Internet and IP Protocol

Vinton G. Cerf
Inventor of TCP/IP
(Now Senior Vice President of WorldCom)

Covered Today

- IP address classes
- Subnetting and subnet masks
- CIDR

If time permits:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

Global Adresses

- Why do we need IP addresses?
  - Physical addresses (e.g., Ethernet) are specific to MAC-layer protocols
    - Need a network independent addressing scheme
  - Need a hierarchical addressing scheme (e.g., unlike Ethernet addresses)

Format of the IP Address

- IP address contains
  - Network field
  - Host field
- Design question: How many bits to allocate to each field?
  - Size of network field = maximum number of networks
  - Size of host field = maximum number of hosts per network

Classes of IP Addresses (Network Types)

- Unicast addresses:

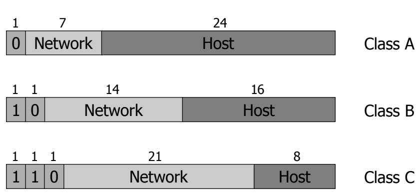

- Note: no unicast address starts with 111 (i.e., has 224 or larger in its most significant byte)

Classes of IP Addresses (Network Types)

- Class A network:
  - 126 networks; 16 million hosts per network
  - #0, and #127 are reserved
- Class B network:
  - 16 thousand networks; 65 thousand hosts per network
- Class C network:
  - 2 million networks; 254 hosts per network
  - #0 and #255 are reserved

Exercise

- Find the IP address class (i.e., network type) of each of the following sites:
  - 171.69.210.245
  - www.cnn.com
  - blue.unix.virginia.edu
  - eecs.mit.edu
  - www.amazon.com
  - mamba.cs.virginia.edu

Other IP Addresses

- Multicast (class D) addresses

1	1	1	1	28	
1	1	1	0	Multicast group ID	Class D

  - The packet is delivered to all destinations listening to the multicast address
- Class E addresses (unused)

1	1	1	1	1	27	
1	1	1	1	0	Reserved for future use	Class E

- Broadcast address: 255.255.255.255

IP Forwarding Tables

- Based on network addresses

Destination	Next hop (port)
Network X	E
Network Y	C
Network Z	C
Network W	A

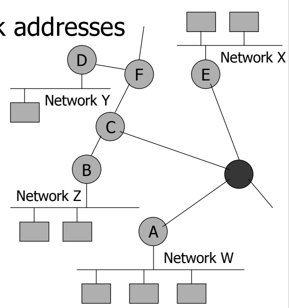

IP Forwarding

```
Lookup packet's destination network_num
if (network_num = one of my interfaces)
then
  Deliver to destination over that interface
else
  if (network_num is in my forwarding table)
  then
     Deliver to Next Hop router
  else
     Deliver to default router
```


Are IP Address Classes a Good Idea?

- Imagine a currency with denominations $1 (class C), $250 (class B), and $62,000 (class A).
- Class C is too small, class A is too big, so most use class B, but don't fully utilize it.
- Internet runs out of IP addresses!
- One temporary solution: subnets

Subnets

- The idea is to share the same IP network number among multiple subnets
- Subnets of a network should reside in the same general locale (e.g., college campus, corporate location, …)
- Routers on an IP network know the local subnets
- Remote routers need to know only the network address

Subnetting

- Subnet mask determines the subnet field within the IP address

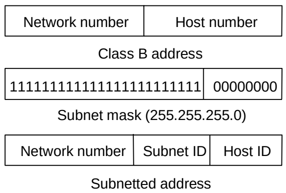

Figure 4.24
Peterson & Davie

Subnet Example

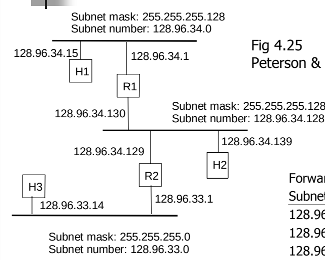

Fig 4.25
Peterson & Davie

Forwarding table at router R1

Subnet Number	Subnet Mask	Next Hop
128.96.34.0	255.255.255.128	interface 0
128.96.34.128	255.255.255.128	interface 1
128.96.33.0	255.255.255.0	R2

Table 4.10 Peterson & Davie

Hierarchical IP Forwarding with Subnets

- The scheme allows a 3-level hierarchy
- Routing table entries fall into three categories:
  - (remote network) : forwarding address
  - (my network, subnet) : forwarding address
  - (my network, my subnet, host) : forwarding address

1992 – The Collapse of the Internet (that never happened)

Quoted from IETF RFC 1517 as causes for an impending Internet collapse:

- Exhaustion of the class-B network address space. One fundamental cause of this problem is the lack of a network class of a size that is appropriate for a mid-sized organization. Class-C, with a maximum of 254 host addresses, is too small, while class-B, which allows up to 65534 addresses, is too large to be densely populated. The result is inefficient utilization of class-B network numbers.
- Routing information overload. The size and rate of growth of the routing tables in Internet routers is beyond the ability of current software (and people) to effectively manage.
- Eventual exhaustion of IP network numbers.

Classless Inter-Domain Routing (CIDR)

- Generalizes the subnet architecture
  - Imposes no restrictions on classes
  - Allows hierarchies of arbitrary length
  - Performs longest prefix matching at routers
- Example:

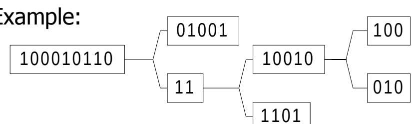

- Match
  - 10001011 01110010 01001110 01001001 ?

Midterm

Average grade: 46.7/50 (A)

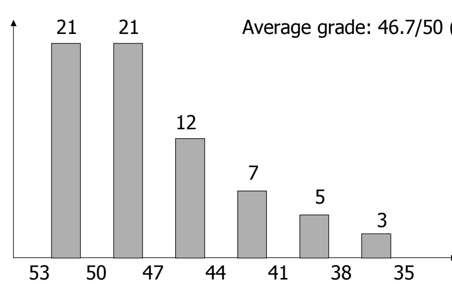

Review of Subnetting

- Problem 16, p. 356:
  - 128.96.39.0 255.255.255.128 Port 1
  - 128.96.39.128 255.255.255.128 Port 2
  - 128.96.40.0 255.255.255.128 Port 3
  - 196.4.153.0 255.255.255.192 Port 4
  - Other (default) Port 5
- Where do the following packets go?
  - 128.96.39.10
  - 128.96.40.12
  - 128.96.40.151
  - 192.4.153.17
  - 192.4.153.90

The Routing Problem

- Routing is a graph-theoretic problem
  - What is the least cost path from source to destination?

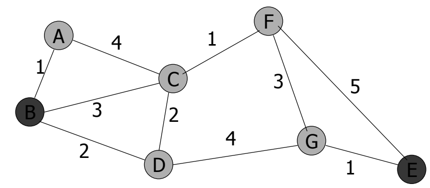

Distance Vector Routing (RIP)

- Each node maintains a distance table with entries of the form:
  - `(Destination, Cost, NextHop)`
- Nodes exchange updates with directly connected neighbors (triggered versus periodic updates)
- Each update is a list of pairs:
  - `(Destination, Cost)`
- Update routing table if receive a "better" route

Historical Perspective

- Initial introduction:
  - IETF RFC 1058, 1988 – Routing Information Protocol (RIP)
- IETF RFC 1723, 1994 – RIP v2
  - Authentication
  - Subnetting/Supernetting
  - Multicast
- RIPng for IPv6, 1997
- Today: IETF RIP Working Group
  - RIP disussion: ietf-rip@baynetworks.com

Initial State

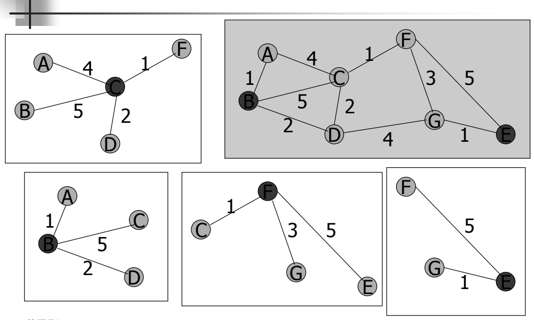

Routing Updates

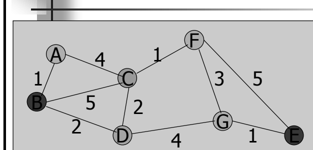

Example:

B 's initial view

D 's update

B 's new view =
min (B's current view, D's update + distance to D)

C 's update

B 's new view =
min (B's current view, C's update + distance to C)

A	B	C	D	E	F	G
1	*	5	2	-	-	-
-	2	2	*	-	-	4
1	*	4 D	2	-	-	6 D
4	5	*	2	-	1	-
1	*	4 D	2	-	5 D	6 D

Routing Table Information

- After first set of message exchanges
  - All neighbors two hops away are known
  - Optimal 2-hop routes are known
- After second set of message exchanges
  - All neighbors three hops away are known
  - Optimal 3-hop routes are known
- After the n^{th} set of message exchanges?

The Count to Infinity Problem

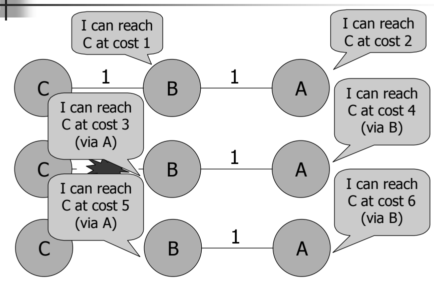

Avoiding Count to Infinity

- Bounded hop count
  - Stop counting early (e.g., at count 16)
- Split horizon
  - Don't send to node A updates you derived from A's reported distances
- Split horizon with poison reverse
- Remaining problem:
  - Cycles with 3 or more nodes.

Split Horizon

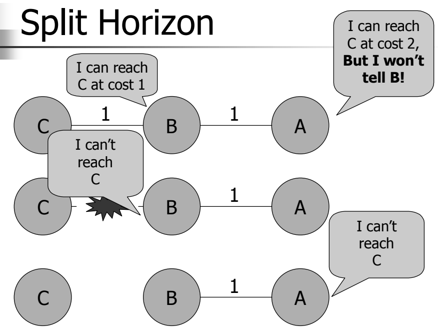

Split Horizon with Poison Reverse

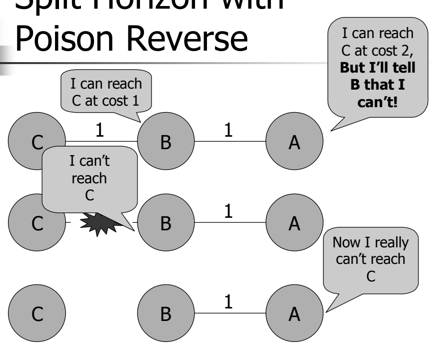

Link State Routing (OSPF)

- Current version: OSPF v2 (RFC 2328, 1998)
- Strategy
  - send to all nodes (not just neighbors) information about directly connected links (not entire routing table)
- Link State Packet (LSP)
  - id of the node that created the LSP
  - cost of link to each directly connected neighbor
  - sequence number (SEQNO)
  - time-to-live (TTL) for this packet

Dijkstra's Algorithm

- Finds the shortest path from a node s to every other node in a graph

Marked = {*s*}
For each node *n* (other than *s*) in the graph
    Cost(*n*) = link(*s*,*n*)
While some nodes remain unmarked
    Find the minimum distance unmarked node from *s*, say *w*
    Marked = Marked + *w*
    For each unmarked node, n
        C(*n*) = min {C(*n*), C(*w*) + link (*w*,*n*)}

CS457 IP Layer

Example:

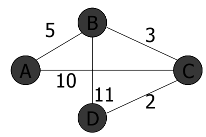

Confirmed	Tentative
D	B(11), C(2)
D, C(2)	$B(5_C)$, $A(12_C)$
D, C(2), $B(5_C)$	$A(10_C)$
D, C(2), $B(5_C)$, $A(10_C)$	

CS457 IP Layer

More Routing Information

- IETF homepage
  - www.ietf.org
- Discussion lists
  - ospf@discuss.microsoft.com
  - ietf-rip@baynetworks.com

CS457 IP Layer