

High speed computer networks

Exterior routing protocols

Sosina M.
Addis Ababa institute of technology (AAiT)
2012 E.C.

Intra AS routing

- Routing algorithms
  - Distance vector
  - Link state
- Protocols
  - RIP and OSPF
    - ✓ Optimal paths from source to destination that are internal to the same AS
- Not suitable for inter AS routing
  - Scalability
    - ✓ LS – flooding
    - ✓ DV – instability
  - Metric

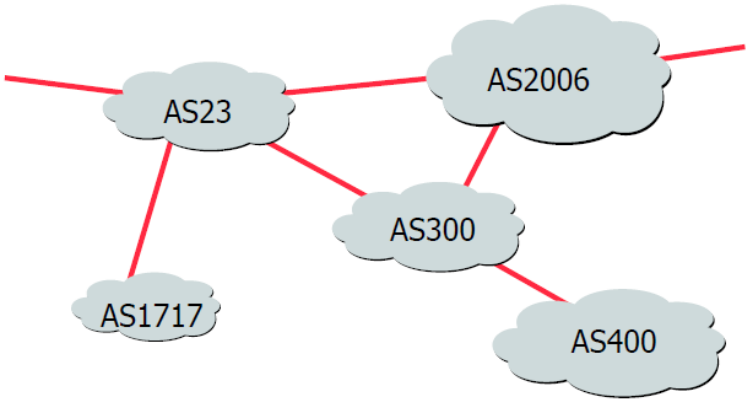

Path-vector routing

- Does not include a distance or cost estimate

- Provides path information
  - Lists the ASs that must be traversed to reach the destination
  - Enables policy routing
    - ✓ E.g., avoid transiting a particular AS

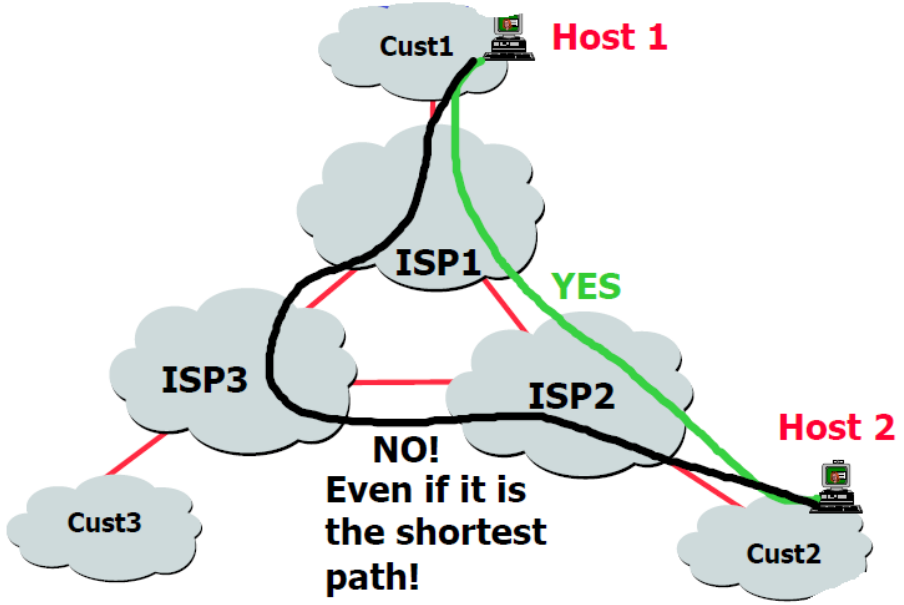

Boarder gateway protocol (BGP)

- Exterior routing protocol for the internet
- Allows gateway routers in different AS to cooperate in the exchange of routing information

- Three Functional procedures
  - Neighbor acquisition
  - Neighbor reachability
  - Network reachability

Neighbor acquisition

- Routers in different ASs wishing to exchange route information
- **Neighbor acquisition -** two neighboring routers in different AS agree to exchange routing information regularly
- Procedures
  - One router sends a request message to the other
  - The other router may either accept or refuse the offer

- *How one router knows the address and how it decides that it needs to exchange routing information with that particular router*?

Neighbor and Network reachability

- After establishing neighbor reachability
- **Neighbor reachability -** to maintain the relationship
  - A periodic keepalive message
- **Network reachability**
  - Each router maintains a database of the networks that it can reach and the preferred route for reaching each network
  - The router issues a broadcast update message when change is made to the database

BGP messages

- Sent over TCP connection
- **Open message**
  - The first message after TCP connection establishment
- **Update message**
  - Used to transfer routing information between BGP peers
    - ✓ Advertise a feasible route
- **Keepalive message**
  - Sent periodically (every 60 s)
  - Wait for a hold-time period (180 s) before declaring the peer is unreachable
- **Notification message**
  - Sent when there is a fatal error condition

BGP messages

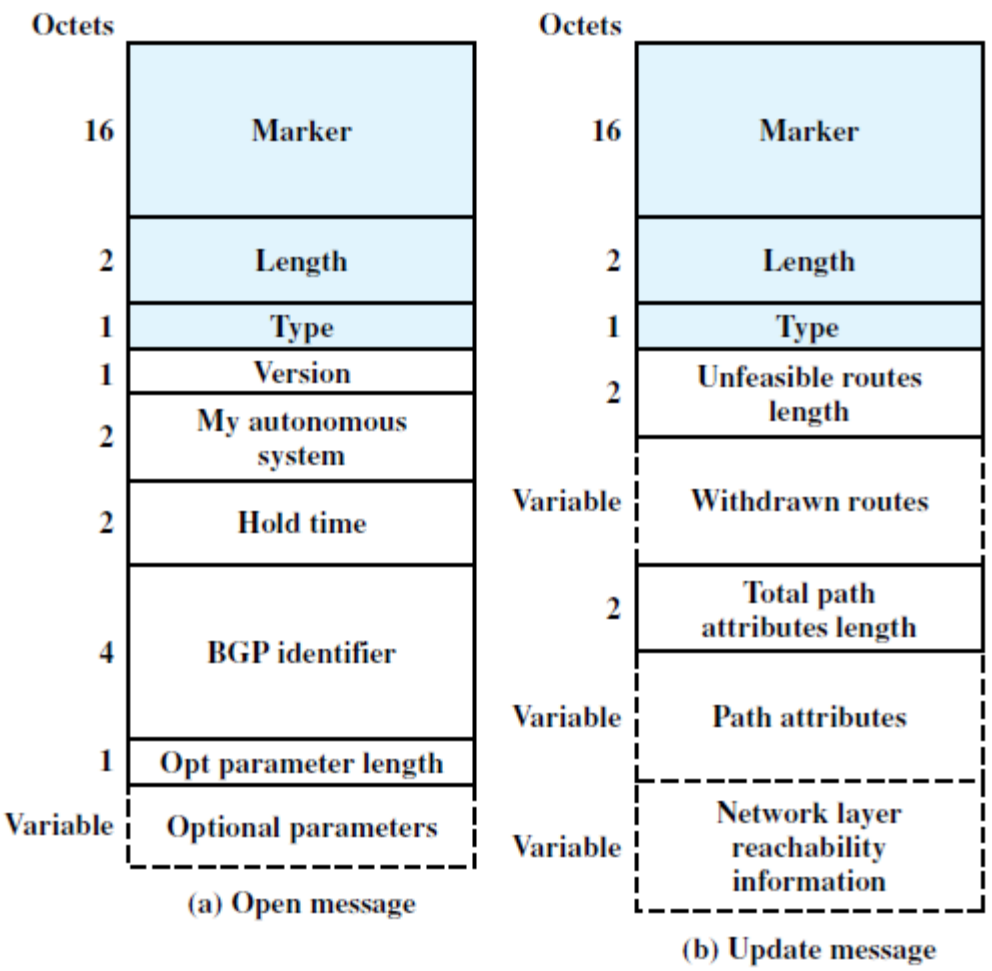

- ***Marker***: included for compatibility
- ***Length***: total length of the message
- ***Type***: type code of the message (open, update, notification, keepalive)

BGP: connection establishment

- A router first opens a TCP connection to a neighbor router
  - BGP runs over TCP on a well-know port (179)
- Sends an open message
  - AS, IP address, Hold time parameter
- Recipient
  - Calculate hold time (min[it hold timer, hold timer in the open message])
  - Reply with keep alive message
    - ✓ consists simply of the header

BGP: update

- Update message communicates two types of information
  - Information about a single route
    - ✓ Network Layer Reachability Information
    - ✓ Total Path Attributes Length
    - ✓ Path Attributes
  - A list of routes previously advertised by this router that are being withdrawn
  - May contain one or both types of information

BGP: update

- ***Network Layer Reachability Information(NLRI)***
  - list of identifiers (IP) of networks that can be reached by this route
- ***Path attributes***
  - Origin –an interior or exterior router
  - AS path – Sequence of ASs that are traversed for this route
    - ✓ Prevent looping
    - ✓ Enables routing policies implementation
  - Next Hop**:** The IP address of the next-hop router along the path to the destination
  - Multi_Exit_disc : Information about routers internal to AS
  - Local_Pref : Tell other routers within AS degree of preference for a particular route
  - Atomic_Aggregate, Aggregator : Uses subnet addresses in tree view of network to reduce information needed in NLRI

BGP: update

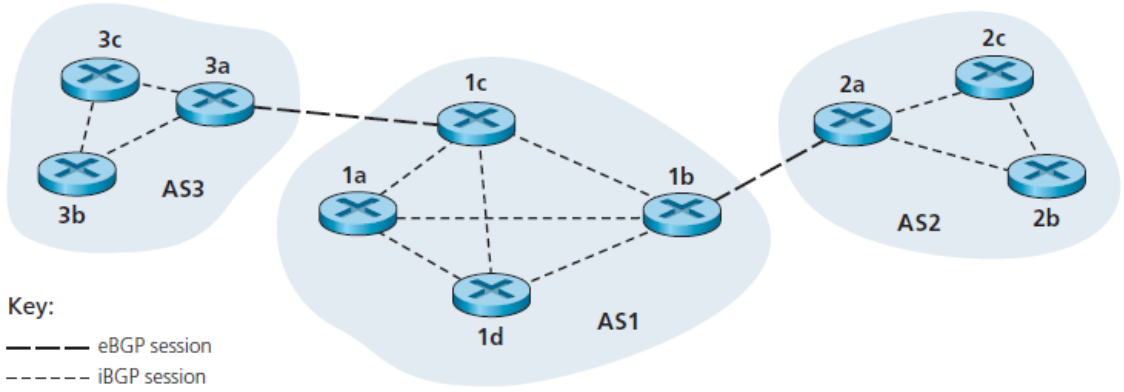

- E.g., 138.16.64/24 is first advertised from AS2 to AS1
  - ✓ AS1 advertises the prefix to AS3: AS-PATH = AS2 AS1

- Withdrawal of one or more routes
  - route is identified by the IP address of the destination network

BGP: notification

- Generated when an error condition is detected
  - **Message header error:** authentication and syntax errors
  - **Open message error:** errors and options not recognized in an Open or unacceptable Hold Time
  - **Update message error:** syntax and validity errors in an Update message
  - **Hold timer expired:**
  - **Finite state machine error:** any procedural error
  - **Cease:** Used by a router to close a connection with another router in the absence of any other error

BGP sessions

- There are two types of BGP sessions
  - eBGP: between BGP speaking routers in different AS
    - ✓ Each AS may have more than one router
  - iBGP: between BGP routers in the same AS
    - ✓ Used to exchange information about external routes

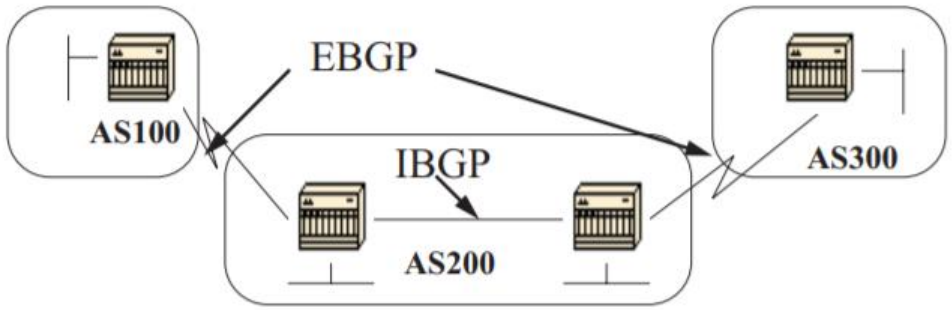

Multicast routing

Multicasting versus unicasting

❑ *Unicasting*

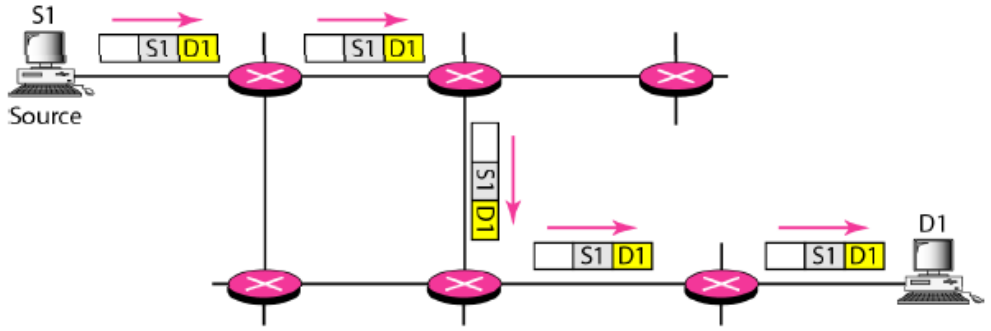

❑ *Multicasting*

- One source and a group of destination

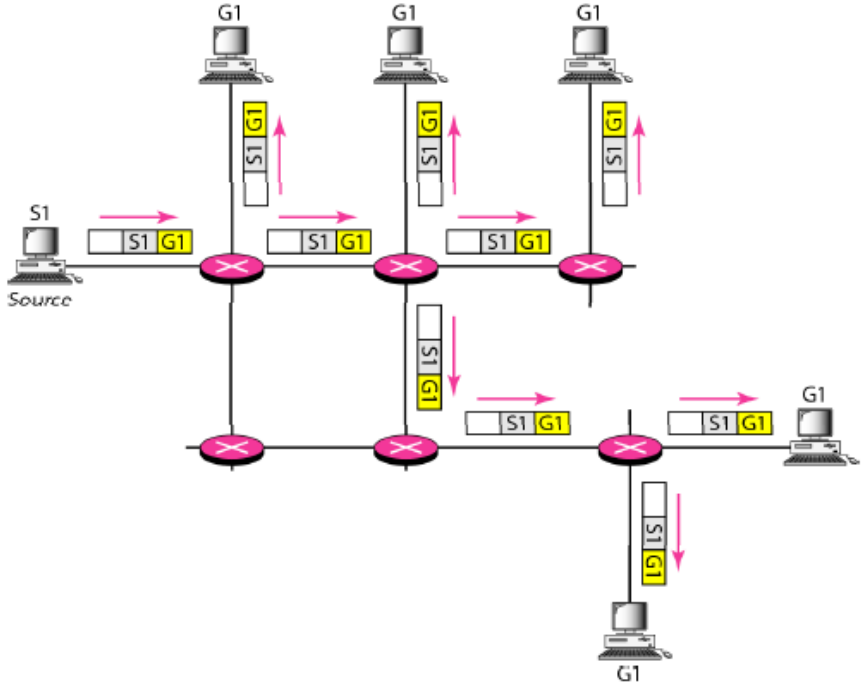

Multicasting versus multiple unicasting

- Multiple unicasting
  - Several packets start from the source
    - ✓ Not efficient
    - ✓ May create long delays, particularly with a large group

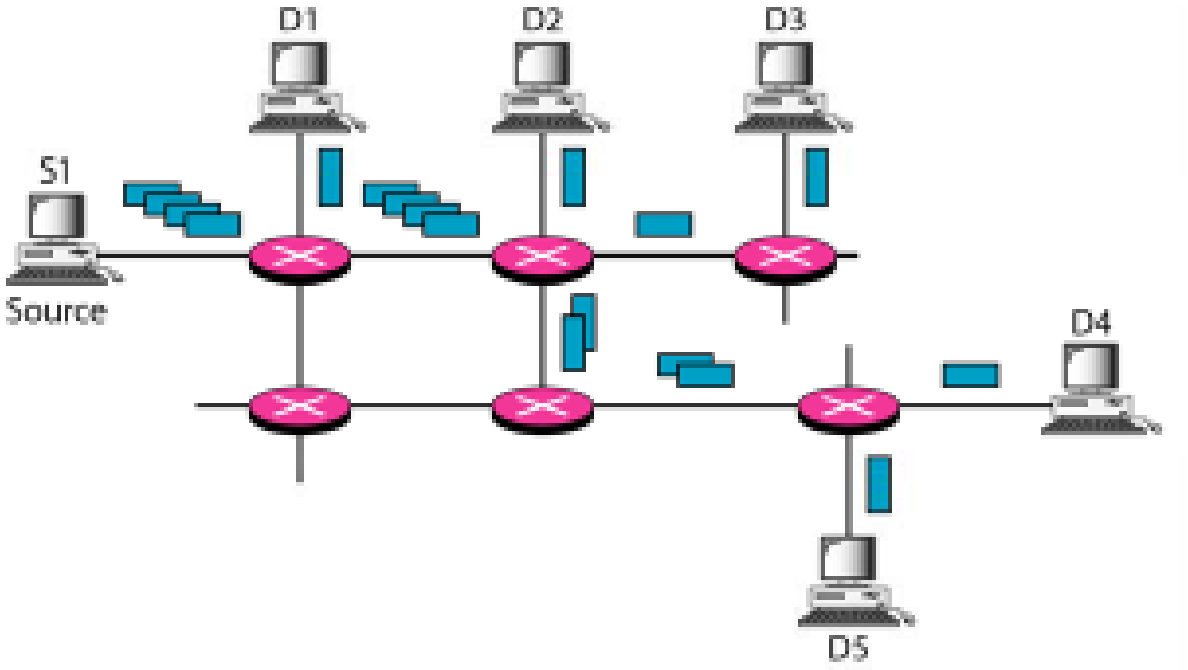

Application of multicasting

- ***Access to distributed database***
  - user's request is multicast to all the database locations, and the location that has the information responds
- ***Information Dissemination***
  - For example, a software update can be sent to all purchasers of a particular software package
- ***Dissemination of News***
- ***Teleconferencing***
  - The individuals attending a teleconference all need to receive the same information at the same time

Multicast routing

- Optimal routing: shortest path tree
  - The root of the tree is the source, and the leaves are the potential destinations
  - The path from the root to each destination is the shortest path
- unicast routing: the router knows the shortest path for each destination

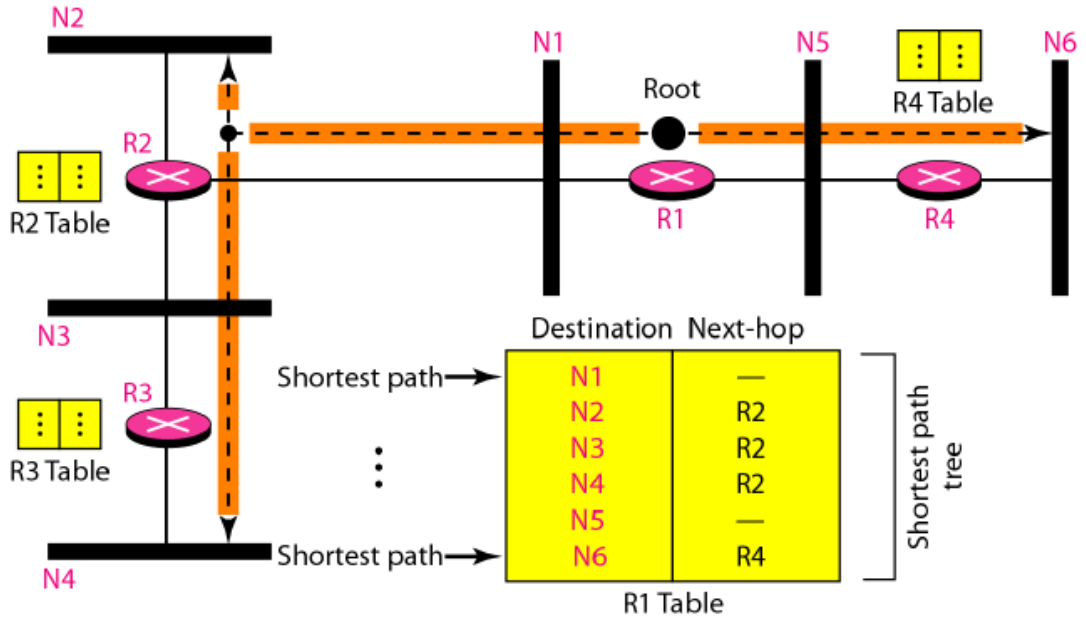

Multicast routing

- Forwarding of a single packet to members of a group requires a shortest path tree
  - If we have *n* groups, we may need *n* shortest path trees
- **Source-based tree approach**
  - In the source-based tree approach, each router needs to have one shortest path tree for each group

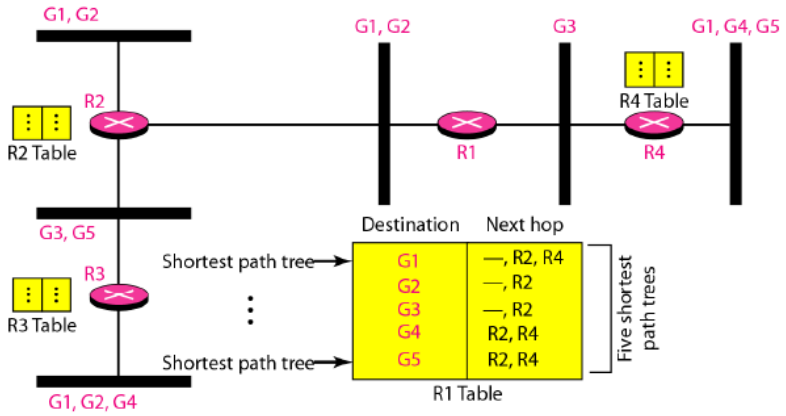

Complex if we have hundreds or thousands of groups

Multicast routing

Group-Shared Tree

- Only one designated (core) router takes the responsibility of distributing multicast traffic
- If a router receives a multicast packet, it encapsulates the packet in a unicast packet and sends it to the core router
- The core router removes the multicast packet from its capsule, and consults its routing table to route the packet

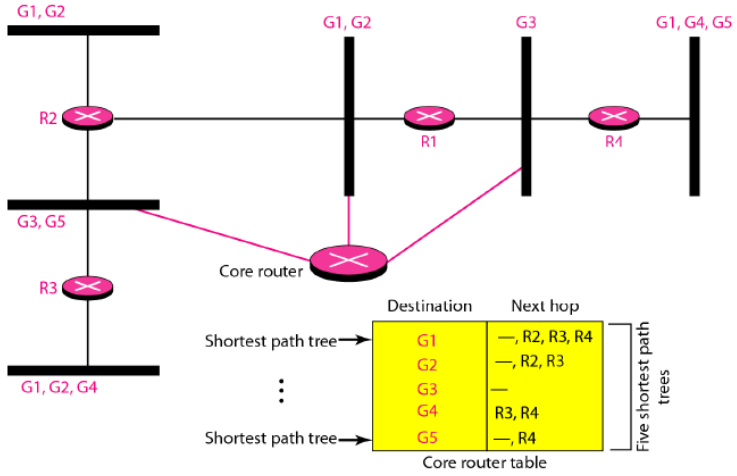

Multicast routing protocols

- ***Multicast link state routing: MOSPF***
  - uses the source-based tree approach
  - Builds a tree that contains all the hosts belonging to a group, the unicast address of the host is used in the calculation
  - For efficiency, the router calculates the shortest path trees on demand (when it receives the first multicast packet)

Multicast routing protocols

❑ *Multicast Distance Vector: DVMRP*

- uses source-based trees, but the router never actually makes a routing table
- a process based on four decision-making strategies
  - ✓ ***Flooding broadcast packet***
    - Every network with active members receives the packet.
    - However, so will networks without active members
    - creates loops in the systems
  - ✓ ***Reverse Path Forwarding (RPF)***
    - is a modified flooding strategy
    - To prevent loops, only one copy is forwarded; the other copies are dropped
    - The router receives a packet and extracts the source address (a unicast address)
    - It consults its unicast routing table as though it wants to send a packet to the source address
    - The router forwards the packet if it has traveled from the shortest path; it discards it otherwise

Reverse path forwarding

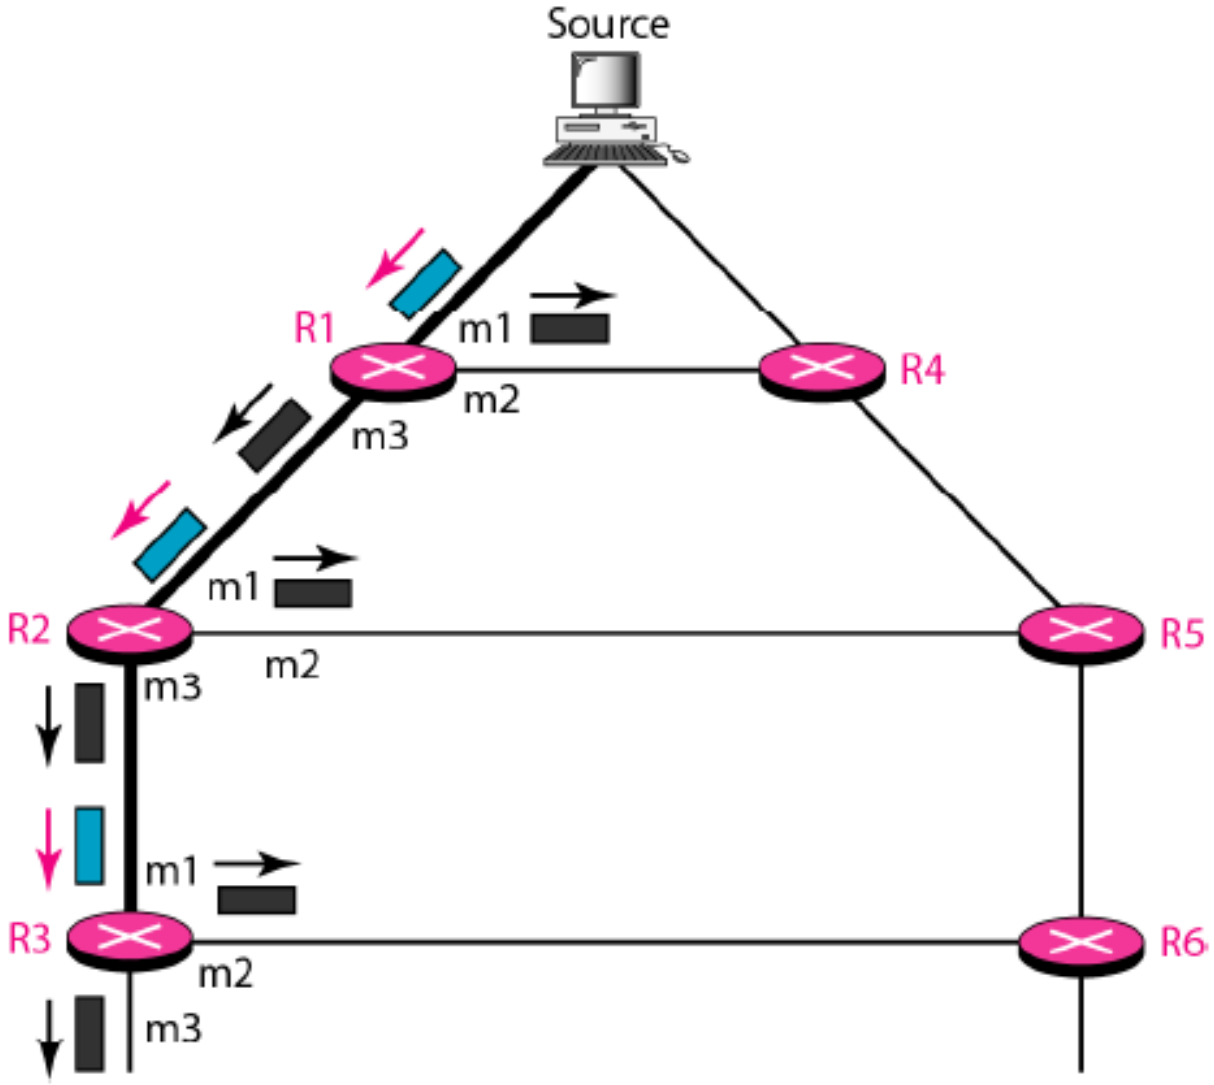

Multicast Distance Vector: DVMRP

- ***Reverse path broadcasting (RPB)***
  - RPF does not guarantee that each network receives only one copy

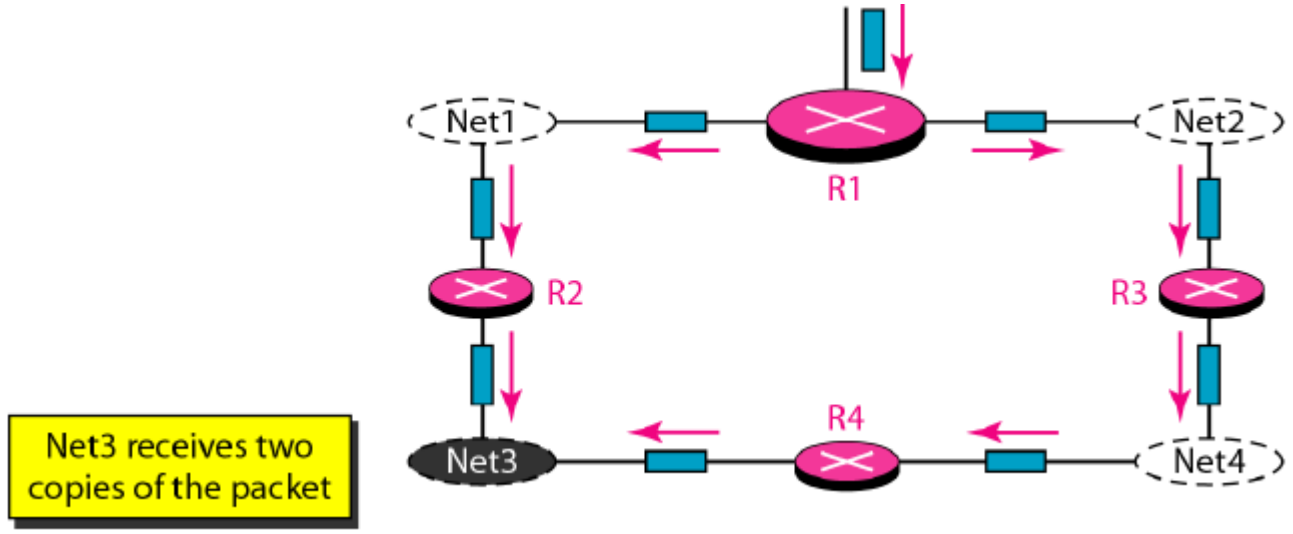

  - A network can receive a multicast packet from a particular source only through a designated parent router

Reverse path broadcasting (RPB)

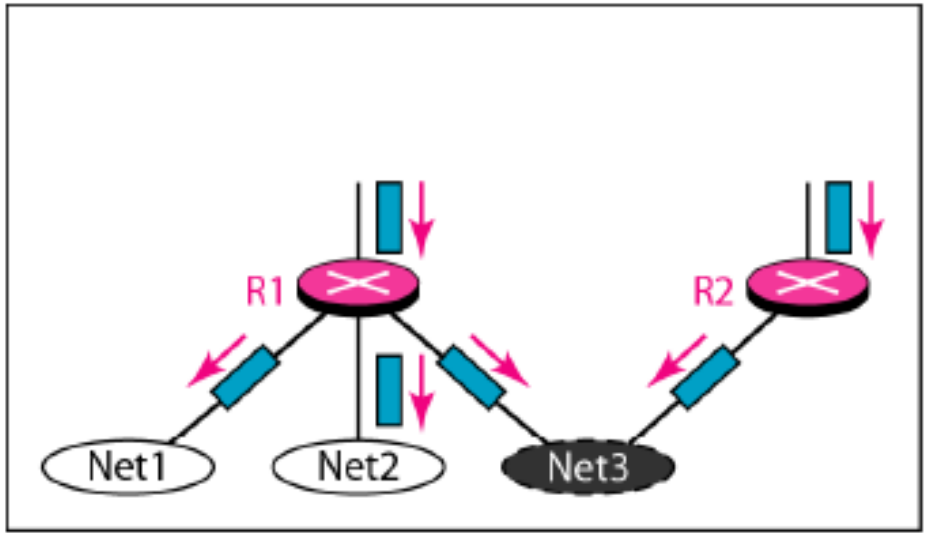

a. RPF

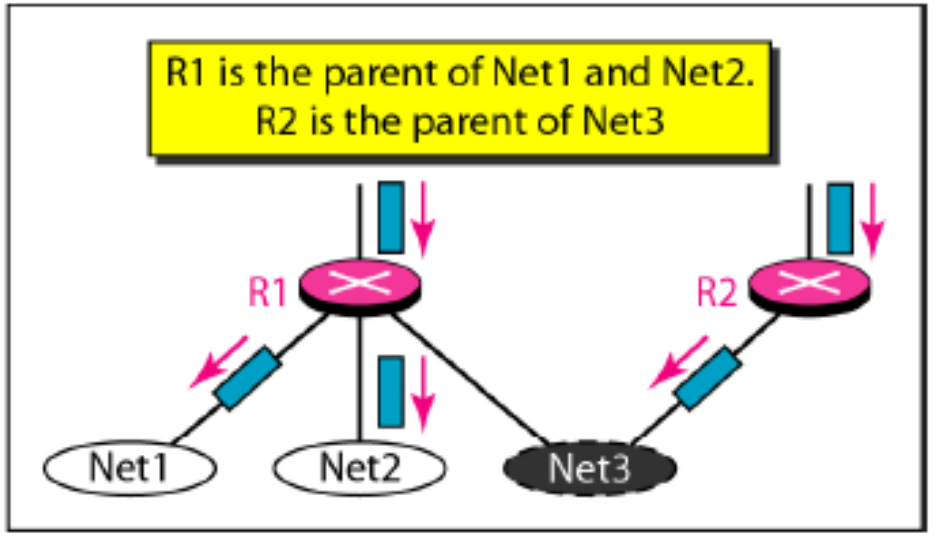

b. RPB

Multicast Distance Vector: DVMRP

- ***Reverse path multicasting (RPM)***
  - To convert broadcasting to multicasting, the protocol uses two procedures, pruning and grafting