Computer Network Programming

IP Overview

Dr. Sam Hsu
Computer Science & Engineering
Florida Atlantic University

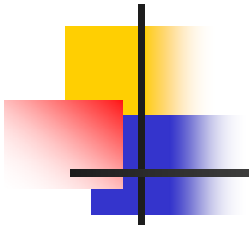

IP Overview

- Introduction
- IP Datagram Format
- Type of Service
- Datagram Fragmentation
- IP Options
- IP Checksum
- IP Addresses
- Domain Name System
- IP Routing
- Mapping Addresses

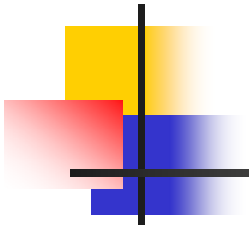

Internet Protocol (1/2)

- IP (Internet Protocols) is the workhorse protocol of the TCP/IP protocol suite
  - TCP/IP is said to be an IP-based technology.
- IP provides an unreliable, connectionless datagram delivery service
  - Unreliable since there is no guarantee that an IP datagram will get to its destination successfully.
  - However, IP will make an earnest attempt to deliver packets, known as a *best-effort* service.

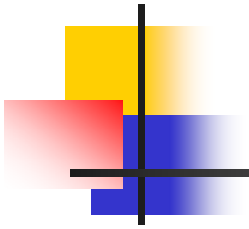

Internet Protocol (2/2)

- IP has three important functionalities:
  - Specification of packet formats.
  - Routing.
  - Packet delivery.

Datagram Format

General form of an IP datagram

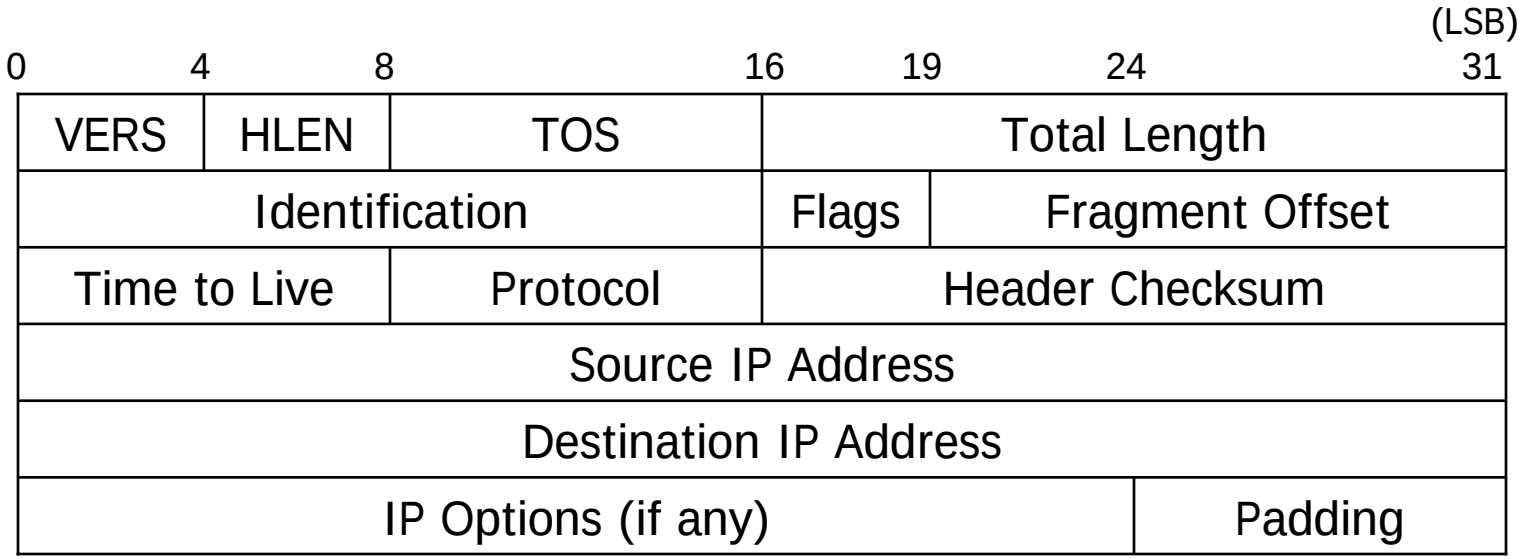

Format of an IP datagram header

Type of Service (TOS) (1/3)

Original subfields of the TOS field

D: Minimize delay
T: Maximize throughput
R: Maximize reliability

Type of Service (TOS) (2/3)

- The TOS field was redefined By the IETF in 1998 as the *differentiated services* (DS) field.

Structure of the DS field

DSCP: differential services codepoint

Type of Service (TOS) (3/3)

- The DSCP subfield is divided into three administrative groups, known as *pools*, for the purpose of codepoint assignment and management.

Pool	Codepoint	Assignment Policy
1	xxxxx0	Standard action
2	xxxx11	Experimental or local use
3	xxxx01	Experimental or local use (for now)

The three administrative pools of DSCP
(***x*** refers to either ‘0’ or ‘1’)

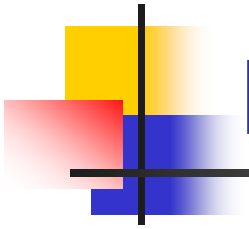

Datagram Fragmentation (1/2)

- An IP datagram may be fragmented due to limitations of the underlying network.
  - Different networks have different *maximum transfer unit*s (MTUs)
    - MTUs are in bytes.

Datagram Fragmentation (1/3)

- An IP datagram may be fragmented due to limitations of the underlying network.
  - Different networks have different *maximum transfer unit*s (MTUs)
    - MTUs are in bytes.
  - An example

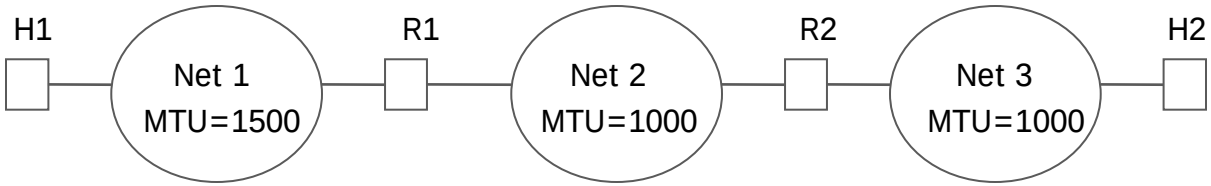

Datagram Fragmentation (2/3)

- Each fragment has the same format as the original datagram.
- All fragments travel as separate datagrams all the way to the ultimate destination where they will be reassembled.
- An example

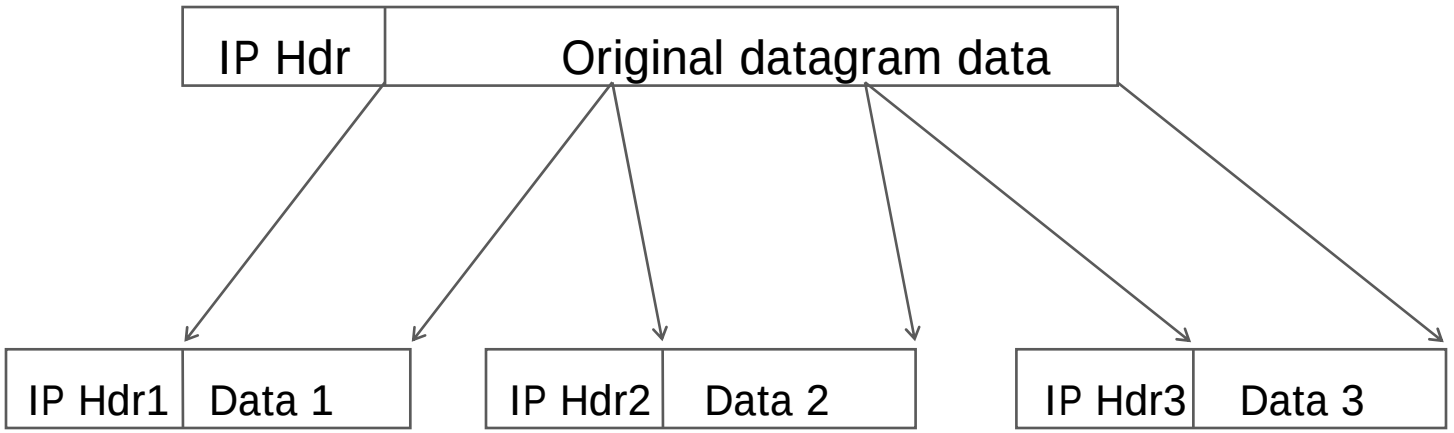

Datagram Fragmentation (3/3)

- MTUs for some networks

Protocol	MTU
Hyperchannel	65,535
Token ring (16 Mbps)	17,914
Token ring (4 Mbps)	4,464
FDDI	4,352
Ethernet	1,500
X.25	576
PPP	296

Ref: *TCP/IP Protocol Suite*, 2nd ed., Behrouz Forouzan, McGraw Hill, 2003, p. 199

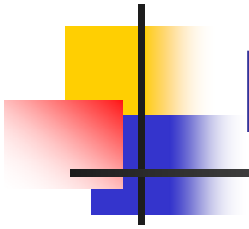

Flags

- There are three subfields in the *Flags* field.

U	DF	MF

U: Unused

DF: Do not fragment

MF: More fragments

IP Options (1/3)

- Used primarily for network testing and debugging.
- Length varies depending on options selected.
- Zero bits padded at the end to ensure header length to be in multiples of 32 bits.

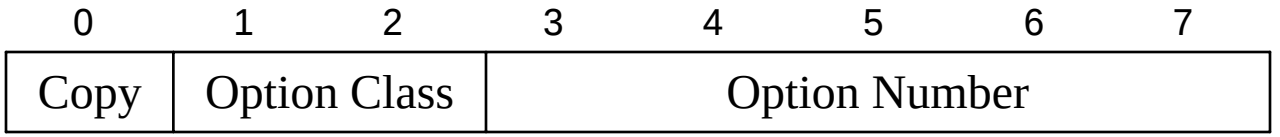

The three subfields of an IP option code octet

IP Options (2/3)

Copy	Meaning
0	Copy option to 1st fragment only
1	Copy option to all fragments

Meanings of the copy bit

Option Class	Meaning
00	Datagram or network control
01	Reserved for future use
10	Debugging and measurement
11	Reserved for future use

IP option classes

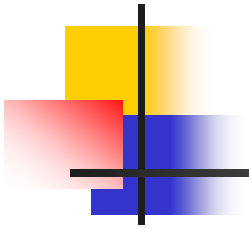

IP Options (3/3)

Option Number	Meaning
00000	End of option
00001	No operation
00011	Loose source route
00100	Timestamp
00111	Record route
01001	Strict source route

Some IP option numbers

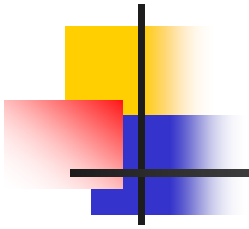

IP Checksum

- Is used to ensure integrity of header values.
  - Formed by treating the header as a sequence of 16-bit integers, in network standard byte order, adding them together using 1's complement arithmetic, and then taking 1's complement of the result.
    - Network standard byte order uses *big-endian* in sending integers; that is, the most significant byte in an integer is sent first.
- Error checking for the data area is left to the higher layer protocol.

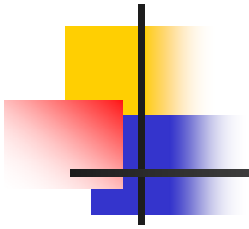

Internet Addresses (1/2)

- A means to identify hosts on the Internet.
- Also known as IP addresses.
- Each host on the Internet is assigned a 32-bit unique address
  - An IP address is assigned to a single host only.
  - A host may have more than one IP address (multi-homed host).

Internet Addresses (2/2)

- Five classes of IP addresses:

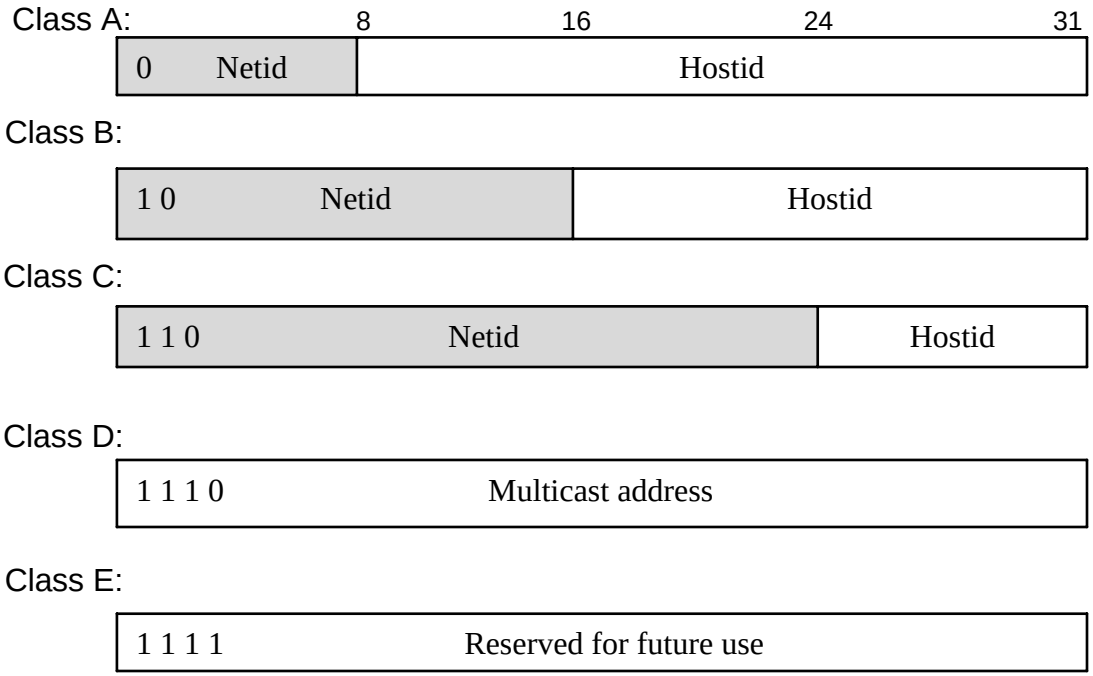

The five classes of IP addresses

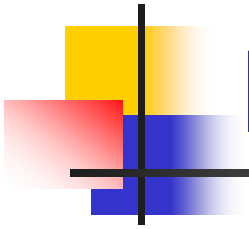

Dotted Representation

- Internet addresses are represented in the form of four integers separated by decimal points.
  - For readability by human.
  - For example:
    - FAU machine = 131.91.*deptid.hostid*

      131.91.128.74 (FAU Web Server)

      131.91.96.106 (FAU CSE Dept.)

The Domain Name System (1/2)

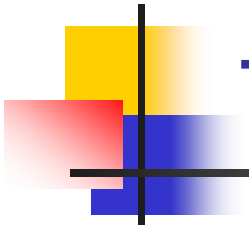

- A high-level naming scheme.
- A sequence of characters grouped into sections delimited by decimal points.
  - Labeled in a meaningful way
  - FAU machine = *hostid.deptid.*fau.edu
- A hierarchical naming system
  - Written in the local-most level first and the top-most level last fashion.

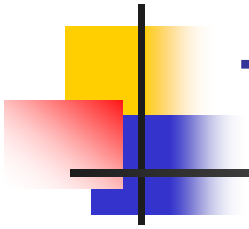

The Domain Name System (2/2)

- Top-level domain names
  - Generic domains: for use by the Internet public (.**aero, .biz, .cat, .com, .coop, .edu, .gov, .info, .jobs, .mobi, .int, .mil, .museum, .name, .net, .org, .pro,** & **.travel**).
  - Country code domains: for use by each individual country (**.uk**, **.de**, **.jp**, **.us**, etc.)
  - Infrastructure domain (**.arpa**)

Ref: http://www.iana.org/domain-names.htm
http://www.norid.no/domreg.html

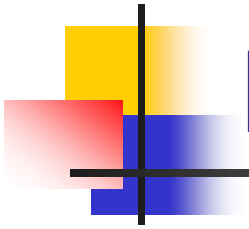

More on IP Addresses (1/3)

- An IP address encodes both a network and a host on that network. It does not specify an individual machine, but a connection to a network.
  - Referred to as a network connection.
- If a host moves from one network to another, its IP address must change.
  - A weakness in Internet addressing.
- For a multi-homed host, knowing one address of it may not be sufficient to reach it when some network(s) are unavailable.
  - Another weakness in Internet addressing.

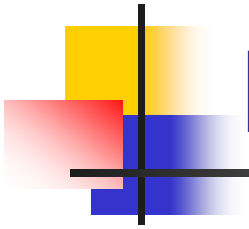

More on IP Addresses (2/3)

- An IP address can refer to a network as well as to an individual host.
- By convention, the network address has *hostid* with all bits ***0***, and a broadcast address has *hostid* with all bits ***1***
  - All ***0***s mean “this”.
  - All ***1***'s mean “all”.
- Class A network address ***127.x.x.x*** is reserved for loopback, and is designed for testing and IPC on the local machine.

More on IP Addresses (3/3)

- Some special address conventions

Netid	Hostid	Types of Address	Purpose
all 0s	all 0s	source host	used during bootstrap
all 0s	specific	destination host	a specific host on local net
specific	all 0s	network	identifies a network
specific	all 1s	directed broadcast	broadcast on specified net
all 1s	all 1s	limited broadcast	broadcast on local net
127	any	loopback	testing

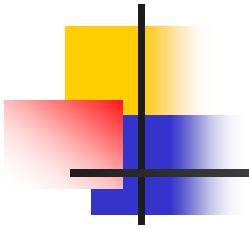

Internet Addressing Authorities

- *Netid*s may be obtained either via a domain name registrar authorized by the *Internet Corporation for Assigned Names and Numbers* (ICANN), or an Internet service provider (ISP).
  - A service used to be handled by *the Internet Assigned Numbers Authority* (IANA).
  - Had been handled by Jon Postel until late 1998.
- *Hostid*s are assigned by the local authority.

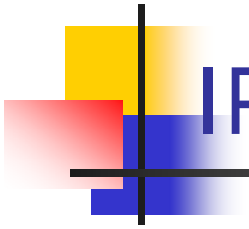

IP Address for Private Networks (1/2)

- What IP addresses should an organization choose for internal use only?
  - An organization may like to use TCP/IP protocols internally only without being connected to the Internet.
- Three options:
  - To apply for a unique network address for internal use only.
    - Problem: Classes A & B addresses are hard to get today.
  - To select an IP address arbitrarily.
    - Problem: What if the organization decides to get connected to the Internet later?

IP Address for Private Networks (2/2)

- To use one of the following addresses, officially assigned to private networks:

Address ranges		
10.0.0.0	to	10.255.255.255
172.16.0.0	to	172.31.255.255
192.168.0.0	to	192.168.255.255

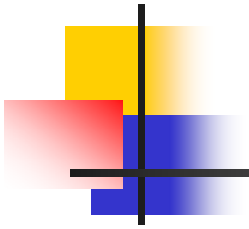

IP Routing (1/4)

- Routing refers to the process of selecting a path over which to send packets to the destination.
  - If both source and destination attach directly to the same physical network,
    - Sender maps the destination IP address onto its physical address and uses the network hardware to deliver the datagram directly.
  - Else if source and destination are on two different physical networks,
    - Sender needs to pass the datagram to a router for delivery.

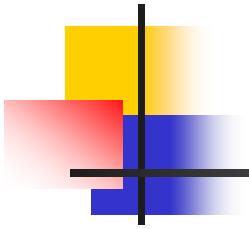

IP Routing (2/4)

- Datagrams pass from router to router until they reach a router that can deliver the datagram directly.
- Routing also applies between different subnets of the same network.
- IP routing is table-driven.

IP Routing (3/4)

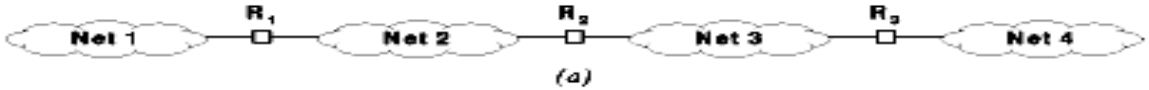

(a)

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

(b)

(a) An example internet with three routers connecting four physical networks, and (b) the conceptual routing table found in router R2. Each entry in the table lists a destination network and the next hop along a route to that network.

Ref: *Computer Networks and Internets, 2nd ed.*, Douglas Comer, Prentice Hall, 1999, p. 243.

IP Routing (4/4)

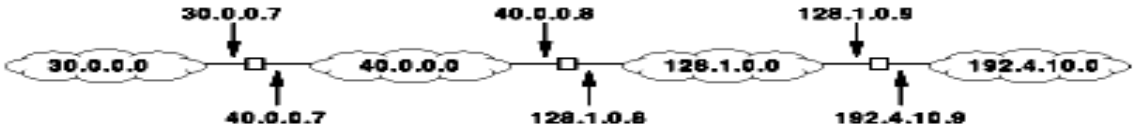

(a)

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)

(a) An internet of four networks and three routers with an IP address assigned to each router interface, and (b) the routing table found in the center router. Each entry in the table lists a destination, a mask, and the next hop used to reach the destination .

Ref: *Computer Networks and Internets, 2nd ed.*, Douglas Comer, Prentice Hall, 1999, p. 271.

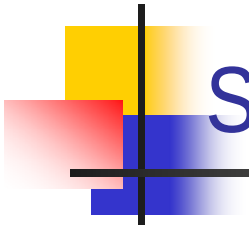

Some Routing Table Terminology (1/2)

- Next-hop routing
  - The routing table holds only the address of the next hop (router), instead of the complete path to the destination.
- Network-specific routing
  - Only the network address (*netid*) is entered into the routing table.
    - All hosts on the same network have only one entry in the table

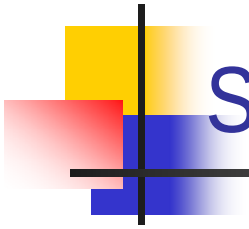

Some Routing Table Terminology (2/2)

- Host-specific routing
  - The complete IP address, down to *hostid*, for the destination host is given in the table.
  - To have a better control over routing; may be for security, debugging purposes, etc.

Default Routing

- A catchall mechanism that matches *none of the above*.
- Must be the very last entry in the table.
  - An example:

    The following routing table in router R_1, using the first network configuration *(a)* on **Slide 31**.

Destination	Next hop
Net_1	Deliver Directly
Net_2	Deliver Directly
Net_3	R_2
Default	R_2

IP Routing Algorithm

Algorithm
RouteDatagram (Datagram, RoutingTable)
Extract destination IP address, *D*, from the datagram and compute the network prefix, *N*;

If *N* matches any directly connected network address, deliver datagram to destination *D* over that network (This involves resolving *D* to a physical address, encapsulating the datagram, and sending the frame.)
else if the table contains a host-specific route for *D*, send datagram to next-hop specified in table
else if the table contains a route for network *N*, send datagram to next-hop specified in table
else if the table contains a default route, send datagram to the default router specified in table
else declare a routing error.

- The algorithm IP uses to forward a datagram. Given an IP datagram and a routing table, this algorithm selects the next hop to which the datagram should be sent. All routes must specify a next hop that lies on a directly connected network.

Ref: *Internetworking with TCP/IP Volume I*, 4th ed., Douglas Comer, Prentice Hall, 2004, p. 122.

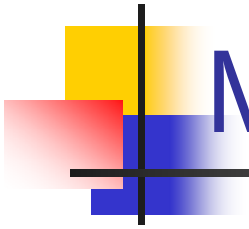

Mapping Addresses (1/2)

- Mapping domain names to IP addresses
  - Between high-level domain names and low-level IP addresses.
  - For local host itself
    - Done at system startup.
  - For hosts within the same subdomain
    - Local table lookup.
  - Others
    - May need to go through a nameserver(s).

Mapping Addresses (2/2)

- Mapping IP addresses to physical addresses
  - Application programs always use IP addresses.
  - Hosts and routers use physical network addresses to deliver datagrams across underlying physical networks.
  - IP addresses are independent of physical network addresses.
    - Need to map IP addresses to physical addresses, known as the *address resolution* problem.
  - A host obtains its IP and physical addresses at system startup
    - IP address is usually kept in a disk file.
    - Physical address is provided by the network interface card.

ARP/RARP Protocols (1/3)

A host uses the Address Resolution Protocol (ARP) to learn the physical addresses of other machines.

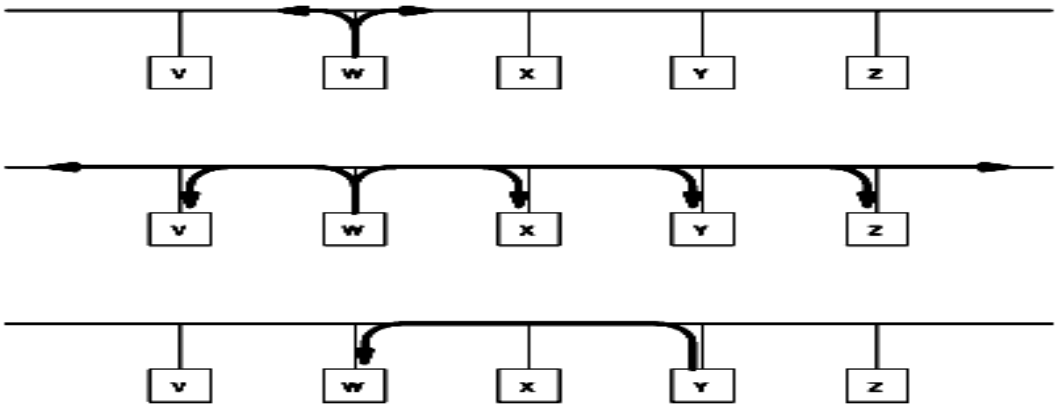

An ARP message exchange. (a) Computer *W* begins to broadcast an ARP request that contains computer *Y*'s IP address. (b) all computers receive the request, and (c) computer *Y* sends a response directly to *W*.

Ref: *Computer Networks and Internets, 2nd ed.*, Douglas Comer, Prentice Hall, 1999, p. 258.

ARP/RARP Protocols (2/2)

- ARP cache timeout
  - An ARP cache contains a set of recently obtained IP-to-physical address bindings.
    - Is used to avoid successive ARP requests/responses for the same destinations.
    - For the purpose of reducing communication cost.
  - Entries in an ARP cache are timed.
    - Each ARP cache entry is associated with a timer, typically 20 minutes.
    - An entry is removed if it has not been referenced before its timer expires.
- A diskless host uses the Reverse Address Resolution Protocol (RARP) to obtain its IP address from its server.

ARP/RARP Protocols (3/3)

ARP/RARP protocol format

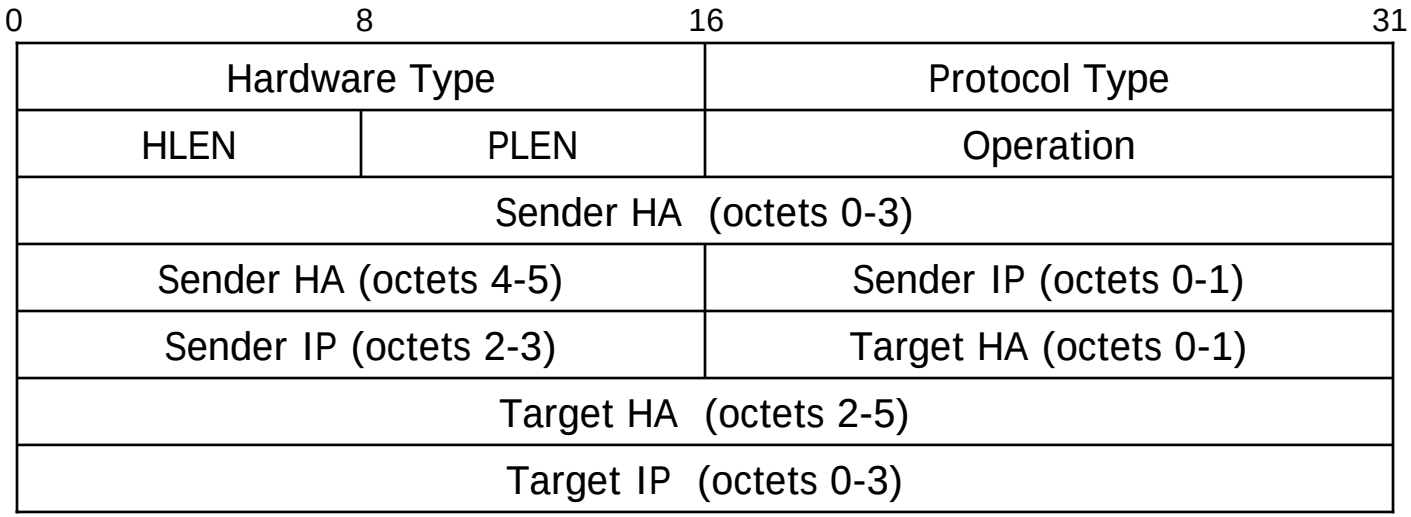

An example of the ARP/RARP message format when used for IP-to-Ethernet address resolution. The length of address fields depends on the hardware and protocol address lengths, which are 6 octets for an Ethernet address and 4 octets for an IP address.

Ref: *Internetworking with TCP/IP Volume I*, 4th ed., Douglas Comer, Prentice Hall, 2000, p. 85.

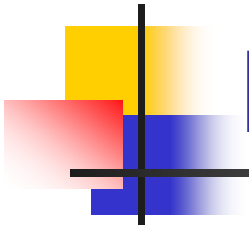

Reading Assignment

- Read Chapters 3-5.

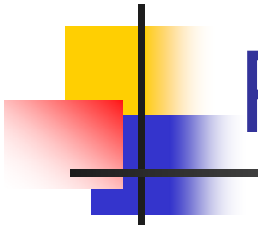

For More Information

- RFC 791 – Internet Protocol, Sep-01-1981.
- RFC 815 – IP datagram reassembly algorithms, Jul-01-1982.
- RFC 894 – Standard for the transmission of IP datagrams over Ethernet networks, Apr-01-1984.
- RFC 1034 – Domains names – concepts and facilities, Nov-01-1987.
- RFC 1042 – Standard for the transmission of IP datagrams over IEEE 802 networks, Feb-01-1988.
- RFC 1071 – Computing the Internet checksum, Sep-01-1988.
- RFC 1624 – Computation of the Internet Checksum via Incremental Update, May 1994.
- RFC 1918 – Address Allocation for Private Internets, February 1996.
- RFC 2474 – Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, December 1998.
- RFC 2517 – Building Directories from DNS: Experiences from WWWSeeker, February 1999.
- RFC 3007 – Secure Domain Name System (DNS) Dynamic Update. November 2000.
- RFC 3172 – Management Guidelines & Operational Requirements for the Address and Routing Parameter Area Domain ("arpa"). September 2001.
- RFC 3232 – Assigned Numbers: RFC 1700 is Replaced by an On-line Database, January 2002.
- RFC 3330 – Special-Use IPv4 Addresses. IANA, September 2002.

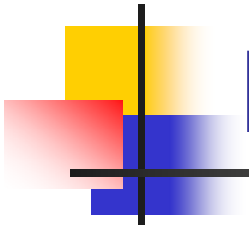

For More Information

- RFC 826 – Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware, Nov-01-1982.
- RFC 903 – Reverse Address Resolution Protocol, Jun-01-1984.
- RFC 1166 – Internet number, Jul-01-1990.
- RFC 1390 – Transmission of IP and ARP over FDDI Networks, January 1993.
- RFC 1433 – Directed ARP, March 1993.
- RFC 1931 – Dynamic RARP Extensions and Administrative Support for Automatic Network Address Allocation, April 1996.
- RFC 2625 – IP and ARP over Fibre Channel, June 1999.
- RFC 2835 – IP and ARP over HIPPI-6400 (GSN), May 2000..