CSCE 515: Computer Network Programming

------ IP routing

Wenyuan Xu

Department of Computer Science and Engineering
University of South Carolina

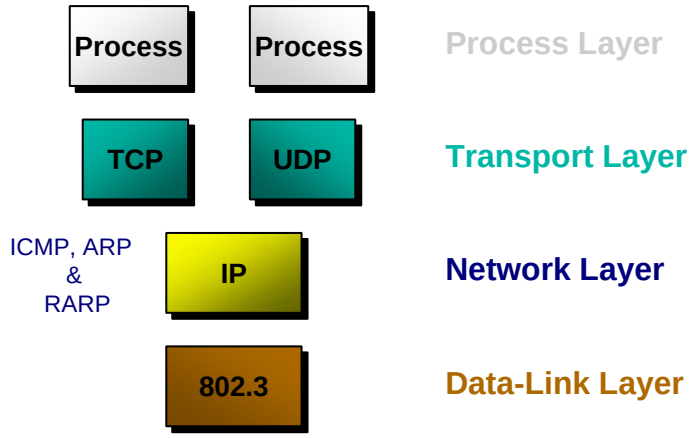

IP - Network Layer

- Provide delivery of packets from one host in the Internet to any other host in the Internet, even if the hosts are on different networks
- Problems:
  - Heterogeneity (addressing, packet size and format, routing)
  - Handling this problem and efficiently route packets across several networks
  - Loops, oscillations, islands
- Solutions
  - Gateways to handle communication between networks
  - Gateways, routers, nodes, links, networks

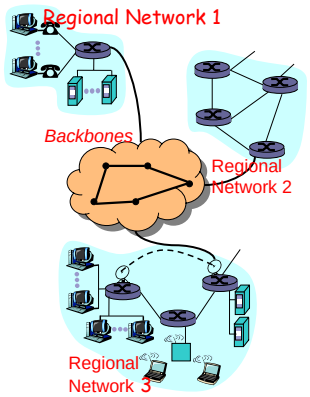

Internet Protocol (IP)

- Provide *unreliable* and *connectionless* datagram delivery service
- Internet packets are called “datagrams” and may be up to 64 kilobytes in length (although they are typically much smaller e.g. 1500 bytes)
- Upper layer data (TCP, UDP, ICMP, IGMP, etc.) are transmitted as IP datagrams
- Q: What’s the advantage of connectionless delivery?

Advantages of Connectionless

- Host software is much simpler at the network layer.
- Transport layer already provides connection-oriented, should not repeat the work
- Many applications do not require sequential delivery of packets (example: packet voice).
- It is better to provide degraded service to everyone than to limit network access.
- Server (or router) could become overloaded managing too many connections.

IP Routing

- Q: How do you get a packet from one network to another?

IP Routing

- A: with a router (or a series of routers)

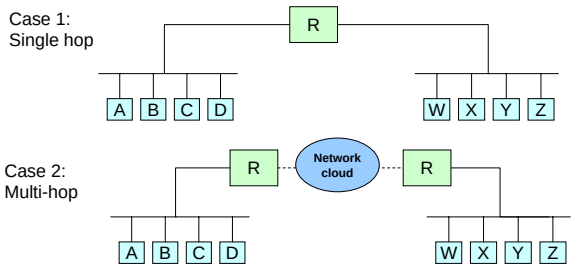

IP Routing

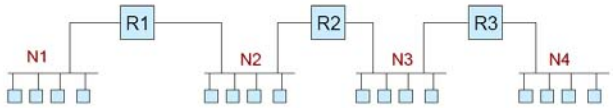

Routing table @ R2

Dest	Next hop
N1	R1
N2	Deliver directly
N3	Deliver directly
N4	R3/default

Actual routing able contains IP addresses, Flags indicating type of entries, net mask etc.

IP Routing

- Forwarding:
  - When each packet arrives, looking up the outgoing line to use for it in the routing table
  - Done on a hop-by-hop basis
  - If destination is directly connected or on a shared network, send IP datagram directly to destination
  - Otherwise send datagram to a router
- Routing updates
  - filling in and updating the routing tables

Routing Table

- Each entry contains following information
  - Destination IP address
  - IP address of next-hop router
  - Specification of network interface
  - Flag
    - **U:** the route is up and operational.
    - **H** :this is a route to a specific host (most routes are to networks).
    - **G**: the route uses an external gateway.

Destination	Gateway	Flags	Ref	Use	Interface
192.168.0.0	129.252.130.203	UG	1	0	
129.252.130.0	129.252.130.106	U	1	68	eri0
224.0.0.0	129.252.130.106	U	1	0	eri0
default	129.252.130.1	UG	1	135	
127.0.0.1	127.0.0.1	UH	1	0	lo0

Host route determination

- Longest prefix match with destination address and entry in the routing table
- First, search for a matching host address
  - Flag H is set
- Second, search for a matching network address
  - Need to know the number of bits to use for network ID
- Third, search for a default entry
  - Execute `netstat -rn` on your machine and find the contents of the routing table
  - Default entry allows for a single entry for a list of entries that have the same next-hop value

IP Routing

- Forward datagrams generated either on local host or on some other hosts toward their ultimate destination
- Static routing: when network is small, single connection point to other networks, no redundant route existent
- Dynamic routing: use routing daemon to run routing protocol in order to communicate with other routers

`ifconfig` Command

- Available at `/usr/sbin`
- Configure or query a network interface used by TCP/IP
- Support address families other than IP address
- `-a` option to display report of all interfaces

`netstat` Command

- Available at `/usr/bin`
- Display network status
- `-a` option to display state of all sockets, all routing table entries, or all interfaces
- `-i` option to display interface information
- `-n` option to print IP addresses instead of host names
- `-r` option to display routing table
- `-M` option to display multicast routing table

The *four for*mats of IP Addresses

- 32 bits long: 129.252.138.8

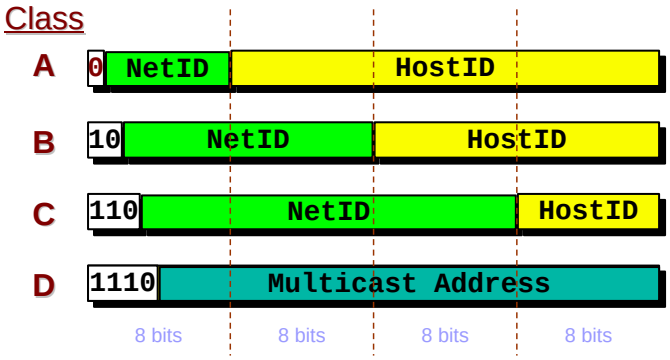

IP Addresses

- Subnet Addressing
  - To make better use of class A and class B addresses, divide host ID into subnet ID and host ID

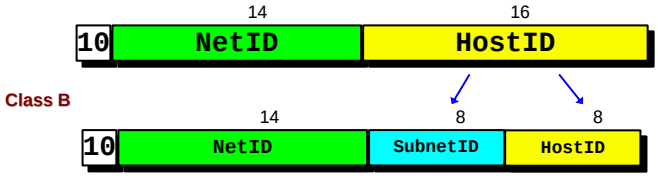

Subnet Mask

- 32-bit value containing “1” bits for network ID and subnet ID, and “0” bits for host ID

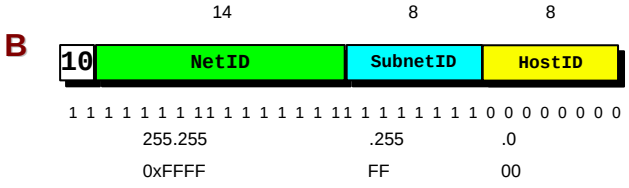

Subnet Masks

- Assume UP addresses A and B share subnet mask M.
- Are IP addresses A and B on the same subnet?
  - 1. Compute (A and M)
  - 2. Compute (B and M)
  - 3. if (A and M) == (B and M) then A and B are one the same subnet.
- Example: A and B are class B addresses
  - A = 165.230.82.52
  - B = 165.230.24.93 *Same network?*
  - M = 255.255.255.0 *Same subnet?*

Variable length subnetting

- Subnet masks allow power of 2 subnets
- Use a hierarchy of routers to allow subnets to be divided with different subnet masks
- Another approach:
  - Variable length subnet masks
  - Allow a subnet to be defined by more than two masks
  - The router applies the masks one after another

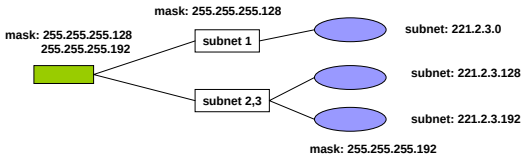

Routing Table

Address	Mask	Interface
221.2.3.0	255.255.255.128 255.255.255.192	L1
221.2.3.128	255.255.255.128 255.255.255.192	L2
221.2.3.129	255.255.255.128 255.255.255.192	L3

Question

- If an ISP has a 203.6.8.0 Network, he wants to give assign to 5 customers who need 60, 60, 60, 30, 30 hosts.
  - What should the subnet mask be?
  - What is the address range in each subnet?
  - What should the routing table entries be?

CIDR

- CIDR (classless Inter domain routing)
- Too many small networks requiring multiple class C addresses
- Running out of class B addresses, not enough nets in class A
- Assign contiguous block of class C addresses
- Use CIDR address mask to aggregate
- Ex 192.17.0.0, 255.255.248.0, Send it to R3

CIDR

- Original addressing schemes (class-based):
  - 32 bits divided into 2 parts:
  - Class A
  - Class B
  - Class C

- Class C address has max of 254 hosts
  - Not enough for many organizations
  - Too many class C addresses → huge routing tables
- Classless Inter domain routing (CIDR)
- CIDR introduced to solve 2 problems:
  - exhaustion of IP address space
  - size and growth rate of routing table

Supernetting

- Example: an organization needs 500 addresses.
  - A single class C address not enough (256 hosts).
  - Instead a class B address is allocated. (~64K hosts). a huge waste.
- CIDR allows multiple Class C addresses to be assigned to an organization but still occupy one entry in the routing table
- <192.5.48.0,2> this is used to specify that 2 network addresses 192.5.48.0 and 192.5.49.0 is allocated to an organization
- Typically the starting address with a CIDR mask that indicates the common most significant bits for the ranges is used to specify the block of addresses
  - /notation
  - 192.5.48.0/23 → 192.5.48.0 and 192.5.49.0 are assigned

Address Arithmetic: Address Blocks

- Address format <IP address/prefix P>.
  - The prefix denotes the upper P bits of the IP address.
  - Can be used to specify arbitrary blocks of addresses
- The <address/prefix> pair defines an address block:
  - Examples:
  - 200.15.0.0/16 => [200.15.0.0 - 200.15.255.255]
  - 192.24.0.0/13 => [192.24.0.0 - 192.31.255.255]

192.**00011**000. 0000 0000. 0000 0000

13th bits fixed — Variable

Reducing Routing Table Size

Without CIDR:

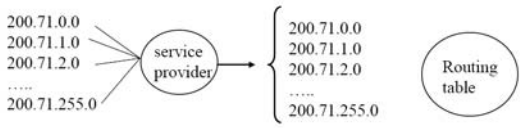

With CIDR:

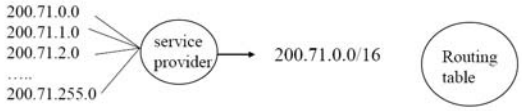

CIDR: Classless Inter-Domain Routing

- Address format <IP address/prefix P>.
  - The prefix denotes the upper P bits of the IP address.
  - Can be used to specify arbitrary blocks of addresses
- Say an ISP has 192.5.48.0, 192,5.49.0,192.5.50.0,192.5.51.0 then the IP address advertised will be 192.5.48.0/22
- An ISP can obtain a block of addresses and partition this further to its customers
- Say an ISP has 200.8.4/24 address (256 addresses)
- He has another customer who needs only 4 addresses from 200.8.4.24 then that block can be specified as 200.8.4.24/30

Other Developments: NAT

- NAT- Network address translation
- Hosts need not have unique global IP address
  - Hosts are assigned private addresses
  - 10.0, 172.16 to 172.31 and 192.168. Are allocated for private hosts (Hmmm, what if those addresses appear on the Internet themselves?)
- These hosts connected to a NAT gateway which has a public IP addresses
- Packets from private hosts are replaced with source address of NAT gateway, use port# to uniquely do the reverse translation.

Assignment & Next time

- Reading:
  - TI Ch 3 **;
  - Variable Length Subnet Masking*
  - On the Assignment of Subnet Numbers *
- Next Lecture:
  - ICMP, Routing Principles