARP standard defines two basic message types
  - A request and a response
- ARP request message containing the IP address is placed in a hardware frame and broadcast to all computers
- A response contains both IP and hardware addresses but it is not broadcast

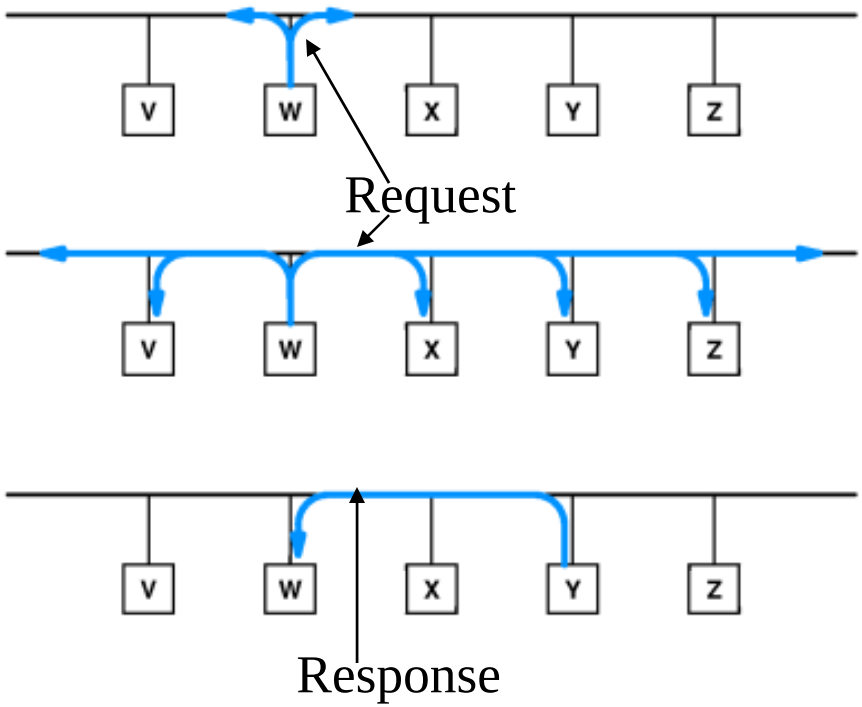

ARP Message Format

- The ARP standard
  - Describes the general form for ARP messages
  - Specifies how to determine the details for each type of network address
- ARP is almost always used to bind a 32-bit IP address to a 48-bit Ethernet address

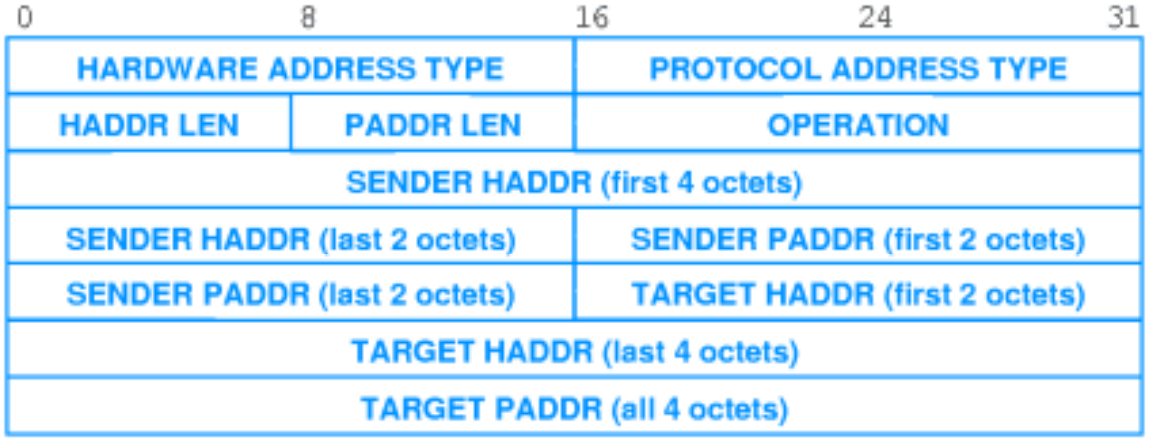

ARP Message and Frames

- Encapsulation
  - Placing a message inside a frame for transport
- ARP is encapsulated directly in a hardware frame
- The type field in the frame header specifies that the frame contains an ARP message.
  - Does not distinguish between a request and a response

Dest. Address	Source Address	Frame Type	Data In Frame
		806	complete ARP message

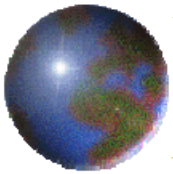

Caching ARP Responses

- Three packets traverse the network for each ARP transmission
- To reduce network traffic, ARP software extracts and saves the information
- ARP manages the table as a cache
- ARP uses the binding(if present) without transmitting a request
- If binding is not present
  - ARP broadcasts a request
  - Waits for a response
  - Updates the cache
  - Proceeds to use the binding

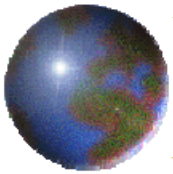

Processing ARP Message

- Receiver must perform two basic steps
  - Extract the sender's address binding and check its presence
  - Determines whether message is a request or a response
- After the computer replies to an ARP request, the computer extracts the sender's address binding
- Optimization is done because
  - Most computers communication involves two-way traffic
  - A computer cannot store an arbitrary number of address bindings

Layering and Addressing

- Address resolution is associated with network interface layer
- Address resolution software hides the details of physical addressing
- Applications and higher-layers of protocol software are built to use protocol addresses only

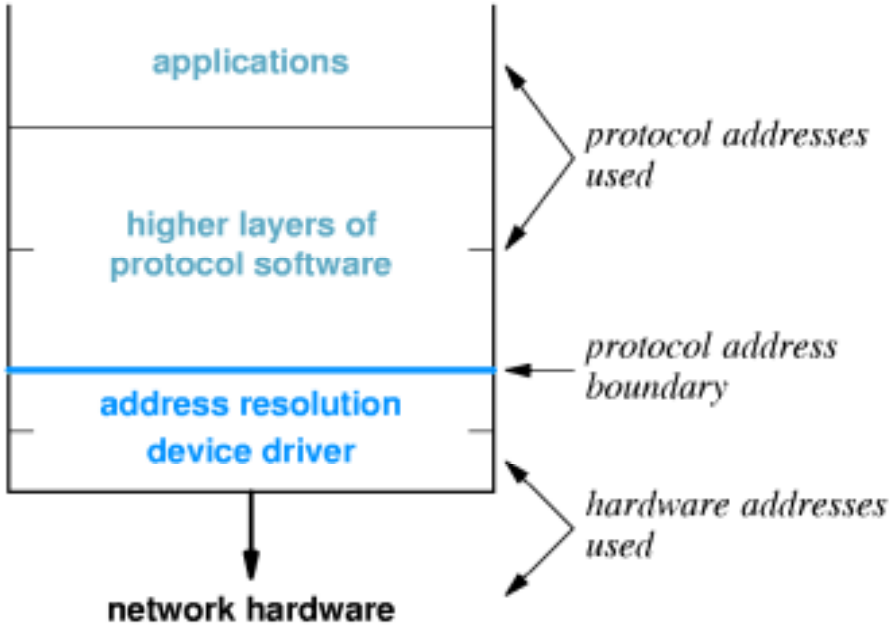

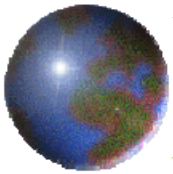

Virtual Packets

- TCP/IP designers include protocol for both connectionless and connection-oriented services
- Applications program remain unaware of the underlying physical networks
- Router forwards each packet from one network to another
- No fixed frame format because
  - Routers can connect heterogeneous networks
- Universal-virtual packet
  - An internet packet format independent of the hardware

The IP Datagram

- A packet sent across a TCP/IP internet
- Each datagram consists of a header followed by data
- Source and destination addresses in the datagram header are IP addresses
- The size of the datagram is variable
  - Makes IP adaptable to a variety of applications

IP Datagram

Forwarding an IP Datagram

- Datagrams traverse from source to destination through routers
- Each IP router keeps information in a routing table
- Each destination listed in a routing table is a network, not an individual host

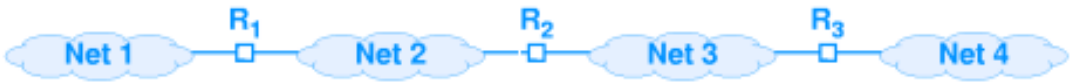

Routing table for R2

Destination	Next Hop
net 1	R_1
net 2	deliver direct
net 3	deliver direct
net 4	R_3

IP Addresses and Routing Table Entries

- In practice, an IP routing table is complex and contains
  - First, the Destination field in each entry contains the network prefix of the destination address
  - Second, an address mask specifying which bits of the destination correspond to network prefix
  - Third, next-hop specifying IP address of the router

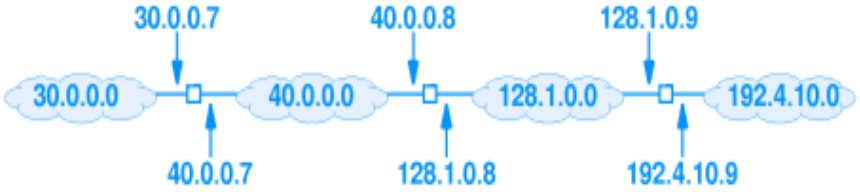

(a)

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

(b)

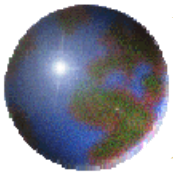

Routing Table Entries

- Routing or Forwarding
  - The process of using a routing table to select a next-hop
- The mask field provides the network part of an address during lookup
- Software computes the Boolean *and* of the mask and the datagram destination address
  - if((Mask [i] & D)== Destination[i])forward to NextHop[i];
- The destination address in the datagram header always refers to the ultimate address
- Although the datagram is forwarded to another router, datagram header retains destination address.

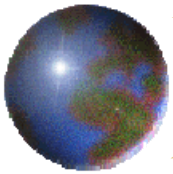

Best-Effort Delivery

- IP is designed to operate over all types of network hardware
- Uses the term best-effort to describe the service it offers
- IP cannot handle the following problems
  - Datagram duplication
  - Delayed of out-of-order delivery
  - Corruption of data
  - Datagram loss
- Additional layers of protocol software are needed to handle each of these errors

IP Datagram Header Format

- Each field in an IP datagram header has a fixed size

VERS- IP version

H.LEN- Header length

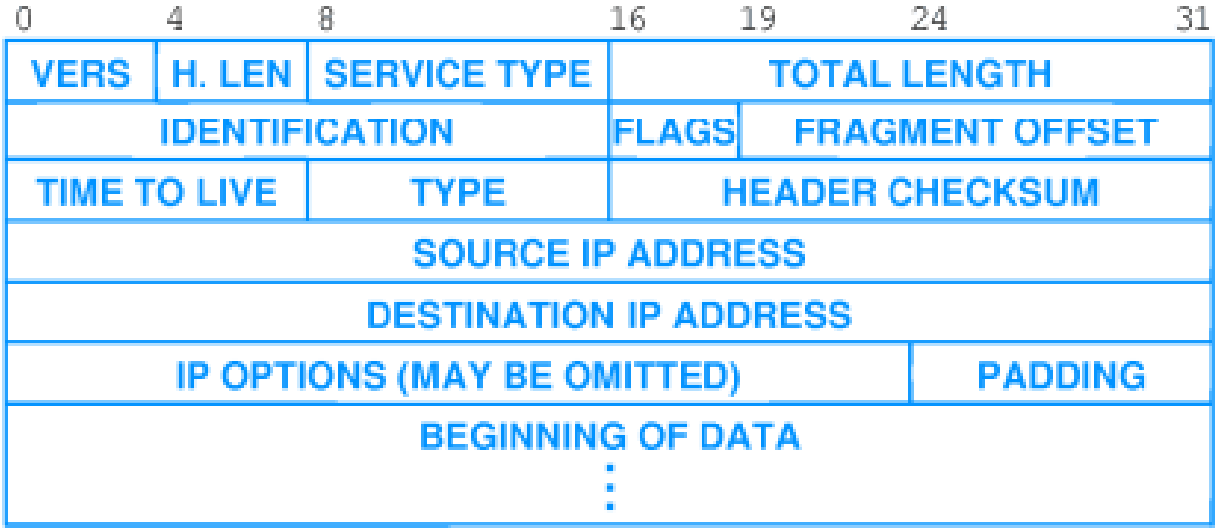

Datagram Transmission and Encapsulation

- Network hardware does not understand datagram format or Internet Addressing
- Encapsulation
  - The entire datagram is placed in the data area of the frame
- The destination address in the frame is the address of the next-hop
- The address is obtained by translating the IP address to an equivalent hardware address

Transmission Across an Internet

- When a datagram arrives in a network frame
  - Receiver extracts the datagram
  - Discards the frame header

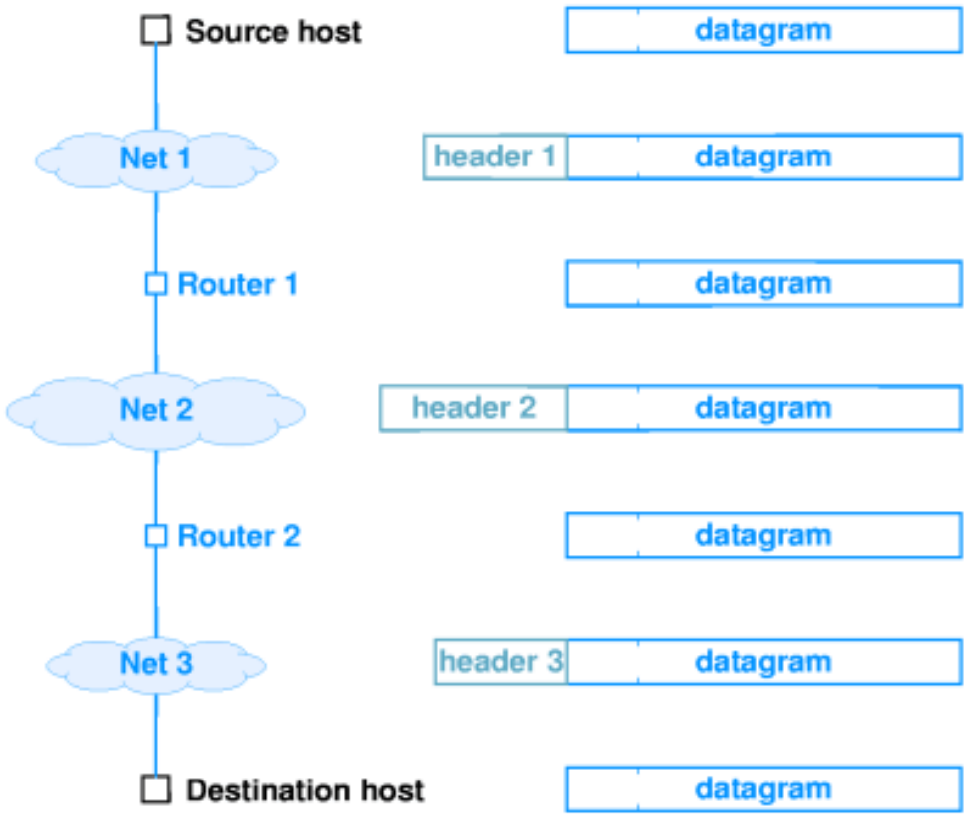

MTU and datagram Size

- Maximum transmission unit (MTU)
  - The maximum amount of data that a frame can carry
- For encapsulation, datagram must be smaller of equal to the network MTU
- Fragmentation
  - An IP router divides a larger datagram into smaller pieces called fragments

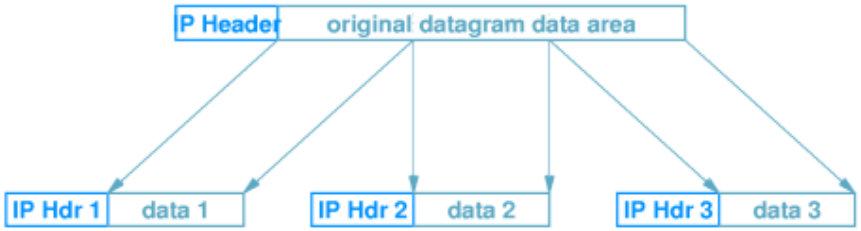

Reassembly

- The process of creating a copy of the original datagram from fragments
- All fragments have the same destination address as the original datagram
- The ultimate destination host reassembles fragments because
  - It reduces the amount of state information in routers
  - It allows routers to change dynamically

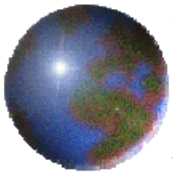

Identifying a datagram

- IP does not guarantee delivery because
  - Individual fragments can be lost
  - Fragments can arrive out-of-order
- Sender places a unique identification number in each outgoing datagram
- Receiver uses the identification number and IP source address to determine the datagram to which the fragment belongs
- The FRAGMENT OFFSET field tells a receiver how to order fragments within a given datagram

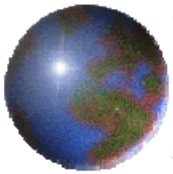

Fragment Loss

- An encapsulated datagram or fragment can be lost or delayed
- The receiver holds fragments until all frames arrive
- IP specifies a maximum time to hold fragments
- IP's reassembly timer is all-or-nothing
  - Either all frames arrive, or
  - IP discards complete datagram
- It is possible to further
  - Fragment a fragment
- IP does not distinguish between original segments and sub-fragments

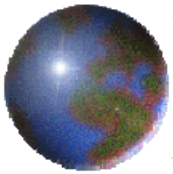

The Success of IP

- The current version is successful because
  - It handles heterogeneous networks
  - It accommodates changes in hardware technology
  - It handles extreme increases in scale
- The motivation for change is
  - Limited address space (only 32-bits)
  - Service for new internet applications (audio and video)
  - More complex addressing and routing capabilities

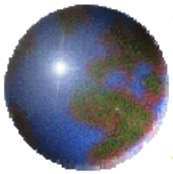

IPv6

- Current IP version is IPv4
- New version became known as IPv6
- IPv6 retains many design features of IPv4
  - Connectionless
  - Basic datagram features like destination address, independent routing and maximum number of hops
- The new features of IPv6 are
  - Address size: Each IPv6 address contains 128 bits
  - Header format: Completely different from IPv4

IPv6 (Cont.)

- Extension headers: IPv6 encodes information into separate headers
- Support for audio and video: IPv6 includes a mechanism that allows a sender and receiver to establish a high-quality path
- Extensible protocol: IPv6 does not specify all possible protocol features

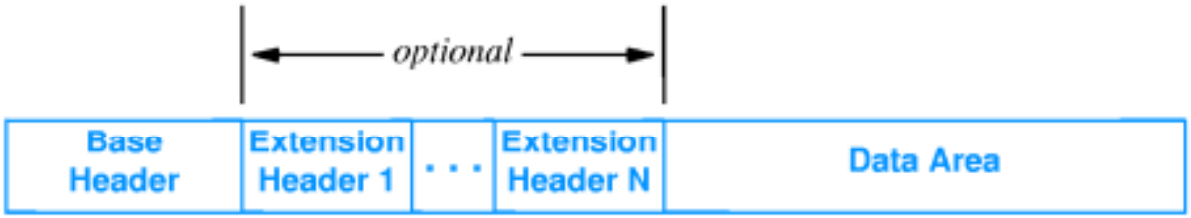

IPv6 Base Header Format

- Twice as large as an IPv4 header but contains less information

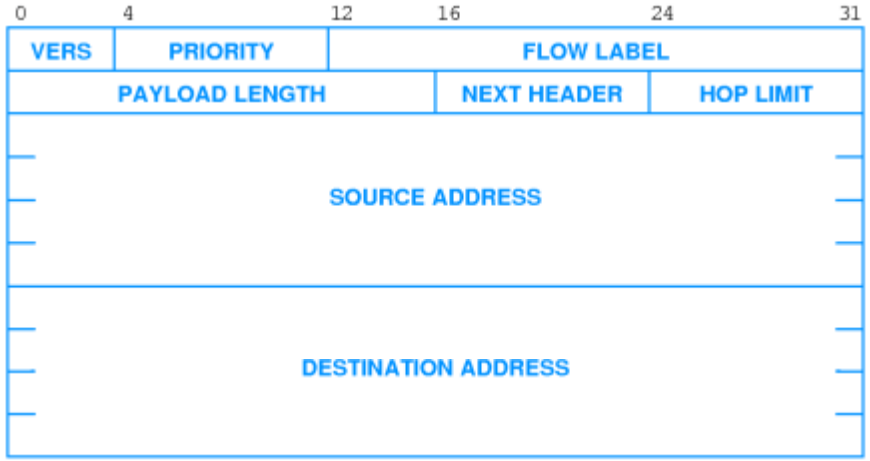

(a)

(b)

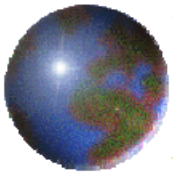

Ipv6 and Multiple Headers

- The standard specifies a unique value for each possible header type
- A receiver uses the NEXT HEADER field to determine what follows
  - If value corresponds to data, the receiver passes the datagram to software
- IPv6 software knows the end of header because
  - Some header types have fixed size
  - For variable size extension headers, the header must contain sufficient information to determine where the header ends

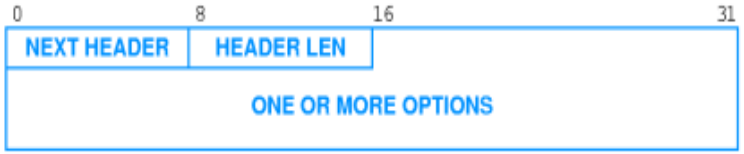

Fragmentation, Reassembly and Path MTU

- IPv6 places fields in a separate fragment extension header
- The presence of a header identifies the datagram as a fragment
- A sending host is responsible for fragmentation
- The host learns the MTU along the path to the destination

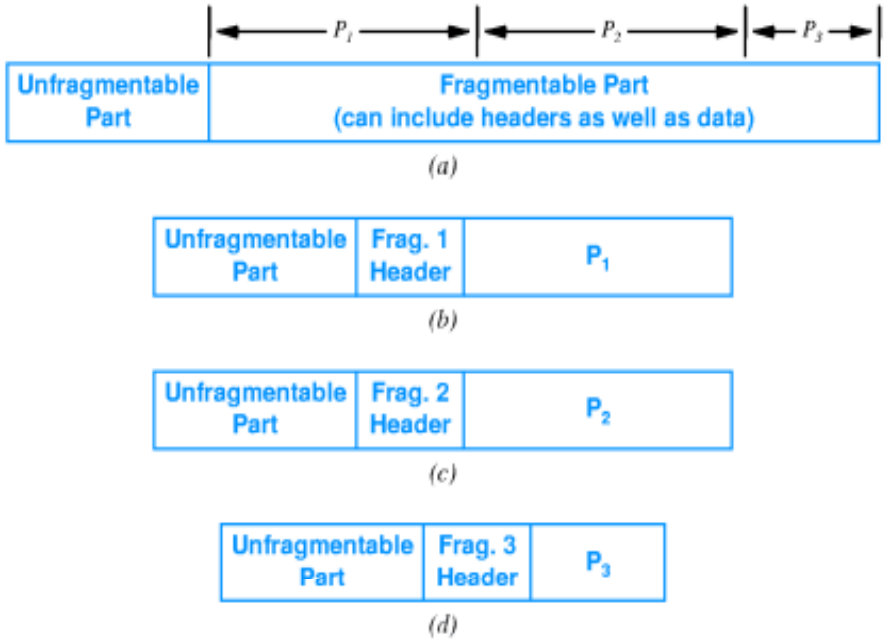

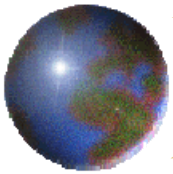

Purpose of Multiple Headers

- Ipv6 uses separate extension headers because it
  - Improves economy and extensibility
- Partitioning the datagram functionality into separate headers saves space
- Reducing datagram size also reduces the bandwidth consumed
- Extensibility: Adding a new feature to a protocol
- Existing protocol headers can remain unchanged
- A new NEXT HEADER type is defined as well as a new header format

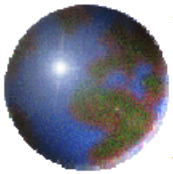

Ipv6 Addressing

- IPv6 assigns a unique address for each connection between a computer and a physical router
- Addresses do not have defined classes
- Each IPv6 address is one of three basic types
- UNICAST: corresponds to a single computer
- MULTICAST: corresponds to a set of computers, possibly at many locations
- ANYCAST: corresponds to a set of computers that share a common address prefix

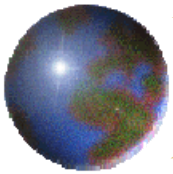

Ipv6 Colon Hexadecimal Notation

- Colon hexadecimal notation
  - A compact syntactic form in which each group of 16 bits is written in hexadecimal with a colon separating groups
- Ex: 69DC:8864:FFFF:FFFF:0:1280:8C0A:FFFF
- Zero compression
  - Replaces sequences of zeros with two colons
- Ex : FFOC:0:0:0:0:0:0:B1 is written as FFOC::B1
- Any IPv6 address that begins with 96 zero bits contains an IPv4 address in the low-order 32 bits

Best-Efforts Semantics and Error Detection

- IP defines a best effort communication service in which datagrams can be
  - Lost, duplicated, delayed or delivered out-of-order
- IP attempts to avoid errors and to report problems
- A header checksum is used to detect transmission errors
- Checksum is verified to ensure that the header arrived intact
- If a checksum error occurs, datagram must be discarded immediately without further processing

Internet Control Message Protocol

- A protocol that IP uses to send error messages
- IP uses ICMP when it sends an error message, and ICMP uses IP to transport messages

Type	Name
0	Echo Reply
1	Unassigned
2	Unassigned
3	Destination Unreachable
4	Source Quench
5	Redirect
6	Alternate Host Address
7	Unassigned
8	Echo
9	Router Advertisement
10	Router Selection
11	Time Exceeded
12	Parameter Problem
13	Timestamp
14	Timestamp Reply
15	Information Request
16	Information Reply
17	Address Mask Request
18	Address Mask Reply
19	Reserved (for Security)
20-29	Reserved (for Robustness Experiment)
30	Traceroute
31	Datagram Conversion Error
32	Mobile Host Redirect
33	IPv6 Where-Are-You
34	IPv6 I-Am-Here
35	Mobile Registration Request
36	Mobile Registration Reply
37-255	Reserved

List of ICMP messages

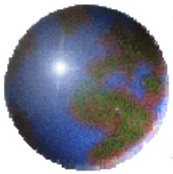

ICMP Error Messages

- Source quench
  - Sent by a router that has no more buffer spaces available
- Time exceeded
  - Send by router if it reduces the TIME TO LIVE field to zero
  - Sent by host if reassembly timer expires
- Destination unreachable
  - Router determines that a datagram cannot be delivered to its final destination
- Redirect
  - Router asks the host to change its route
- Parameter problem
  - One of the parameters specified is incorrect

ICMP Message Transport

- ICMP message is placed in the data area of the IP datagram
- ICMP messages are created in response to a datagram
- Either the datagram has a problem or it carries an ICMP request message to which a router replies
- If ICMP error message causes an error, no error message is sent

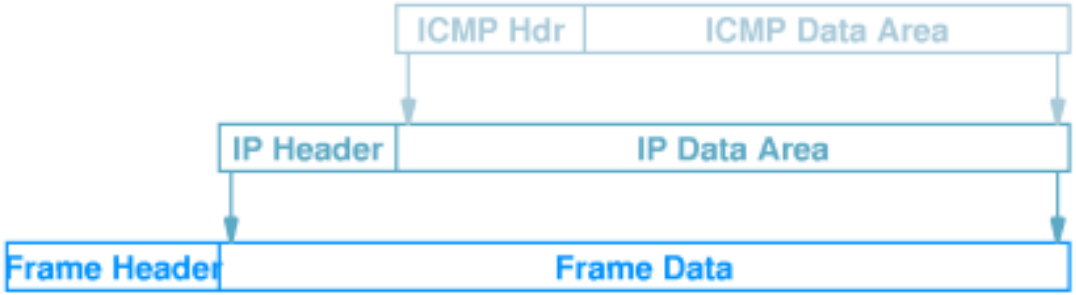

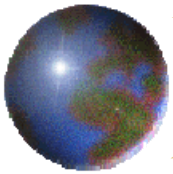

Using ICMP Messages

- Ping uses the ICMP echo request and echo reply
- traceroute uses the ICMP to construct a list of all routers along a path to a given destination
- Traceroute sets the TIME TO LIVE values to extract the IP address of the routers
- Traceroute faces many problems
  - Datagrams can be lost, duplicated or delivered out-of-order
  - Routers can change dynamically
- Traceroute uses UDP(User Defined Protocol) when TIME TO LIVE is large enough to reach the destination host
- Path MTU can be determined from ICMP error messages and then datagram size is fixed.