PART XV

ROUTING: INSIDE AN AUTONOMOUS SYSTEM (RIP, OSPF, HELLO)

Static Vs. Dynamic Interior Routes

- Static routes
  - Initialized at startup
  - Never change
  - Typical for host
  - Sometimes used for router
- Dynamic router
  - Initialized at startup
  - Updated by route propagation protocols
  - Typical for router
  - Sometimes used in host

Illustration Of Topology In Which Static Routing Is Optimal

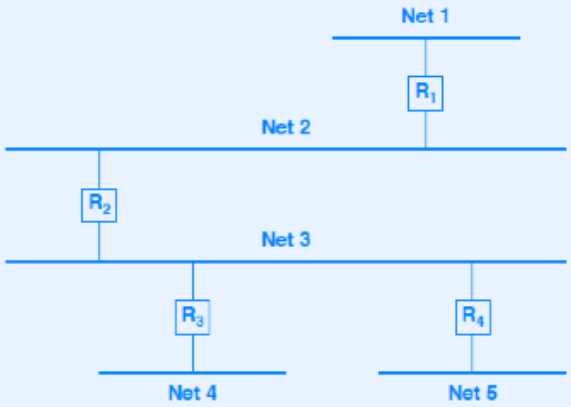

- Only one route exists for each destination

Illustration Of Topology In Which Dynamic Routing Is Needed

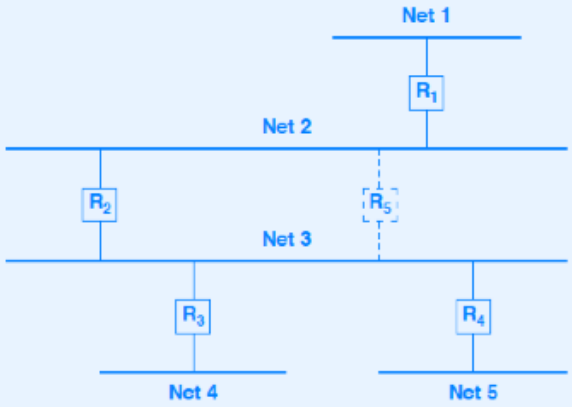

- Additional router introduces multiple paths

Exchanging Routing Information Within An Autonomous System

- Mechanisms called interior gateway protocols, IGPs
- Choice of IGP is made by autonomous system
- Note: if AS connects to rest of the world, a router in the AS must use an EGP to advertise network reachability to other autonomous systems.

Example Of Two Autonomous Systems And the Routing Protocols Used

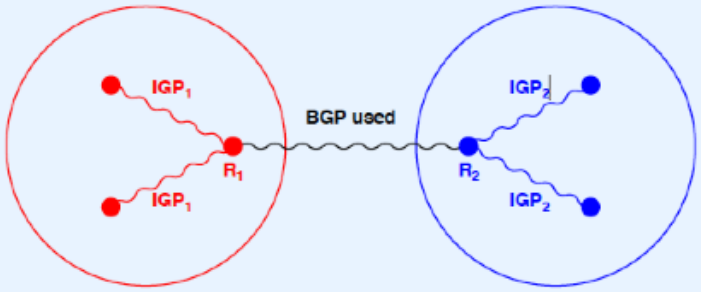

Example IGPs

- RIP
- HELLO
- OSPF

Routing Information Protocol (RIP)

- Implemented by UNIX program *routed*
- Uses hop count metric
- Distance-vector protocol
- Relies on broadcast
- Assumes low-delay local area network
- Uses split horizon and poison reverse techniques to solve inconsistencies
- Current standard is RIP2

Two Forms Of RIP

- Active
  - Form used by routers
  - Broadcasts routing updates periodically
  - Uses incoming messages to update routes
- Passive
  - Form used by hosts
  - Uses incoming messages to update routes
  - Does not send updates

Illustration Of Hosts Using Passive RIP

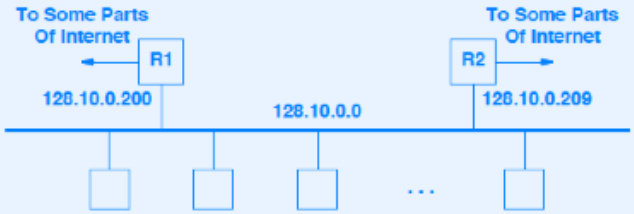

- Host routing table initialized to:

Destination	Route
128.10.0.0	direct
default	128.10.0.200

- Host listens for RIP broadcast and uses data to update table
- Eliminates ICMP redirects

RIP Operation

- Each router sends update every 30 seconds
- Update contains pairs of
  (destination address, distance)
- Distance of 16 is *infinity* (i.e., no route)

Slow Convergence Problem (Count To Infinity)

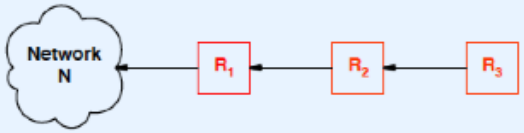

Routers with routes to network N

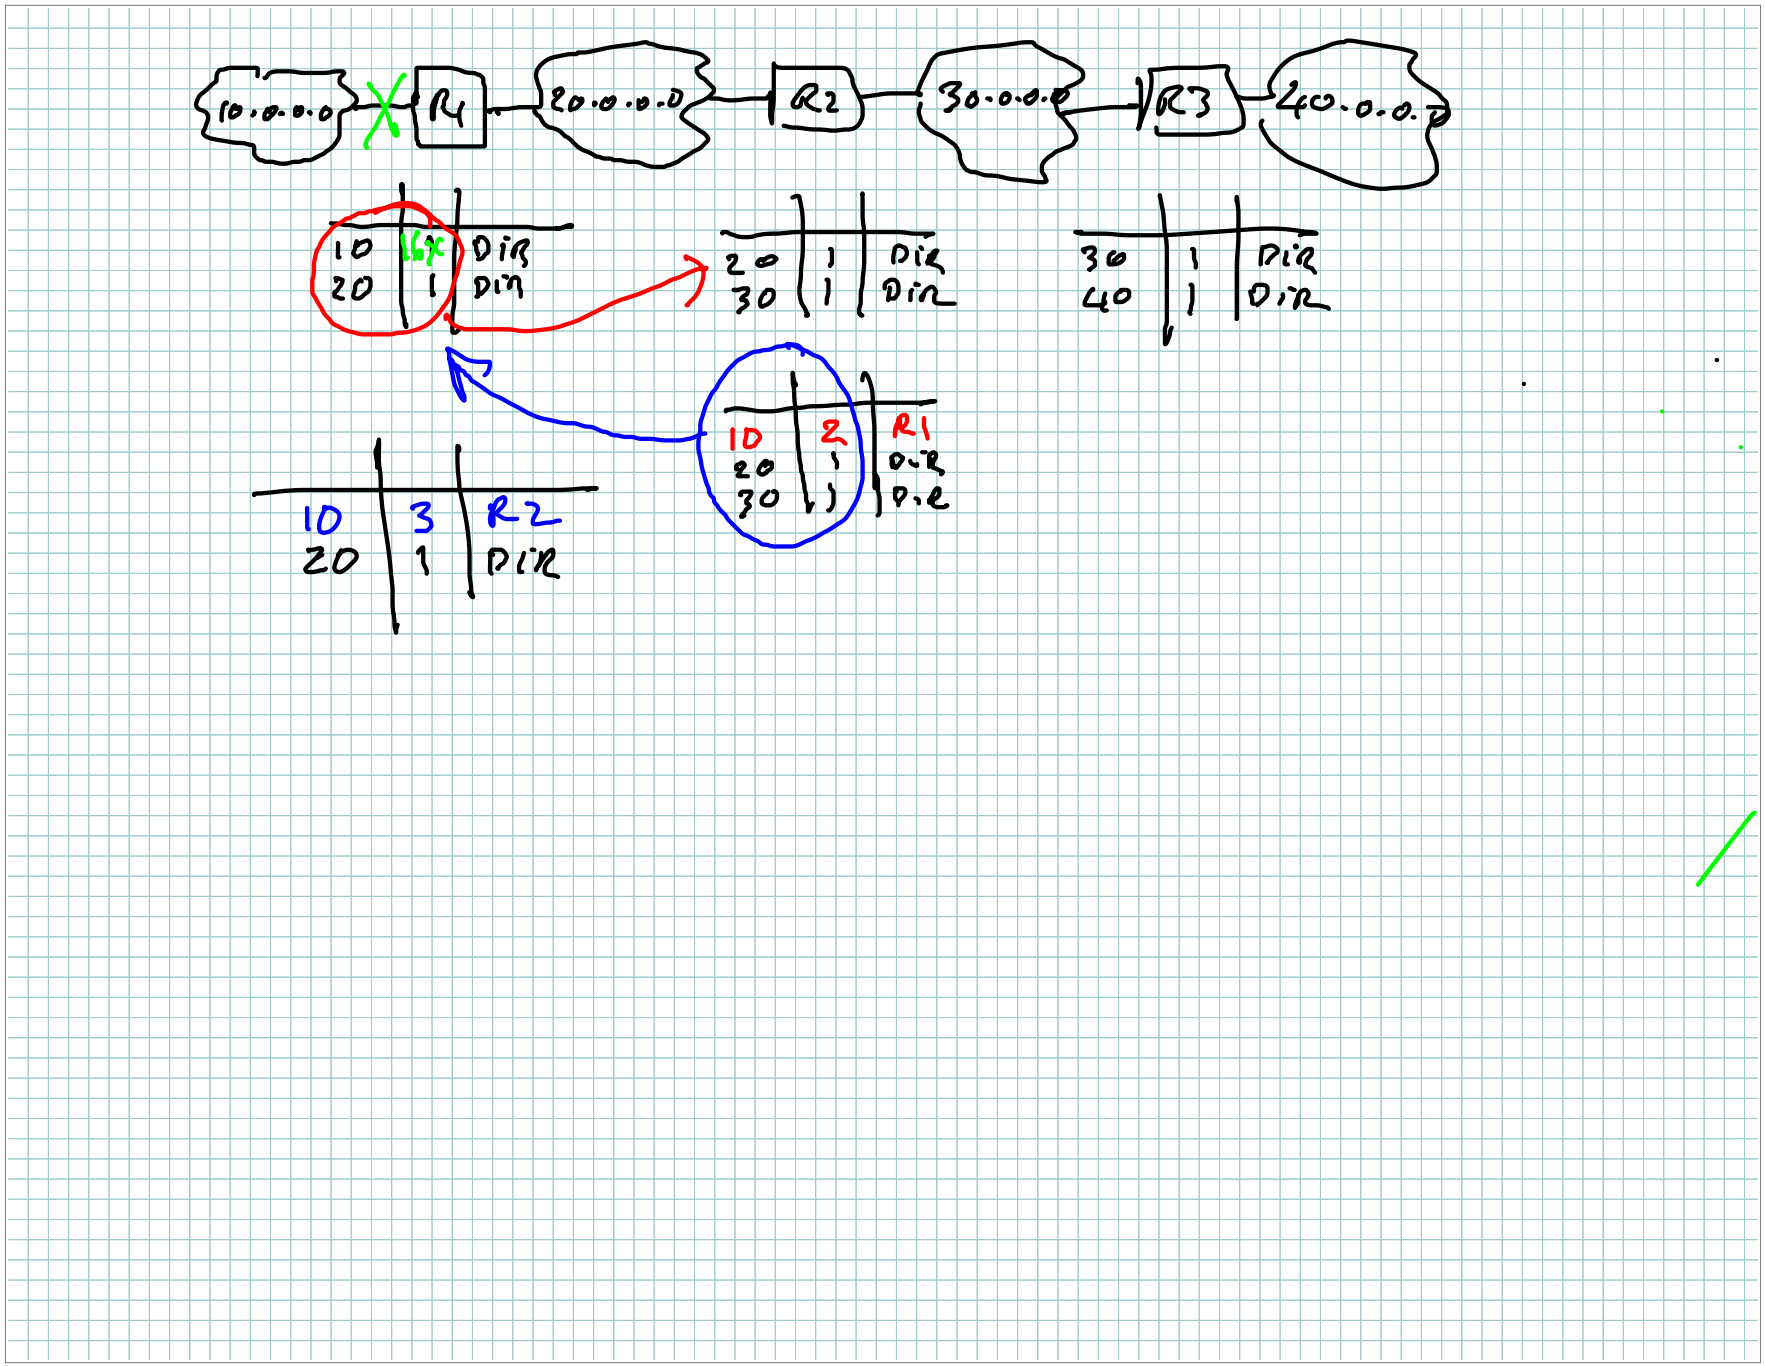
10.0.0.0
R1
20.0.0.0
R2
30.0.0.0
R3
40.0.0.0
10 16 Dir
20 1 Dir
20 1 Dir
30 1 Dir
30 1 Dir
40 1 Dir
10 2 R1
20 1 Dir
30 1 Dir
10 3 R2
20 1 Dir

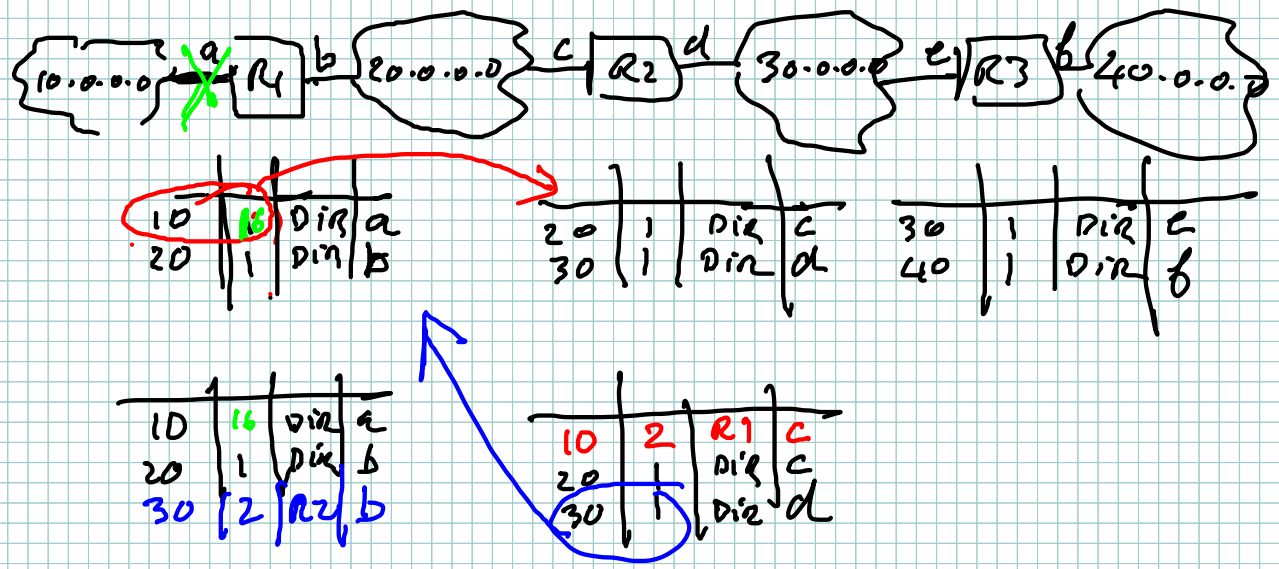
10.0.0.0
a
R1
b
20.0.0.0
c
R2
d
30.0.0.0
e
R3
f
40.0.0.0
10 | 16 | Dir | a
20 | 1 | Dir | b
20 | 1 | Dir | c
30 | 1 | Dir | d
30 | 1 | Dir | e
40 | 1 | Dir | f
10 | 16 | Dir | a
20 | 1 | Dir | b
30 | 2 | R2 | b
10 | 2 | R1 | c
20 | 1 | Dir | c
30 | 1 | Dir | d

Slow Convergence Problem (Count To Infinity)

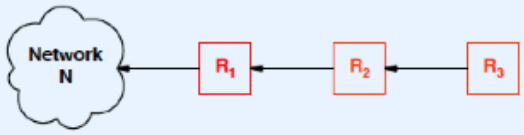

Routers with routes to network N

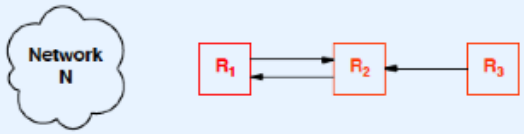

R_1 erroneously routes to R_2 after failure

RIP1 Update Format

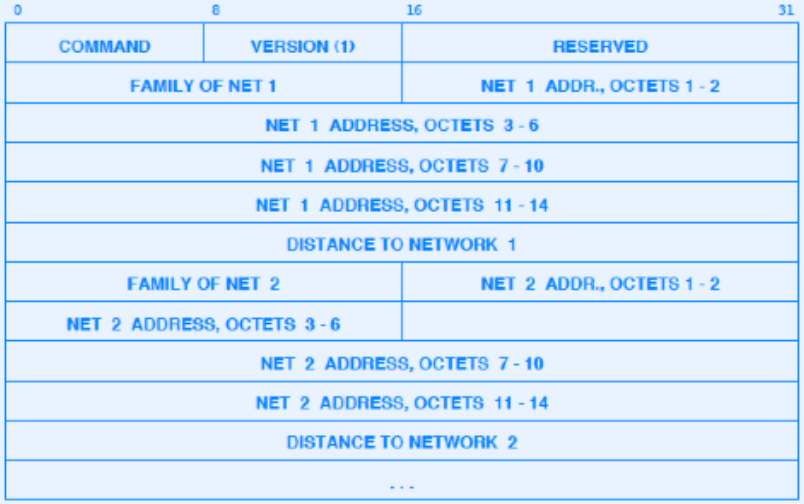

- Uses *FAMILY* field to support multiple protocols
- IP address sent in octets *3 - 6* of address field
- Message travels in UDP datagram

Changes To RIP In Version 2

- Update includes subnet mask
- Authentication supported
- Explicit next-hop information
- Messages can be multicast (optional)
  - IP multicast address is 224.0.0.9

RIP2 Update Format

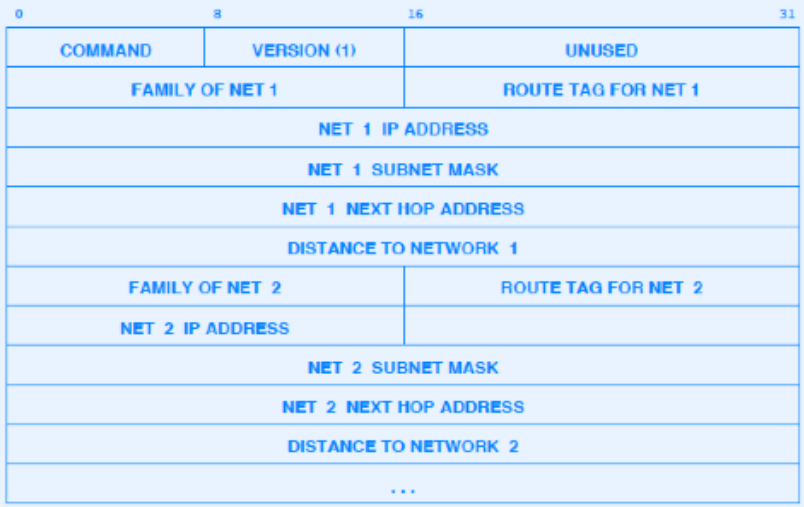

- Packet format is backward compatible
- Infinity still limited to *16*
- RIP2 *can* be broadcast

Measures Of Distance That Have Been Used

- Hops
  - Zero-origin
  - One-origin (e.g., RIP)
- Delay
- Throughput
- Jitter

HELLO: A Protocol That Used Delay

- Developed by Dave Mills
- Measured delay in milliseconds
- Used by NSFNET fuzzballs
- Now historic

How HELLO Worked

- Participants kept track of delay between pairs of routers
- HELLO propagated delay information across net
- Route chosen to minimize total delay

Route Oscillation

- Effective delay depends on traffic (delay increases as traffic increases)
- Using delay as metric means routing traffic where delay is low
- Increased traffic raises delay, which means route changes
- Routes tend to oscillate

Why HELLO Worked

- HELLO used only on NSFNET backbone
- All paths had equal throughput
- Route changes damped to avoid oscillation

Open Shortest Path First (OSPF)

- Developed by IETF in response to vendors' proprietary protocols
- Uses SPF (link-state) algorithm
- More powerful than most predecessors
- Permits hierarchical topology
- More complex to install and manage

OSPF Features

- Type of service routing
- Load balancing across multiple paths
- Networks partitioned into subsets called *areas*
- Message authentication
- Network-specific, subnet-specific, host-specific, and CIDR routes
- Designated router optimization for shared networks
- Virtual network topology abstracts away details
- Can import external routing information

OSPF Message Header

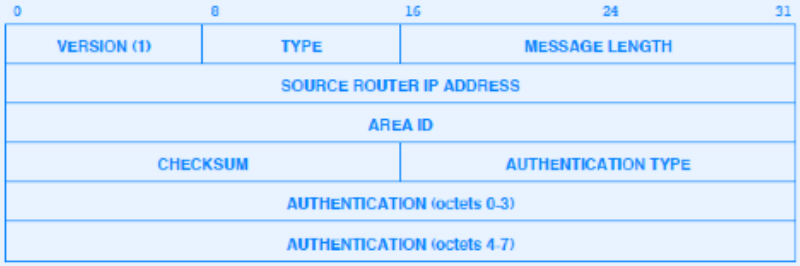

- Each message starts with same header

OSPF Message Types

Type	Meaning
1	Hello (used to test reachability)
2	Database description (topology)
3	Link status request
4	Link status update
5	Link status acknowledgement

OSPF HELLO Message Format

- Used to test reachability

OSPF Database Description Message Format

- Fields starting at *LINK TYPE* are repeated

Values In The LINK Field

Link Type	Meaning
1	Router link
2	Network link
3	Summary link (IP network)
4	Summary link (link to border router)
5	External link (link to another site)

OSPF Link Status Request Message Format

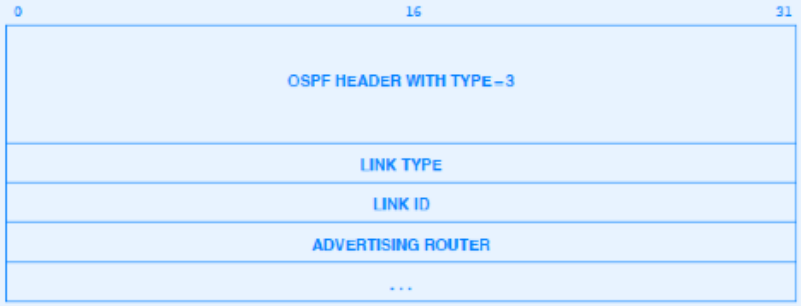

OSPF Link Status Update Message Format

Header Used In OSPF Link Status Advertisements

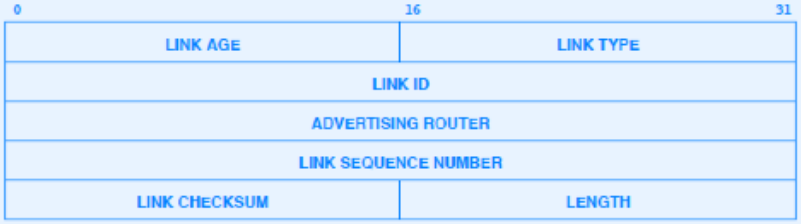

- Four possible formats follow
  - Links from a router to given area
  - Links from a router to physical net
  - Links from a router to physical nets of a subnetted IP network
  - Links from a router to nets at other sites

Discussion Question

- What are the tradeoffs connected with the issue of routing in the presence of partial information?

Summary

- Interior Gateway Protocols (IGPs) used within an AS
- Popular IGPs include
  - RIP (distance vector algorithm)
  - OSPF (link-state algorithm)