DATA COMMUNICATOIN NETWORKING

Instructor: Ouldooz Baghban Karimi

Course Book & Slides:
Computer Networking, A Top-Down Approach
By: Kurose, Ross

Course Overview

- **Basics of Computer Networks**
  - Internet & Protocol Stack
  - Application Layer
  - Transport Layer
  - Network Layer
  - Data Link Layer

- **Advanced Topics**
  - Case Studies of Computer Networks
  - Internet Applications
  - Network Management
  - Network Security

Broadcast Routing

- Deliver packets from source to all other nodes
- Source duplication is inefficient
- Source duplication: How does source determine recipient addresses?

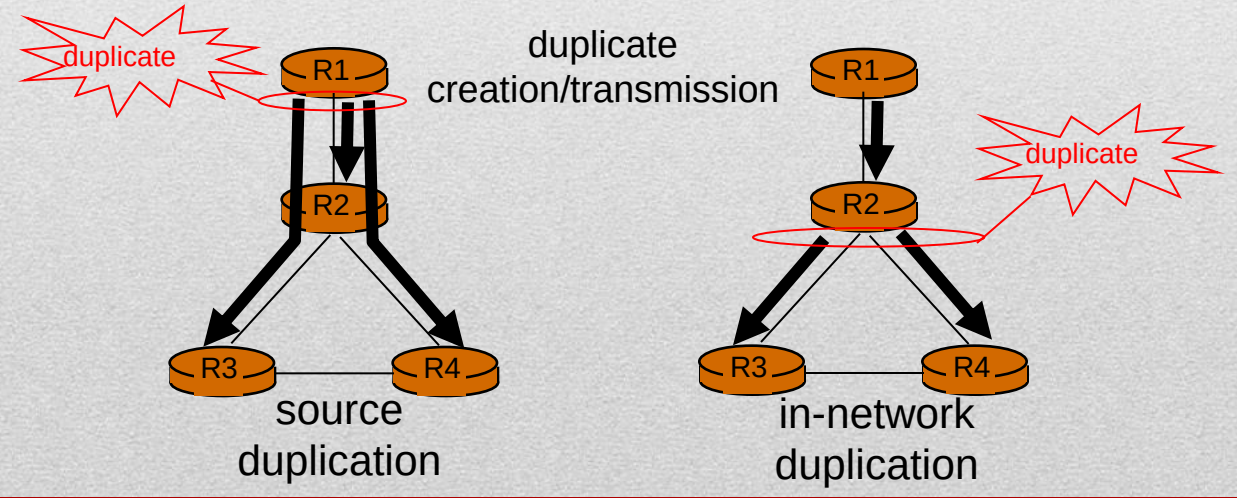

In Network Duplication

- **Flooding**
  - When node receives broadcast packet, sends copy to all neighbors
  - Problems: cycles & broadcast storm
- **Controlled flooding**
  - Node only broadcasts packet if it has not broadcast same packet before
  - Node keeps track of packet ids already broadcasted
  - Or reverse path forwarding (RPF): only forward packet if it arrived on shortest path between node and source
- **Spanning tree**
  - No redundant packets received by any node

Spanning Tree

- First construct a spanning tree
- Nodes then forward/make copies only along spanning tree

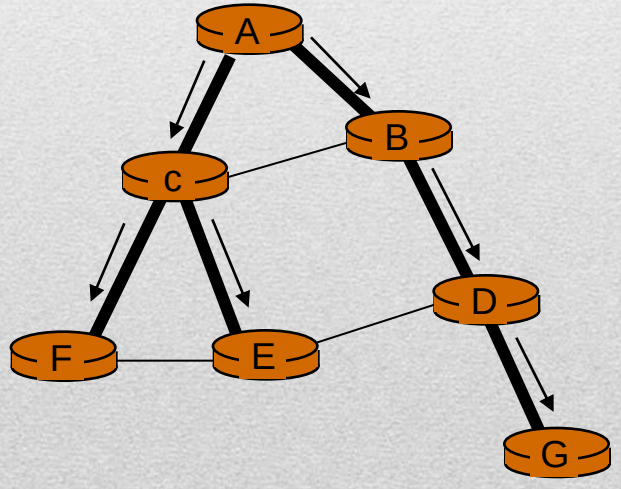

(a) broadcast initiated at A

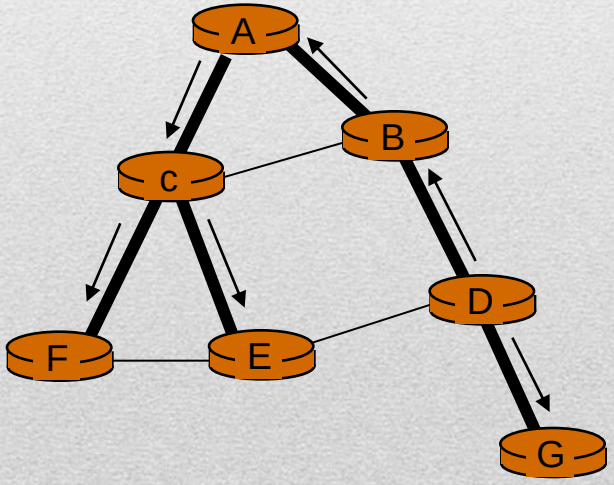

(b) broadcast initiated at D

Spanning Tree: Creation

- Center node
- Each node sends unicast join message to center node
  - Message forwarded until it arrives at a node already belonging to spanning tree

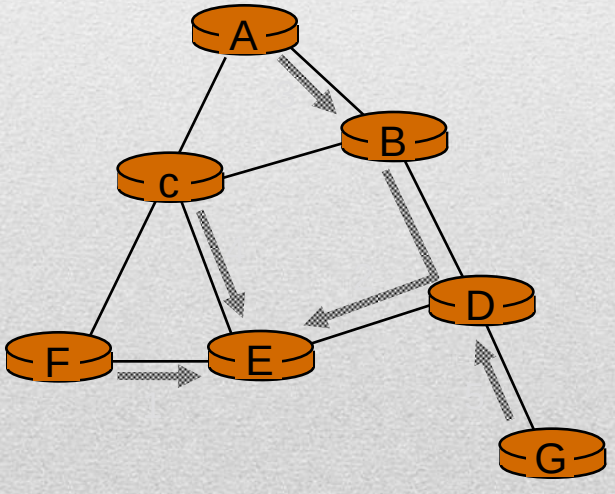

(a) stepwise construction of spanning tree (center: E)

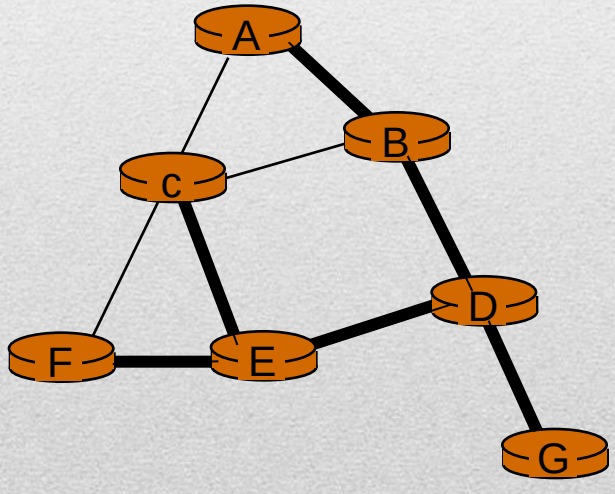

(b) constructed spanning tree

Multicast Routing

Goal: find a tree (or trees) connecting routers having local multicast group members

- Tree: not all paths between routers used
- Shared-tree: same tree used by all group members
- Source-based: different tree from each sender to receivers

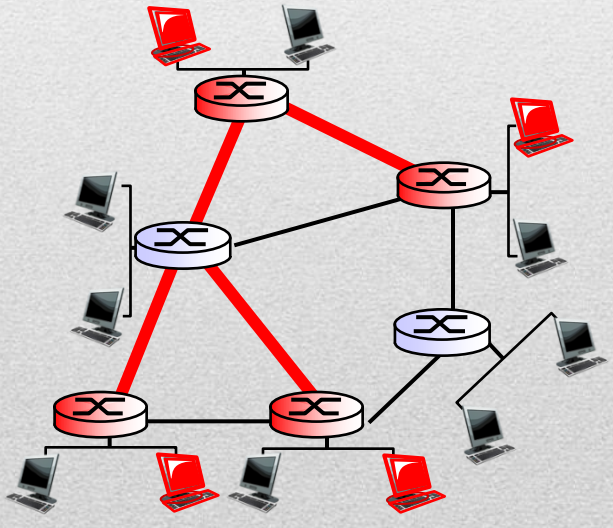

shared tree

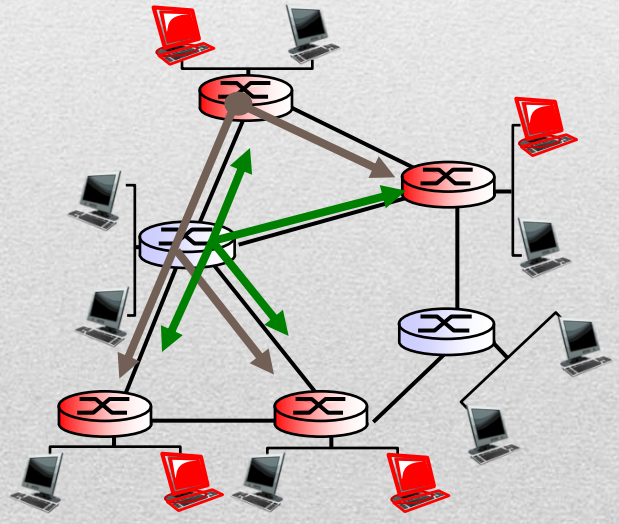

source-based trees

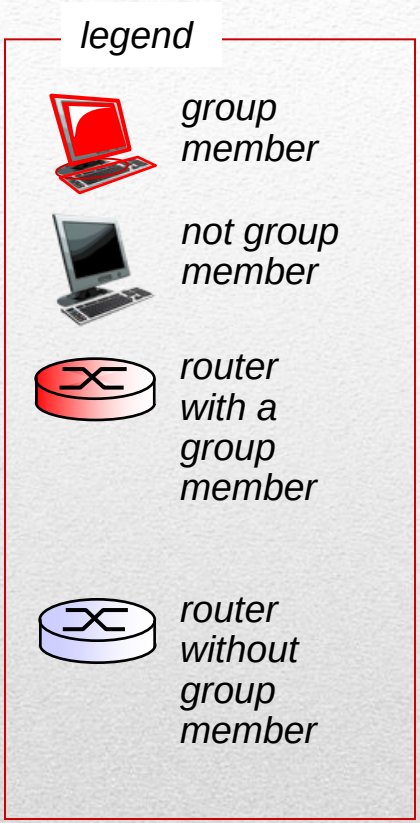

Building Multicast Trees

- **Approaches**
  - *Source-based tree*: one tree per source
    - Shortest path trees
    - Reverse path forwarding
  - *Group-shared tree*: group uses one tree
    - Minimal spanning (Steiner)
    - Center-based trees

Building Multicast Trees

- Multicast forwarding tree: tree of shortest path routes from source to all receivers
  - Dijkstra's algorithm

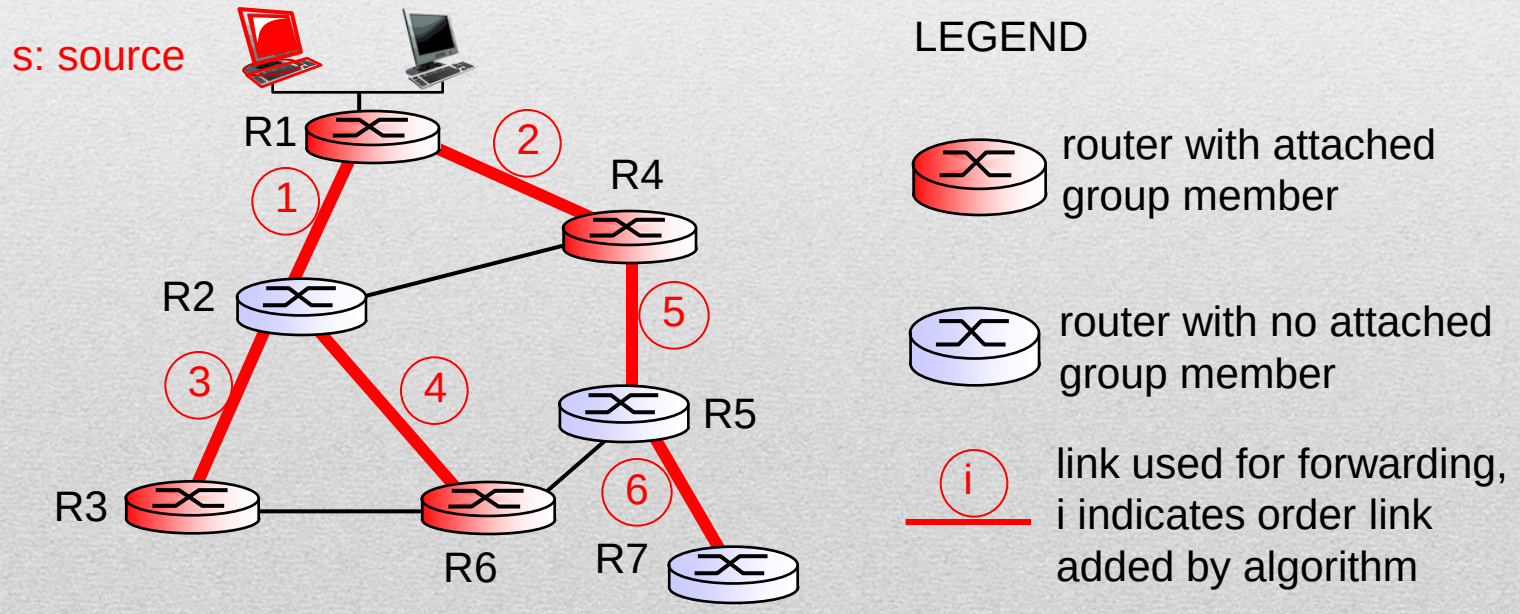

Reverse Path Forwarding

- Rely on router's knowledge of unicast shortest path from it to sender
- Each router has simple forwarding behavior

if (multicast datagram received on incoming link on shortest path back to center)

then flood datagram onto all outgoing links

else ignore datagram

Reverse Path Forwarding

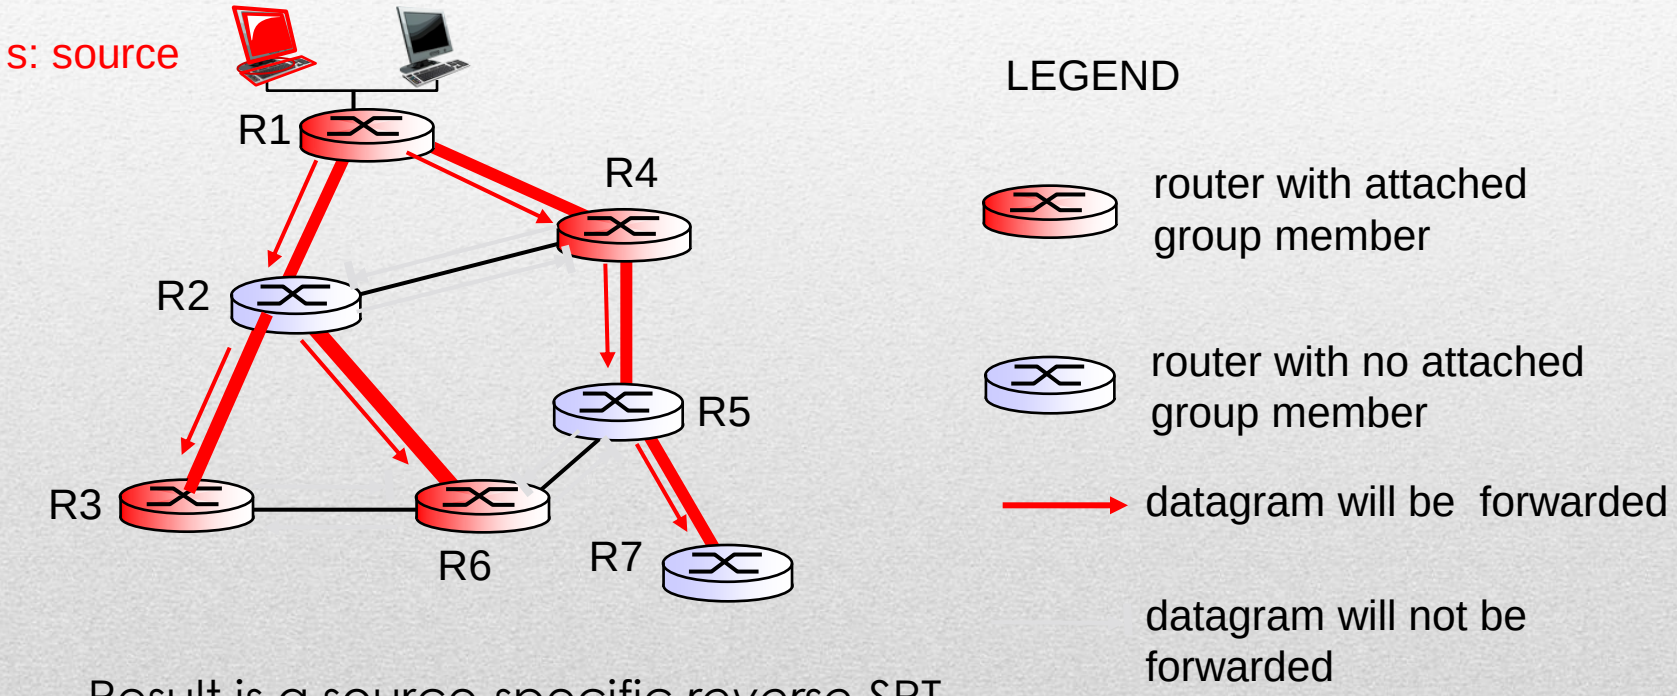

- Result is a source-specific *reverse* SPT
  - May be a bad choice with asymmetric links

Reverse Path Forwarding: Pruning

- Forwarding tree contains sub-trees with no multicast group members
  - No need to forward datagrams down sub-tree
  - Prune messages sent upstream by router with no downstream group members

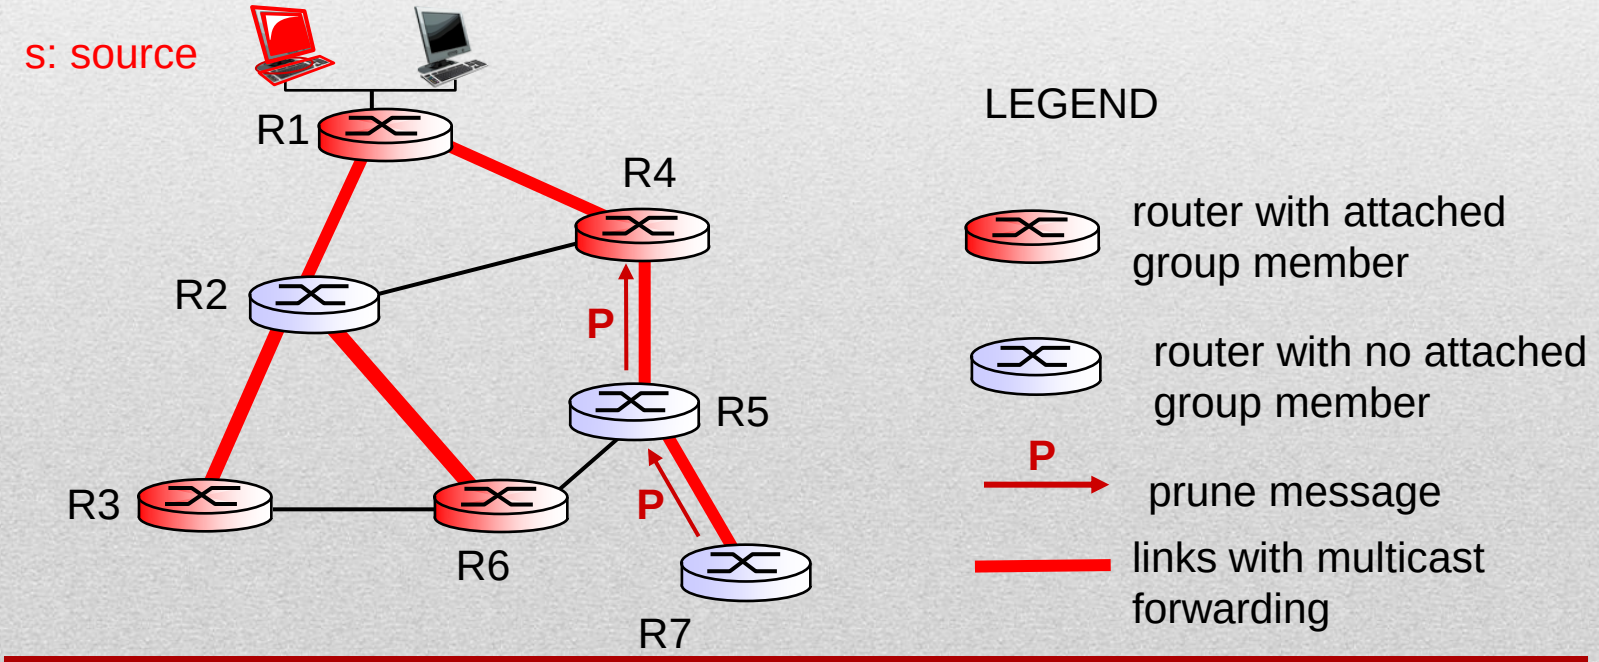

Shared Tree: Steiner Tree

- **Steiner tree:** minimum cost tree connecting all routers with attached group members
- Problem is NP-complete
- Excellent heuristics exists
- Not used in practice:
  - Computational complexity
  - Information about entire network needed
  - Monolithic: rerun whenever a router needs to join/leave

Center-Based Trees

- Single delivery tree shared by all
- One router identified as center of tree
- To join
  - Edge router sends unicast join-message addressed to center router
  - Join-message "processed" by intermediate routers and forwarded towards center
  - Join-message either hits existing tree branch for this center, or arrives at center
  - Path taken by join-message becomes new branch of tree for this router

Center-Based Trees

Suppose R6 chosen as center:

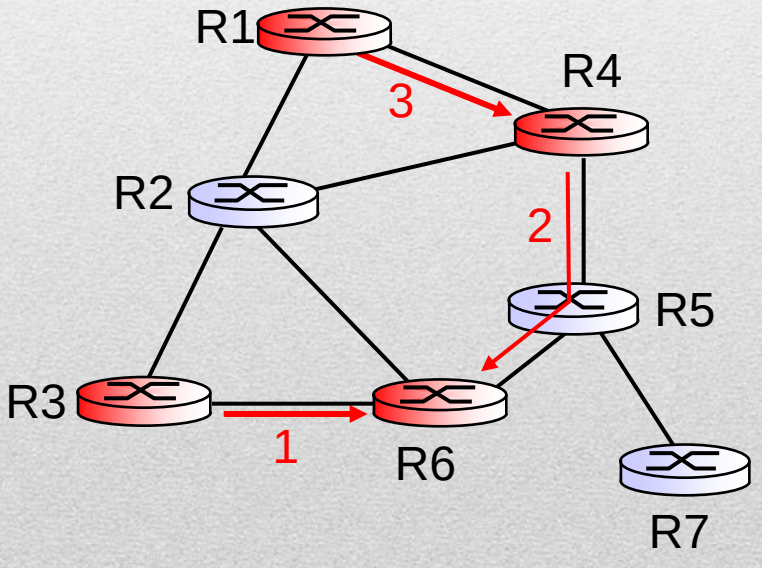

LEGEND

router with attached group member

router with no attached group member

1 path order in which join messages generated

DVMRP: Internet Multicast

- DVMRP: distance vector multicast routing protocol, RFC1075

- Flood and prune: reverse path forwarding, source-based tree
  - RPF tree based on DVMRP's own routing tables constructed by communicating DVMRP routers
  - No assumptions about underlying unicast
  - Initial datagram to multicast group flooded everywhere via RPF
  - Routers not wanting group: send upstream prune messages

DVMRP: Internet Multicast

- Soft state: DVMRP router periodically (1 min.) "forgets" branches are pruned:
  - Multicast data again flows down un-pruned branch
  - Downstream router: re-prune or else continue to receive data
- Routers can quickly re-graft to tree
  - Following IGMP join at leaf
- Odds and ends
  - Commonly implemented in commercial router

Tunneling

Q: how to connect islands of multicast routers in a sea of unicast routers?

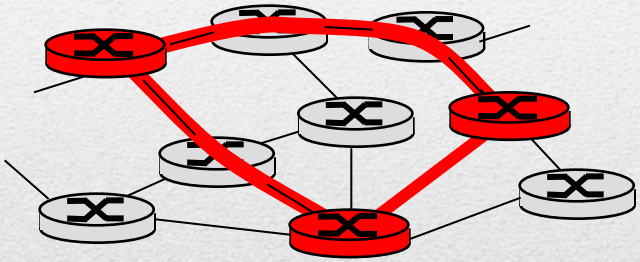

physical topology

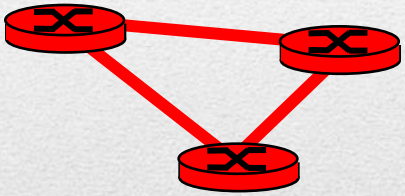

logical topology

- Multicast datagram encapsulated inside "normal" (non-multicast-addressed) datagram
- Normal IP datagram sent through "tunnel" via regular IP unicast to receiving multicast router (recall IPv6 inside IPv4 tunneling)
- Receiving multicast router un-encapsulates to get multicast datagram

PIM: Protocol Independent Multicast

- Not dependent on any specific underlying unicast routing algorithm (works with all)
- Two different multicast distribution scenarios :

Dense:

- Group members densely packed, in "close" proximity.
- Bandwidth more plentiful

Sparse:

- Number of networks with group members small interconnected networks
- Group members widely dispersed
- Bandwidth not plentiful

PIM: Protocol Independent Multicast

Consequences of sparse-dense dichotomy:

Dense

- Group membership by routers assumed until routers explicitly prune
- Data-driven construction on multicast tree (e.g., RPF)
- Bandwidth and non-group-router processing profligate

Sparse

- No membership until routers explicitly join
- Receiver- driven construction of multicast tree (e.g., center-based)
- Bandwidth and non-group-router processing conservative

PIM: Dense Mode

- Flood-and-prune RPF: similar to DVMRP but…
- Underlying unicast protocol provides RPF info for incoming datagram
- Less complicated (less efficient) downstream flood than DVMRP reduces reliance on underlying routing algorithm
- Has protocol mechanism for router to detect it is a leaf-node router

PIM: Sparse Mode

- Center-based approach
- Router sends join message to rendezvous point (RP)
  - Intermediate routers update state and forward join
- After joining via RP, router can switch to source-specific tree
  - Increased performance: less concentration, shorter paths

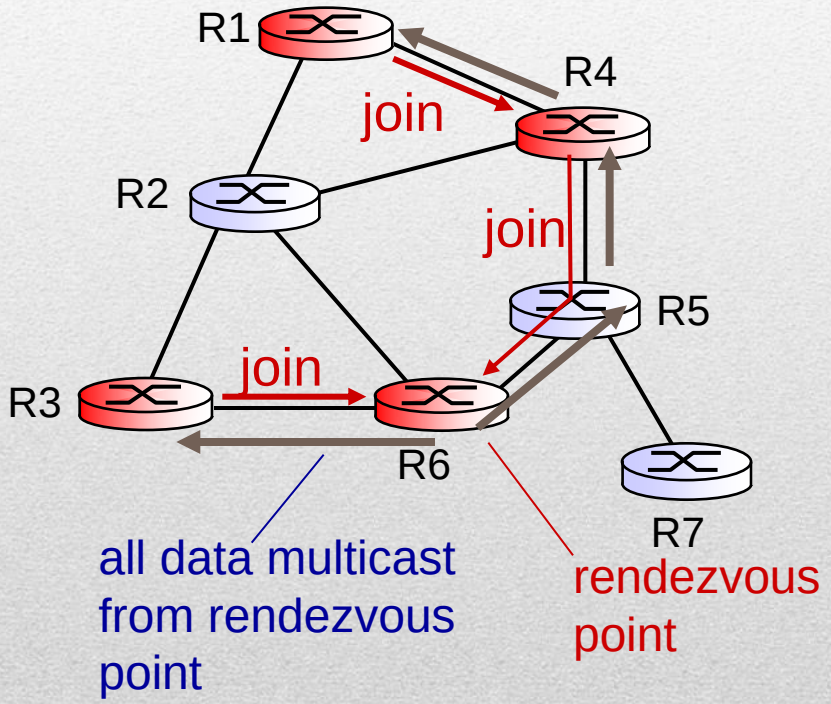

PIM: Sparse Mode

Sender(s)

- Unicast data to RP, which distributes down RP-rooted tree
- RP can extend multicast tree upstream to source
- RP can send stop message if no attached receivers
  - "no one is listening!"