CSE331: Introduction to Networks and Security

Lecture 8

Fall 2006

Announcements

- Reminders:
  - Project I is due on Monday, Sept. 25th.
  - Homework 1 is due on Friday, Sept. 29th.

Internet Protocol Interoperability

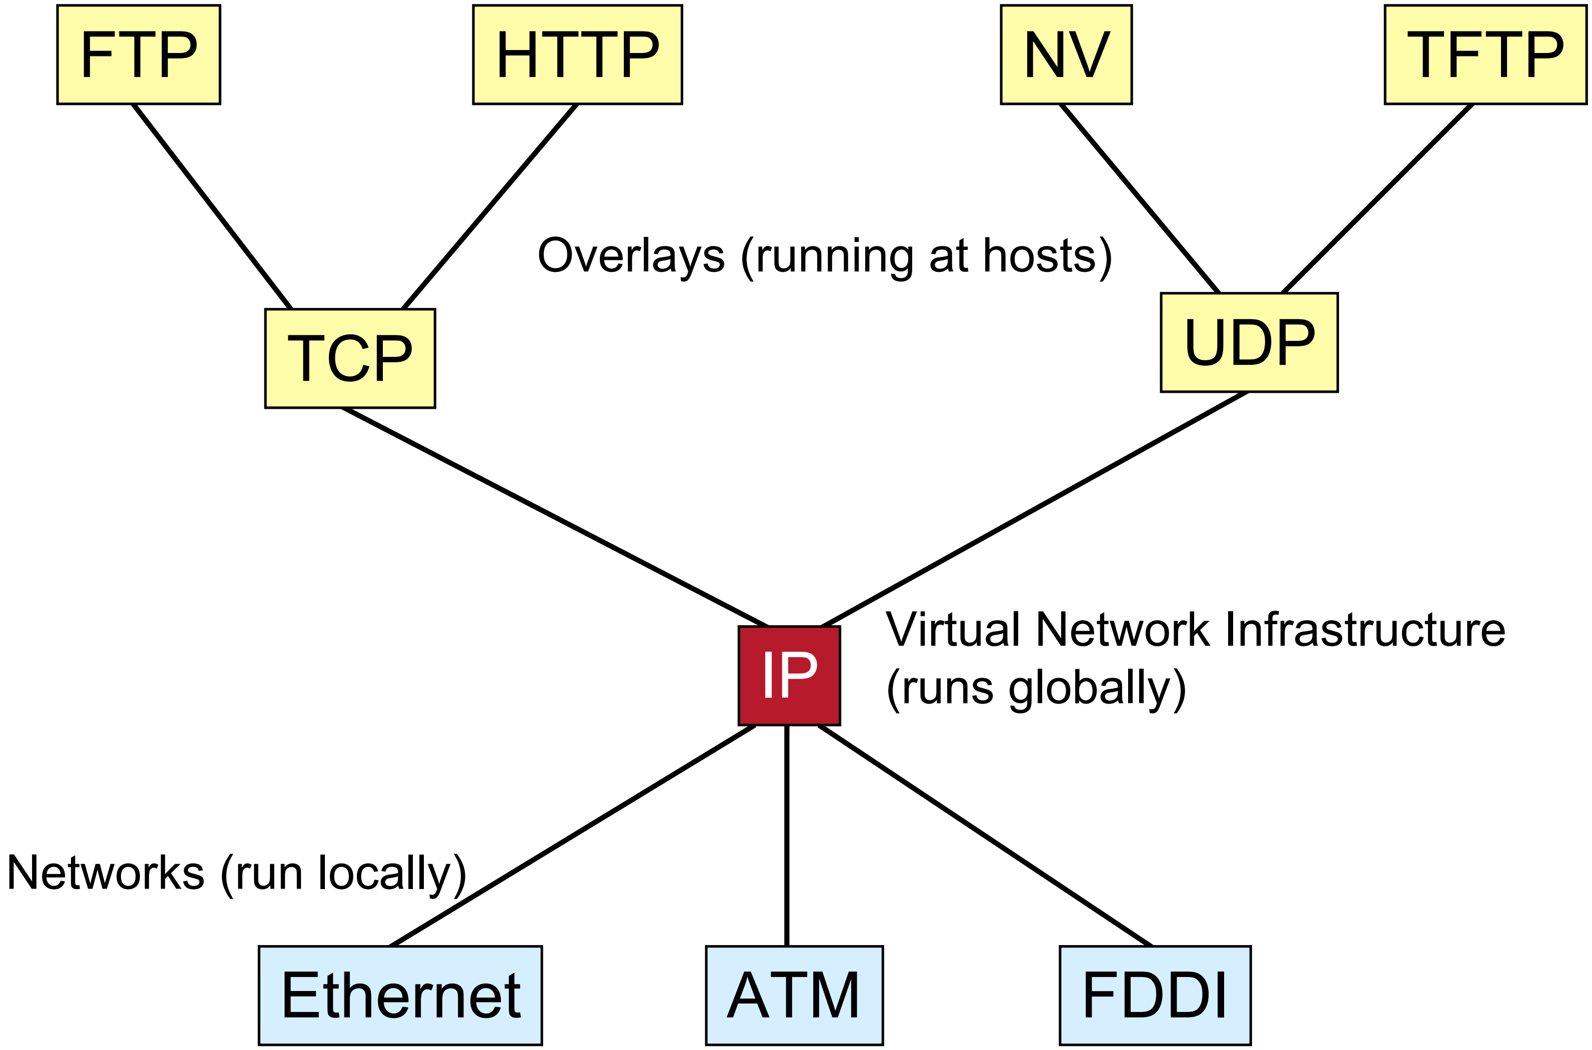

IPv4 Packet Format

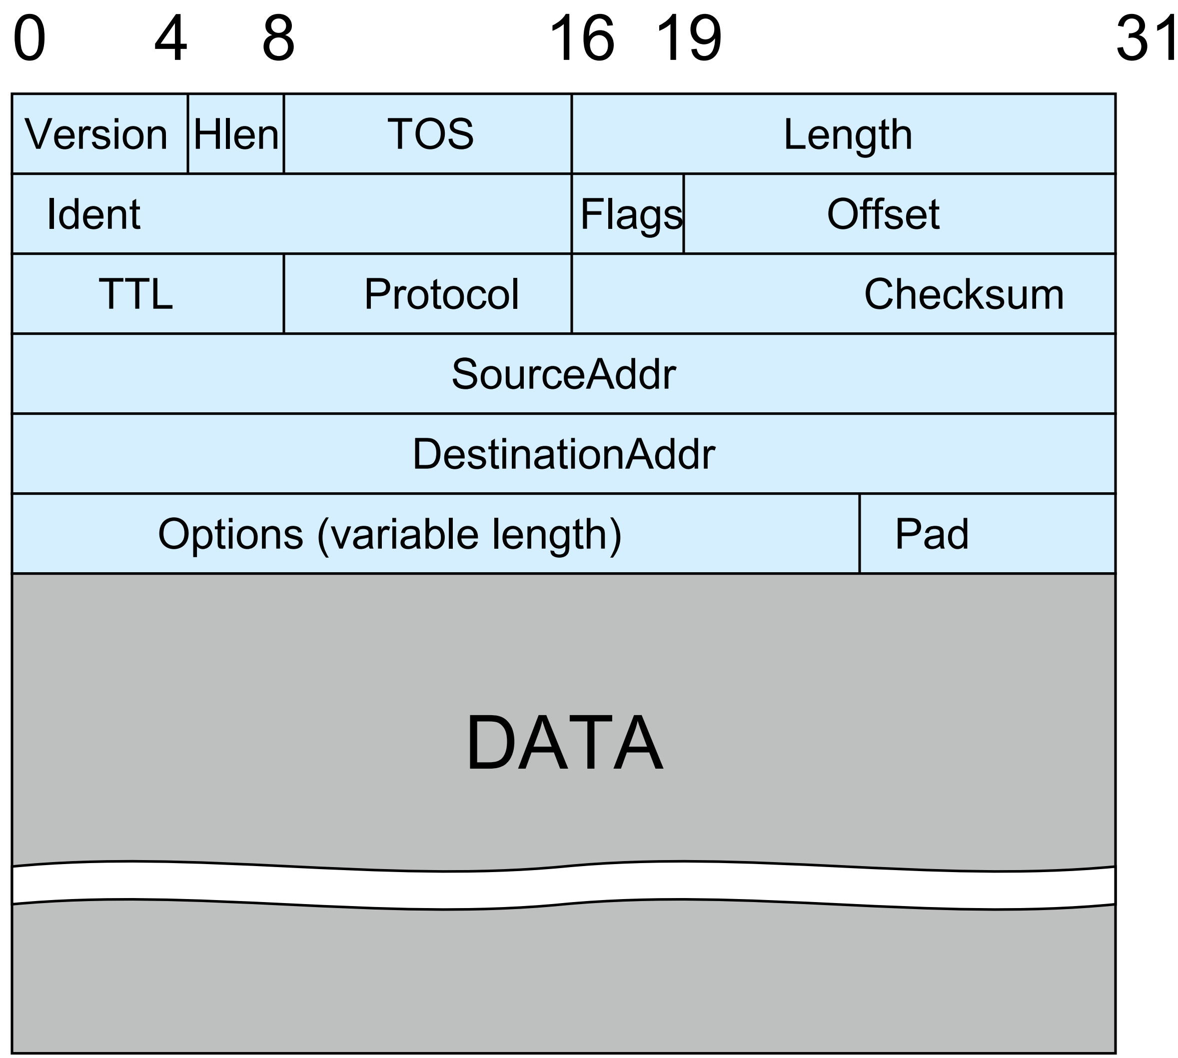

Scaling Problems

- Not enough network numbers.
  – Class C network with 2 nodes wastes 253 IP addresses
  – Class B network with ~300 nodes wastes 64,000 IP addresses
  – Only 2^{14} ~ 16,500 class B networks
- Routing information too cumbersome.
  – More networks means larger routing tables

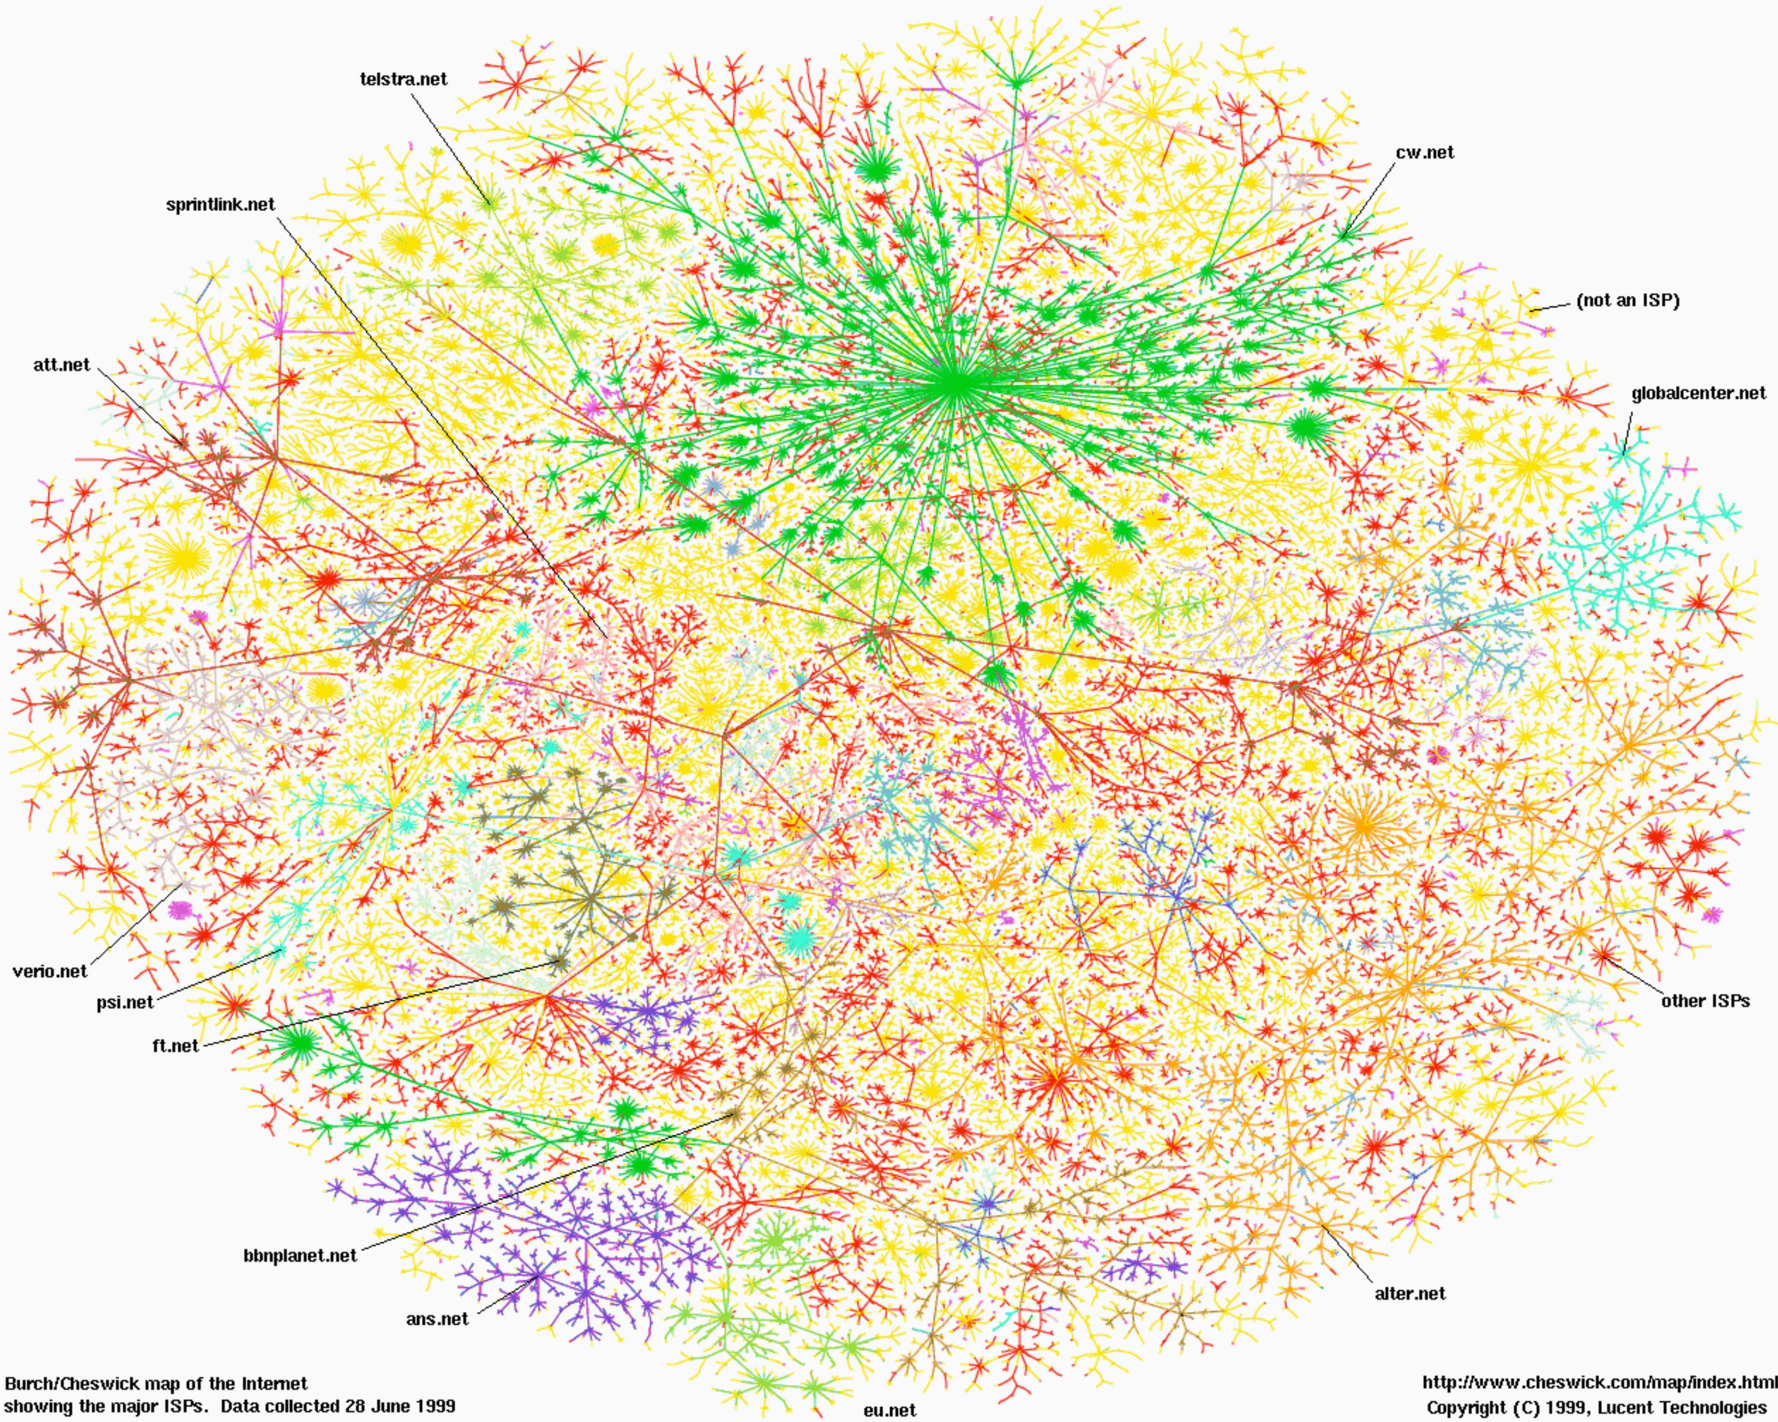

Burch/Cheswick map of the Internet
showing the major ISPs. Data collected 28 June 1999

http://www.cheswick.com/map/index.html

Subnetting

- Idea: One IP network number allocated to several physical networks.
  - The multiple physical networks are called *subnets*
  - Should be close together (why?)
  - Useful when a large company (or university!) has many physical networks.

IP addresses

- Hierarchical, not flat as in Ethernet

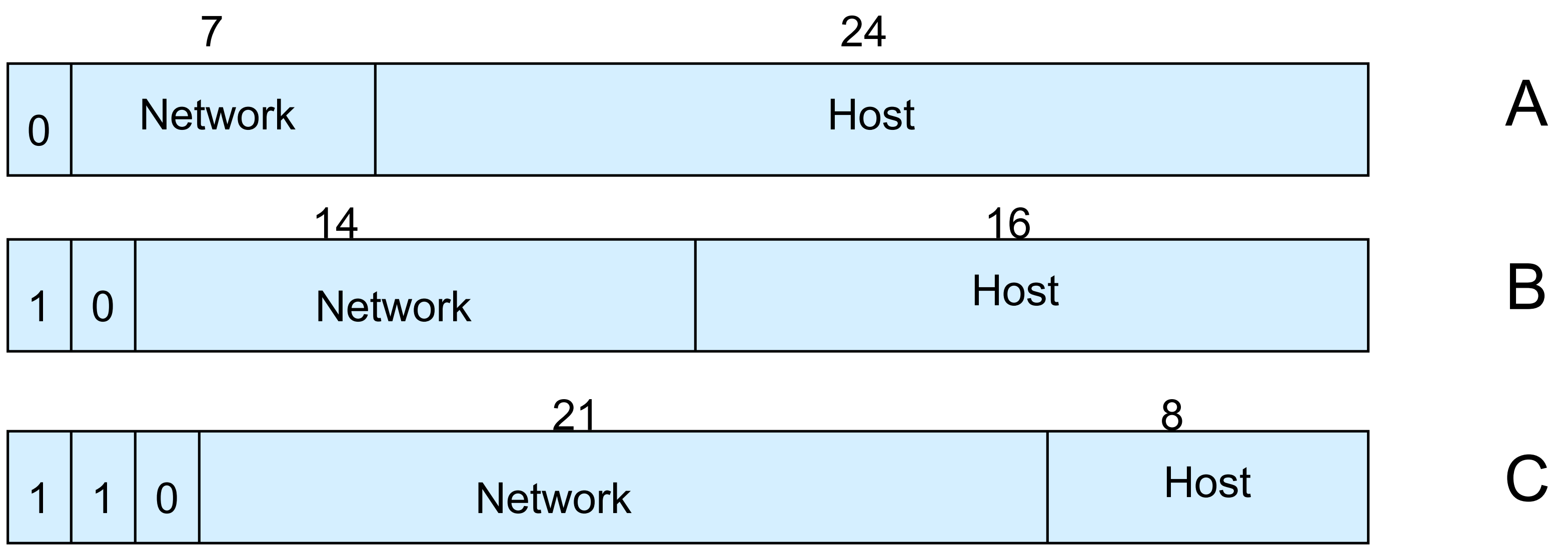

- Written as four decimal numbers separated by dots: 158.130.14.2

Subnet Numbers

- *Subnetting*
  - All nodes are configured with *subnet mask*
  - Allows definition of a *subnet number*
    - All hosts on a physical subnetwork share the same *subnet number*

Subnet Mask (255.255.255.0)

111111111111111111111111	00000000

Subnetted Address:

Network number	Subnet ID	Host ID

Example of Subnetting

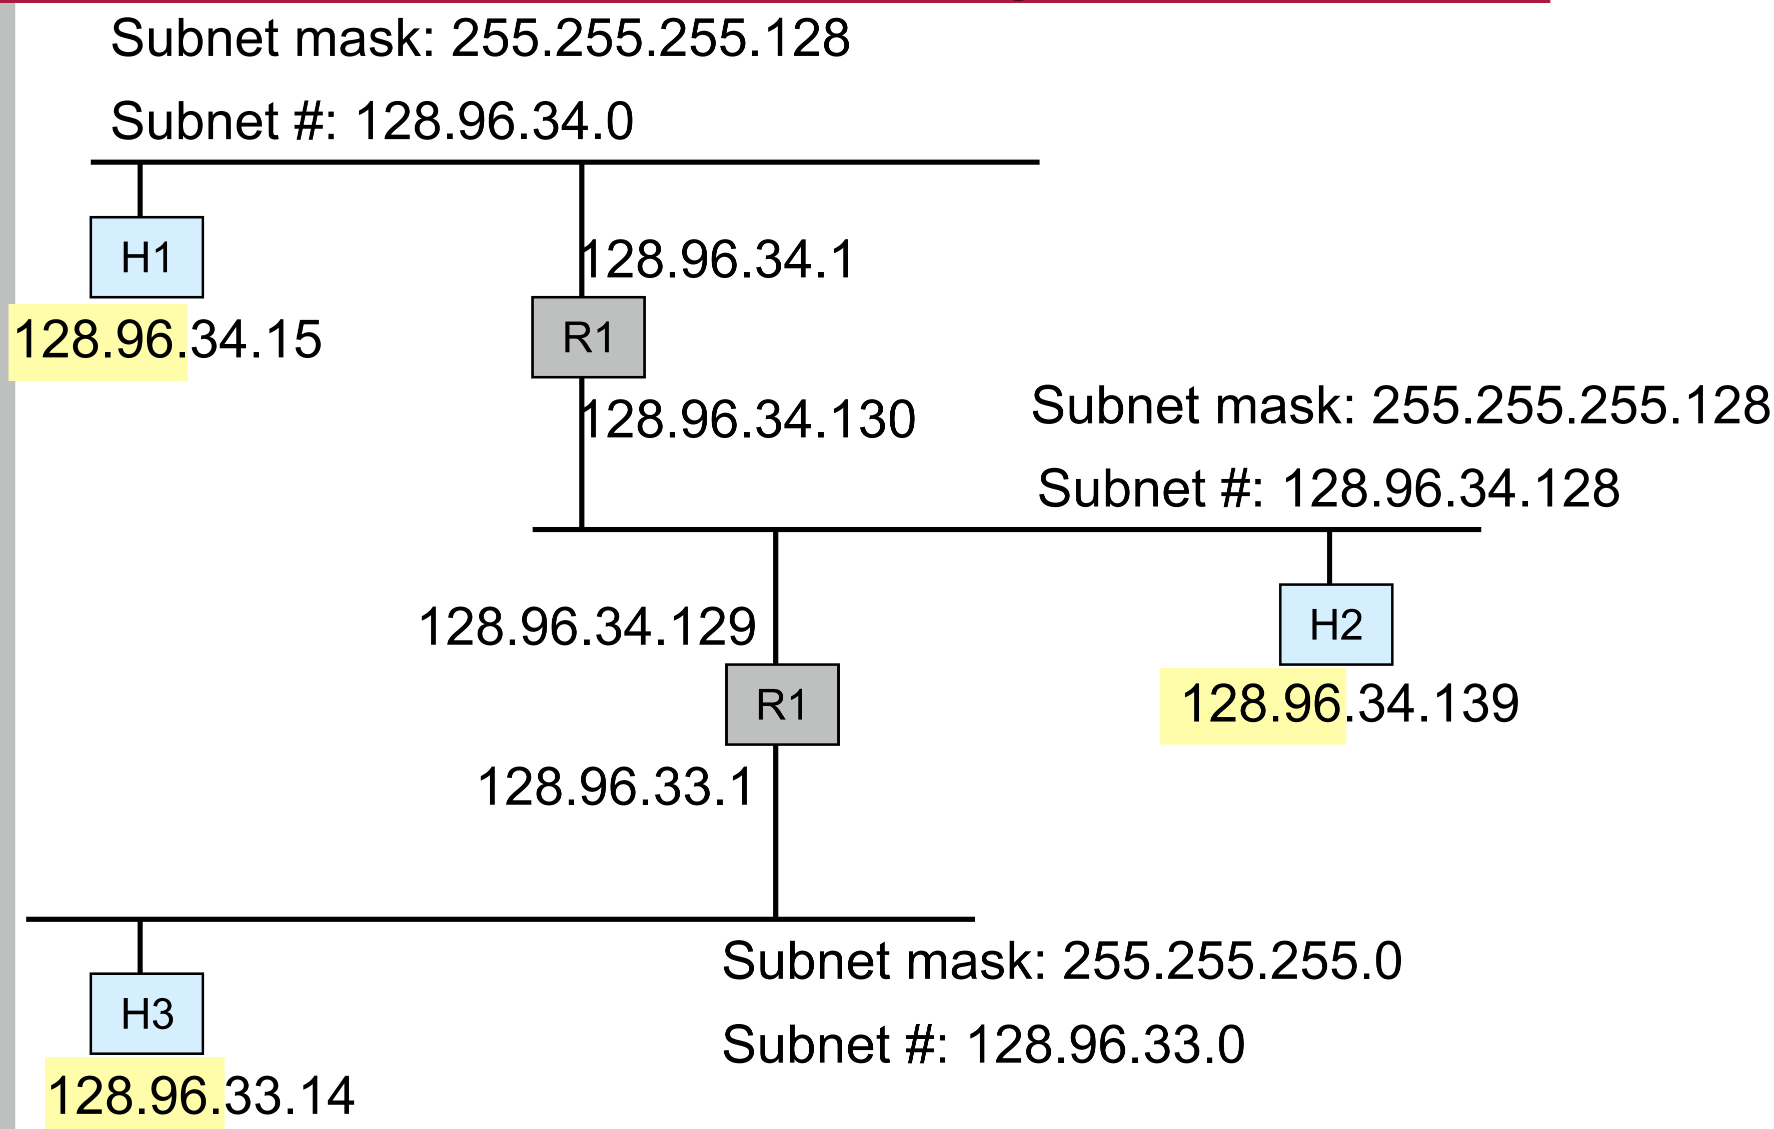

Subnets, continued

- Mask is bitwise-ANDed with address
- This is done at routers
- Router tables in this model:
  - <Subnet #, Subnet Mask, NextHop>
- Subnetting allows a set of physical networks to look like a single logical network from elsewhere

Forwarding Algorithm

```
D = destination IP address
for each forwarding table entry
(SubnetNumber, SubnetMask, NextHop)
   D1 = SubnetMask & D
   if  D1 = SubnetNumber
      if NextHop is an interface
         deliver datagram directly to destination
      else
         deliver datagram to NextHop (router)
```

Intradomain Routing

- RIP - Routing Information Protocol
  - Uses distance vector algorithm
  - Limited to small nets; <15 hops
- OSPF - Open Shortest Path First
  - Augmented version of link-state
  - Augmentation includes authentication, load-balancing, and defined areas

Spanning Trees (Abstractly)

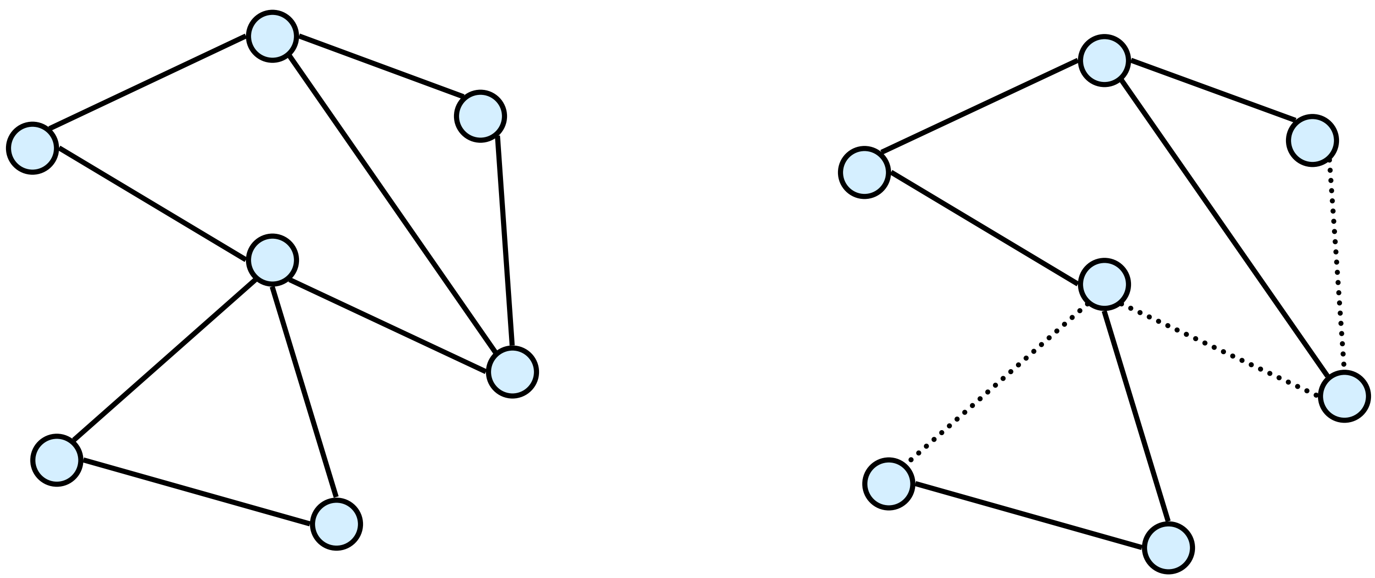

- Given a connected graph G
- A *spanning tree* is an acyclic, connected subgraph of G that contains all the nodes.

Spanning Tree Algorithm (Abstractly)

- Pick a root node
- Compute shortest paths to root
- Need to break ties

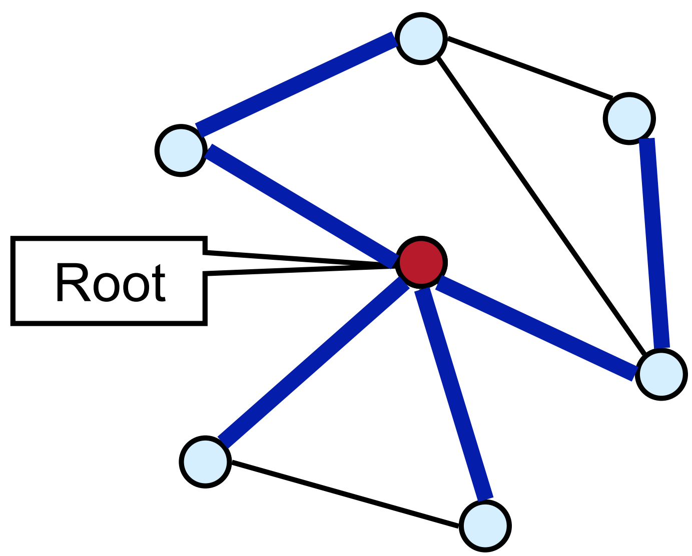

Distance Vector Algorithm (RIP)

- Similar to the Spanning Tree Algorithm
  – Except that information about distance to ALL nodes is forwarded (not just info. about root.)
  – Sometimes called Bellman-Ford algorithm

- Each node constructs a *Distance Vector*
  – Contains distances (costs) to reach all other nodes
  – Initially:
    - Distance to neighbors = 1
    - Distance to others = ∞
  – Routing table reflects node's beliefs

Example Network Graph

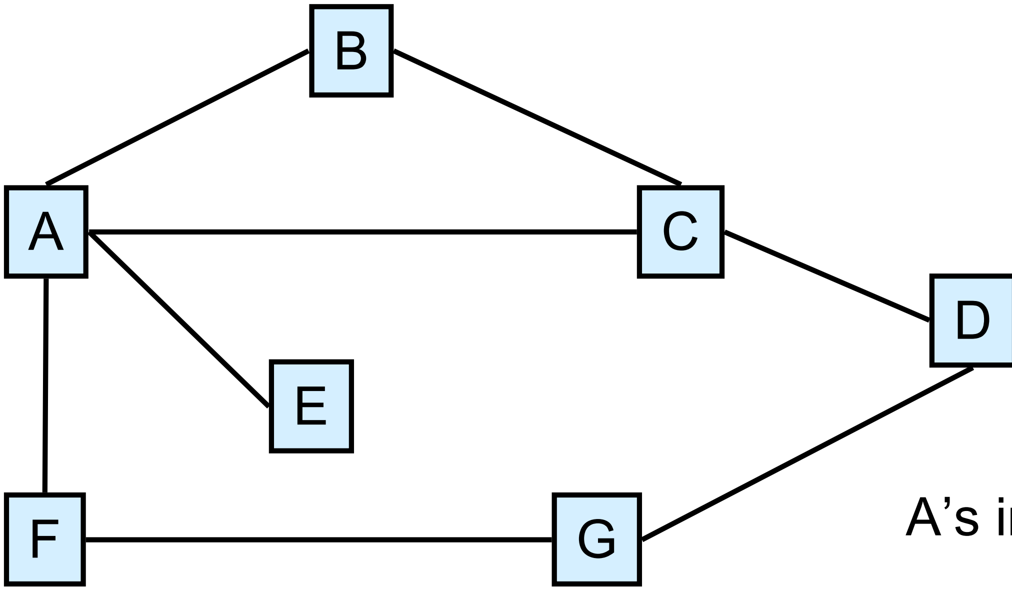

A's initial information

Dest.	Cost	NextHop
B	1	B
C	1	C
D	∞	-
E	1	E
F	1	F
G	∞	-

Iteration Steps

- Each host sends its DV to its neighbors
- Neighbors can update their distance vectors and routing information accordingly.
  - As in spanning tree, the nodes ignore worse information
  - Update any better routes
- If host changed its tables, send new DV to neighbors
- After a few iterations, routing information *converges*

Example Iteration Steps

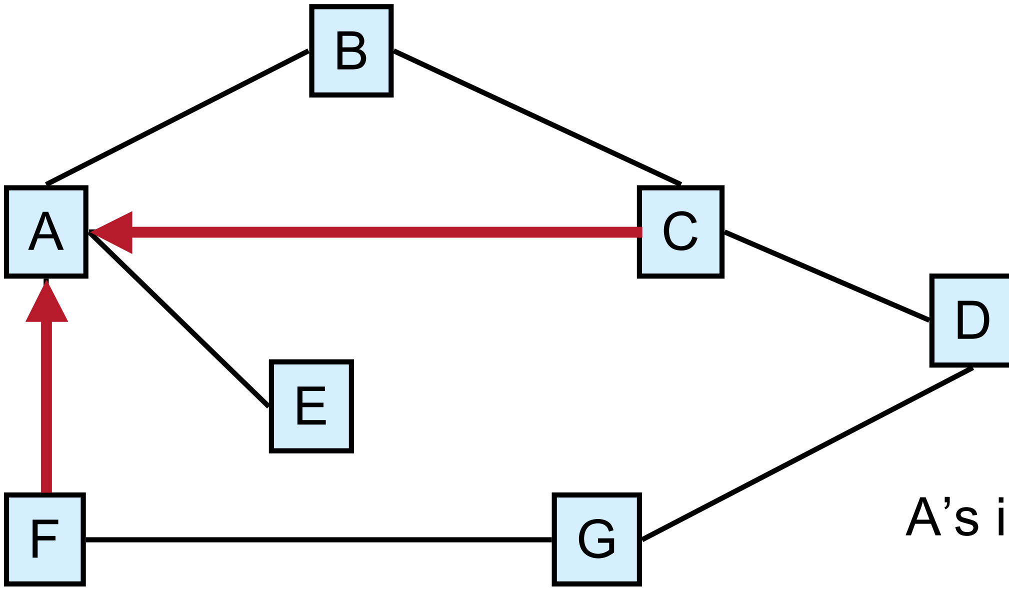

A's initial information

Dest.	Cost	NextHop
B	1	B
C	1	C
D	~~∞~~ 2	~~-~~ C
E	1	E
F	1	F
G	~~∞~~ 2	~~-~~ F

F sends A its DV.
- A discovers that G can be reached in two hops through F

C sends A its DV.
- A discovers that D can be reached in two hops through C

Details

- Note: No single host has all routing information.
- When to send update vectors?
  - When your routing table changes (triggered)
  - Periodically (“I’m alive!”)
- Detecting link/node failure
  - (1) Periodically exchange “I’m alive!” messages.
  - (2) Timeout mechanism