CSE331: Introduction to Networks and Security

Lecture 9

Fall 2004

Recap

- Spanning tree algorithm
- IPv4

- Today:
  - Subnetting
  - ARP / ICMP / DNS
  - Routing

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

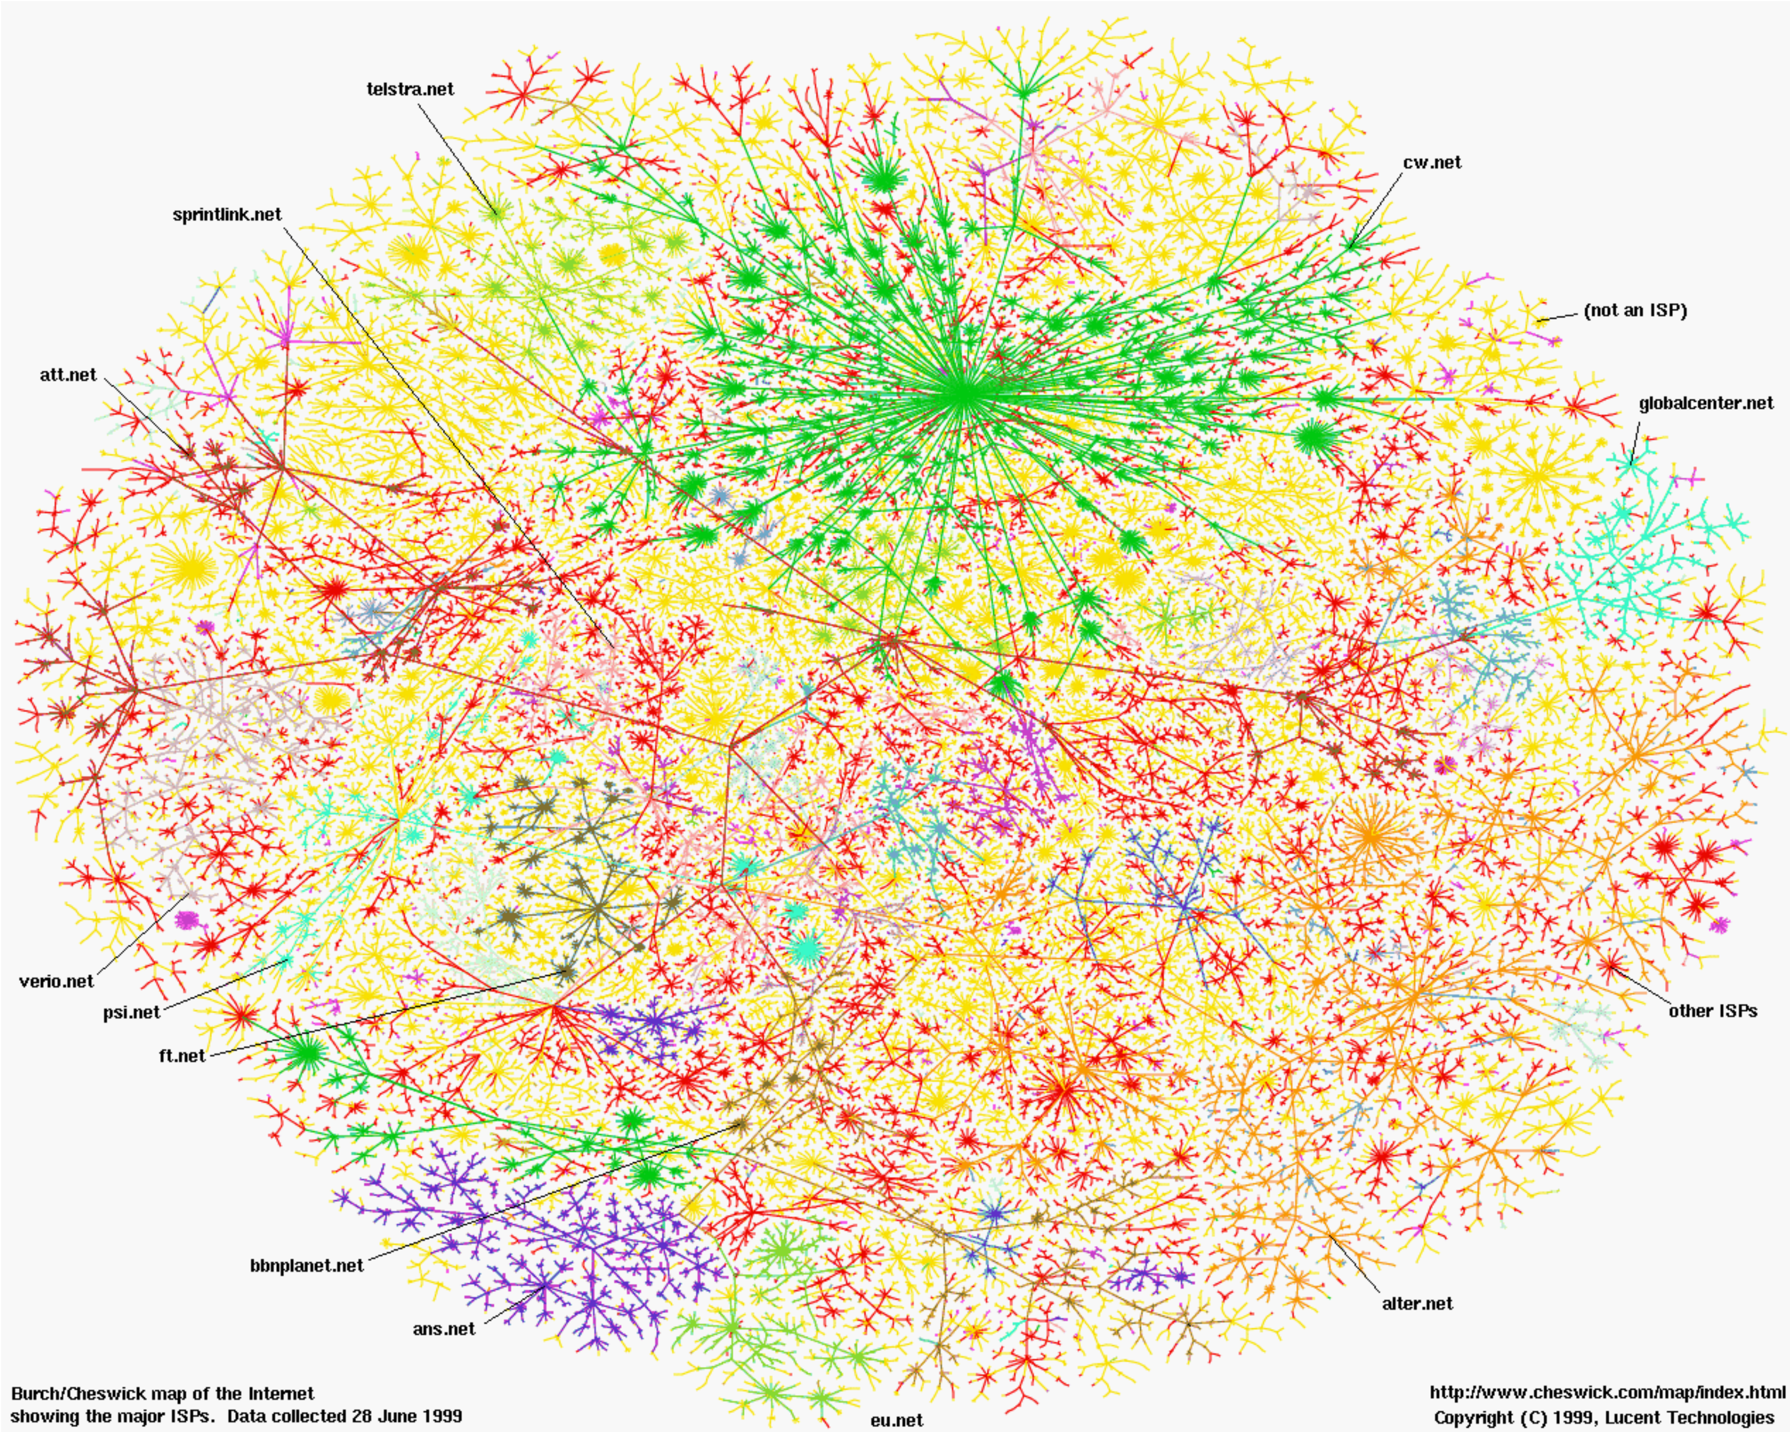

telstra.net
cw.net
sprintlink.net
(not an ISP)
att.net
globalcenter.net
verio.net
psi.net
ft.net
other ISPs
bbnplanet.net
ans.net
alter.net
eu.net
Burch/Cheswick map of the Internet
showing the major ISPs. Data collected 28 June 1999
http://www.cheswick.com/map/index.html
Copyright (C) 1999, Lucent Technologies

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

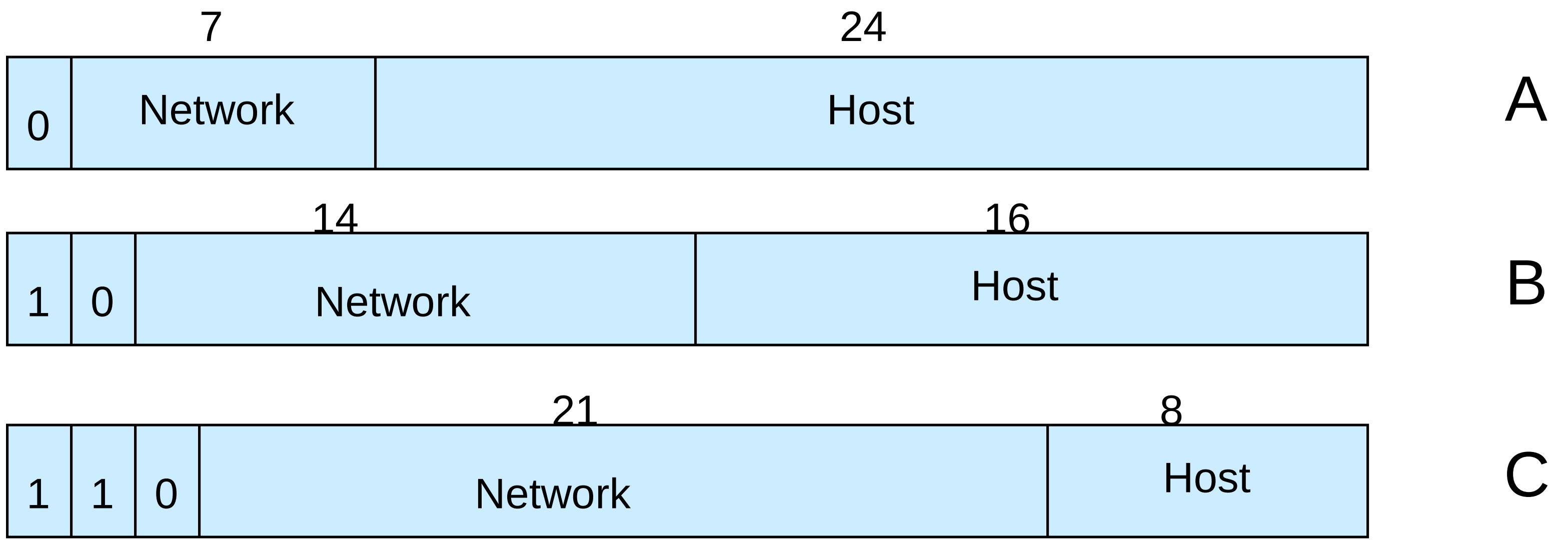

- Written as four decimal numbers separated by dots: 158.130.14.2

Subnet Numbers

- Solution: *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

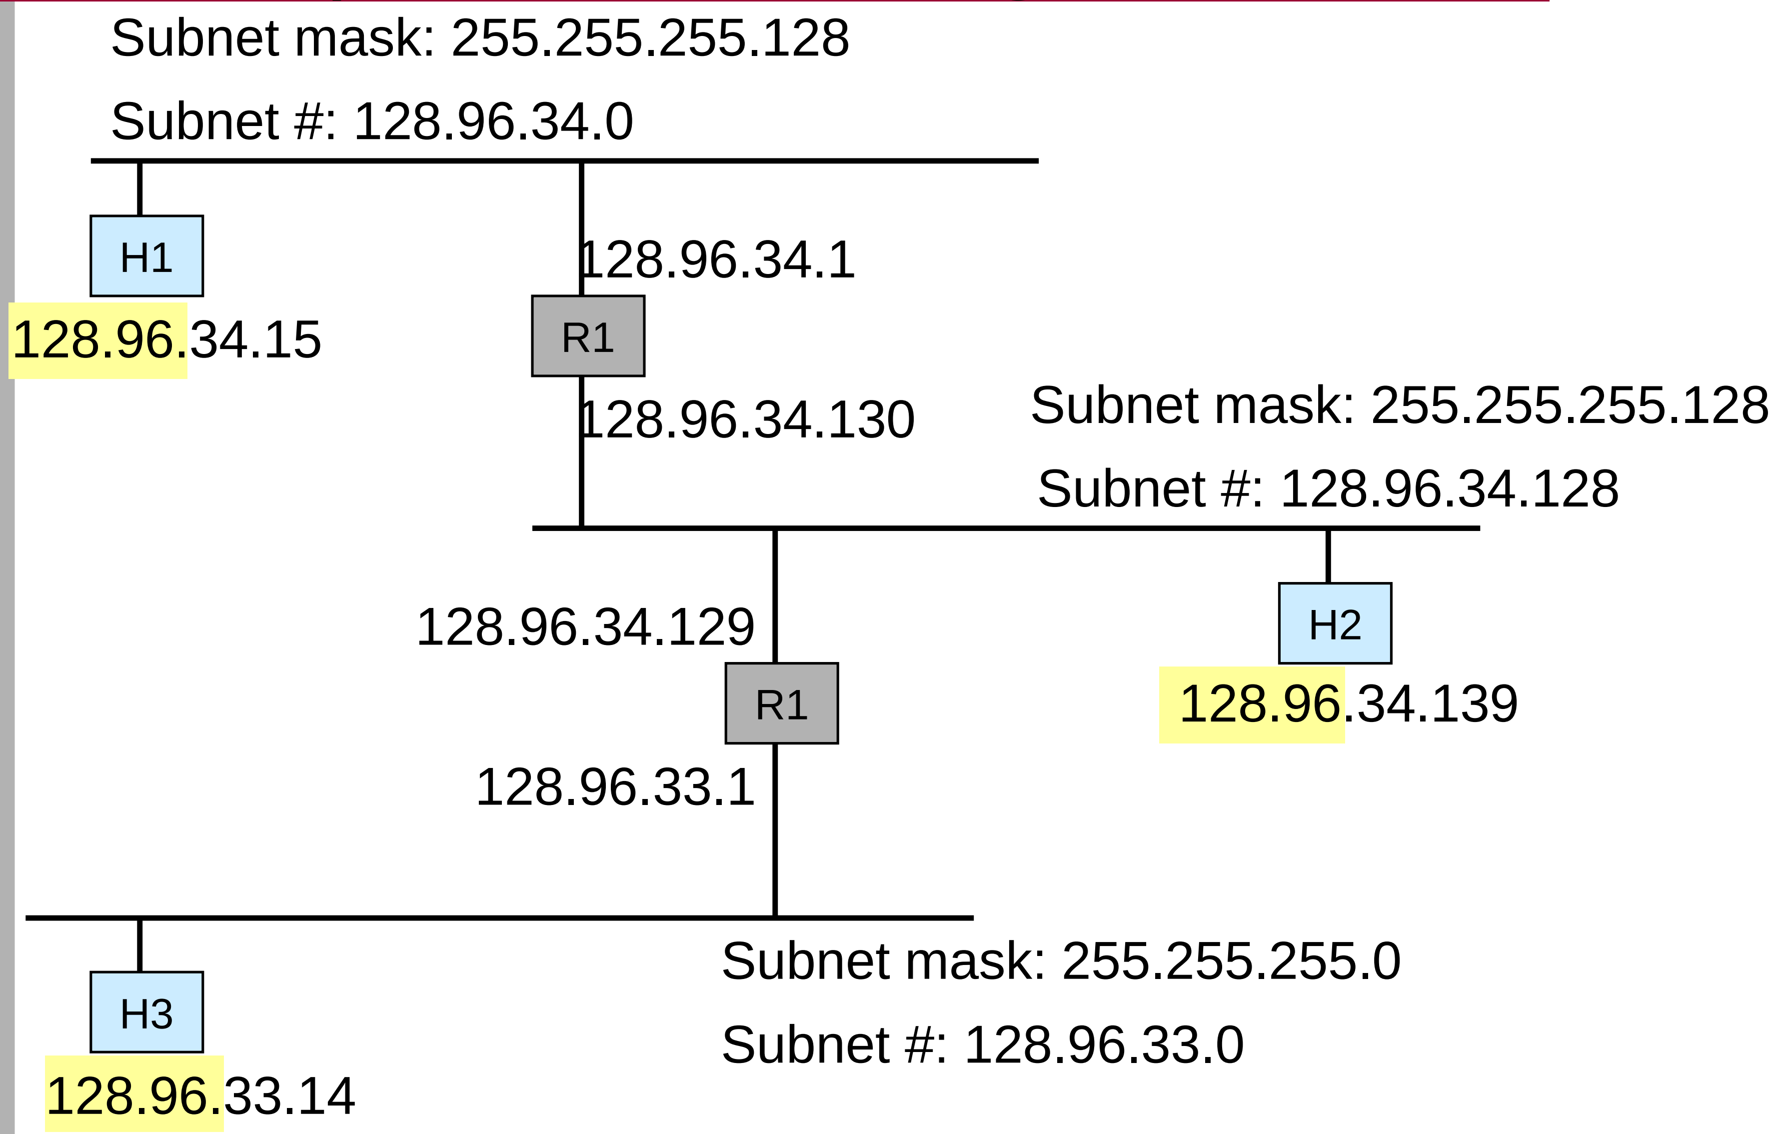

Subnets, continued

- 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)
```

ARP - Address Resolution Protocol

- Problem:
  - Need mapping between IP and link layer addresses.

- Solution: ARP
  - Every host maintains IP–Link layer mapping table (cache)
  - Timeout associated with cached info (15 min.)
- Sender
  - Broadcasts “Who is IP addr X?”
  - Broadcast message includes sender’s IP & Link Layer address
- Receivers
  - Any host with sender in cache “refreshes” time-out
  - Host with IP address X replies “IP X is Link Layer Y”
  - Target host adds sender (if not already in cache)

ICMP: Internet Control Message Protocol

- Collection of error & control messages
- Sent back to the source when Router or Host cannot process packet correctly
- Error Examples:
  - Destination host unreachable
  - Reassembly process failed
  - TTL reached 0
  - IP Header Checksum failed
- Control Example:
  - Redirect – tells source about a better route

Domain Name System

- System for mapping mnemonic names for computers into IP addresses.

  zeta.cis.upenn.edu → 158.130.12.244

- Domain Hierarchy
- Name Servers
- Name Resolution

Domain Name Hierarchy

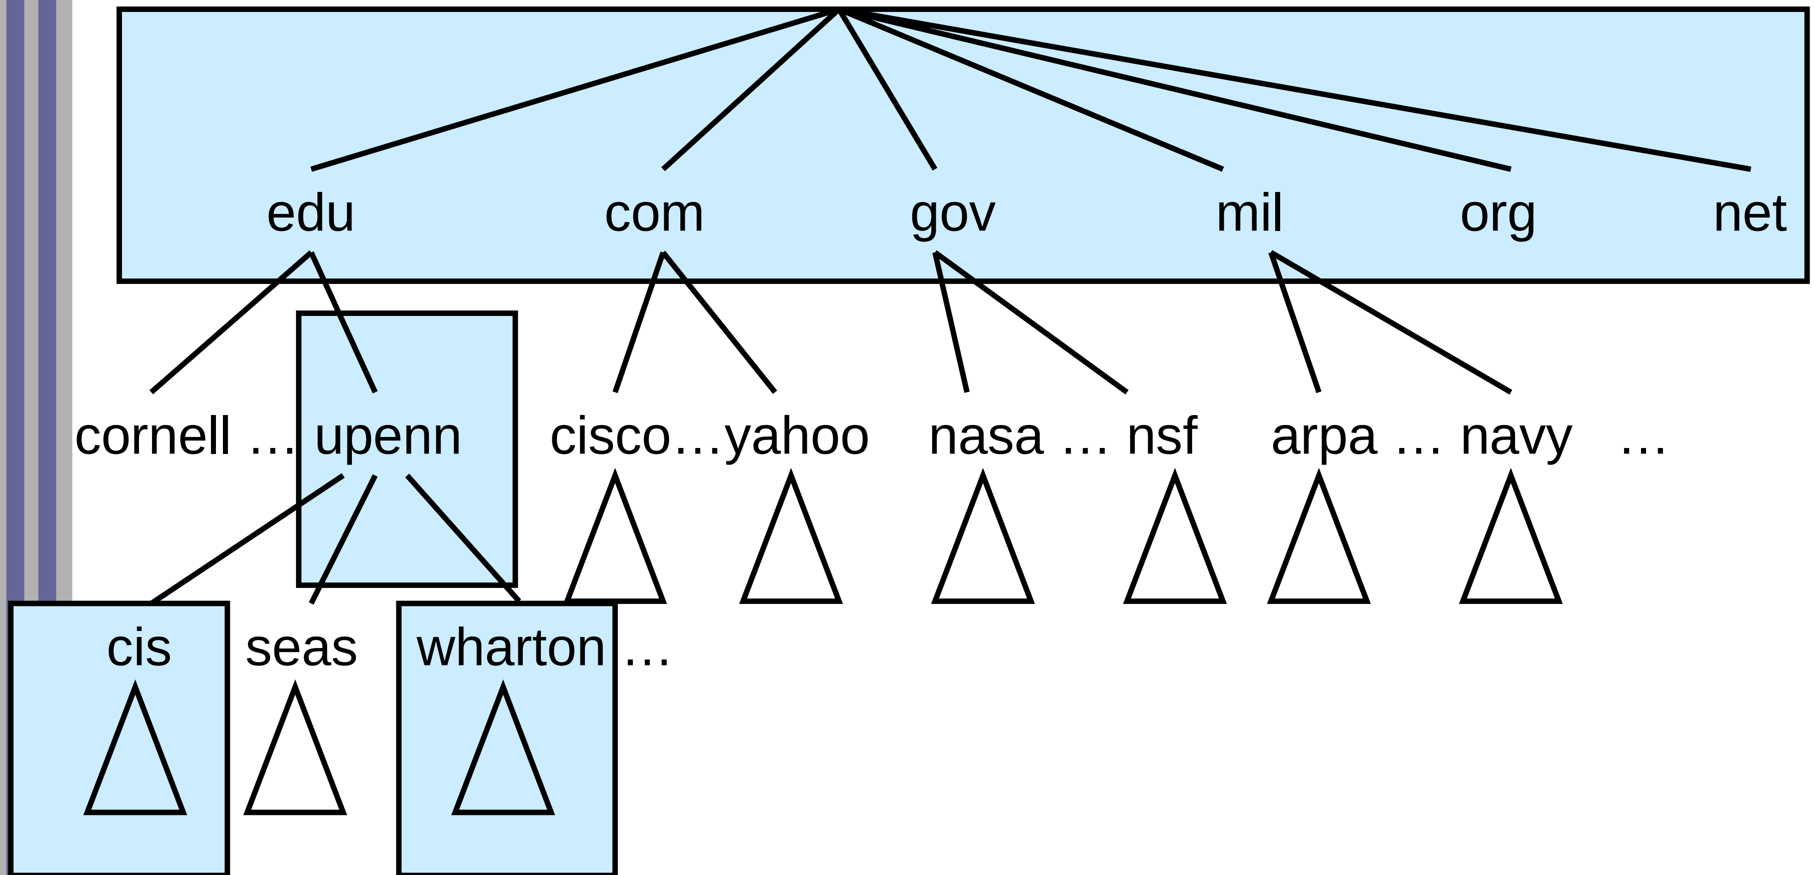

Hierarchy of Name Servers

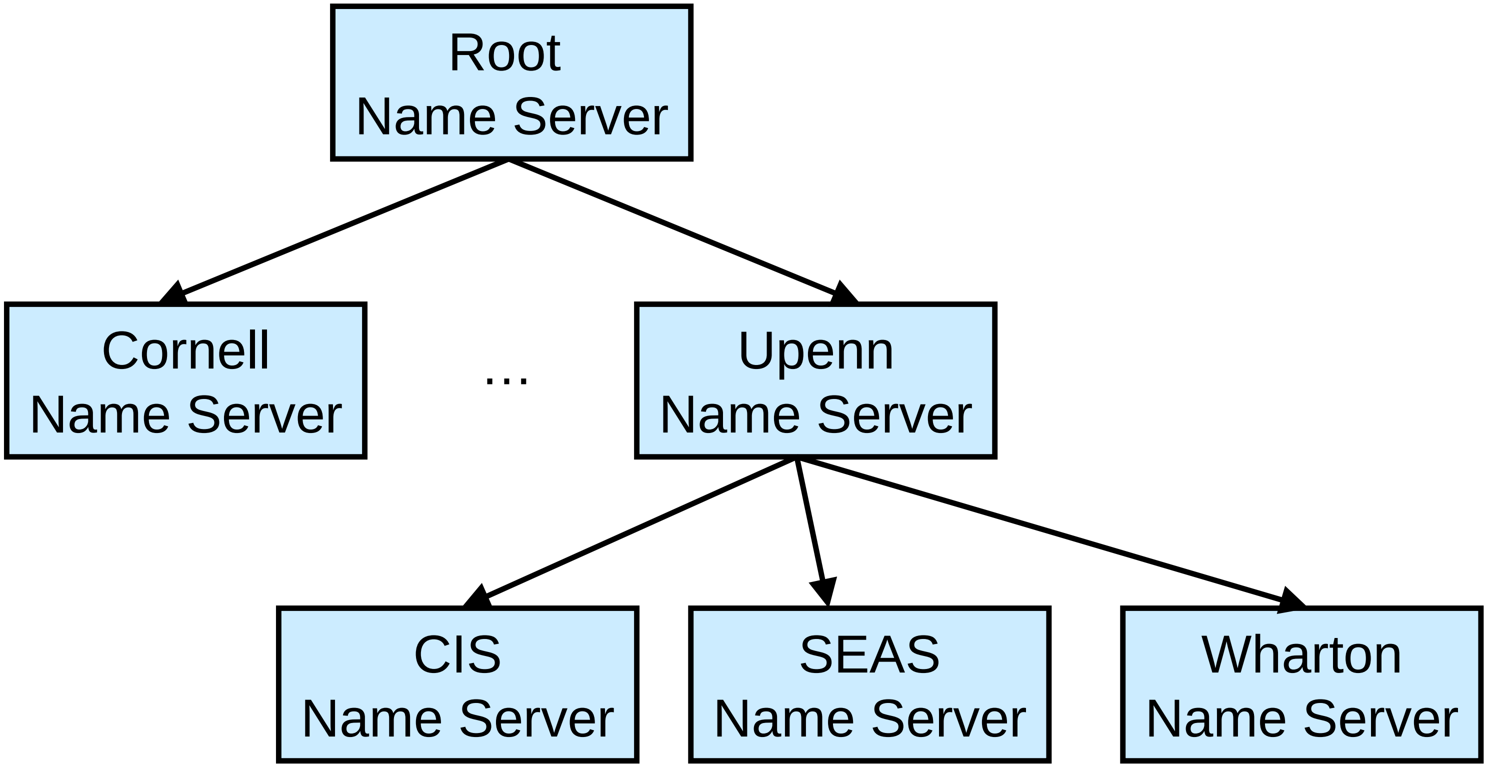

Records on Name Servers

- < Name, Value, Type, Class >
- Types
  - A Host to address mappings
  - NS Name server address mappings
  - CNAME   Aliases
  - MX Mail server mappings
- Class IN for IP addresses

Name resolution

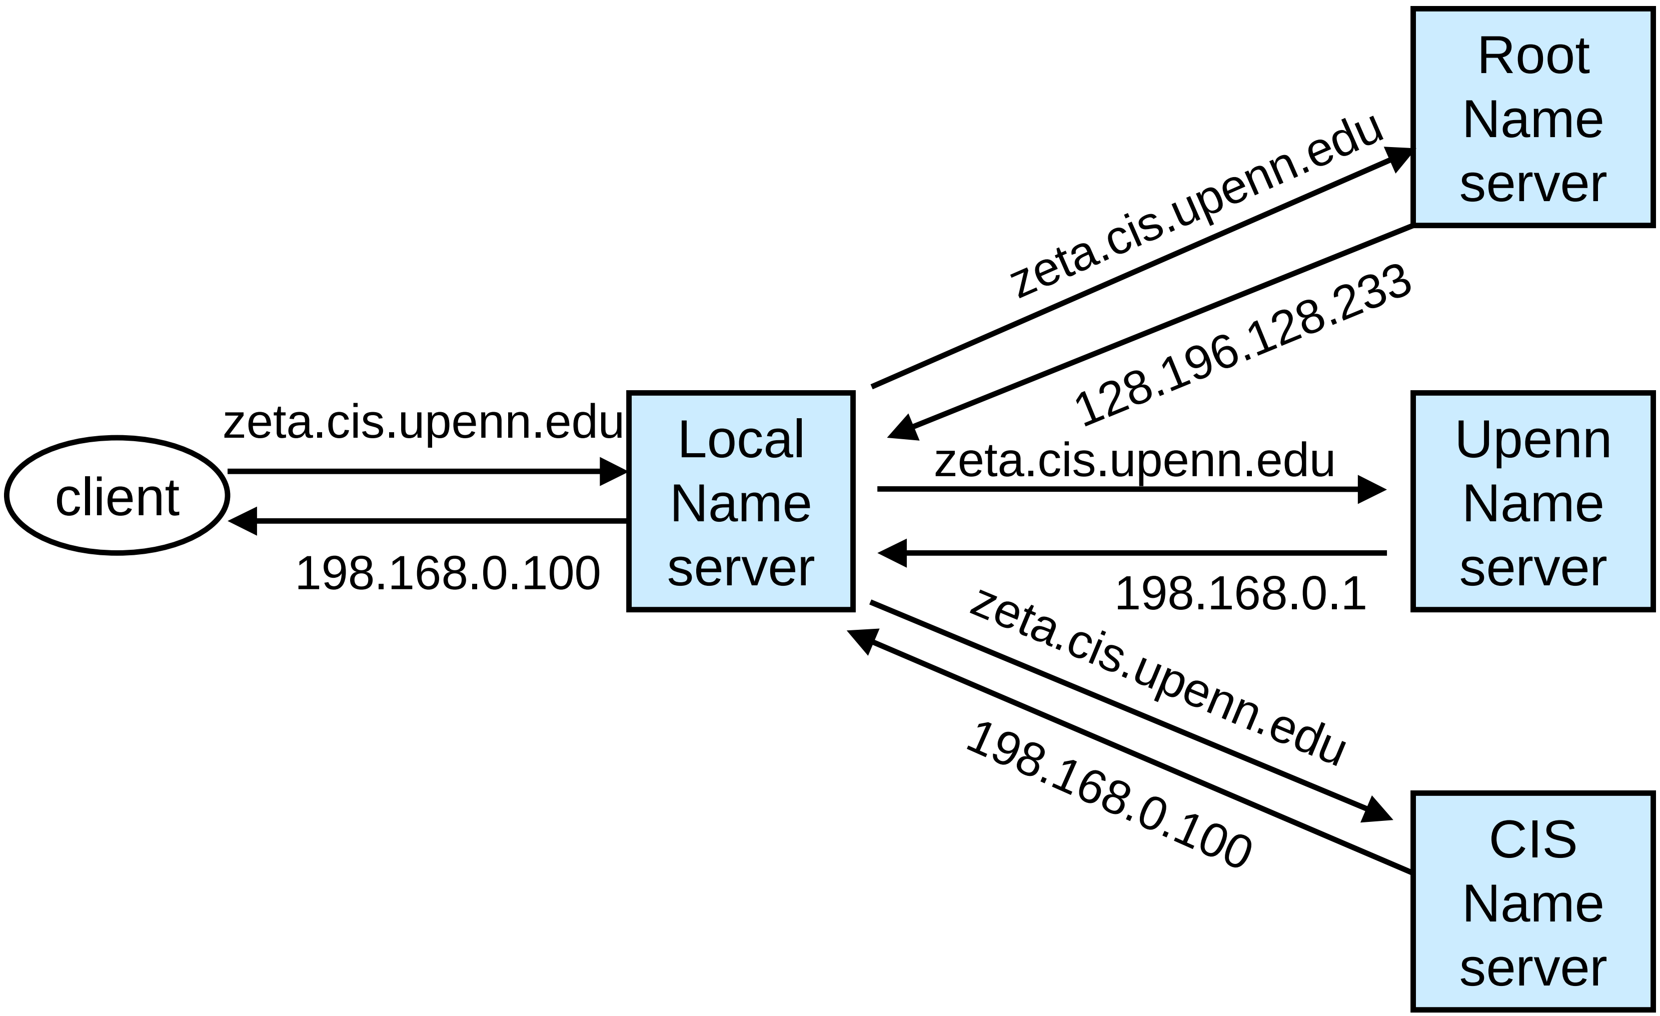

IP Routing

- Begin by partitioning problem:
- Intradomain Routing
  – Inside *administrative domains* (AD's)
- Interdomain Routing
  – Between administrative domains (e.g., companies)
  – Exterior Gateway Protocol (EGP)
  – Border Gateway Protocol (BGP) [Replaced EGP]

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

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

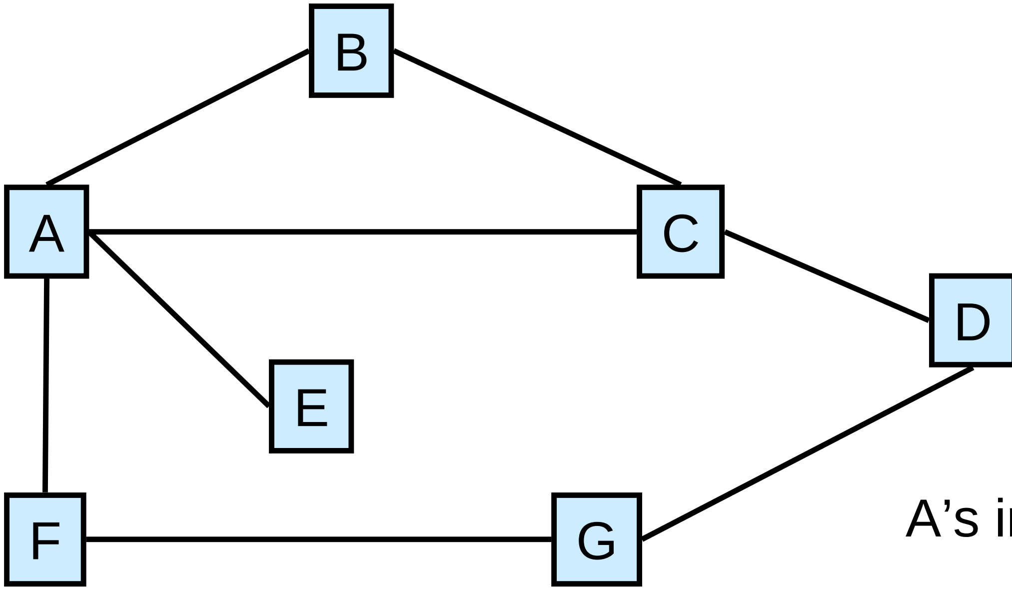

A's initial information

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

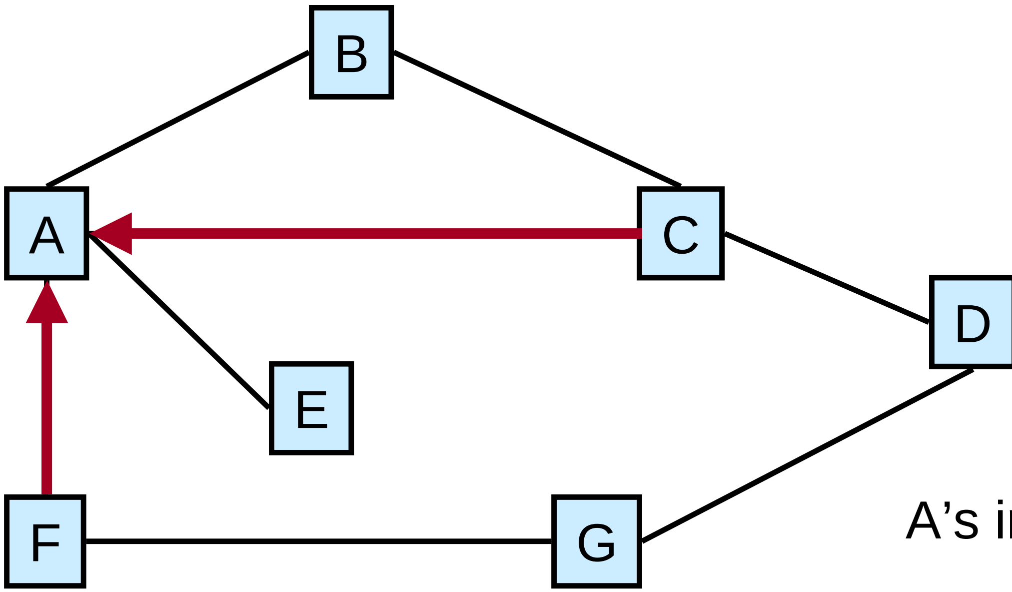

A's initial information

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