ECE/COMPSCI 356 Computer Network Architecture

Lecture 17: Multicast

Neil Gong

neil.gong@duke.edu

Slides credit: Xiaowei Yang, PD

Overview

- Multicast routing protocols

Communication models

- Unicast
  - One-to-one
  - Unicast routing
- Multicast
- Broadcast

Needs for Multicast

- One-to-many
  - Radio station broadcast
  - Transmitting news, stock-price
  - Software updates to multiple hosts
- Many-to-many
  - Multimedia teleconferencing
  - Online multi-player games
  - Distributed simulations

Needs for Multicast

- Without support for multicast
  - A source needs to send a separate packet with the identical data to each member of the group
    - This redundancy consumes more bandwidth
    - Redundant traffic is not evenly distributed, concentrated near the sending host
  - Source needs to keep track of the IP address of each member in the group
    - Group may be dynamic
- To support many-to-many and one-to-many, IP provides an IP-level multicast

Overview

- IP multicast model is based on multicast groups
  - Each group has its own IP multicast address
  - Hosts that are members of a group receive copies of any packets sent to that group's multicast address
  - A host can be in multiple groups
  - A host can join and leave groups

Overview

- One-to-many multicast
  - Source specific multicast (SSM)
  - A receiving host specifies both a multicast group and a specific sending host
- Many-to-many model
  - Any source multicast (ASM)

Overview

- A host signals its desire to join or leave a multicast group by communicating with its local router using a special protocol
  - In IPv4, the protocol is Internet Group Management Protocol (IGMP)
  - In IPv6, the protocol is Multicast Listener Discovery (MLD)

- The router has the responsibility for making multicast behave correctly with regard to the host

Multicast Routing

- A router's unicast forwarding tables
  - Nexthop is a link
- Multicast forwarding tables
  - Nexthop includes multiple links
- Unicast forwarding tables collectively specify a set of paths
- Multicast forwarding tables collectively specify a set of trees
  - Multicast distribution trees

Multicast distribution trees

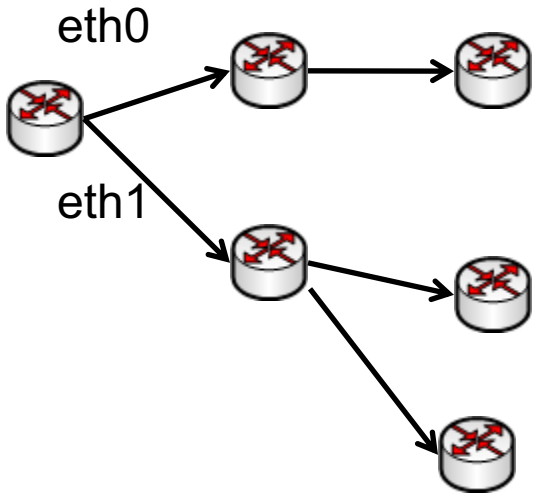

- Multicast distribution trees: multiple outgoing interfaces for a multicast destination address

Multicast Routing Protocols

- Reverse Path Broadcast (RPB)
- Protocol Independent Multicast (PIM)

Reverse Path Broadcast (RPB)

- Using existing distance vector routing protocol
- Each router already knows that shortest path to source S goes through link L.
- When receive multicast packet from S, forward on all outgoing links (except the one on which the packet arrived), iff packet arrived from L.

Reverse Path Broadcast (RPB)

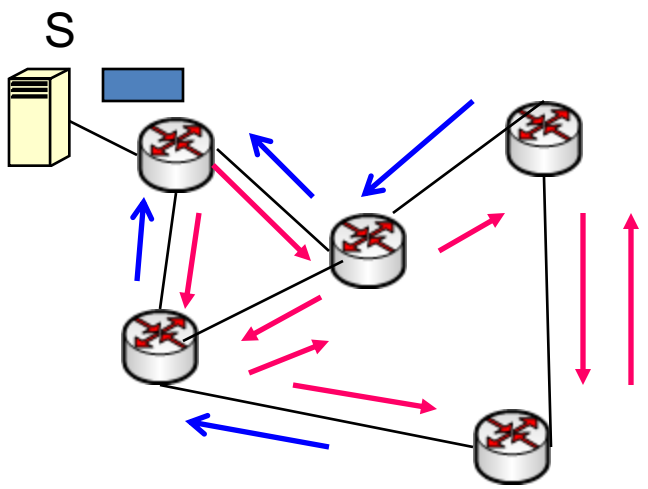

- If packet comes from link L, and next hop to S is L, broadcast to all outgoing links except the incoming one

Problems with RPB

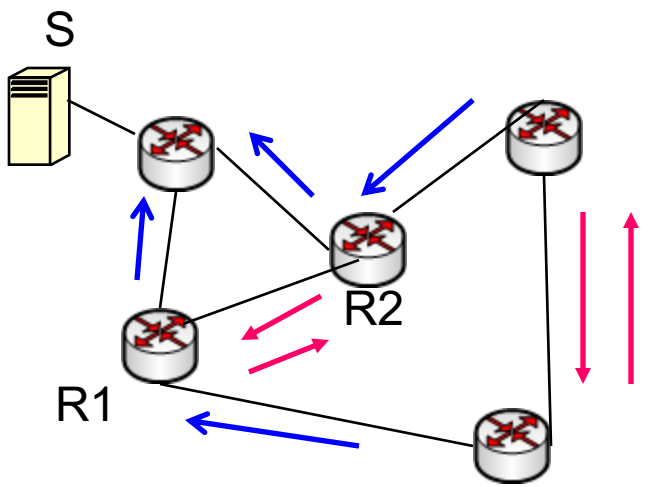

- Problems
  - multiple routers on a LAN → receiving multiple copies of packets
  - Not all hosts are in the multicast group. Broadcast is a waste

Designated router election

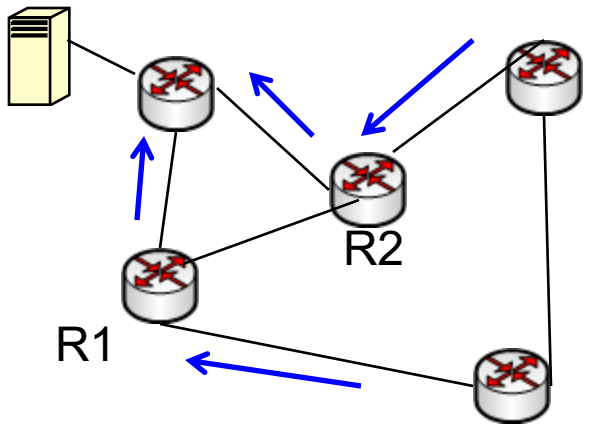

- Address the duplicate broadcast packet problem
- Routers on the same LAN elect a parent that has the shortest distance to S
  - Parent is one with the shortest path
    - Routers can learn this from DV routing messages
  - If tie, elect one with the smallest router ID

Pruning networks

- Goal: Prune networks that have no hosts in group G to save cost

- Step 1: Determine if LAN is a *leaf* with no members in G
  - leaf if parent is only router on the LAN
  - determine if any hosts are members of G using IGMP

- Step 2: Propagate “no members of G here” information

A pruning example

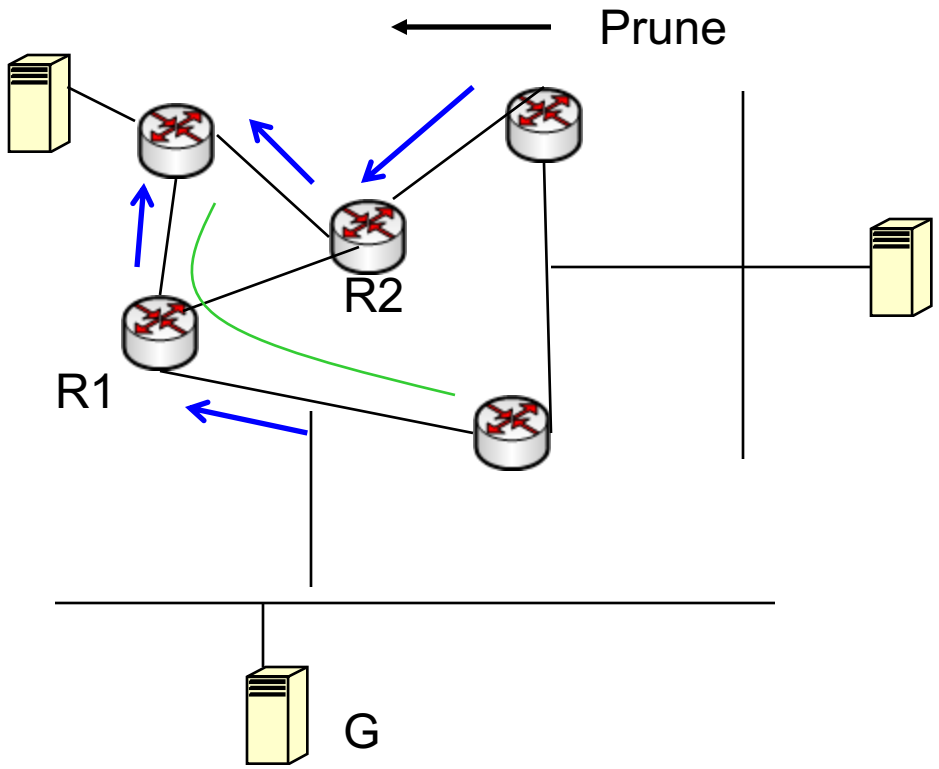

Protocol Independent Multicast (PIM)

- Problem with RPB
  - Broadcast is inefficient if few nodes are interested
  - Most routers must explicitly send prune messages
  - Dependent on routing protocols

- Solution
  - PIM
    - Dense mode (DM): flood & prune similar to RPB
    - Sparse mode (SM): send join messages to rendezvous point (RP)
    - Not dependent on any unicast routing protocol, unlike RPB

PIM-SM

1. Routers explicitly join a shared distribution tree

2. Source-specific trees are created later for more efficient distribution if there is sufficient traffic

Join

- PIM-SM assigns each group a special router known as the rendezvous point (RP)
- A router that has hosts interested in a group G sends a Join message to RP
- A router looks at the join message and create a multicast routing entry (*,G) pointing to the incoming interface. This is called an all sender forwarding entry
- It propagates join to RP

PIM-SM

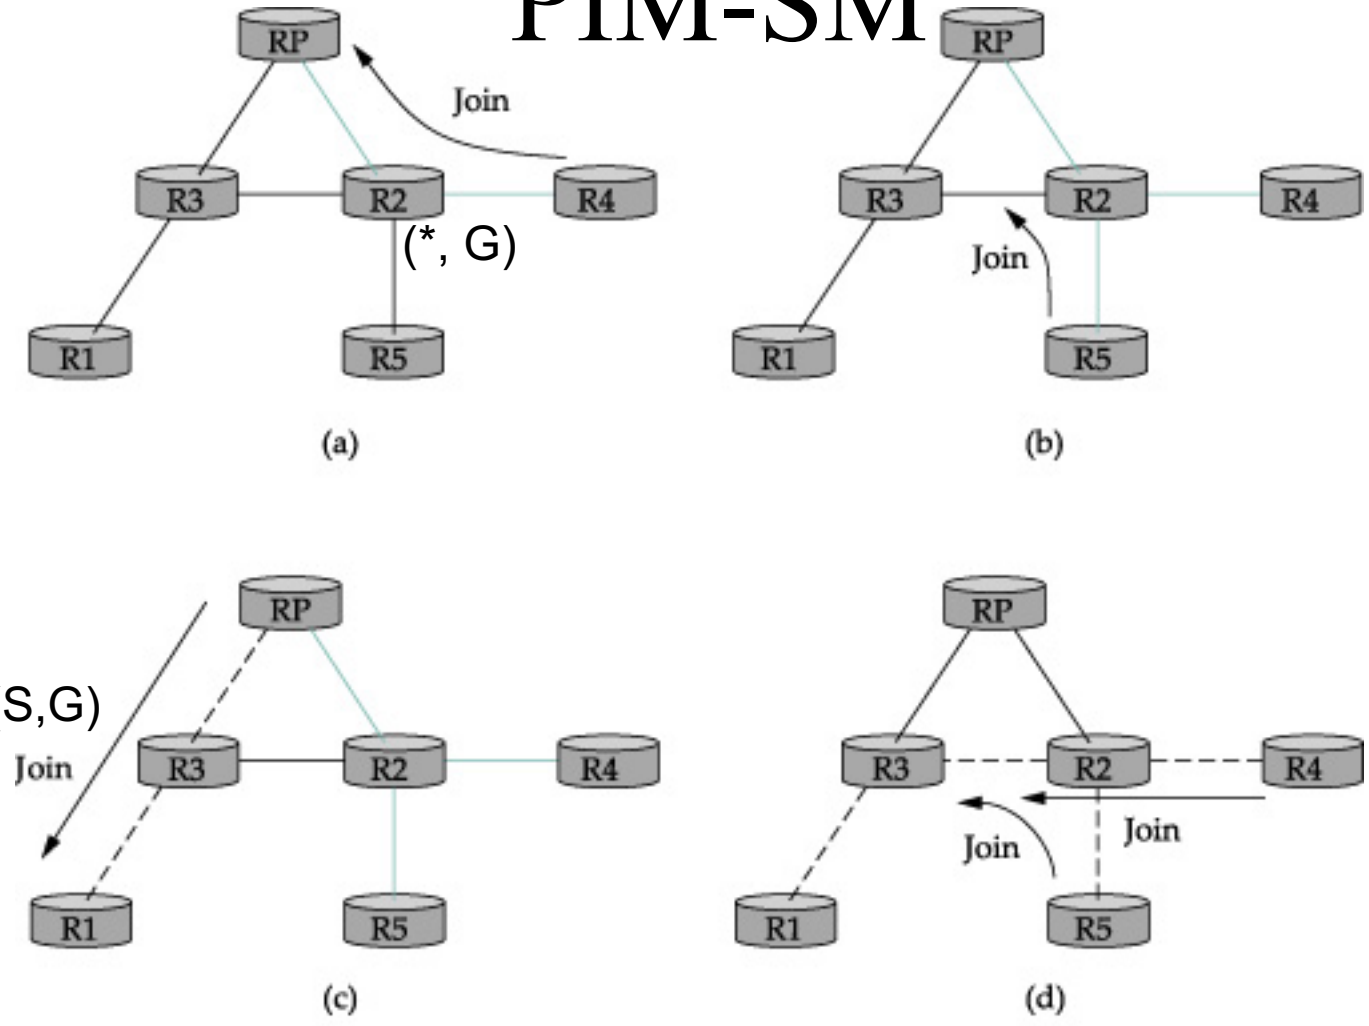

- (a): R4 joins the multicast group
- (b): R5 joins the group
  - The Join message travels to R2

Forwarding along a shared tree

- If a source S wishes to send to the group
  - S sends a packet to its designated router (R1) with the multicast group as the destination address
  - R1 encapsulates the packet into a PIM register message, unicast it to RP
    - IP Tunnel
- RP decapsulates it and forwards to the shared trees

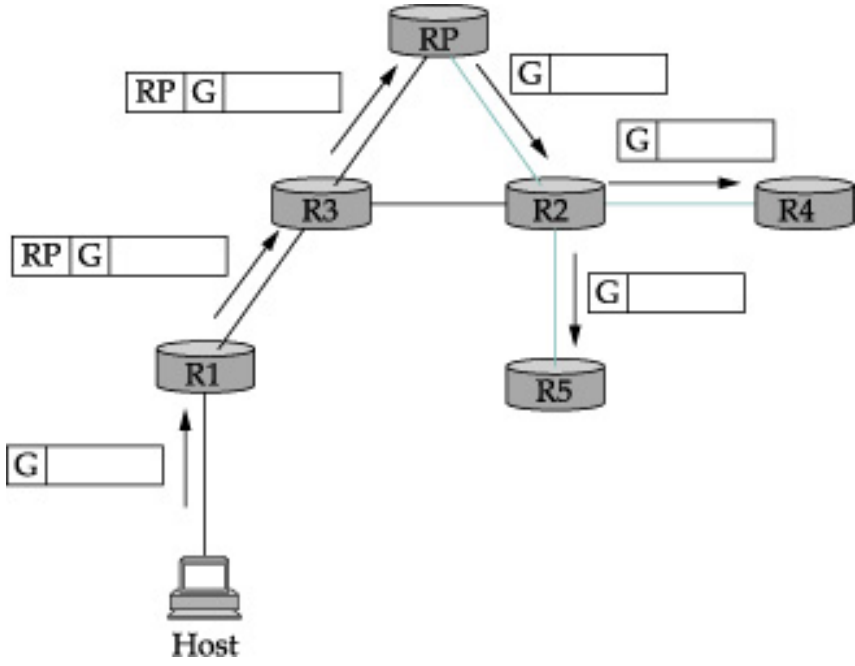

Avoiding IP tunnel

- Problems
  - Encapsulation is inefficient
- Solution:
  - RP sends Join message to source S
  - R3 now knows the group (S,G)
    - Send packets without tunnel

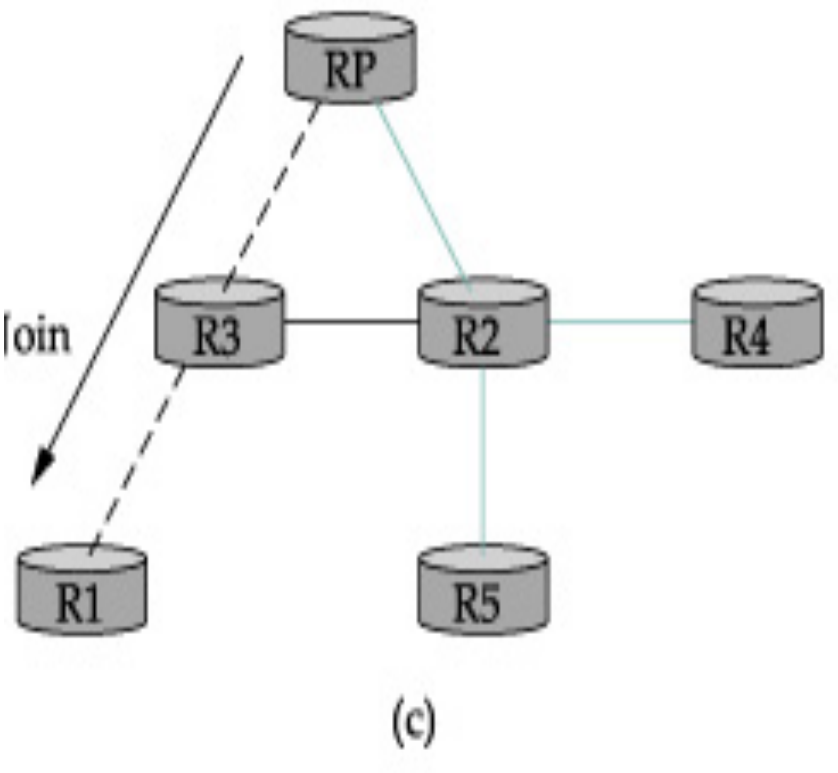

Source specific tree

- Problem: shared trees are inefficient as paths could be longer than shortest path

- Solution
  - If S sends at high rates, routers send source-specific Join messages
  - Trees may no longer involve RP

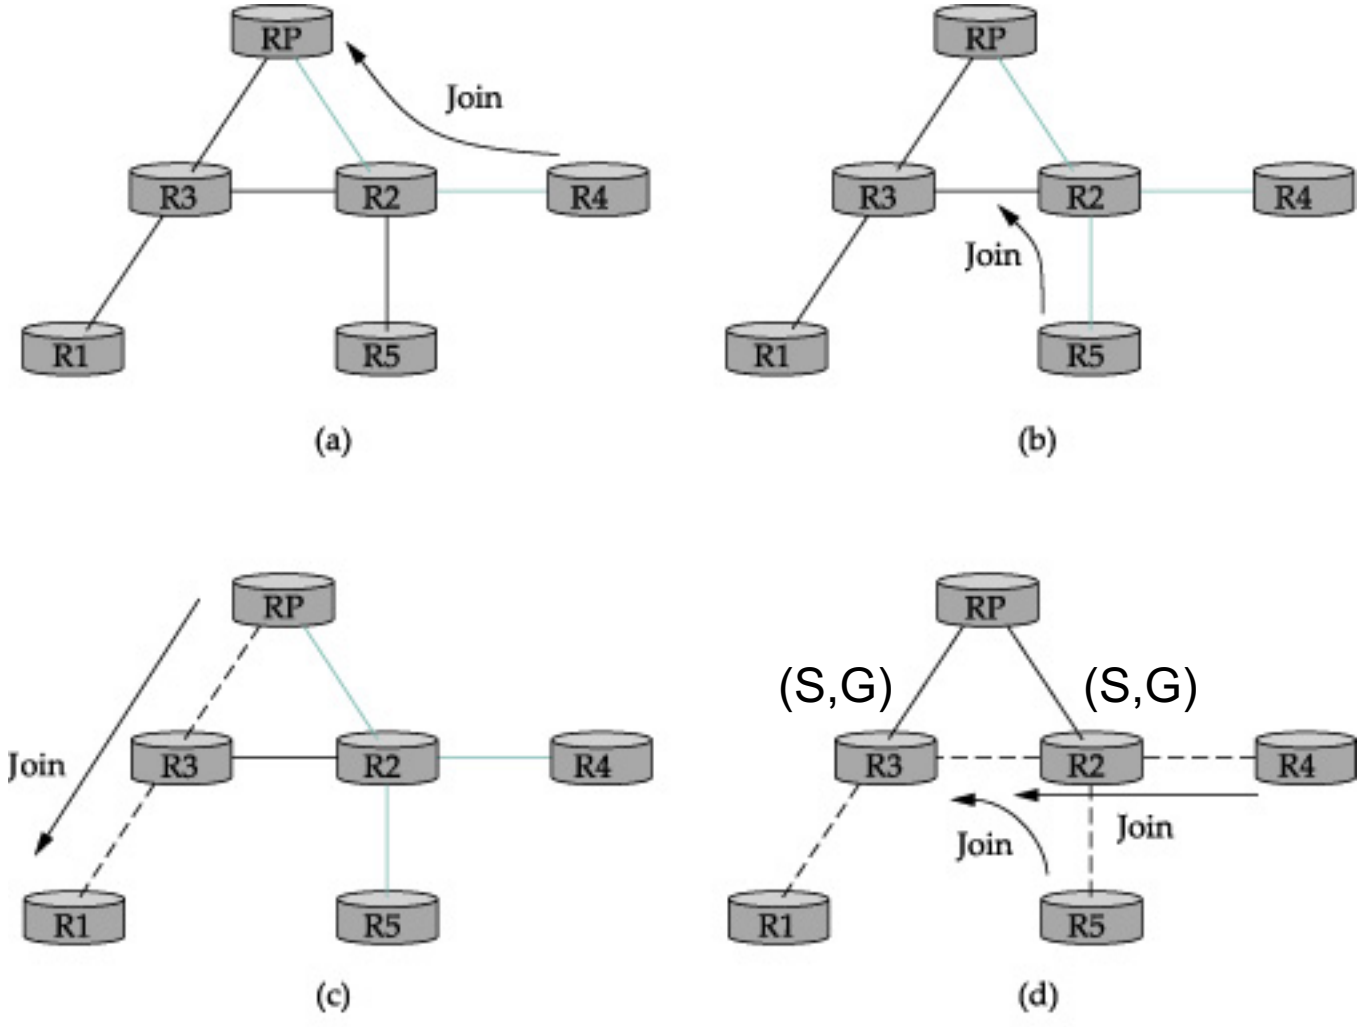

- R1 is the source

Inter-domain multicast

- Challenging for the entire Internet to agree on a single RP for a group G

- Multicast Source Discovery Protocol
  - Hierarchical
    - Similar to intradomain and interdomain routing
  - Intradomain: PIM-SM
  - Interdomain: a distribution tree among all domain's RPs

Inter-domain Multicast

Multicast Source Discovery Protocol (MSDP)

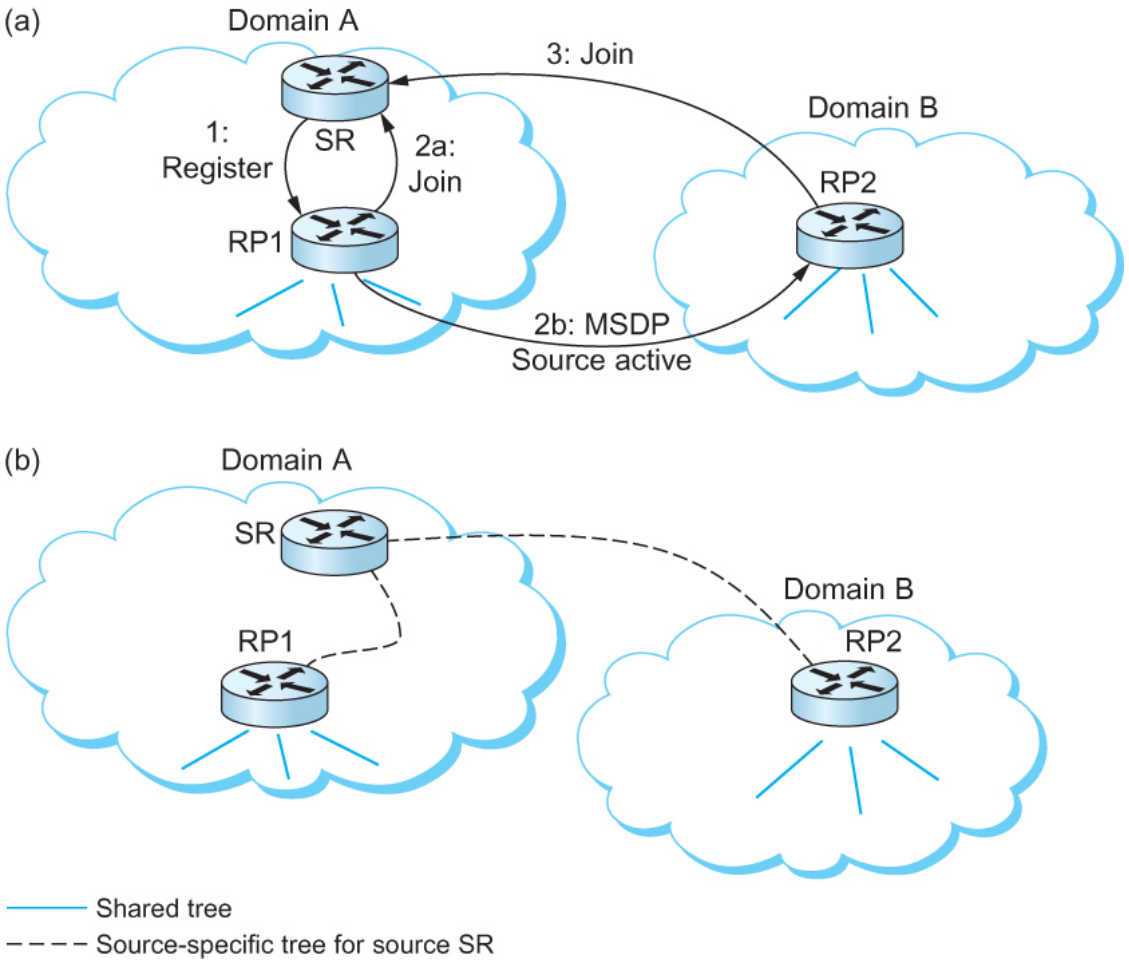

Summary

- Reverse Path Broadcast (RPB)
- Protocol Independent Multicast (PIM)