Chapter 10

Network Layer Overview

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

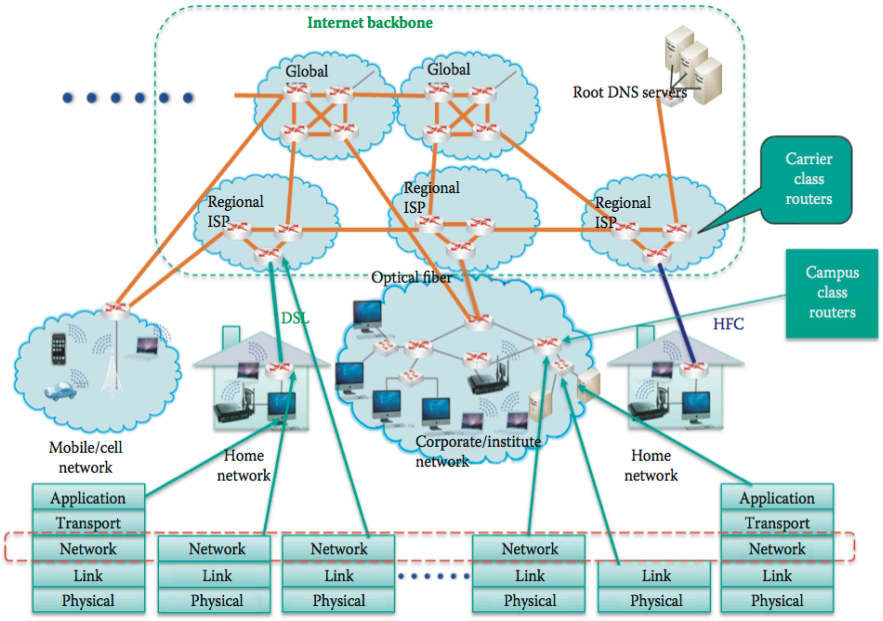

Network Layer

- Route datagram from sending to receiving host
- Source host encapsulates segments (passed down by transport layer) into datagrams
- Destination host delivers segments up to transport layer
- Network layer protocols are built in every host, and router
  - The majority of hosts do not know how to route
  - Client OS does not have routing module
- Router examines header fields in all IP datagrams from one interface and forwards to another interface in accordance with the routing table
  - Routers work together to generate routing tables
- Router understands the network, link and physical layers

Two Key Router Functions

- Routing
  - Generate and maintain routing table
  - Routing algorithms: routers work together to find routes from a subnet to the other subnets
  - Distributed processing motivated by the cold war
  - If there exists a path, routers will put it in the routing table automatically
- Forwarding:
  - A router/layer 3 switch moves a packets from router's input port to appropriate router output in accordance with the routing table

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

Datagram Networks

- Routers: no state information about end-to-end connections
- Packets forwarded using destination host IP address
  - Packets between same source-destination pair may take different paths
  - Packets may not arrive in the original order

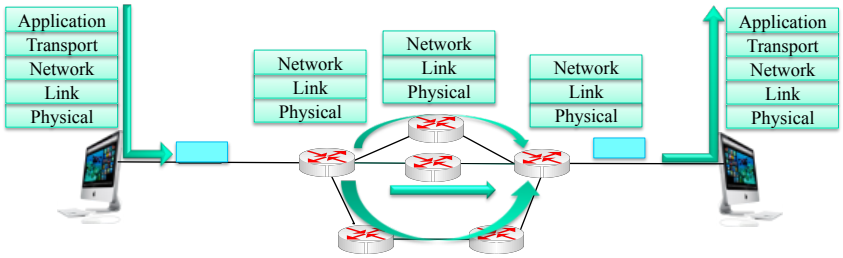

Forwarding Table

Destination Address Range	Router Interface
11001000 11010111 00000000 00000000 through 11001000 11010111 00001111 11111111	0
11001000 11010111 00001000 00000000 through 11001000 11010111 00001000 11111111	1
11001000 11010111 00010000 00000000 through 11001000 11010111 00011111 11111111	2
Default	3

Prefix, or subnet part (network ID + subnet ID)

Host ID

Use Longest Prefix Matching

IP address prefix: 200.215.x.x (C8.D7.x.x)

Prefix for router Interface	Router Interface
11001000 11010111 0000	0
11001000 11010111 00001000	1
11001000 11010111 0001	2
otherwise	3

- Examples
  - Router uses the longest matching prefix in the routing table to select the output interface
    - Destination IP Address: 11001000 11010111 00001000 10100001
      - Route to Interface 1
    - Destination IP Address: 11001000 11010111 00000000 10101010
      - Route to Interface 0

Routing Table Example

Prefix for router Interface	Router Interface
200.215.0/20	0
200.215.8/20	1
200.215.16/24	2
otherwise	3

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

The Internet Network Layer

- Router and network layer functions:

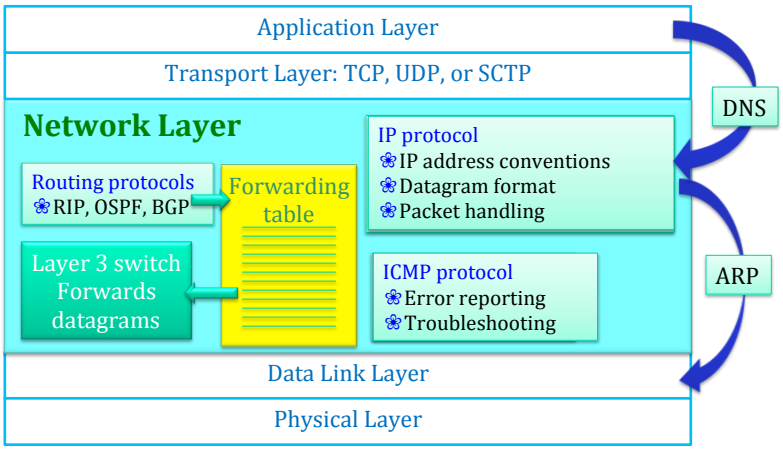

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

IPv4 Header - Context

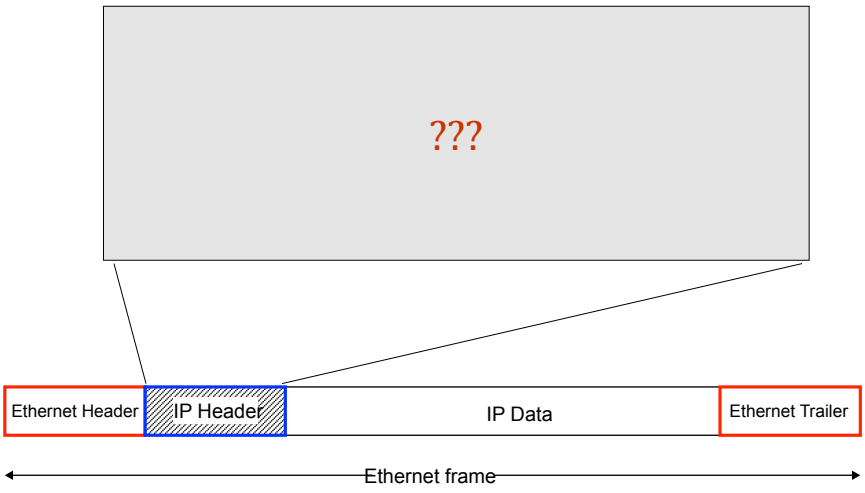

IPv4 Datagram Format

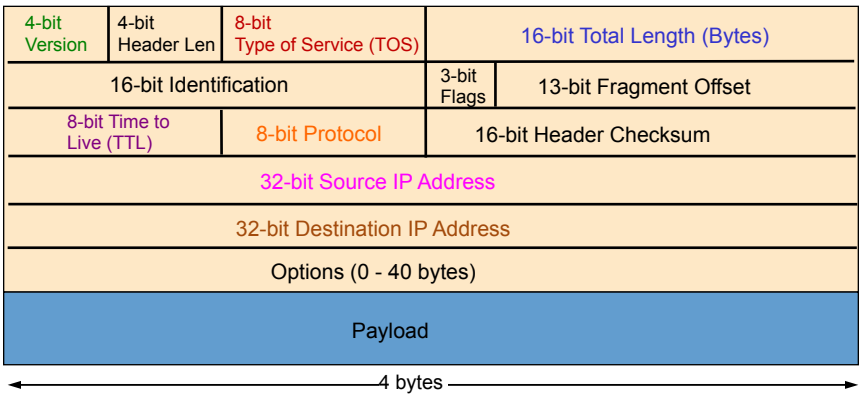

20 bytes ≤ Header Len < 2^4 x 4 bytes = 60 bytes

20 bytes ≤ Total Length < 2^{16} bytes = 65536 bytes

The purpose for each IP header field

IP header field	Description
The *IP protocol version number*	It is 4 bits long and 0100 for IPv"4".
The 4-bit *header length*	The number of 32-bit words in the header (normally 5)
8-bit *type of service*	e.g. voice over IP in order to prioritize traffic.
The *datagram length in the number of bytes*	Including header and data, comprises the remaining 16 bits.
The *ID* and *Flag* fields	They are used to identify and control fragments.
The *fragment offset*	It is used for fragmentation/reassembly.
The *time to live*	It is the allowed maximum number of remaining hops, decremented at each router.
The *upper layer protocol*	It is the one used in the payload of the IP datagram, e.g. TCP has a value of 6 as shown in Table
The *header checksum*	It does error checking on the header.
The options	It follows the *source* and *destination addresses*. The options can be used for source route, i.e. specifying the IP address of each hop, and route record, i.e. recording the IP address of each hop. But they are rarely used for security reasons.

Protocol field

Protocol	Decimal	Hex
TCP	06	0x06
UDP	17	0x11
SCTP	132	0x84
ICMP	01	0x01
IGMP	02	0x02
L2TP	115	0x73
OSPF	89	0x59
ESP	50	0x32
AH	51	0x33

IP Fragmentation & Reassembly

- Network links have MTU (Maximum Transmission Unit)
  - Token ring: 5000 bytes
  - Ethernet: 1500 bytes
- Large IP datagram divided (fragmented) by a router
  - One datagram becomes several datagrams
  - Datagrams are reassembled only at final destination (end computer)
  - IP header bits used to identify, order related fragments

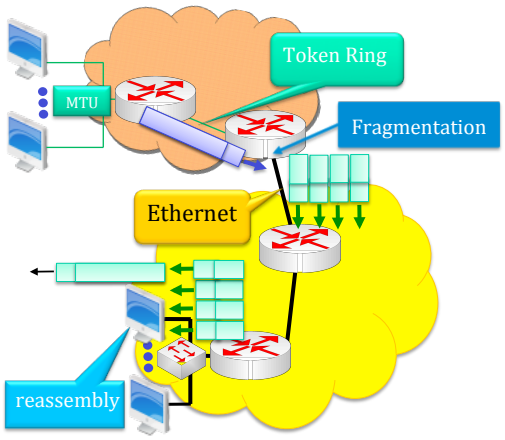

IP Fragmentation and Reassembly (1)

- Example
  - 5000 byte datagram
    (IP header = 20, payload = 4980)
  - MTU = 1500 bytes
    (IP header = 20, payload = 1480)

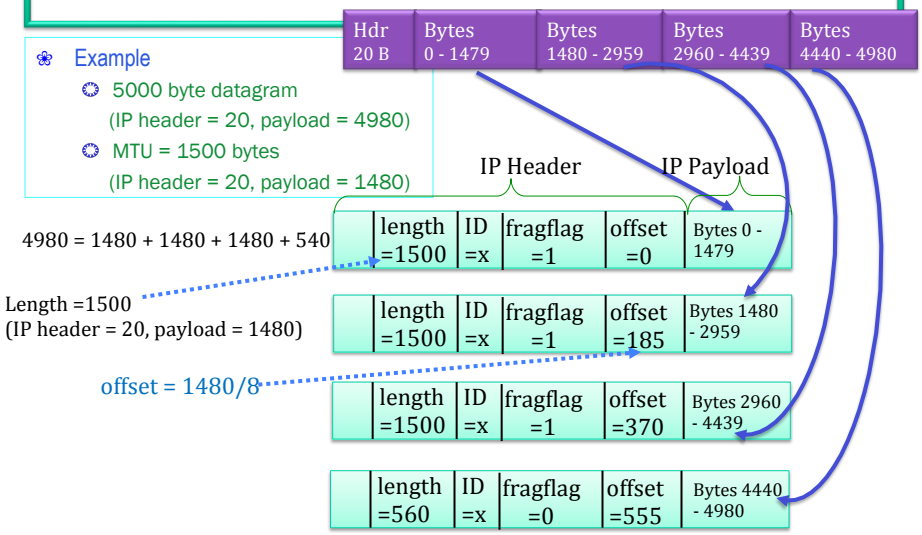

IP Fragmentation and Reassembly (2)

16-bit Identification	0	DF	MF	13-bit Fragment Offset

- The ID is the same in all 4 fragments
- The offset value and the flags used to reassemble the original 5000-byte datagram
- A three-bit Flag field is used to control or identify fragments. They are (in order, from high order to low order):
  - Bit 0: Reserved; must be zero
  - Bit 1: Do not Fragment (DF)
  - Bit 2: More Fragments (MF)
- When a packet is fragmented all fragments, with the exception of the last fragment, have the MF flag set (binary 001)
- If the DF flag is set and fragmentation is required to route the packet, then the datagram will be dropped
  - This is useful when sending packets to a host that does not have sufficient resources to handle fragmentation

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

IPv4 Addressing

- IP address:
  - 32-bit identifier (IPv4) for host, or router interface
- NIC interfaces
  - Router typically has multiple interfaces (at least 2 interfaces)
  - Host typically has one interface
  - One IP address and one MAC address associated with each interface
  - Note: a layer 2 switch has no interface

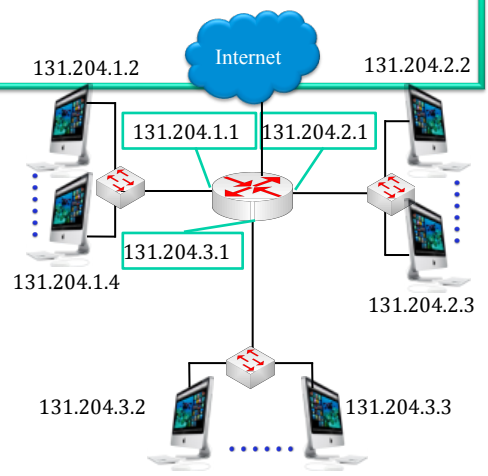

131.204.1.1 = $\underbrace{10000011}_{131}\ \underbrace{11001100}_{204}\ \underbrace{00000001}_{1}\ \underbrace{00000001}_{1}$

Subnet

- Each IP address has:
  - Subnet part (high order bits): aka prefix
  - Host part (low order bits)

  131.204.1.2

- Subnet
  - Hosts that have interfaces with the same subnet part of IP address
  - Hosts can communicate with each other without router

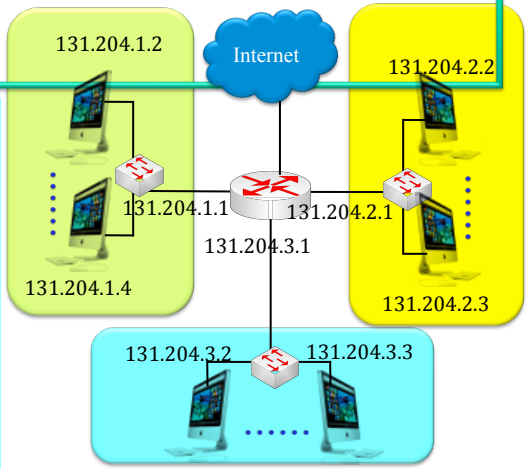

A network consisting of 3 subnets

Subnet

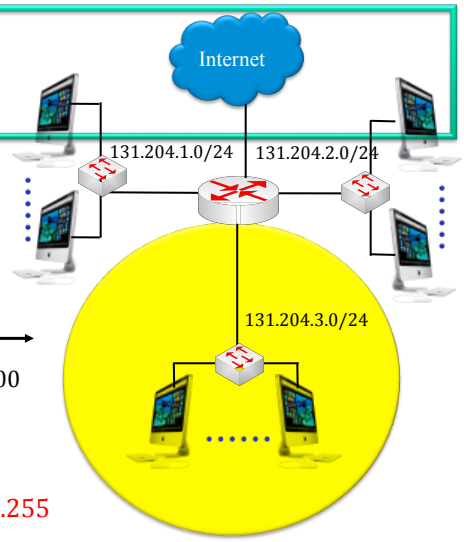

Subnet address: 131.204.3.0/24
Subnet mask: 255.255.255.0

Subnet Part or prefix ← host part →

10000011 11001100 00000011 00000000

Subnet broadcast address: 131.204.3.255

Subnet Part or prefix ← host part →

10000011 11001100 00000011 11111111

Subnets

- 8 subnets
  - 131.204.1.0/24
  - 131.204.2.0/24
  - 131.204.3.0/24
  - 131.204.4.0/24
  - 131.204.5.0/30
    - 2 gateways
  - 131.204.6.0/30
    - 2 gateways
  - 131.204.7.0/30
    - 2 gateways
  - 131.204.8.0/30
    - 2 gateways

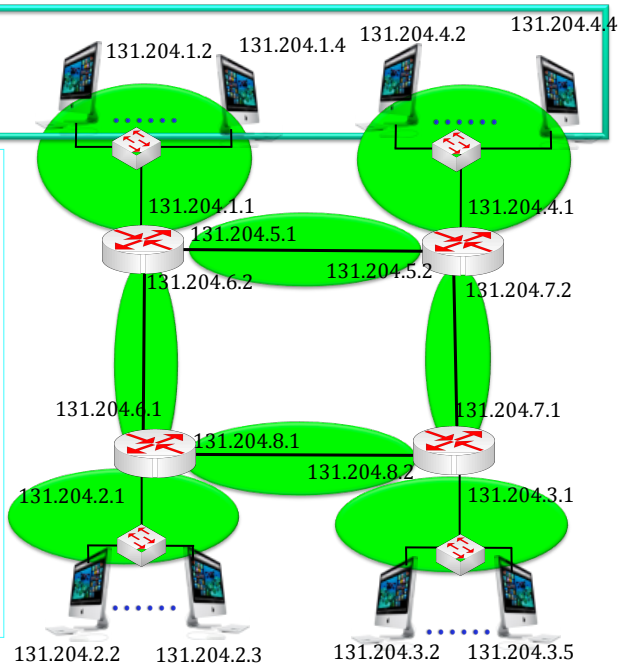

Class A, B, C, and D

- In the olden days, only fixed allocation sizes
  - Class A: 0*
    - Very large /8 blocks (e.g., MIT has 18.0.0.0/8)
  - Class B: 10*
    - Large /16 blocks (e.g,. Princeton has 128.112.0.0/16)
  - Class C: 110*
    - Small /24 blocks (e.g., AT&T Labs has 192.20.225.0/24)
  - Class D: 1110*
    - Multicast groups

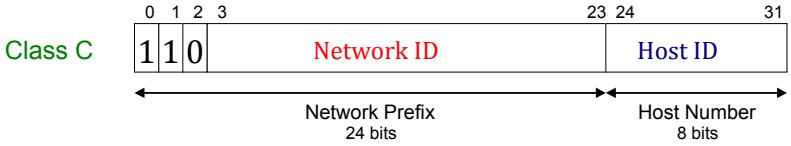

Class A, B, C, and D

- Class A
  - Network ID: 1 byte, MSBit = 0
  - 0<MSByte<127
- Class B
  - Network ID: 2 bytes, MSBits = 10
  - 128≤MSByte≤191
- Class C
  - Network ID: 3 byte, MSBits = 110
  - 192≤MSByte≤223
- Class D
  - MSBits= 1110
  - 224≤MSByte≤239 (Multicast)

Auburn University IP address:

131.204.121.1 ~ 254

Network ID | Subnet ID | Host ID

Subnet part | Host part

Auburn University:

- 131.204.255.255 (broadcast address to every IP address in auburn.edu)
- 131.204.0.0 (refer to Auburn network)

ICANN

- ICANN (Internet Corporation for Assigned Names and Numbers)
  - Allocates IP addresses
  - Manages DNS
  - Assigns domain names, and resolves disputes
- IPv4 address assigned list:
- http://www.iana.org/assignments/ipv4-address-space
  - 008/8 Level 3 Communications, Inc. 1992-12
  - 009/8 IBM 1992-08
  - 012/8 AT&T Bell Laboratories 1995-06
  - 013/8 Xerox Corporation 1991-09
  - 015/8 Hewlett-Packard Company 1994-07
  - 017/8 Apple Computer Inc. 1992-07
  - 018/8 MIT 1994-01

Special Address

- 127.0.0.0 through 127.255.255.255 for loopback purposes
  - Localhost
  - The adapter/NIC intercepts all loopback messages and returns them to the sending application
- Zero Addresses
  - As with the loopback range, the address range from 0.0.0.0 through 0.255.255.255 should not be considered part of the normal Class A range.
  - 0.x.x.x addresses serve no particular function in IP, but nodes attempting to use them will be unable to communicate properly on the Internet
  - 0.0.0.0 means any IP address

Private IP Addresses

- Private Addresses
  - Relieve the shortage of IPv4 addresses
  - The IP standard defines specific address ranges within Class A, Class B, and Class C reserved for use by private networks
  - The table below lists these reserved ranges of the IP address space
- Hosts are effectively free to use addresses in the private ranges if they are not connected to the Internet
- May reside behind firewalls or other gateways that use Network Address Translation (NAT)
- Private IP addresses are blocked by an ISP router's firewall

Class	Private start address	Private finish Address
A	10.0.0.0	10.255.255.255
B	172.16.0.0	172.31.255.255
C	192.168.0.0	192.168.255.255

Classless InterDomainRouting (CIDR)

- Eliminating class limitations of a network ID
- Subnet portion of the address of an arbitrary length
- Address format: a.b.c.d/x, where x is # of bits in subnet part of IP address
- Subnet mask: use all 1's for the subnet part and 0's for the host part

Subnet Part or prefix ← → **host part** ← →

10000011 11001100 0000001 000000001

131.204.2.0/23

- The number of hosts that can be in this subnet: $2^9 - 2$ hosts

CIDR Benefits

- Dramatically reduce the size of routing tables in Internet core routers
  - Most organizations received multiple class C addresses, which are continuous
  - Multiple entries become a single entry in routing tables for one organization
- CIDR is the representation used for configuring network equipment, such as a router and firewall

Aggregate Multiple Class C Nets

- Continuous Class C: 193.1.0.0/24 to 193.1.3.0/24
  - 0: 00000000
  - 3: 00000011

```
193.1. 00000000 00000001
…………………………………………………………………………. .
193.1. 00000011
```

- CIDR: 193.1.0.0/22: 10-bit host ID (freedom bits)
  - 2^{10} – 2 hosts (can be used to assign subnet ID and host ID)
  - Only one entry in the routing table
- Subnet mask 255.255.11111100.0 (255.255.252.0)
  - (replace host ID portion by 0's and subnet part by 1's)
- Broadcast address
  - 193.1.00000011.11111111 = 193.1.3.255 (replace host ID part by 1's)

Supernetted Subnet

- If about 1000 IP addresses are required, 4 Class C networks could be supernetted together:
  - 192.60.128.0 (11000000.00111100.10000000.00000000) Class C subnet address
  - 192.60.129.0 (11000000.00111100.10000001.00000000) Class C subnet address
  - 192.60.130.0 (11000000.00111100.10000010.00000000) Class C subnet address
  - 192.60.131.0 (11000000.00111100.10000011.00000000) Class C subnet address
- 192.60.128.0 (11000000.00111100.10000000.00000000) /22
  - Supernetted Subnet address
- 255.255.252.0 (11111111.11111111.11111100.00000000)
  - Subnet Mask (replace host ID part by 0's and subnet portion by 1's)
- 192.60.131.255 (11000000.00111100.10000011.11111111)
  - Broadcast address (replace host ID part by 1's)

Use of Subnet Mask in Host

- A host uses subnet mask to extract the subnet part of its IP address
- A host uses subnet mask to extract the subnet part of the destination IP address
- Compare to see if both subnet parts are the same
  - If they are the same, use the ARP to obtain the destination MAC address
  - If they are different, then packet must be sent to the gateway (router), and use ARP to obtain the router's MAC address

Example: CIDR and Subnet

131.204.2.00010000/28: IP address representing the subnet

- Gateway:131.204.2.00010001
- Hosts:

  131.204.2.00010010

  .

  .

  .

  131.204.2.00011110

  13 hosts

- Broadcast address: 131.204.2.31

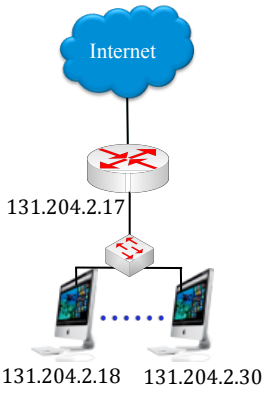

131.204.2.00011111: broadcast

Try It Out …

- Suppose that a small college is allocated the block of addresses 153.104.11.0/24. Assuming the subnet layout below, choose a network address and a network mask for each subnet.

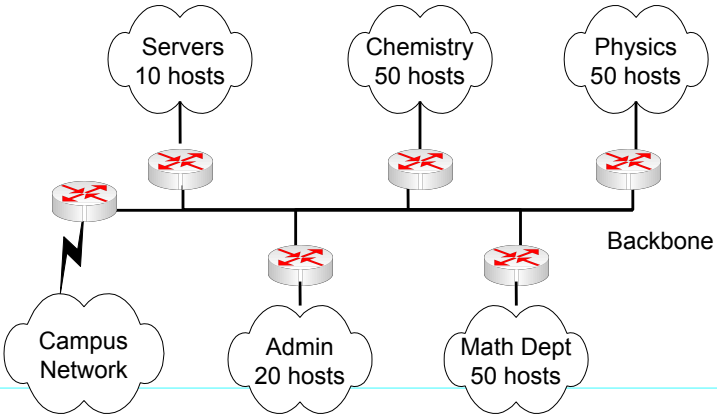

CIDR and Routing Information

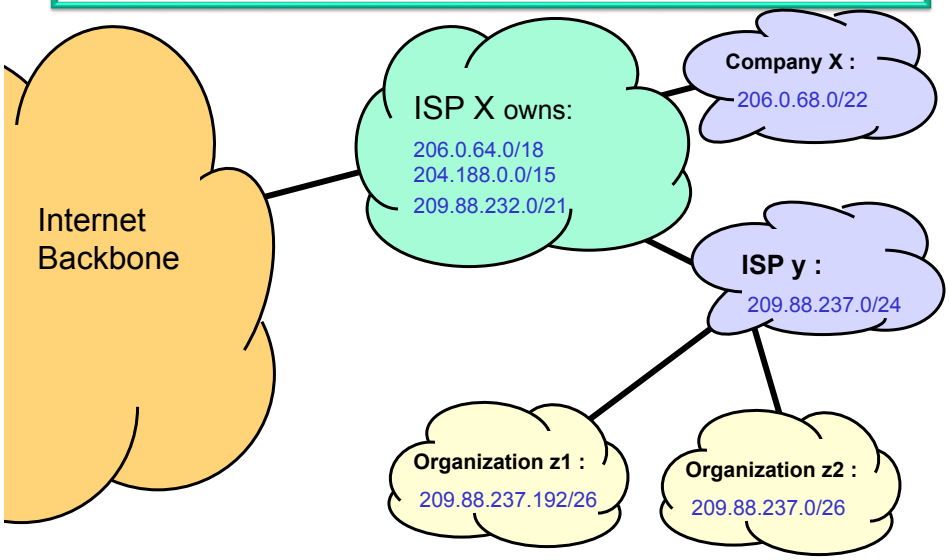

CIDR and Routing Information

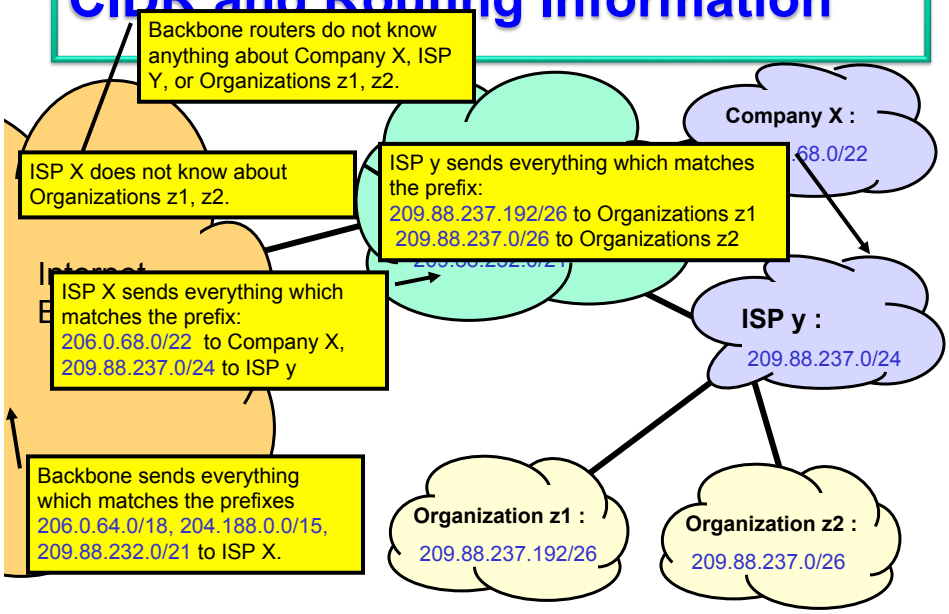

Outline

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

Part 3

Obtaining IP Addresses for My PC

- Hard-coded by hand to a computer
  - Obtain IP address, subnet mask, Gateway IP address, and DNS IP address from ISP or network manager
- DHCP (Dynamic Host Configuration Protocol)
  - Dynamically get IP address from DHCP server
  - "Plug-and-play"

DHCP: Dynamic Host Configuration Protocol

- Function:
  - Allow host to dynamically obtain its IP address from DHCP server when joining network
  - Allow IP address reuse
- Support wired/mobile stations joining network
- Host holds an IP address only while actively connected
- Renew IP address already in use
- DHCP provides:
  - IP address
  - Subnet mask
  - Gateway IP address
  - DNS server IP address

DHCP Client/Server

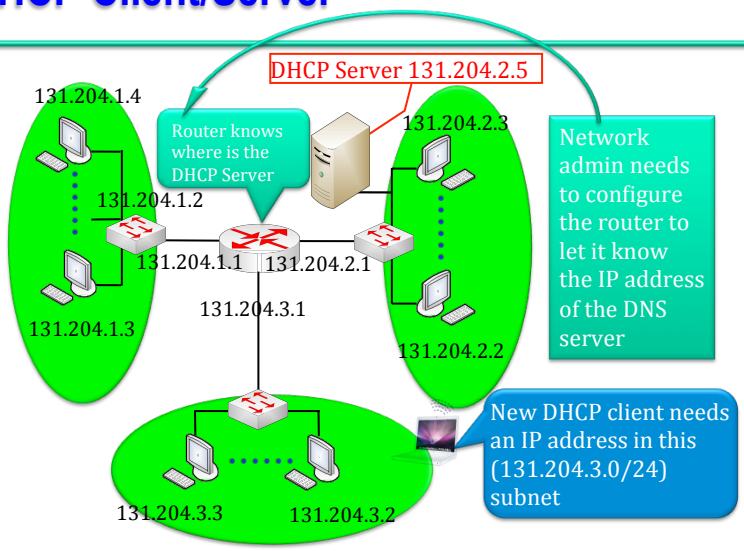

DHCP Client/Server

DHCP procedure:

1. Host broadcasts "DHCP Discover" message
2. DHCP server responds with "DHCP offer" message
3. Host requests IP address: "DHCP request" message
4. DHCP server sends address: "DHCP ACK" message

yiaddr: (offered) your IP address

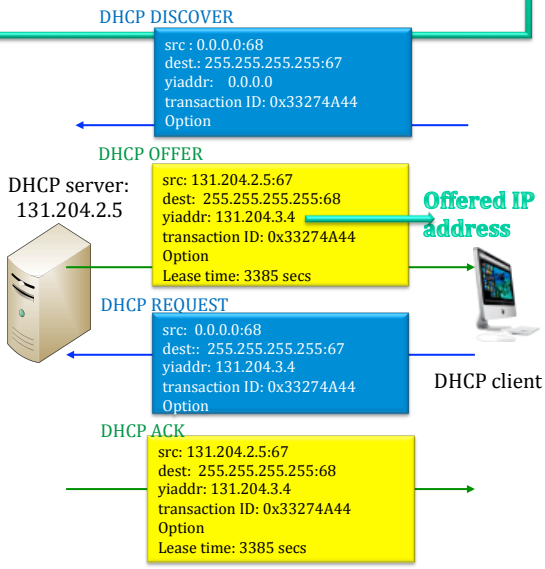

Reuse a previous network address

- Most DHCP servers are configured to let a client reuse a previously allocated network address
  - This can reduce the amount of broadcast traffic resulting from a DHCP DISCOVER message and a DHCP OFFER message
  - If a client remembers and wishes to reuse a previously allocated network address, a client may choose to omit the DHCP discover message
  - The client broadcasts a DHCPREQUEST message on its local subnet. The message includes the client's network address in the 'requested IP address' option
- DHCP addresses have a lease time
  - Typical time is 1-30 days
  - When half the lease time expires, client sends request to the DHCP server to renew the lease
  - If server does not respond, three quarters through the lease time the client broadcasts a new DHCP DISCOVER to locate a new server

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

ARP: Address Resolution Protocol

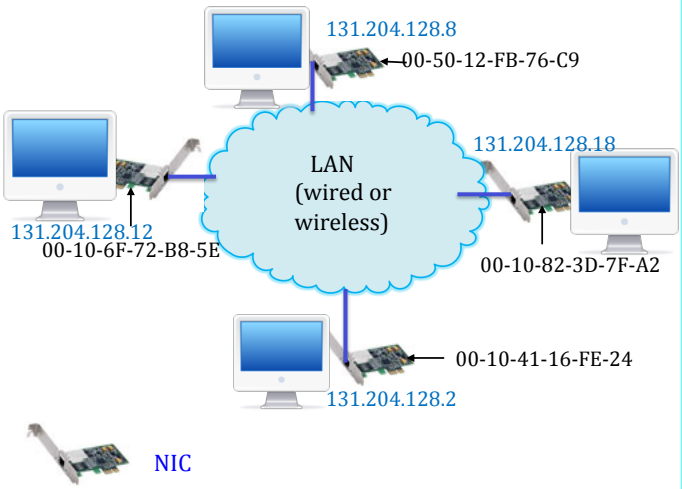

- Each IP station (host, router) on LAN has an ARP table
- ARP table: IP to MAC address mappings for some LAN stations
  - < IP address; MAC address; TTL>
  - TTL (Time To Live): time after which address mapping will be erased (typically 20 min)
- MAC address never leaks outside the subnet

ARP protocol

- Station A wants to send datagram to Station B, and B's MAC address is not in A's ARP table
- Station A broadcasts ARP query packet, containing B's IP address
  - Dest MAC address = FF-FF-FF-FF-FF-FF
  - All machines on LAN receive ARP query
- Station B receives ARP packet, replies to Station A with its MAC address
  - Frame sent to A's MAC address (unicast)
- Station A caches (saves) IP-to-MAC address pair in its ARP table/cache until information times out
- ARP is "plug-and-play":
  - No ARP server
  - Stations create their ARP tables without the involvement of network administrator

ARP protocol

- Every node maintains an ARP table
  - (IP address, MAC address) pair
- Consult the table when sending a packet
  - Map destination IP address to destination MAC address
  - Encapsulate and transmit the data packet
- But, what if the IP address is not in the table?
  - Sender broadcasts: "Who has IP address 1.2.3.156?"
  - Receiver responds: "MAC address 58-23-D7-FA-20-B0"
  - Sender caches the result in its ARP table
  - Receiver caches the sender info in its ARP table
  - Cache table fills up => Least Recently Used policy used

ARP Request

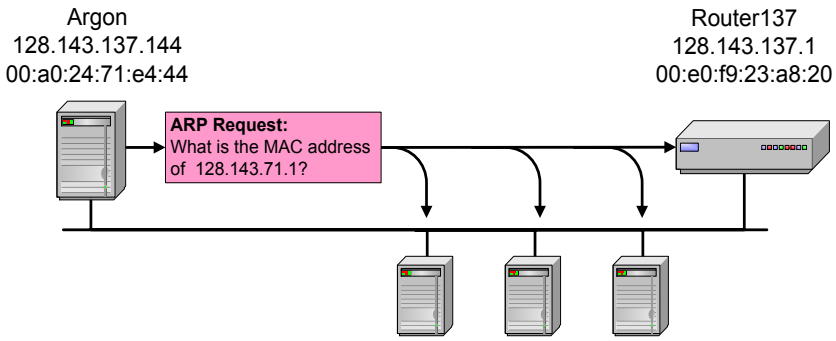

ARP Reply

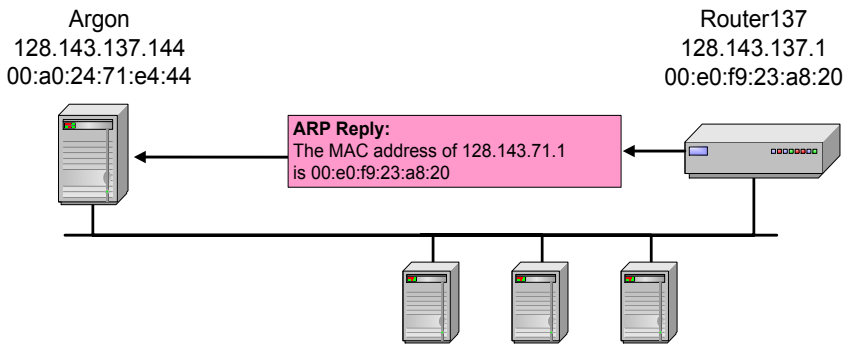

ARP Table

- Use

  arp –a

  to view the ARP cache on your own machine.

- Command available in both Unix and Windows

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

Routing to Another LAN

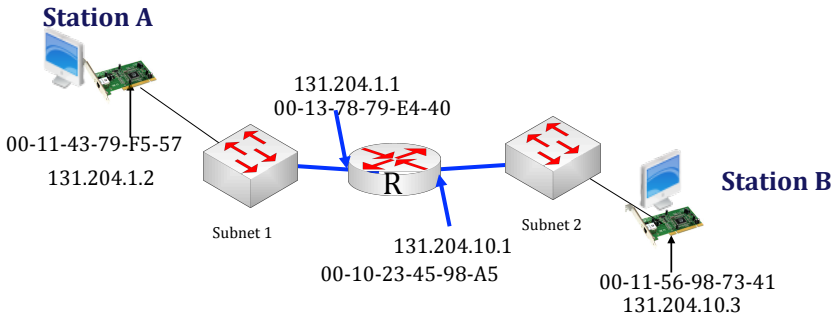

- Walkthrough:
  - Send datagram from Station A to Station B via router R
  - Station A knows B's IP address
- Two ARP tables in router R
  - One for each IP network (LAN)

- Station A creates an IP datagram with source A, destination B
- Station A uses ARP to get R's MAC address for 131.204.1.1 that is obtained by DHCP
- Station A creates a link-layer frame with R's MAC address as dest MAC address, and frame contains A-to-B IP datagram
- A's NIC sends frame
- R's NIC receives frame
- R removes IP datagram from Ethernet frame, sees its destined to Station B
- R forwards the packet to Interface 131.204.10.1 interface based on routing table
- R uses ARP to get B's MAC address
- R creates frame containing A-to-B IP datagram and sends to B using B's MAC address as the destination MAC address

This is a really important example

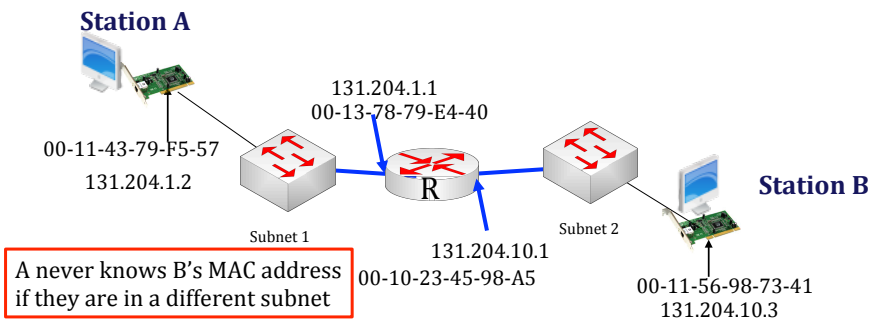

A never knows B's MAC address if they are in a different subnet

Sending Datagram to Another Subnet

- Mission: Sending a datagram from A to B via R

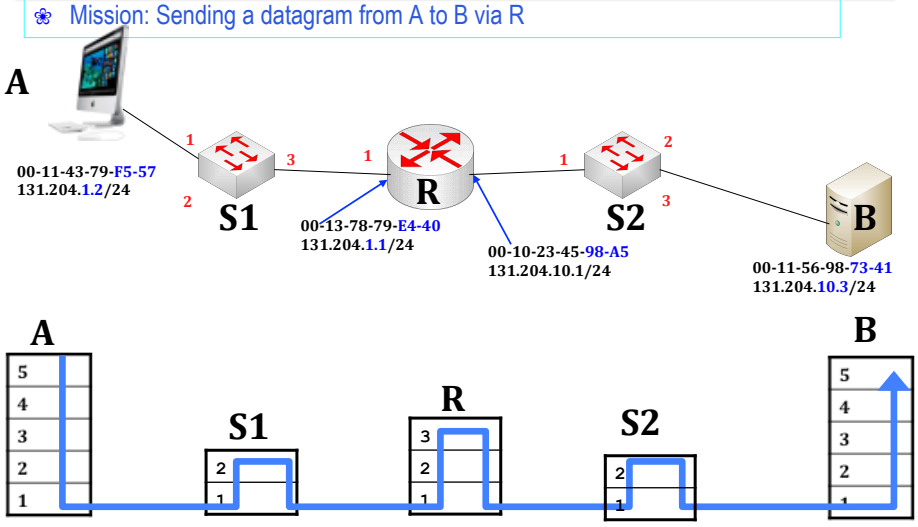

Simplified Data Format

- 1 – Source IP
- 2 – Destination IP
- 3 – Source Mac Address
- 4 – Destination Mac Address

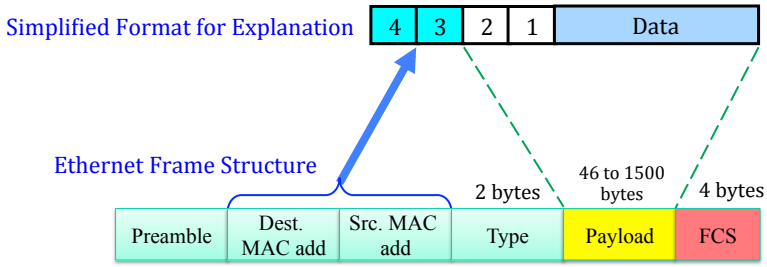

Host A (Network Layer)
- Set ***source & destination IP*** in datagram format
- A datagram will be forwarded to a gateway (router) using the subnet portion of the network ID since the destination host is not on the same subnet

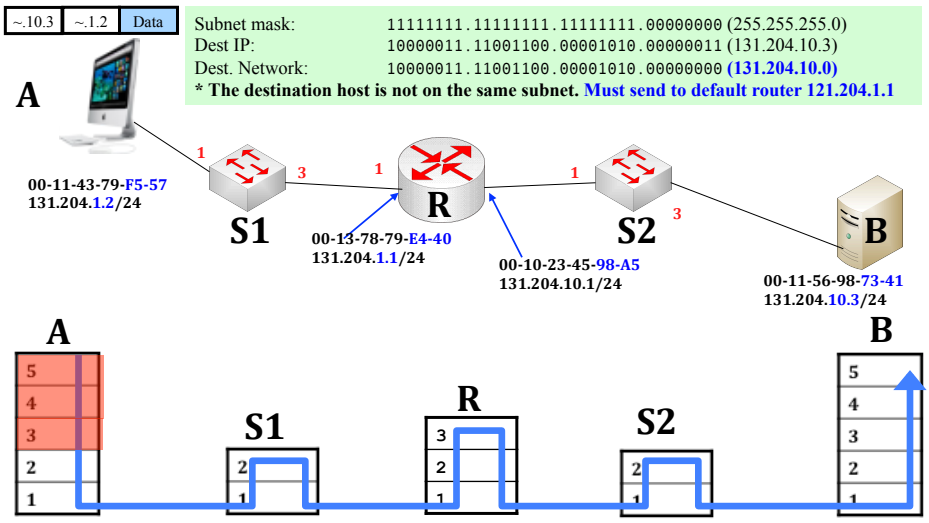

Host A (Data link layer)
- Set ***source & destination Mac address*** in an Ethernet frame by means of an ARP table

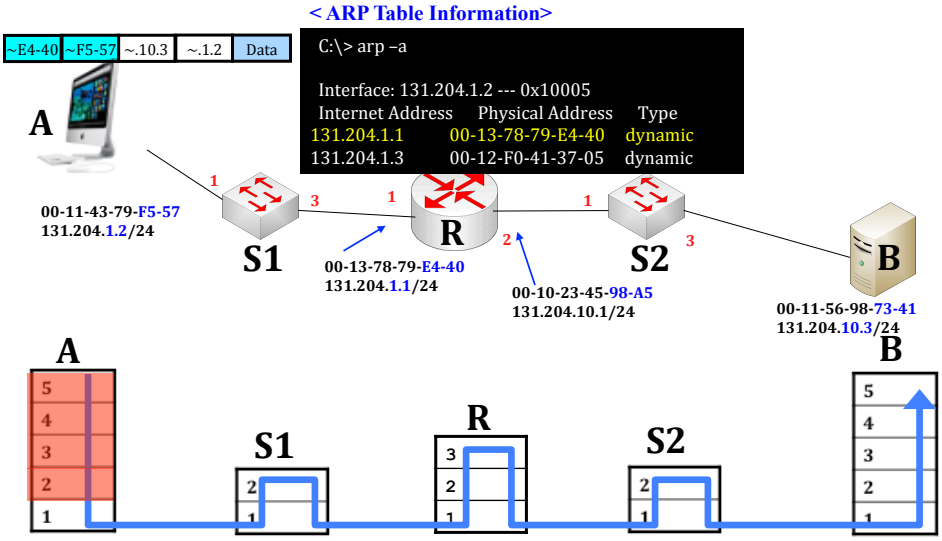

Host A : A's NIC sends an Ethernet frame
Switch S1
- S1's NIC of interface 1 receives an Ethernet frame
- ***Look up a switch table*** & ***Forward a frame*** into a corresponding port

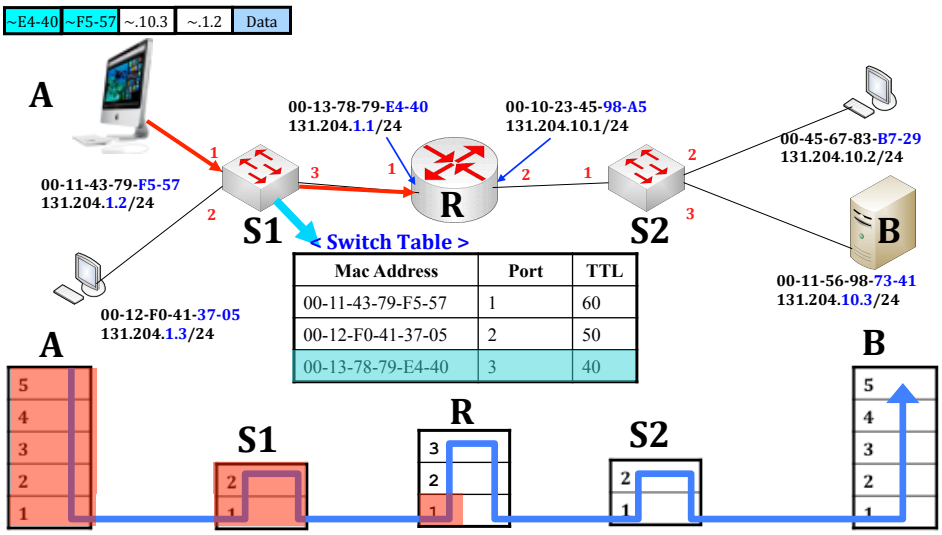

Router R
- Look up a routing table & determine a next forwarding link (Layer 3)
- Change MAC addresses for next hop (Layer 2)

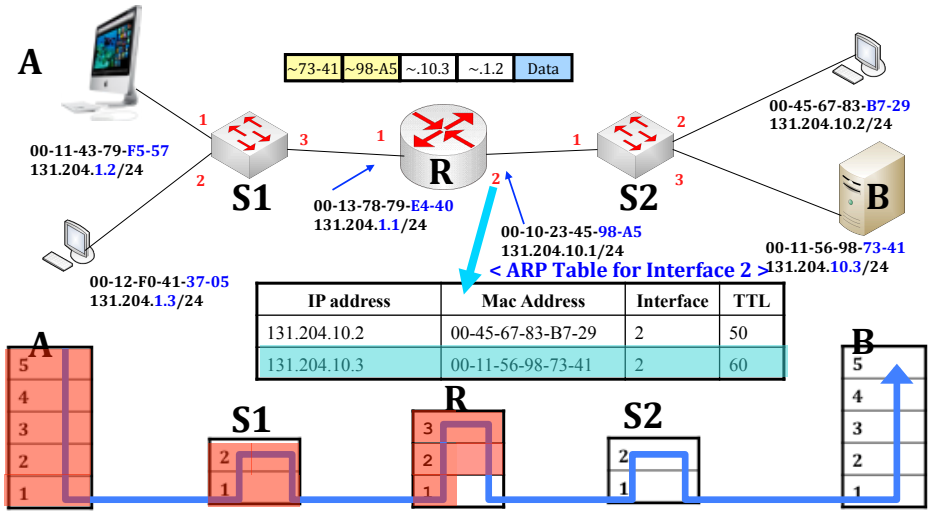

IP address	Mac Address	Interface	TTL
131.204.10.2	00-45-67-83-B7-29	2	50
131.204.10.3	00-11-56-98-73-41	2	60

Router R : R's interface sends an Ethernet frame
Switch S2
- S2's NIC for port 1 receives an Ethernet frame.
- ***Look up a switch table*** & ***Forward a frame*** into corresponding port 3

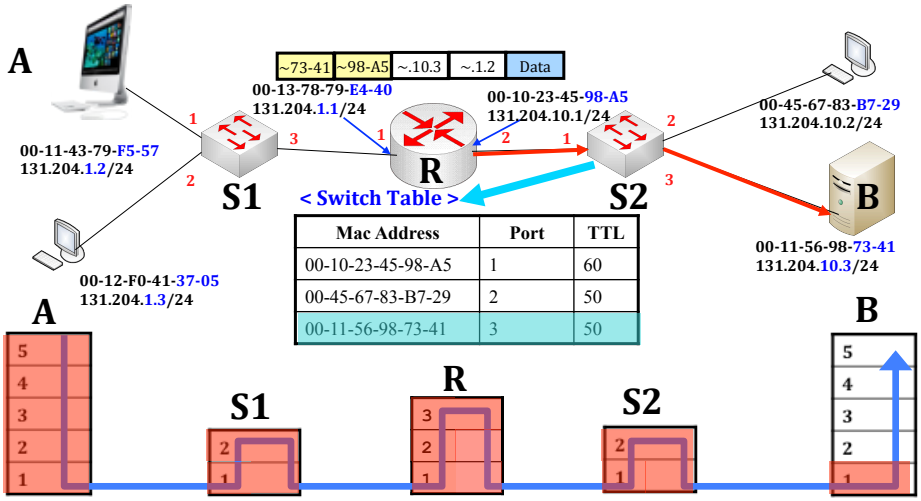

Mac Address	Port	TTL
00-10-23-45-98-A5	1	60
00-45-67-83-B7-29	2	50
00-11-56-98-73-41	3	50

Host B (From layer 1 to layer 5)
- Check if destination addresses are the same as those of itself.
  (Data link layer – MAC Address, Network layer – IP)
- Each layer drops its own header and trailer.

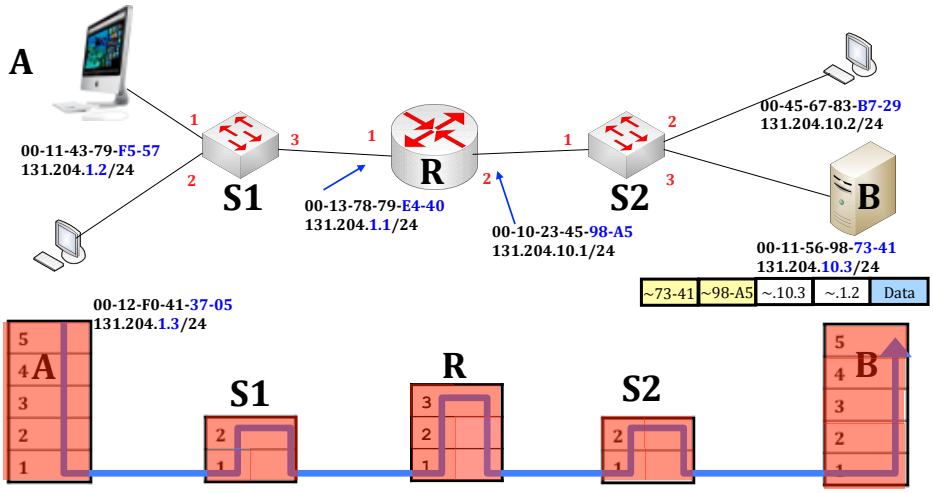

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

Network Address Translation (NAT)

- Goals
  - Private networks use private IP addresses
  - A private network uses just one global IP address provided by ISP to connect to the Internet
  - Can change address of devices in private network without notifying outside world
  - Can change ISP without changing the address of devices in private network
  - Devices inside private network are not explicitly addressable by external network, or visible by outside world (a security plus)
- A NAT device would not advertise private networks to the external/public network; however, the external/public network services may be advertised within the private network

Basic NAT

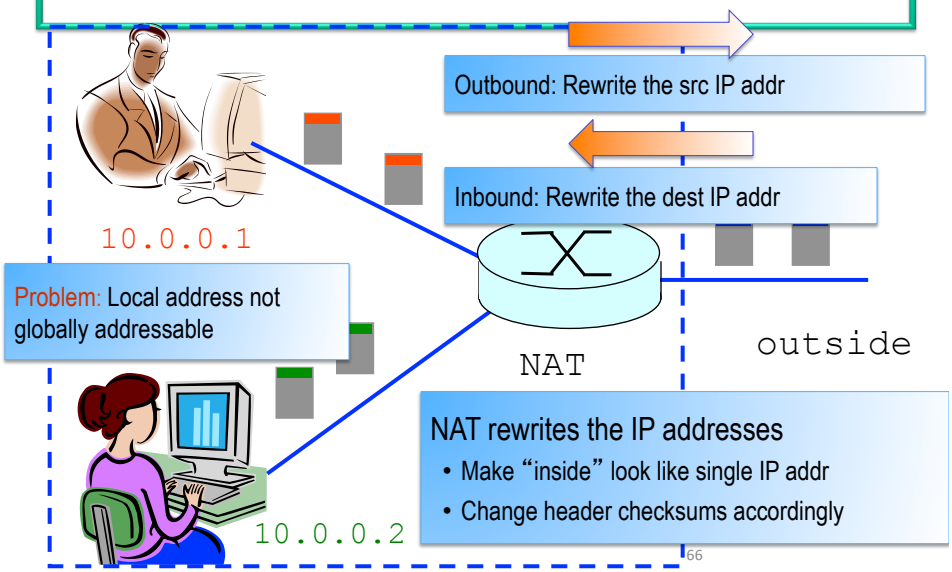

NAPT: Network Address Port Translation

- A NAPT router performs the following for every outgoing datagram
  - Replace (source IP address, port #) with (Public IP address, new port #)
  - Remote clients/servers will respond using (Public IP address, new port #) as destination addr
  - Remember (in NAPT translation table) every (source IP address, port #) to (Public IP address, new port #) translation pair
- A NAPT router performs the following for every incoming datagram
  - Replace (Public IP address, new port #) in destination fields of every incoming datagram with corresponding (source IP address, port #) stored in NAPT table
- 10.0.0.4:5555 represents an IP address and port number separated by a ":"
  - 10.0.0.4:5555 is referred as a transport address in the RFCs

NAPT

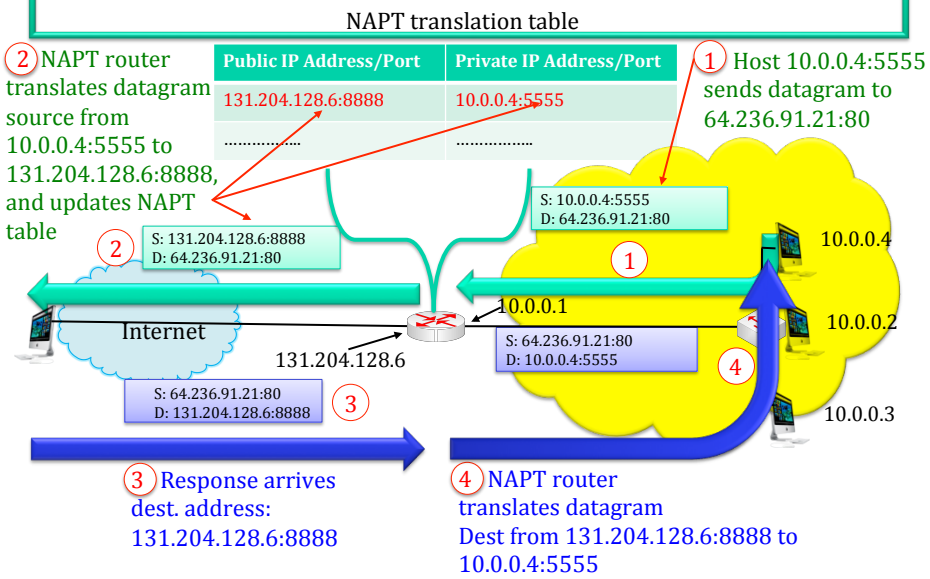

Where is NATP Implemented?

- Home router (e.g., Linksys box)
  - Integrates router, DHCP server, NAT, etc.
  - Use single IP address from the service provider
  - ... and have a bunch of hosts hiding behind it
- Campus or corporate network
  - NAT at the connection to the Internet
  - Share a collection of public IP addresses
  - Avoid complexity of renumbering end hosts and local routers when changing service providers

Practical Objections Against NATP

- Port #s are meant to identify *sockets*
  - Yet, NAT uses them to identify *end hosts*
  - Makes it hard to run a server behind a NAT

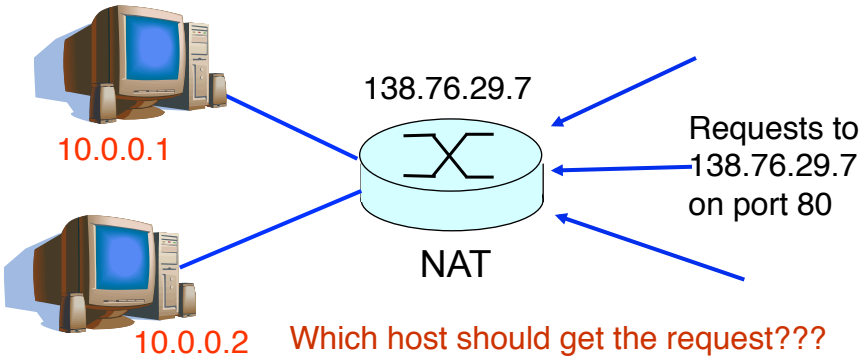

Running Servers Behind NATs

- Running servers is still possible
  - Admittedly with a bit more difficulty
- By explicit configuration of the NAT box
  - E.g., internal service at <dst 138.76.29.7, dst-port 80>
  - ... mapped to <dst 10.0.0.1, dst-port 80>

NAT can Do Load Balancing of Servers

Split load over server replicas
(At the connection level)

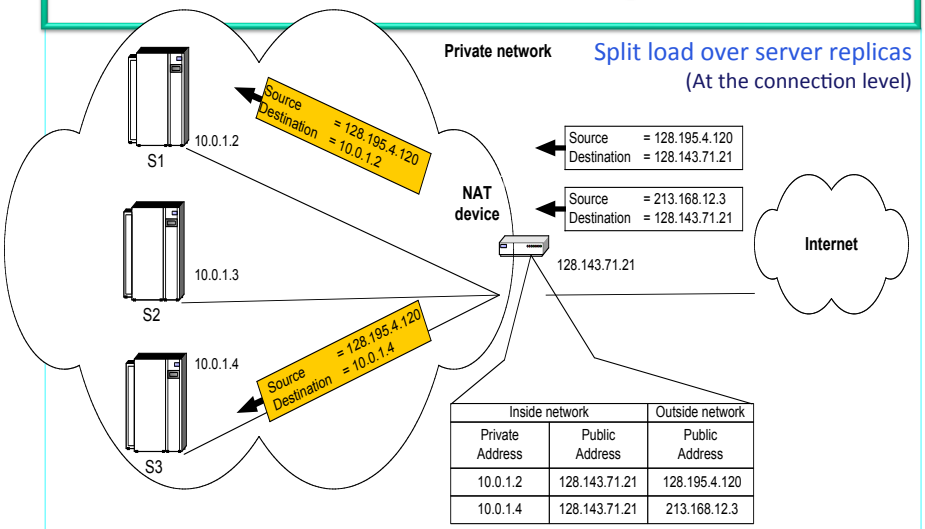

Inside network		Outside network
Private Address	Public Address	Public Address
10.0.1.2	128.143.71.21	128.195.4.120
10.0.1.4	128.143.71.21	213.168.12.3

Apply load balancing policies

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

Error Reporting and Diagnosis

- Examples of errors a router may encounter
  - Router does not know where to forward a packet
  - Packet's time-to-live field expires
- Best effort
  - Router just silently drop packets
- Network diagnosis
  - IP includes basic test and feedback for solving network problems
  - Internet Control Message Protocol (ICMP)
- ICMP runs on top of IP
  - In parallel with TCP and UDP

Internet Control Message Protocol

- Diagnostics
  - Triggered when an IP packet encounters a problem
    - E.g., time exceeded or destination unreachable
  - ICMP packet sent back to the source IP address
    - Includes the error information (e.g., type and code)
    - Excerpt of the original data packet for identification
  - Source host receives the ICMP packet
    - Inspects the excerpt of the packet (e.g., protocol and ports)
    - Inform the socket that should receive the error information

ICMP Message

- ICMP message is carried in IP datagram as payload
- ICMP message: type, code plus first 8 bytes of IP datagram causing error

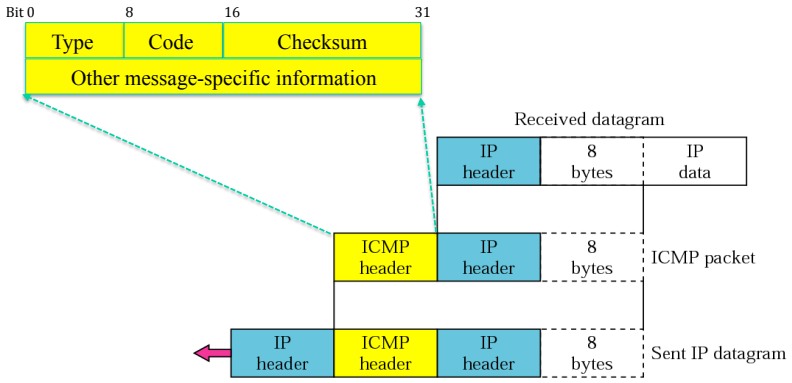

ICMP: Internet Control Message Protocol

Type	Code	Description
0	0	echo reply (ping)
3	0	Destination network unreachable
3	1	Destination host unreachable
3	2	Destination protocol unreachable (the designated transport protocol is not supported)
3	3	Destination port unreachable (the designated protocol is unable to inform the host of the incoming message)
3	6	Destination network unknown
3	7	Destination host unknown
4	0	source quench (congestion control - not used)
8	0	echo request (ping)
9	0	route advertisement
10	0	router discovery
11	0	TTL expired
12	0	bad IP header

Ping

- Uses ICMP Echo request/reply to
  - test destination reachability
  - compute round trip time
  - count the # of hops to destination
- Source sends ICMP echo request message to the destination address
  - echo request packet contains timestamp also
- Destination replies with an ICMP echo reply message containing the data in the original request message
- Source can calculate RTT of packets
- If no echo reply comes back, destination unreachable

Ping (contd.)

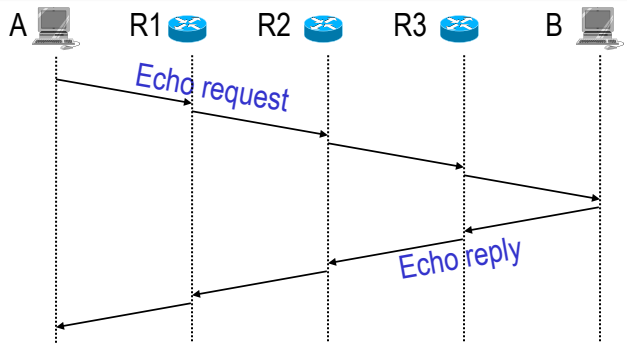

- Sample output:

  Reply from 164.107.144.3: 48 bytes in 47 msec. TTL: 253

ICMP Echo Request / Reply Message

8: Echo request
0: Echo reply

Type: 8 or 0	Code: 0	Checksum
Identifier		Sequence number
Optional data Sent by the request message; repeated by the reply message		

- The identifier and sequence number may be used by the echo sender to aid in matching the replies with the echo requests.

Traceroute Based on ICMP

- Source host sends a series of UDP packets in UNIX OS to destination host
  - Send the first UDP with TTL =1 for three times
  - Send the second UDP with TTL=2 for three times
  - .....
  - Send the ith UDP with TTL=i for three times
  - All UDP packets use a port number that is not opened as a service at the destination host
- When ith datagram arrives at the ith router:
  - The router discards datagram
  - Sends an ICMP message to source host (type 11, code 0: TTL expired)
  - Message includes name of router & IP address

Traceroute Based on ICMP

- When ICMP message arrives, source host calculates RTT (round trip time)
- UDP packet eventually arrives at destination host
- Destination host returns ICMP port unreachable message (type 3, code 3)
- When source host gets this ICMP message, it stops the sending of UDP packets

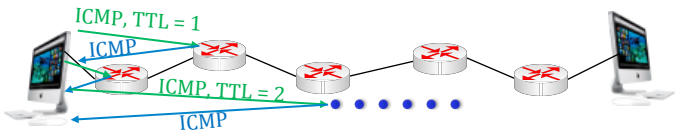

ICMP Risks

- ICMP can be used to scan hosts and available ports (services)
  - Turn off ICMP in Firewalls, including host firewalls
- Traceroute can be used to discover the router interface IP addresses
  - Attacks can be directed toward a particular interface
- To protect hosts and routers, ICMP responses must be turned off
  - Reconnaissance is useless if a host does not respond

Traceroute is a security risk and not allowed in most networks

Outline

Part 3

- Network layer
- Connectionless datagram forwarding
- Network layer functions in the protocol stack
- IPv4 header
- IPv4 address
- DHCP
- ARP
- Routing between LANs
- Network address translation (NAT)
- The Internet control message protocol (ICMP)
- Mobile IP

Mobile IP

- Mobility issues
  - Internet hosts/interfaces are identified by IP address
  - Moving to another network requires different network address
  - Transition should be seamless
- RFC 3344
  - Home agents
  - Foreign agents
  - Foreign-agent registration
  - Care-of-addresses
  - Tunnel/encapsulation (packet-within-a-packet)
  - Agent discovery
  - Registration with home agent
  - Routing

Home and Visited Network

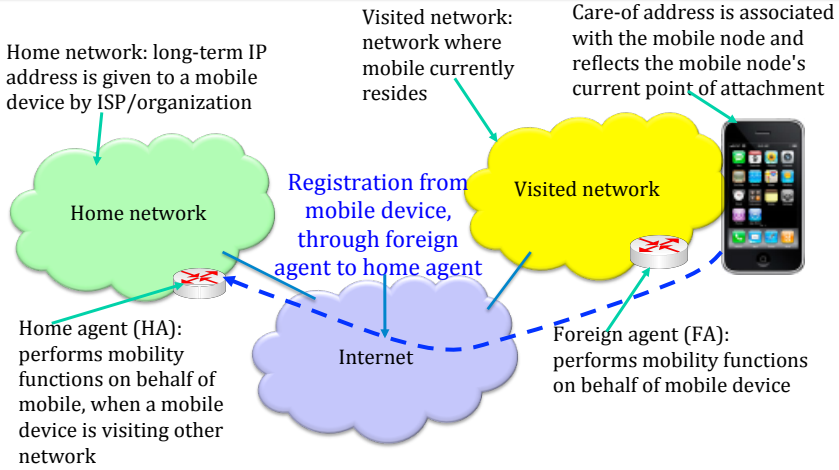

Registration Procedure

- Foreign/home agents advertise service (once every few seconds)
- Mobile device receives the advertisement from FA
- Mobile device sends in the Registration Request to FA
  - UDP packet with Destination Port 434
- FA registers the mobile device to HA by sending a Registration Request
- HA sends Registration Reply to FA
- FA sends Registration Reply to mobile device
- Each mobile node, foreign agent, and home agent must be able to support a mobility security association for a mobile device, indexed by its SPI (security Parameter index) and IP address
  - Authentication between mobile device and FA
  - Authentication between FA and HA
  - Authentication of mobile device/user by HA
- Each registration has a lifetime

Routing

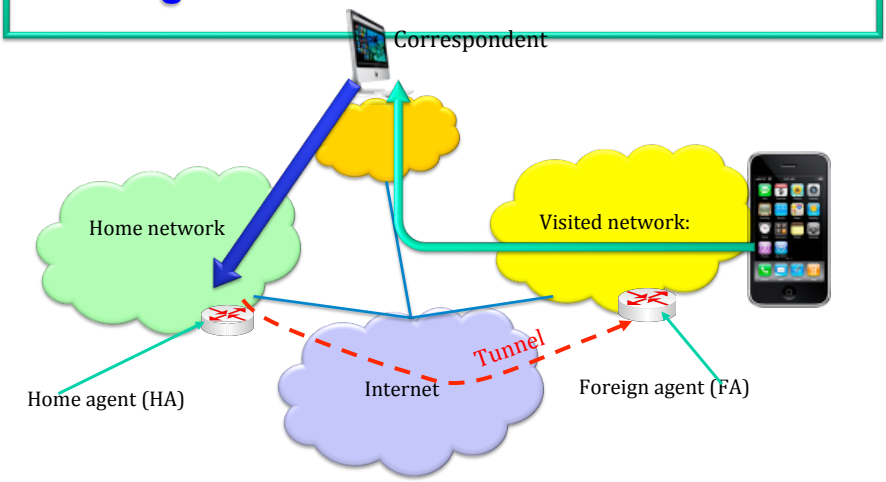

Routing Procedure

- Mobile IP uses protocol tunneling to hide a mobile device's home address from intervening routers between its home network and its current visited network
  - The tunnel terminates at the mobile device's care-of address
  - The care-of address must be an address to which datagrams can be delivered via conventional IP routing
  - At the care-of address, the original datagram is removed from the tunnel and delivered to the mobile device
- Datagrams sent to the mobile devices's home address are
  - Intercepted by its home agent
  - Tunneled by the home agent to the mobile device's care-of address
  - Received at the tunnel endpoint: either at a foreign agent or at the mobile device itself
  - Delivered to the mobile device
- In the reverse direction, datagrams sent by the mobile device are generally delivered to their destination using standard IP routing mechanisms, not necessarily passing through the home agent

Correspondent Node → Mobile Node

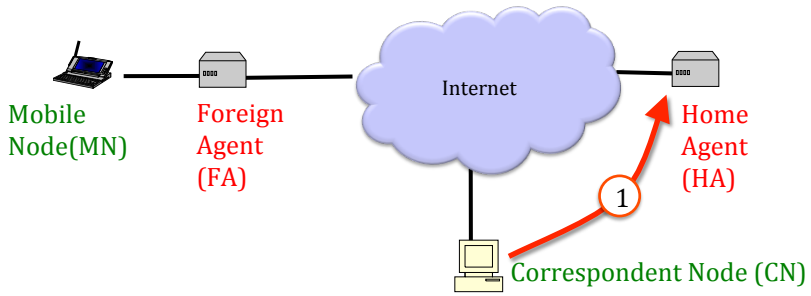

Packet from CN to MN

Source address =
Destination address =
Payload

Correspondent Node → Mobile Node

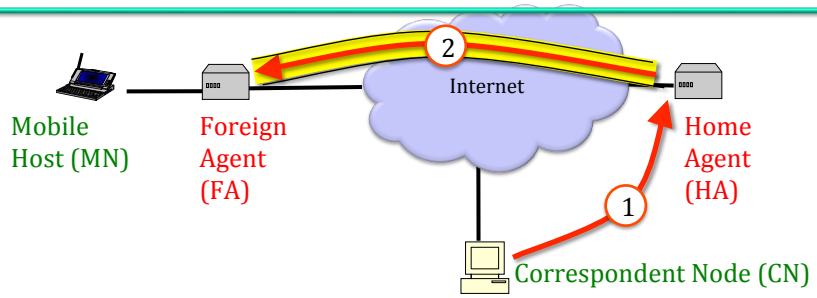

Home agent intercepts packet and tunnels it to FA

Source address =
Destination address =
Source address = address of CN
Destination address = home IP address of MN
Original payload

Correspondent Node → Mobile Node

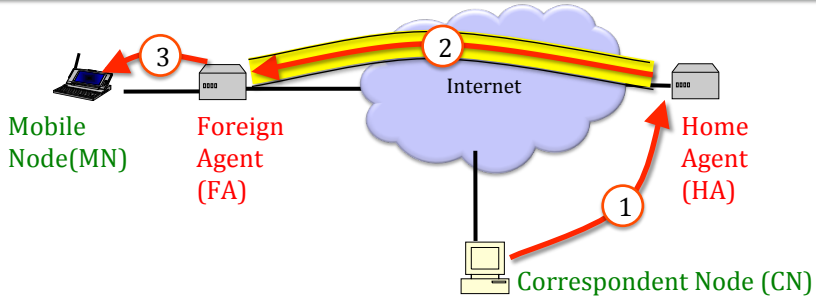

Foreign Agent removes original packet from the tunnel and delivers it to MN over the foreign link.

Source address = address of CN
Destination address = home IP address of MN
Payload

Mobile Node → Correspondent Node

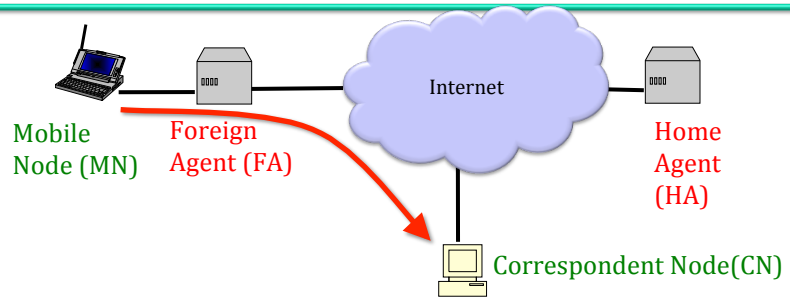

Packet from MN to CN

Source address =
Destination address =
Payload