Routing Review

Autonomous System Concept

- Term Autonomous System (AS) to specify groups of routers
- One can think of an AS as a contiguous set of networks and routers all under control of one administrative authority
  - For example, an AS can correspond to an ISP, an entire corporation, or a university
  - Alternatively, a large organization with multiple sites may choose to define one AS for each site
  - In particular, each ISP is usually a single AS, but it is possible for a large ISP to divide itself into multiple ASs
- The choice of AS size can be made for
  - economic, technical, or administrative reasons

The Two Types of Internet Routing Protocols

- All Internet routing protocols are divided into two major categories:
  - Interior Gateway Protocols (IGPs)
  - Exterior Gateway Protocols (EGPs)
- After defining the two categories
  - we will examine a set of example routing protocols that illustrate each category
- Interior Gateway Protocols (IGPs)
- Exterior Gateway Protocols (EGPs)

The Two Types of Internet Routing Protocols

Interior Gateway Protocols (IGPs)

- Routers within an AS use an IGP exchange routing information
- Several IGPs are available
  - each AS is free to choose its own IGP
- Usually, an IGP is easy to install and operate
- IGP may limit the size or routing complexity of an AS

The Two Types of Internet Routing Protocols

Exterior Gateway Protocols (EGPs)

- A router in one AS uses an EGP to exchange routing information with a router in another AS
- EGPs are more complex to install and operate than IGPs
  - but EGPs offer more flexibility and lower overhead (i.e., less traffic)
- To save traffic
  - an EGP summarizes routing information from an AS before passing it to another AS
- An EGP implements policy constraints
  - that allow a system manager to determine exactly what information is released outside the organization

The Two Types of Internet Routing Protocols

Illustration of How IGPs and EGPs Are Used

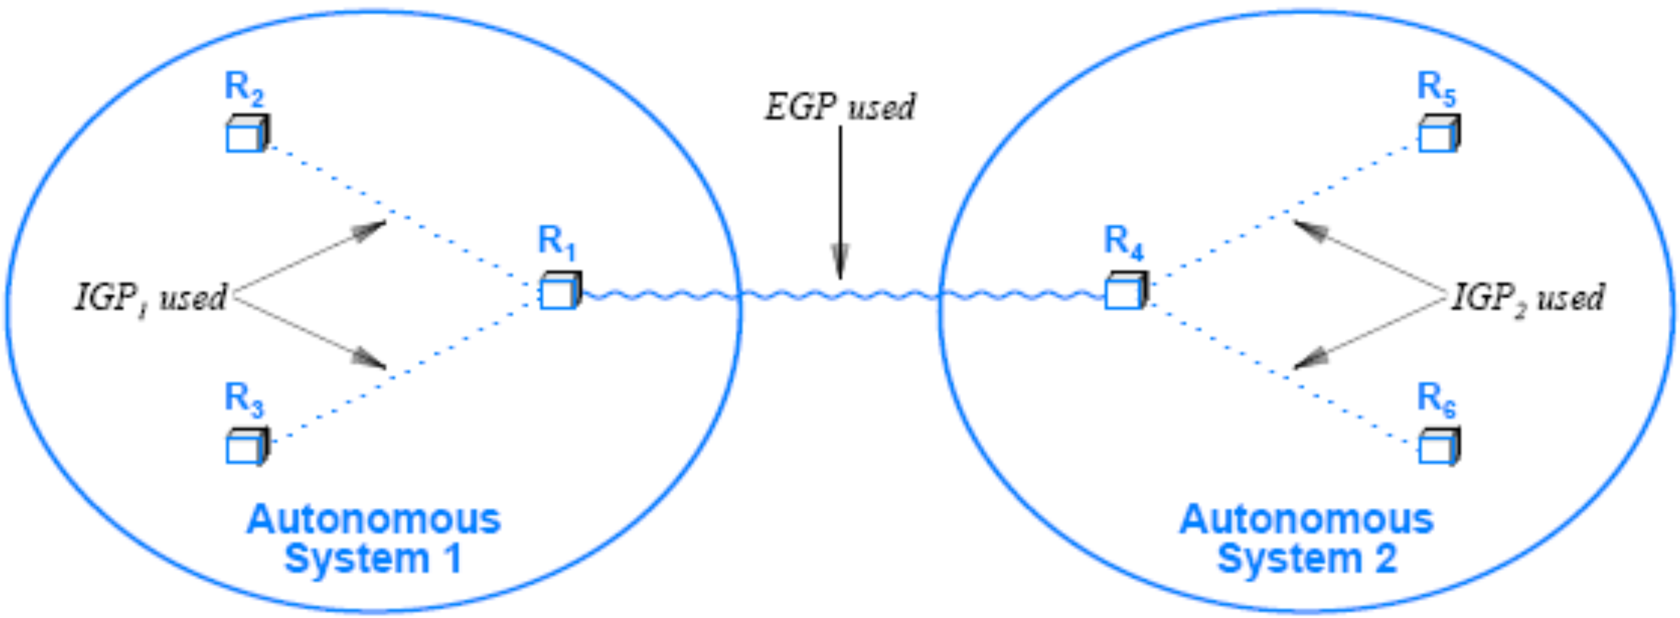

Figure 27.3 Illustration of Internet routing with an IGP used in each autonomous system and an EGP used between autonomous systems.

The Two Types of Internet Routing Protocols

Optimal Routes, Routing Metrics, and IGPs

- Routing software should find all possible paths and then choose one that is optimal
- Although the Internet usually has multiple paths between any source and destination
  - there is no universal agreement about which path is optimal
- Consider the requirements of various applications
  - For a remote desktop application
    - a path with least delay is optimal
  - For a browser downloading a large graphics file
    - a path with maximum throughput is optimal
  - For an audio webcast application that receives real-time audio
    - a path with least jitter is optimal

The Two Types of Internet Routing Protocols

Optimal Routes, Routing Metrics, and IGPs

- IGPs and EGPs differ in an important way with respect to routing metrics:
  - IGPs use routing metrics, but EGPs do not
    - each AS chooses a routing metric and arranges internal routing software to send the metric with each route so receiver can use the metric to choose optimal paths
- Outside an AS, an EGP does not attempt to choose an optimal path
  - Instead, the EGP merely finds a path
- Each AS is free to choose a routing metric
- An EGP cannot make meaningful comparisons
  - Suppose one AS reports the number of hops along a path to destination D and another AS reports the throughput along a different path to D
  - An EGP that receives the two reports cannot choose which of the two paths has least cost because there is no way to convert from hops to throughput
- Thus, an EGP can only report the existence of a path and not its cost

Routes and Data Traffic

- Data traffic for a given destination flows in exactly the opposite direction of routing traffic. For example
  - suppose an AS owned by ISP_1 contains network N
    - Before traffic can arrive destined for N, ISP_1 must advertise a route to N
    - That is, when the routing advertisement flows out, data will begin to flow in

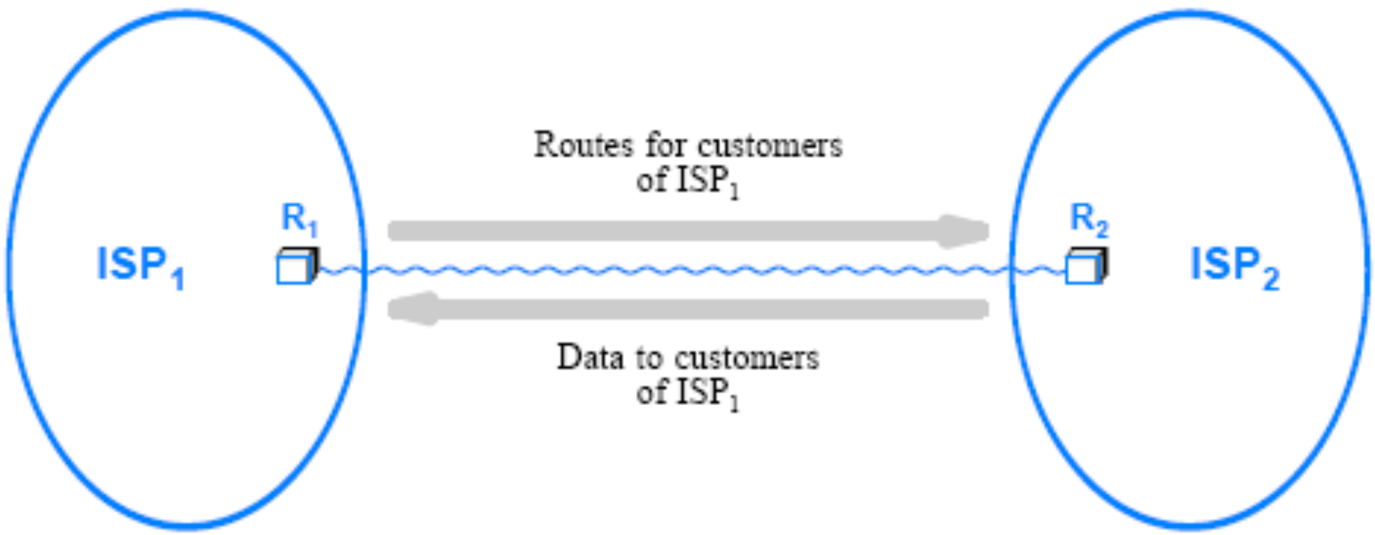

The Border Gateway Protocol (BGP)

- Most widely used EGP in the Internet is BGP
  - Current standard is version 4, abbreviated BGP-4
- BGP has the following characteristics:
- Routing Among AS
  - Because it is intended for use as an EGP
  - BGP provides routing information at the AS level
  - No way for BGP to provide details about the routers within each AS on the path
- Provision for Policies
  - BGP allows the sender and receiver to enforce policies
    - a manager can configure BGP to restrict which routes BGP advertises to outsiders
- Reliable Transport
  - BGP uses TCP for all communication
    - a BGP program on a router in one AS forms a TCP connection to a BGP program on a router in another AS and then sends data across the connection

The Border Gateway Protocol (BGP)

- Facilities for Transit Routing
  - BGP classifies each AS
    - as a transit system if it agrees to pass traffic through to another AS or
    - as a stub system if it does not
  - Traffic passing through on its way to another AS is classified as transit traffic
  - Classification allows BGP to distinguish between ISPs and other AS
  - BGP allows a corporation to classify itself as a stub
    - even if it is multi-homed (i.e., a corporation with multiple external connections can refuse to accept transit traffic)
- BGP provides the glue that holds Internet routing together at the center of the Internet
  - Tier-1 ISPs use BGP to exchange routing information and learn about each other's customers

BGP uses Path Vector Routing

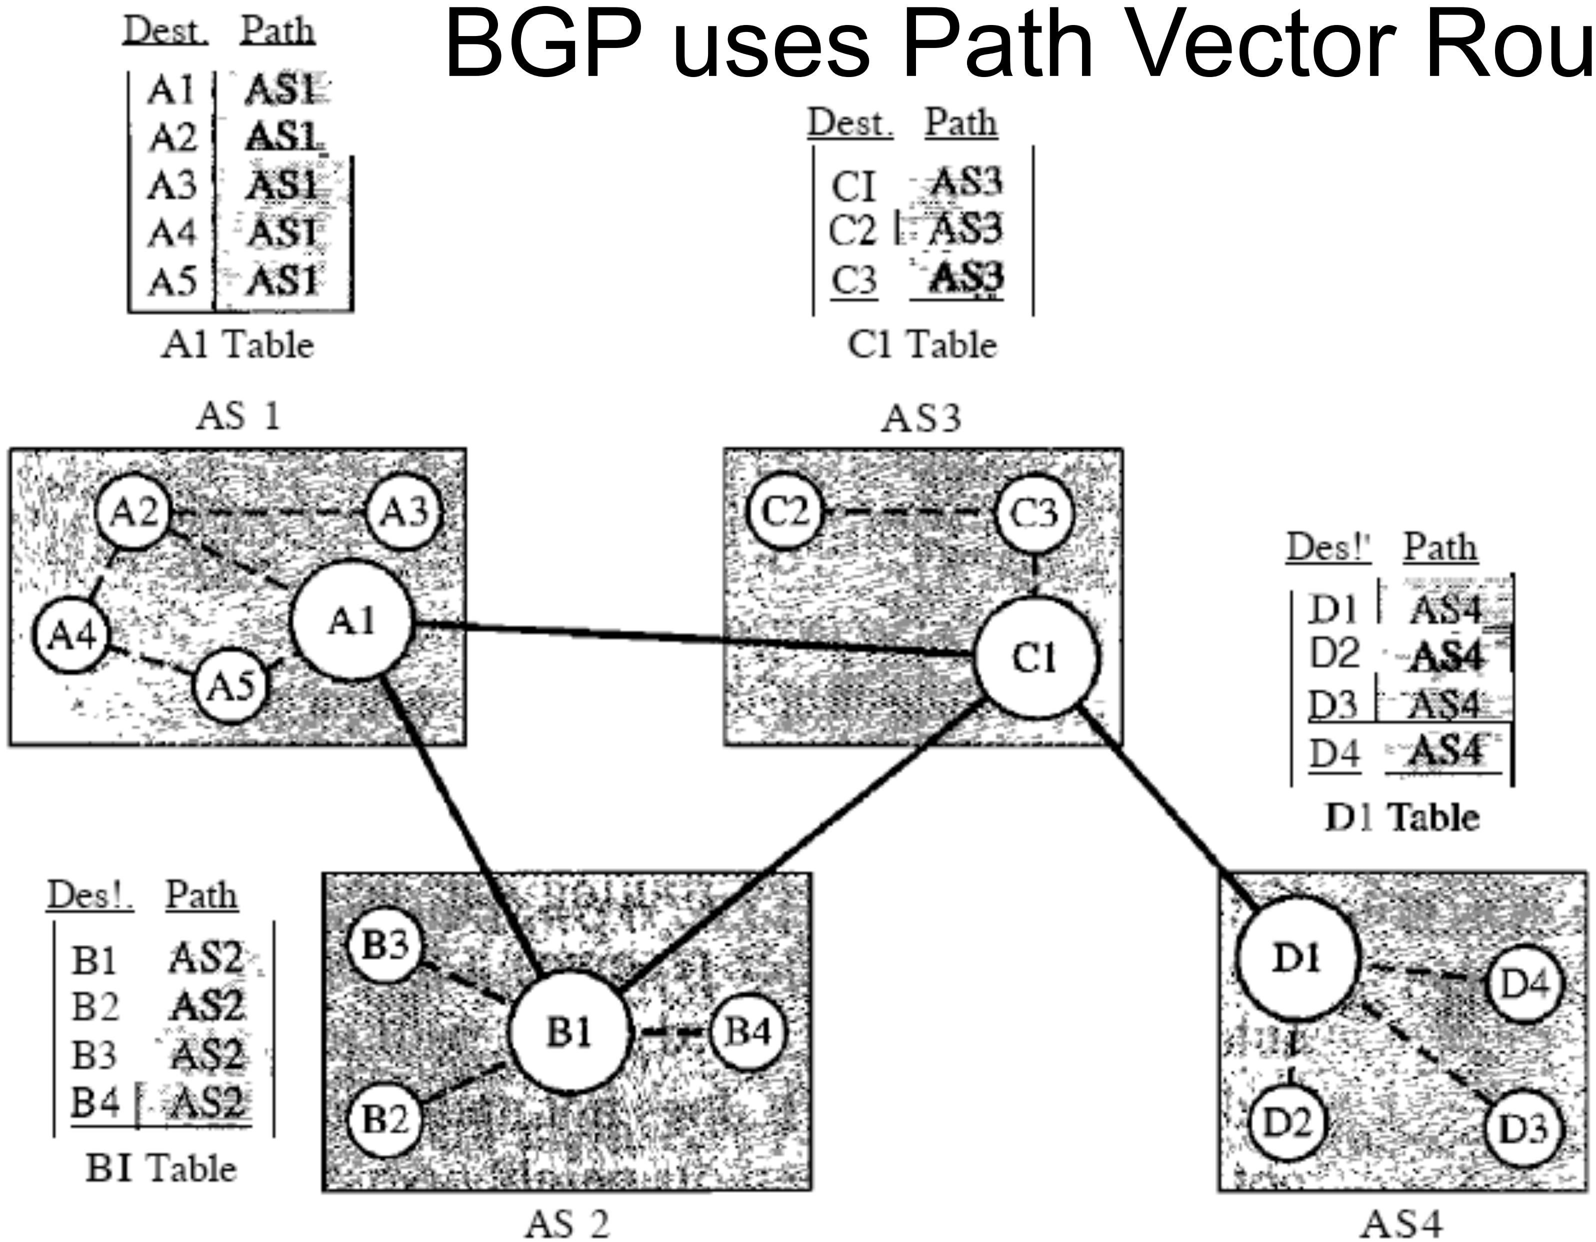

The Routing Information Protocol (RIP)

- RIP was among the first IGP used in the Internet
- RIP has the following characteristics:
- Routing within an AS
  - RIP is designed as an IGP to be used among routers within an AS
- Hop Count Metric
  - RIP measures distance in network hops
    - where each network between the source and destination counts as a single hop
  - RIP counts a directly connected network as one hop away
- Unreliable Transport
  - RIP uses UDP to transfer messages among routers
- Broadcast or Multicast Delivery
  - RIP is intended for use over LAN technologies
    - that support broadcast or multicast (e.g., Ethernet)
  - Version 1 of RIP broadcasts messages
  - Version 2 allows delivery via multicast

The Routing Information Protocol (RIP)

- Support for CIDR and Subnetting
  - RIP version 2 includes an address mask with each destination
- Support for Default Route Propagation
  - RIP allows a router to advertise a default route
- Distance Vector Algorithm
  - RIP uses the distance-vector approach to routing defined in Algorithm 18.3
- Passive Version for Hosts
  - RIP allows a host to listen passively and update its forwarding table
  - Passive RIP is useful on networks where a host selects among multiple routers

The Routing Information Protocol (RIP)

- How RIP propagates routes? (Distance Vector Routing)
  - Each outgoing message contains an advertisement that lists the networks the sender can reach along with a distance to each
  - When it receives an advertisement
    - RIP software uses the list of destinations to update the local forwarding table
- Each entry in a RIP advertisement consists of a pair:

  (destination network, distance)

  where distance is the number of hops to the destination
- When a message arrives
  - if the receiver does not have a route to an advertised destination or
  - if an advertised distance is shorter than the distance of the current route
    - the receiver replaces its route with a route to the sender

The Routing Information Protocol (RIP)

- Advantage of RIP is simplicity, requires little configuration
  - A manager merely
    - starts RIP running on each router in the organization and
    - allows the routers to broadcast messages to one another
  - After a short time
    - all routers in the organization will have routes to all destinations
- RIP also handles the propagation of a default route
  - The organization needs to configure one of its routers to have a default (an organization chooses a router that connects to an ISP)
  - RIP propagates the default route to all other routers in the organization
    - which means that any datagram sent to a destination outside the organization will be forwarded to the ISP

The Routing Information Protocol (RIP) Example
Assume R1 is sending its information to R2

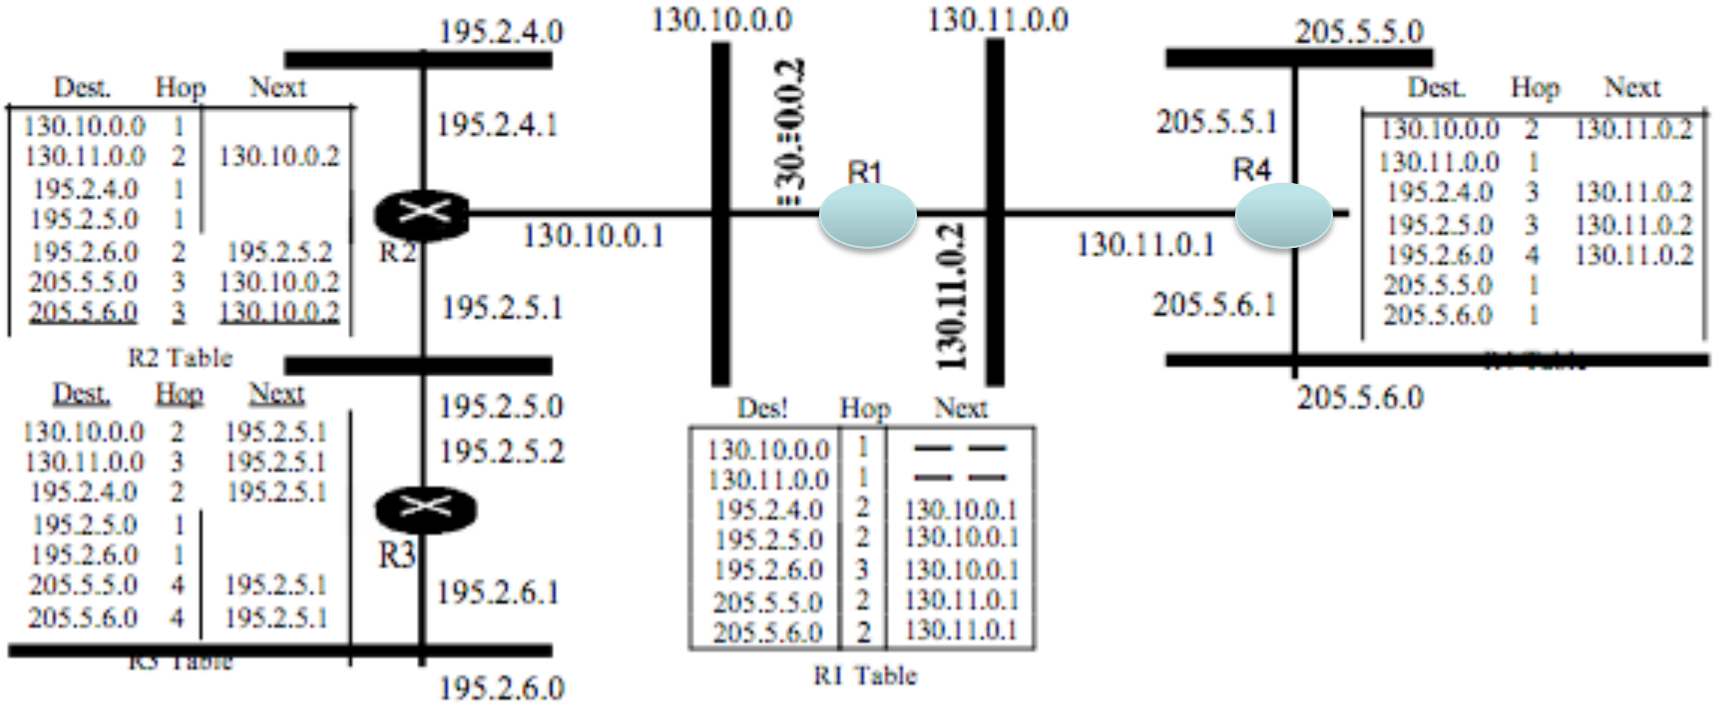

RIP Packet Format

- The RIP message format helps explain how a distance vector routing protocol operates
- To permit RIP to be used with CIDR or subnet addressing
  - an entry contains a 32-bit address mask
- Each entry also has a next hop address
- Two 16-bit fields that identify the entry as an IP address and provide a tag used to group entries together
- Each entry contains 20 octets (bytes)

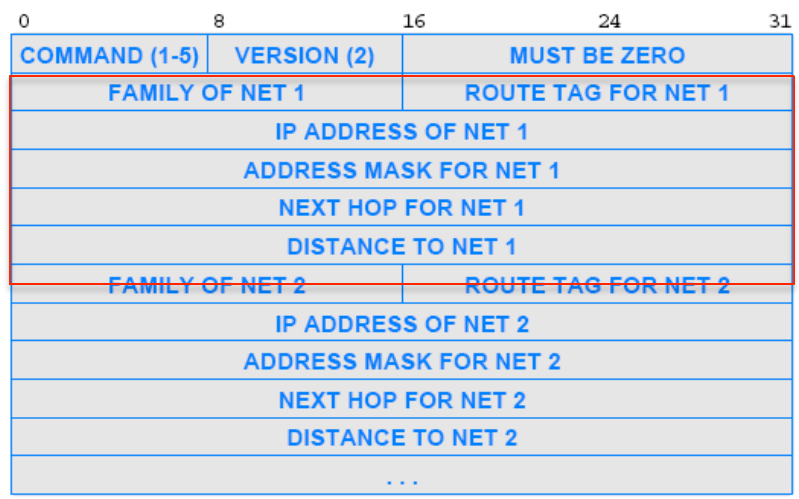

The Open Shortest Path First Protocol (OSPF)

- The RIP message format illustrates a disadvantage of DVR protocols:
  - The size of a message is proportional to the number of networks that can be reached
  - Sending RIP messages introduces delay
  - The delay means that route changes propagate slowly
  - RIP works well among a few routers, but it does not scale well
  - Processing RIP messages consumes many CPU cycles
- To satisfy demand for a routing protocol that can scale to large organizations
  - IETF devised an IGP known as the OSPF
  - The name is derived from the use of Dijkstra's SPF algorithm

The Open Shortest Path First Protocol (OSPF)

- OSPF has the following characteristics:
  - Routing within an AS (OSPF is an IGP)
  - CIDR Support
    - OSPF includes a 32-bit address mask with each address
  - Authenticated Message Exchange
    - A pair of routers using OSPF can authenticate each message
  - Imported Routes
    - OSPF allows a router to introduce routes learned from another means (e.g., from BGP)
  - Link-State Algorithm
  - Support for Metrics
    - OSPF allows an administrator to assign a cost to each route
  - Support for Multi-access Networks
    - LSR is inefficient across a multi-access network, such as an Ethernet, because all routers attached to the network broadcast link status
    - OSPF optimizes by designating a single router to broadcast on the network

An Example OSPF Graph

- Consider the network and associated OSPF graph

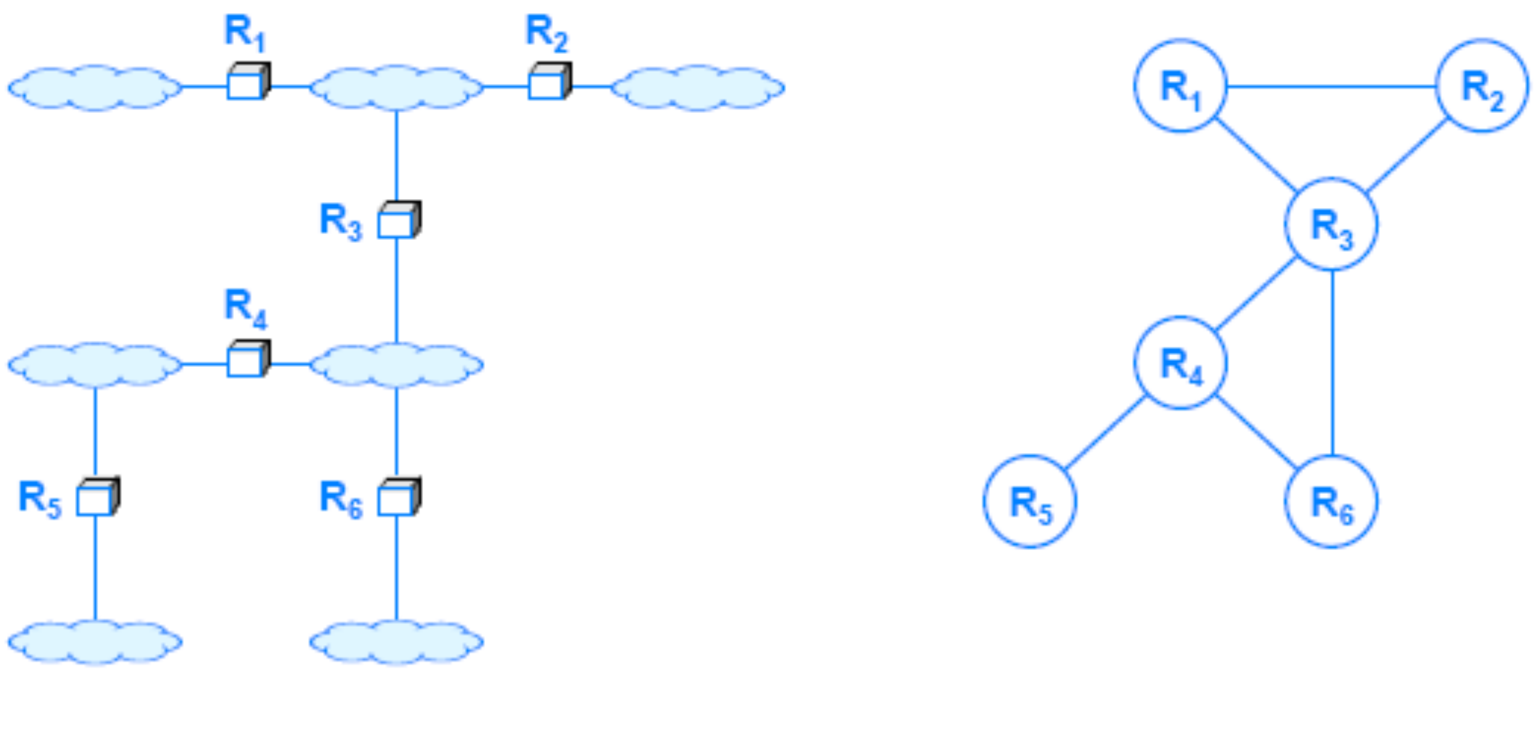

(a) (b)

An Example OSPF Graph

- A typical OSPF graph
  - each node corresponds to a router
  - an edge in the graph is a connection between a pair of routers
- To follow a link-state algorithm
  - each pair of routers connected by a network periodically probe one another and then broadcast a link-state message to other routers
- All routers receive the broadcast message
  - each uses the message to update its local copy of the graph
  - and recomputes shortest paths when the status changes

OSPF Areas

- Hierarchical routing makes OSPF more complex than other routing protocols but also makes it more powerful
- To achieve a hierarchy
  - OSPF allows an AS to be partitioned for routing purposes
  - We can divide routers in an AS into subsets, called areas
  - Each router is configured to know the area boundary
  - Routers within a given area exchange link-state messages periodically
- OSPF allows communication between areas
  - One router in each area communicates with a router in one or more other area(s)
  - The two routers summarize routing information learned from other routers within their respective area
    - and then exchange the summary
  - Instead of broadcasting to all routers in the AS
    - OSPF limits link-state broadcasts to routers within an area
- As a result of the hierarchy
  - OSPF can scale to handle much larger AS than other routing protocols

OSPF Areas

- Hierarchical routing makes OSPF more complex than other routing

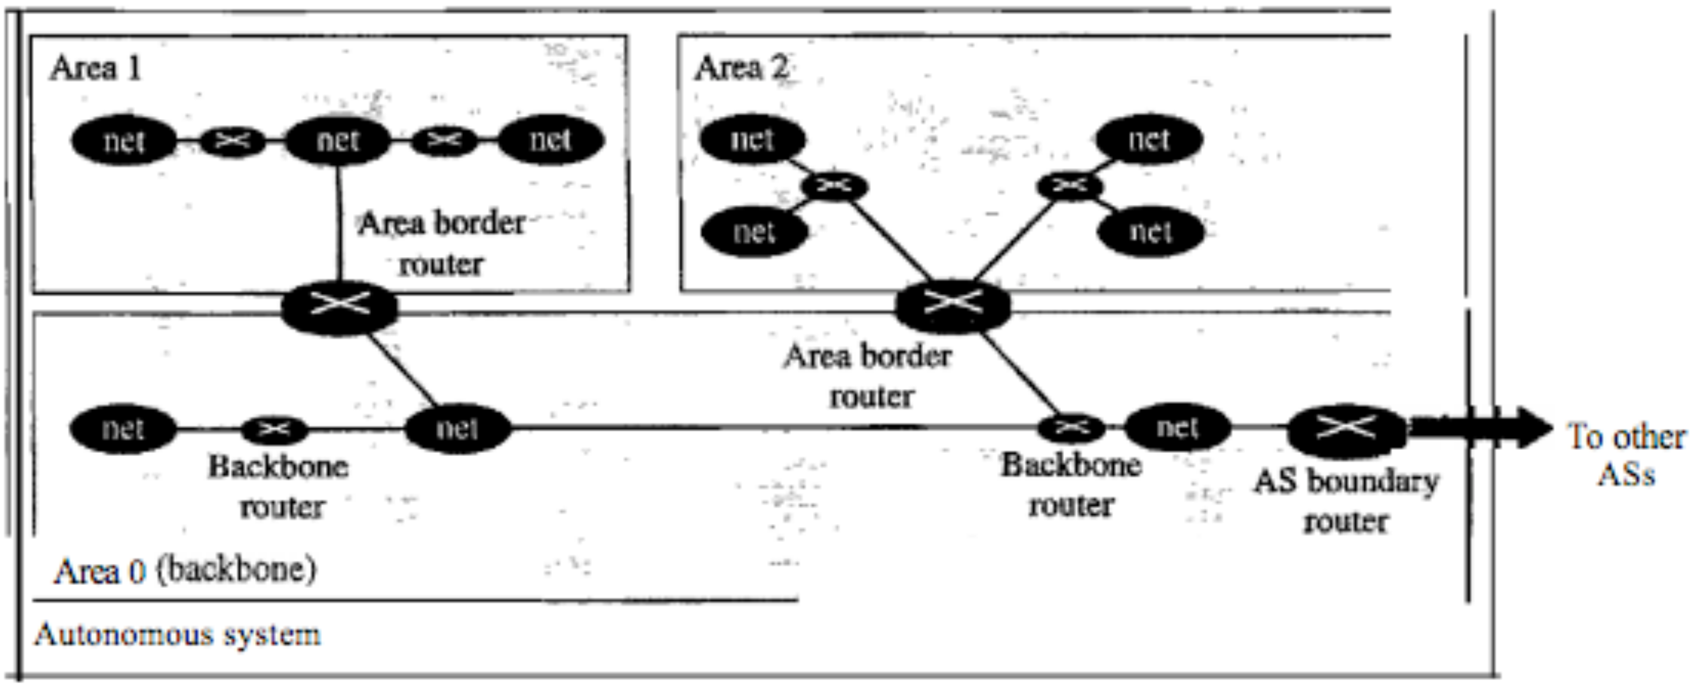

  – Instead of broadcasting to all routers in the AS
    - OSPF limits link-state broadcasts to routers within an area
- As a result of the hierarchy
  – OSPF can scale to handle much larger AS than other routing protocols

Intermediate System - Intermediate System (IS-IS)

- Originally designed by DEC to be part of DECNET V,
- The IS-IS is an IGP
  - naming follows Digital's terminology in which a router was called an IS and a host was called an End System
- IS-IS was created around the same time as OSPF
- The two protocols are similar in many ways
  - Both use the link-state approach
    - employ Dijkstra's algorithm to compute shortest paths
  - Both protocols require two adjacent routers to periodically test the link between them and broadcast a status message

Intermediate System - Intermediate System (IS-IS)

- The chief differences between OSPF and the original IS-IS can be summarized as:
- IS-IS was proprietary (owned by DEC)
  - and OSPF was created as an open standard available to all vendors
- OSPF was designed to run over IP
  - IS-IS was designed to run over CLNS (part of the ill-fated OSI protocol)
- OSPF was designed to propagate IPv4 routes
  - IS-IS was designed to propagate routes for OSI protocols
- Over time, OSPF gained many features
  - As a result, IS-IS now has less overhead
- When the protocols were initially invented
  - OSPF's openness made it much more popular than IS-IS
  - In fact, IS-IS was almost completely forgotten

Intermediate System - Intermediate System (IS-IS)

- As the years progressed
  - OSPF's popularity encouraged the IETF to add additional features
- Ironically, in the early 2000's, ten years after the protocols were designed
  - several things changed to give IS-IS a second chance
- DEC had dissolved, and IS-IS was no longer considered valuable proprietary property
- A newer version of IS-IS was defined to integrate it with IP and the Internet
- Because OSPF was built for IPv4
  - a completely new version had to be developed to handle larger IPv6 addresses
- The largest ISPs have grown to a size where the extra overhead in OSPF makes IS-IS more attractive
- As a result, IS-IS has started to make a comeback

Multicast Routing

IP Multicast Semantics

- So far, we have discussed unicast routing
  - That is, we have considered routing protocols that propagate information about destinations that each have a static address and a location that does not change
- One of the design goals for unicast route propagation is stability
  - continual changes in routes are undesirable
    - because they lead to higher jitter and datagrams arriving out of order
- Once a unicast routing protocol finds a shortest path
  - it usually retains the route until a failure makes the path unusable
- Propagating multicast routing information differs dramatically from unicast route propagation
  - The difference arises because Internet multicast allows dynamic group membership and anonymous senders
  - Dynamic group membership means that an application can choose to participate in a group at any time
    - and remain a participant for an arbitrary duration

Unicasting vs. Multicasting

- What are the issues?

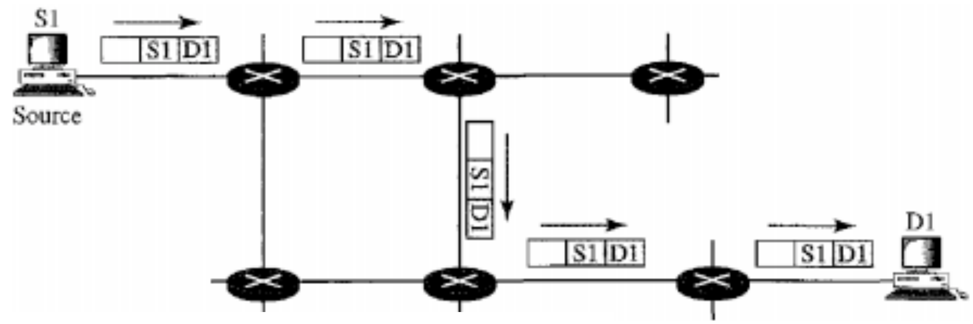

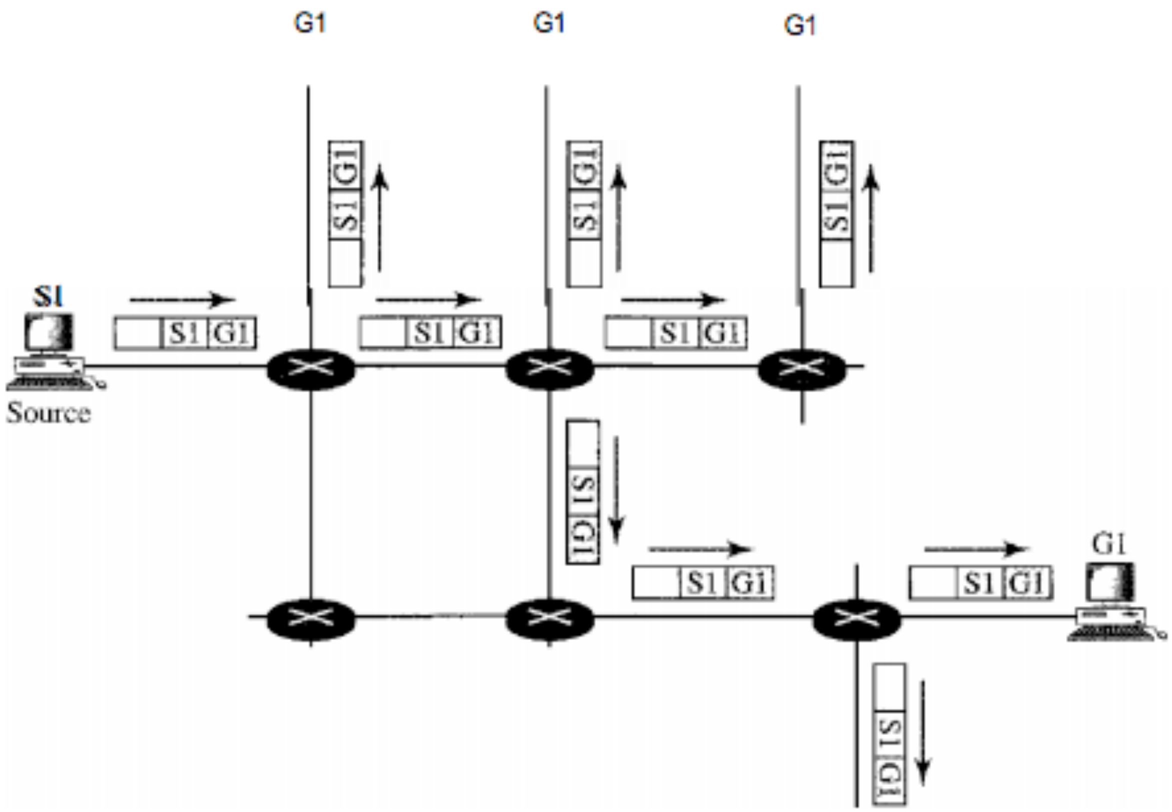

Multicasting
&
Multiple Unicasting

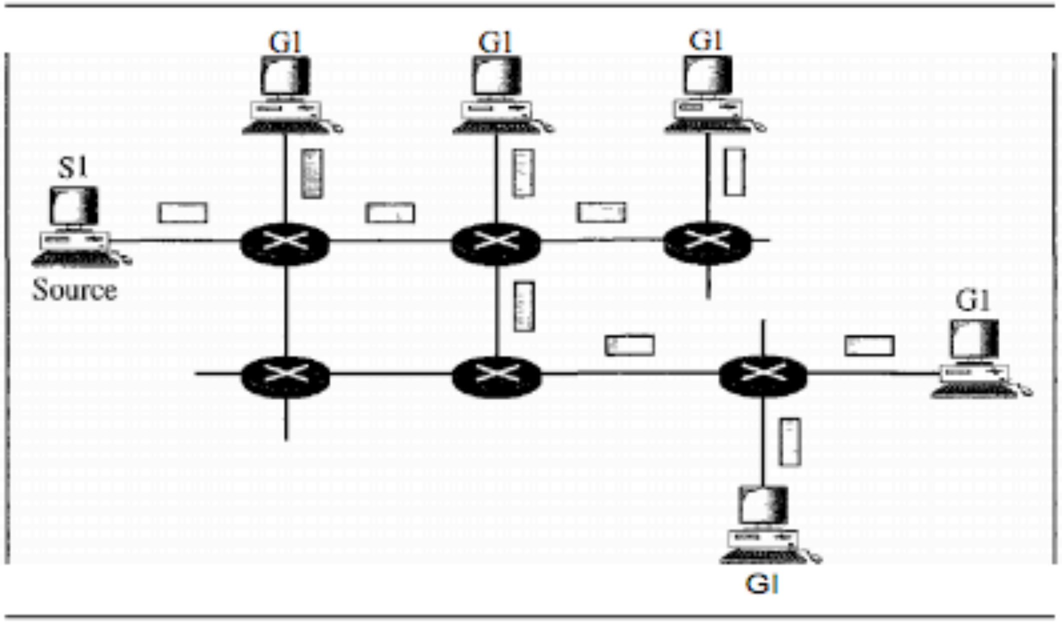

a. Multicasting

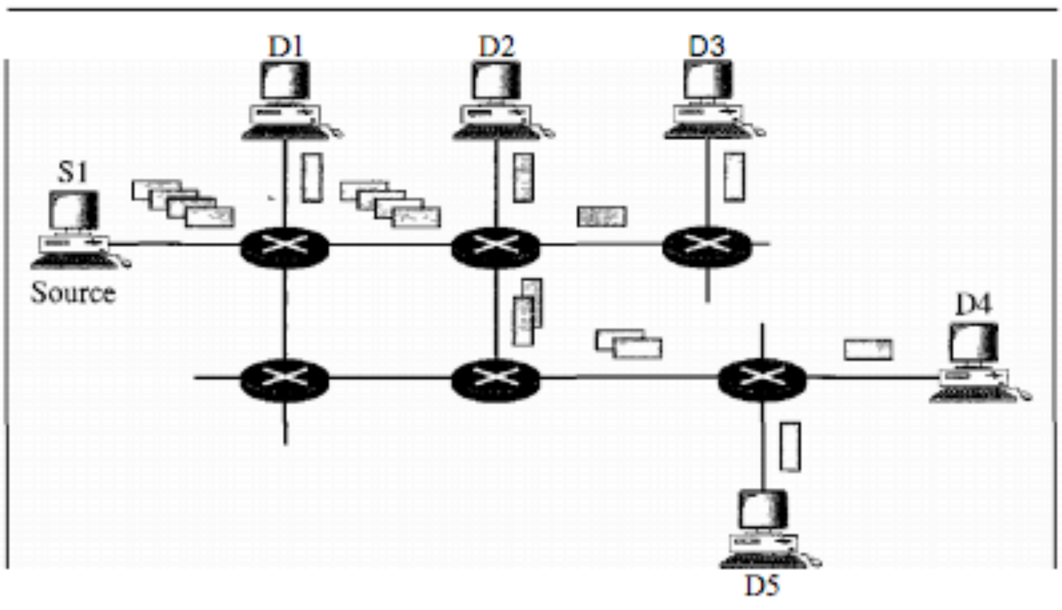

b. Multiple nnicasting

Multicast Routing

IP Multicast Semantics

- An arbitrary computer can Join to a multicast group at any time and begin receiving packets sent to the group
- To join a group, a host informs a nearby router
  - If multiple applications on the same host decide to join a group
    - the host receives one copy of each datagram sent to the group and makes a local copy for each application
- A host periodically sends group membership messages to the local router
- An arbitrary computer can Leave a multicast group at any time
  - Once the last application on the host leaves the group
    - the host informs the local router that it is no longer participating in the group

Multicast Routing

IP Multicast Semantics

- An IP multicast group is anonymous in a number of ways:
- First
  - neither a sender nor a receiver knows (nor can they find out) the identity or the number of group members
- Second
  - Routers and hosts do not know which applications will send a datagram to a group
    - because an arbitrary application can send a datagram to any multicast group at any time
- Membership in a multicast group only defines a set of receivers
  - a sender does not need to join a multicast group before sending a message to the group

Multicast Routing

IGMP

- How does a host join or leave a multicast group?
- A protocol exists that allows a host to inform a nearby router
  - whenever the host needs to join or leave a particular multicast group
  - protocol is known as the Internet Group Multicast (**Management**) Protocol (IGMP)
- The protocol is used only on the network between a host and a router
- The protocol defines the host
  - not the application, to be a group member
- The protocol does not specify anything about applications
- If multiple applications on a given host join a multicast group
  - the host must make copies of each datagram it receives for local applications
- When the last application on a host leaves a group
  - the host uses IGMP to inform the local router that it is no longer a member of the group

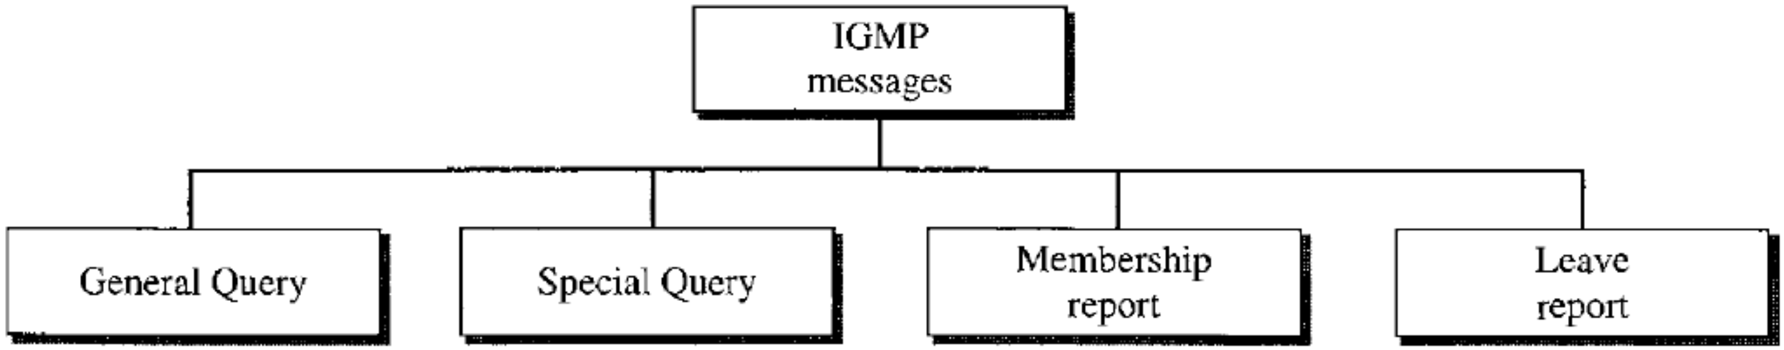

Type	*Value*
General or special query	0x11 or 00010001
Membership report	0x16 or 00010110
Leave report	0x17 or 00010111

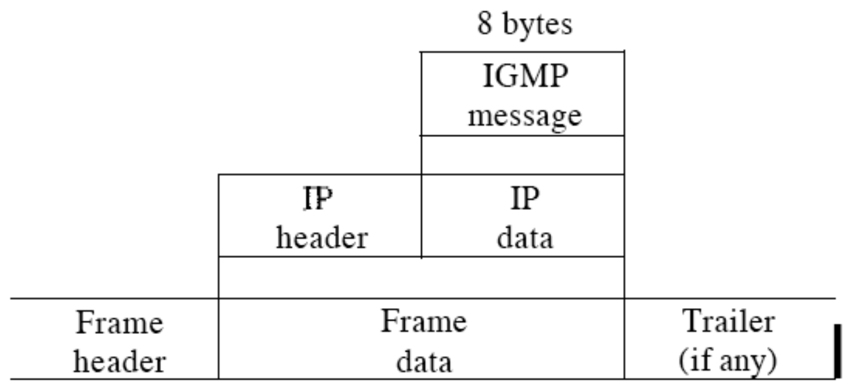

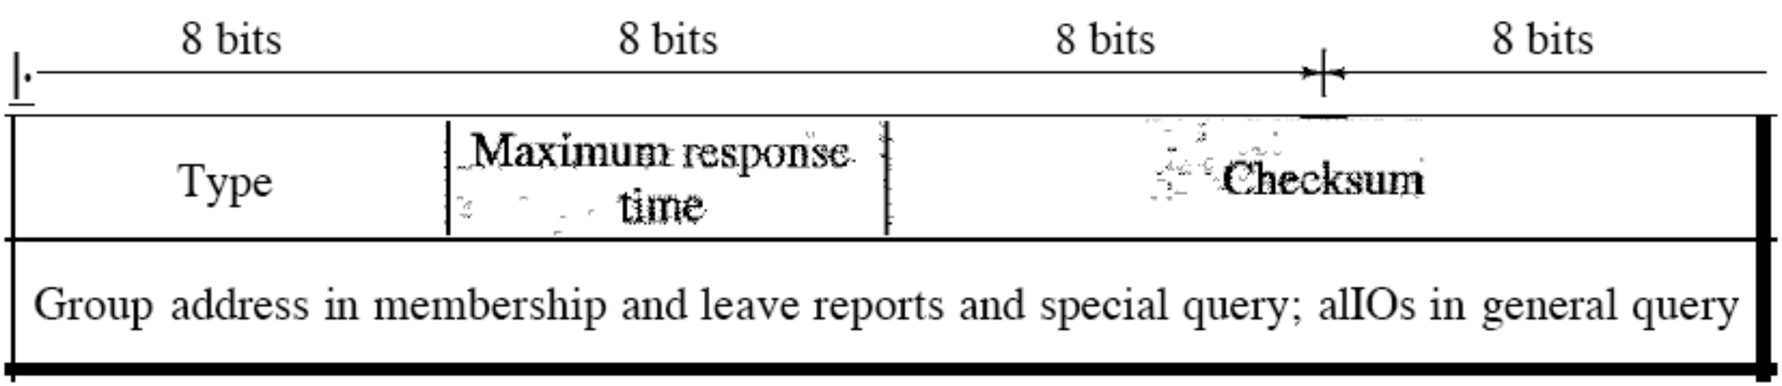

Multicast Routing Protocols

R: Reverse
P: Path
F: Forwarding
B: Broadcasting
M: Multicasting

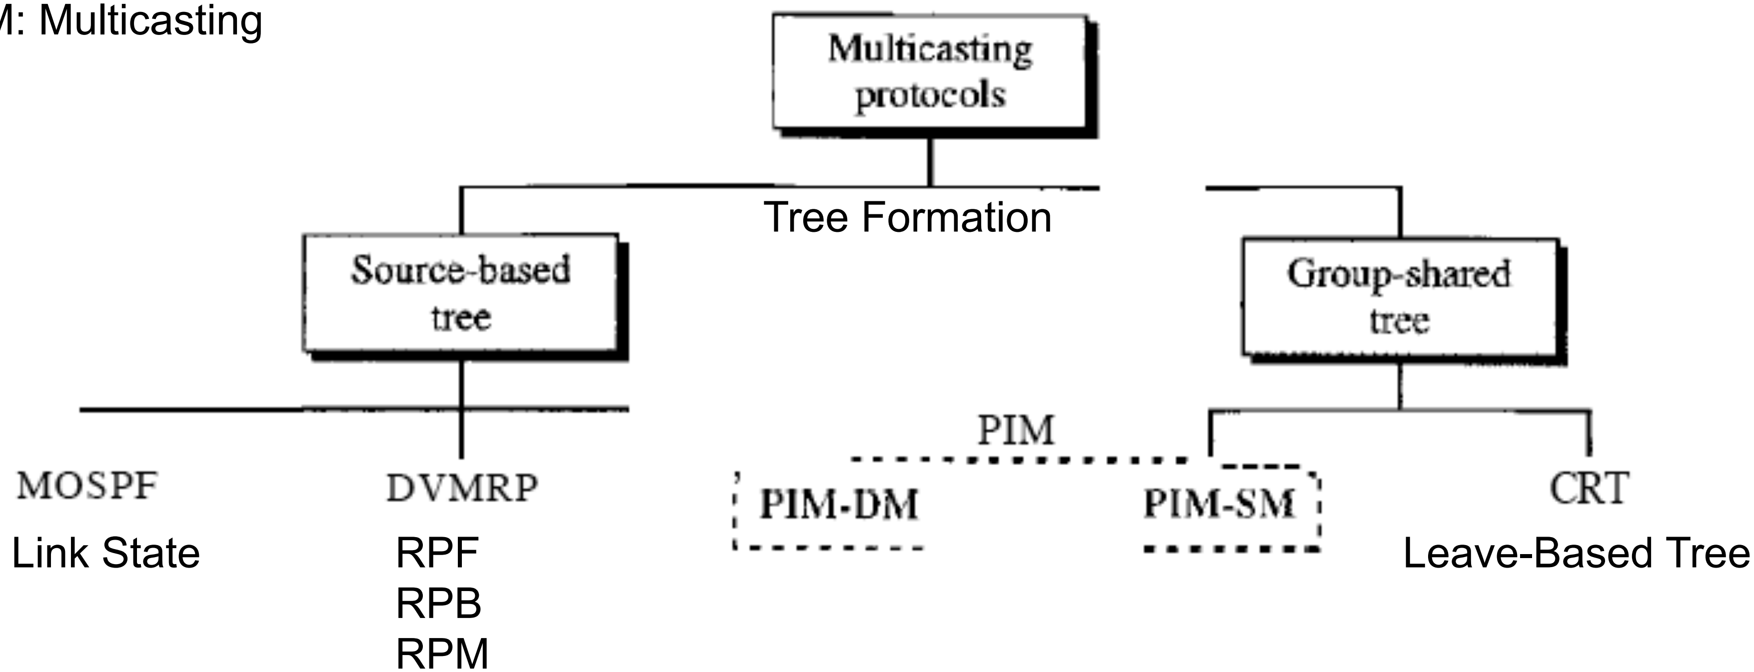

Multicast Routing
Tunneling and Using Unicast Addressing

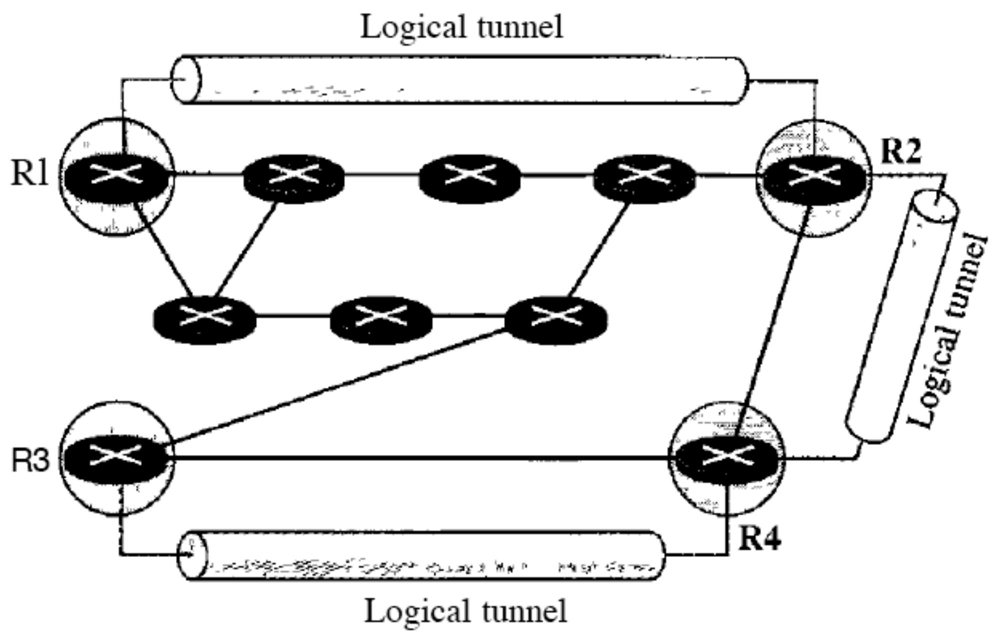

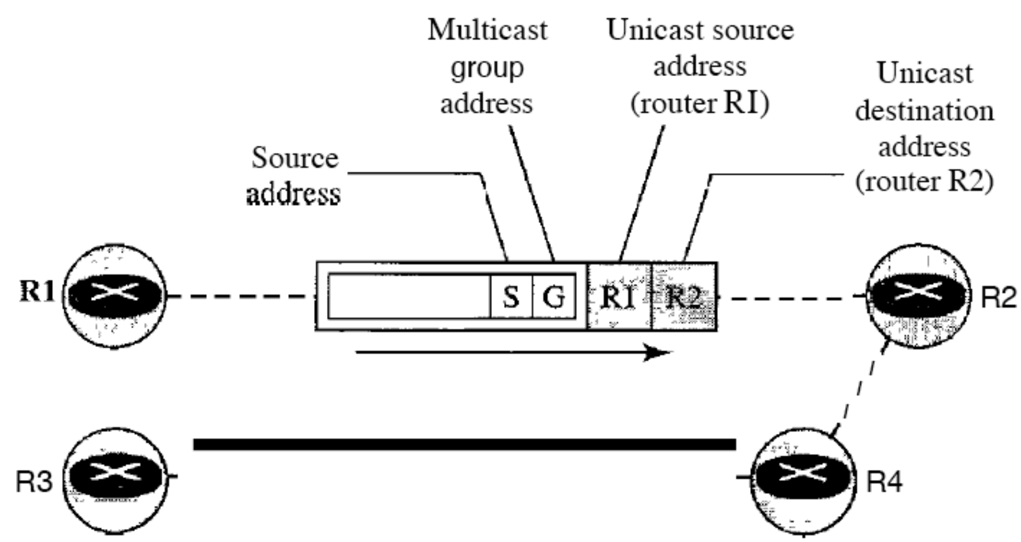

Multicast Routing Protocols

R: Reverse
P: Path
F: Forwarding
B: Broadcasting
M: Multicasting

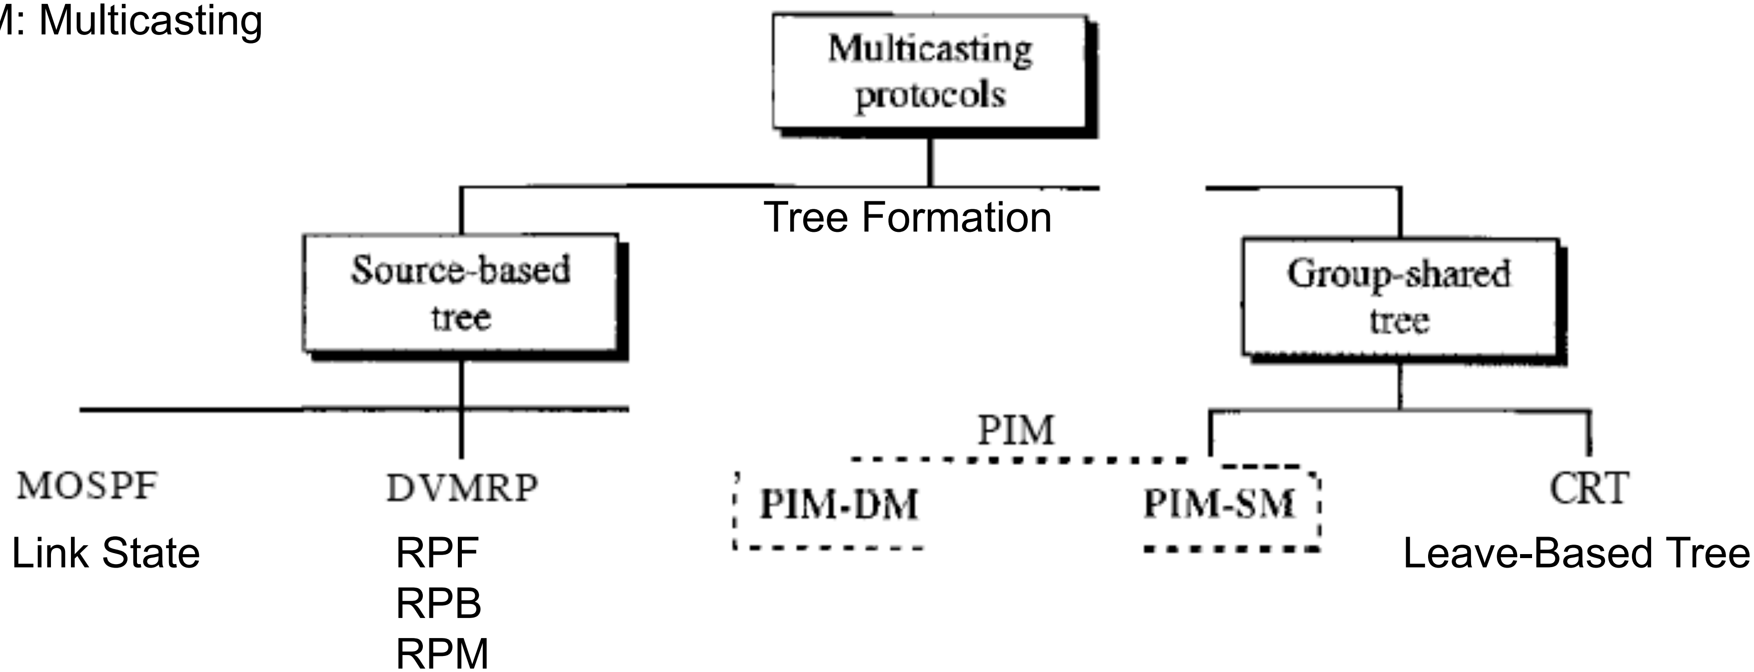

Multicast Routing
DVMRP - RPF

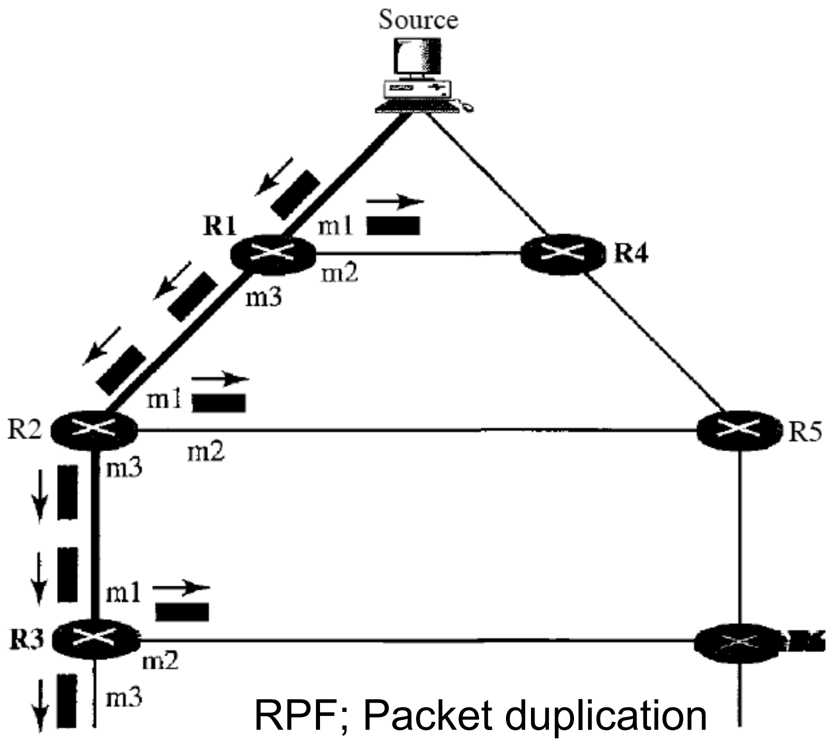

RPF; Packet duplication

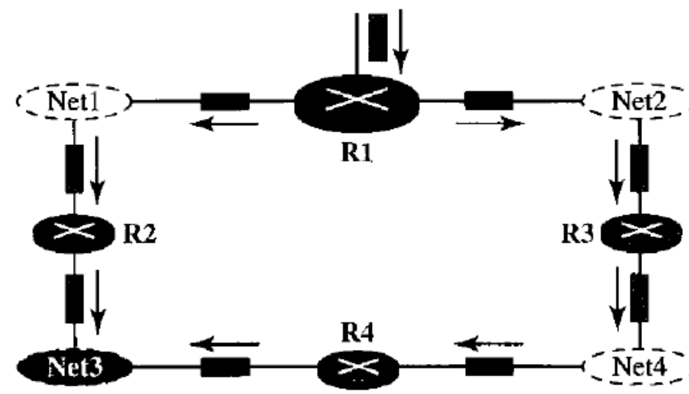

Multicast Routing
DVMRP – RPB
Avoids duplication using parents
Each network has a parent
Parents are selected using shortest path

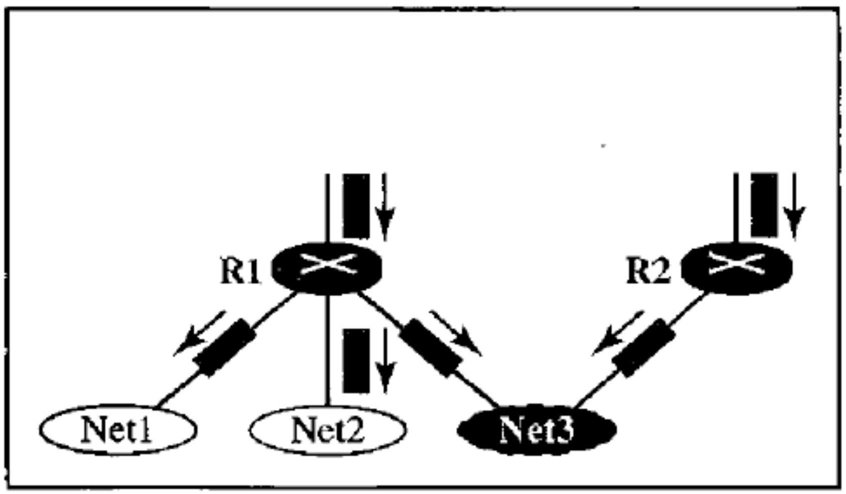

a. RPF

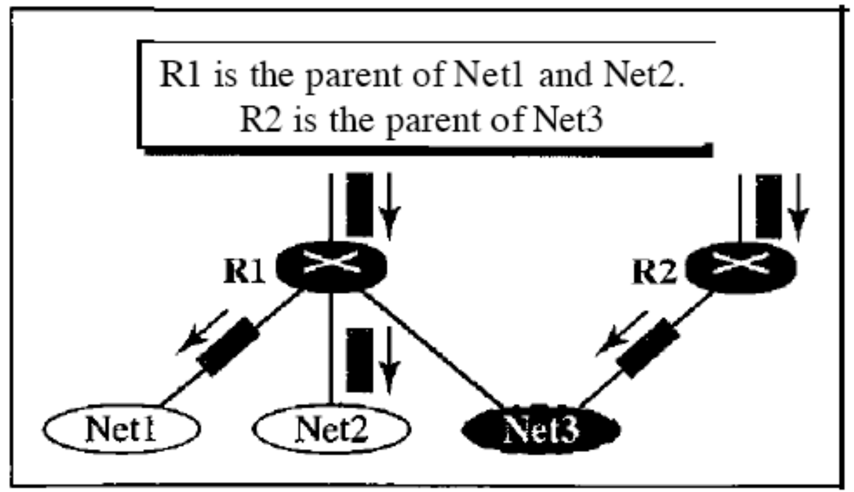

b.RPB

Multicast Routing
DVMRP – RPM
Uses pruning and grafting
Useful for dynamic membership

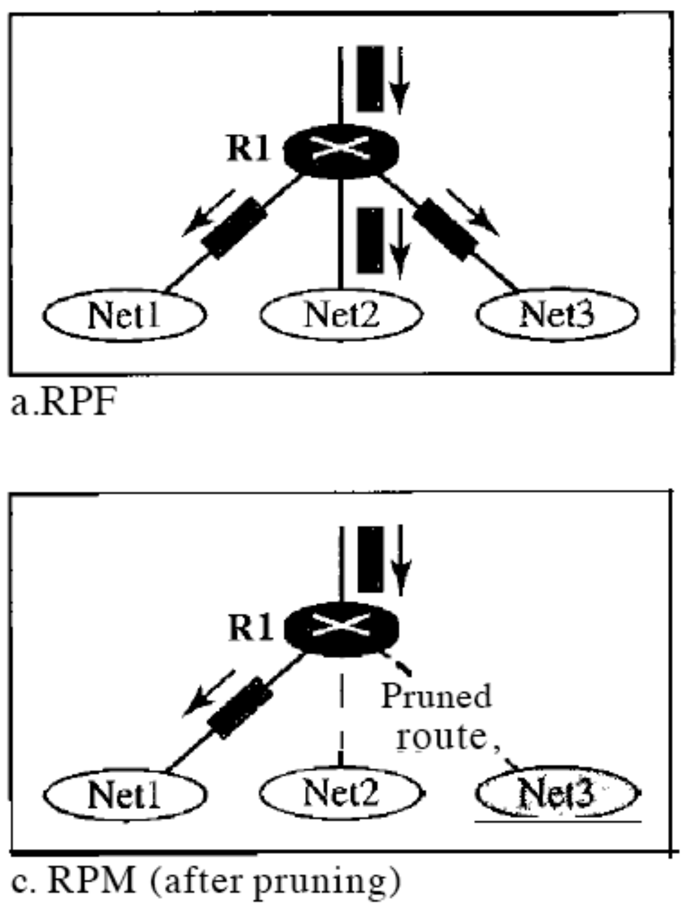

Multicast Routing
MOSPF

- Uses least cost tree

Multicast Routing
CBT

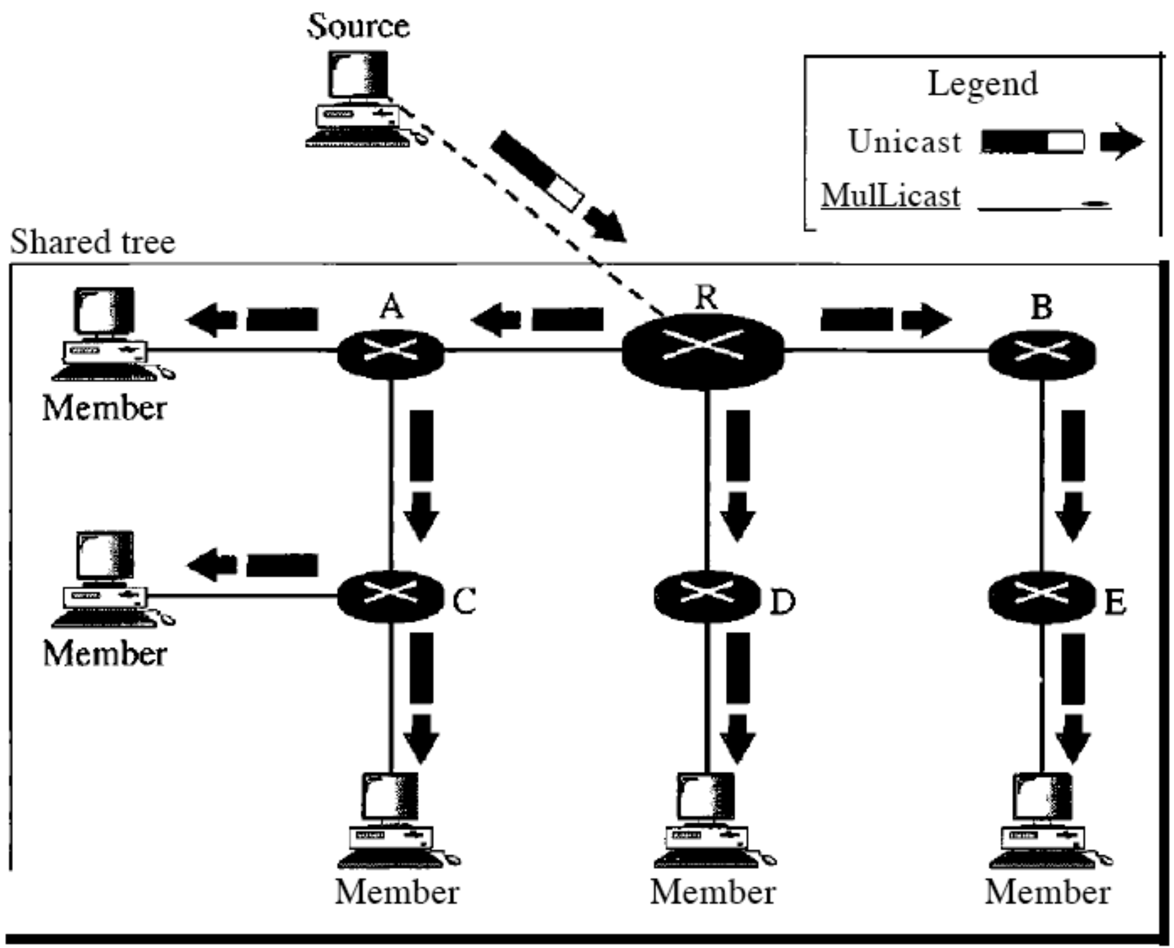

Multicast Routing Protocols