CMPE 150: Introduction to Computer Networks Set 13:

Internet Multicasting

Dr. Soumya Roy (soumya@cse.ucsc.edu)

Homework Assignments

Homework assignment #3

Chapter Four

Due by May 22……

Multicast: one sender to many receivers

- Multicast: act of sending datagram to multiple receivers with single "transmit" operation
  - analogy: one teacher to many students

Uses of Multicasting

- Multimedia “broadcast”: IP radio, webcasts
- Teleconferencing: multiparty videoconferencing with (Mbone, CUSeeMe)
- Database replication
- Distributed computing: immediate updates to all members
- Real-time workgroups: multimedia collaboration
- System management: concerted file updates to group of hosts
- Stock ticker: low-bandwidth quotes to millions of hosts (cf. Pointcast)

Multicast: one sender to many receivers

Question: how to achieve multicast

Multicast via Unicast

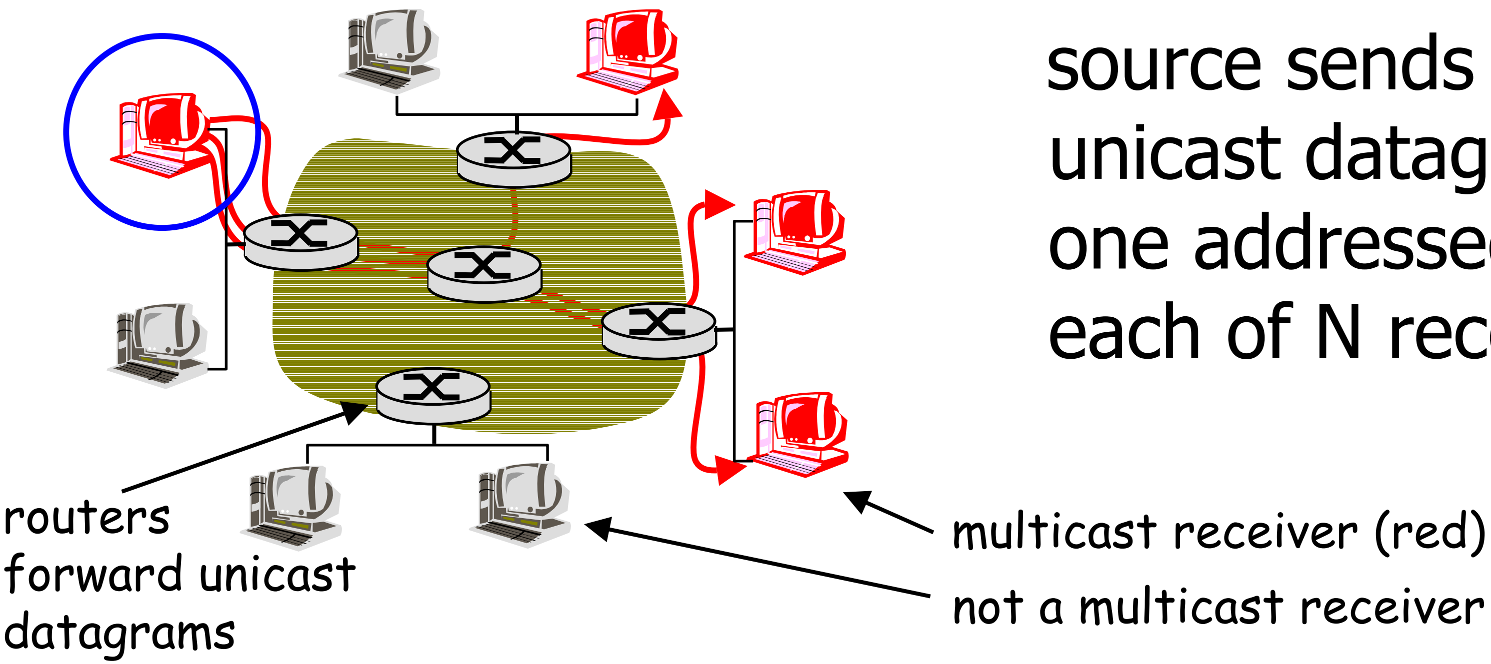

source sends N unicast datagrams, one addressed to each of N receivers

Application Layer Multicast

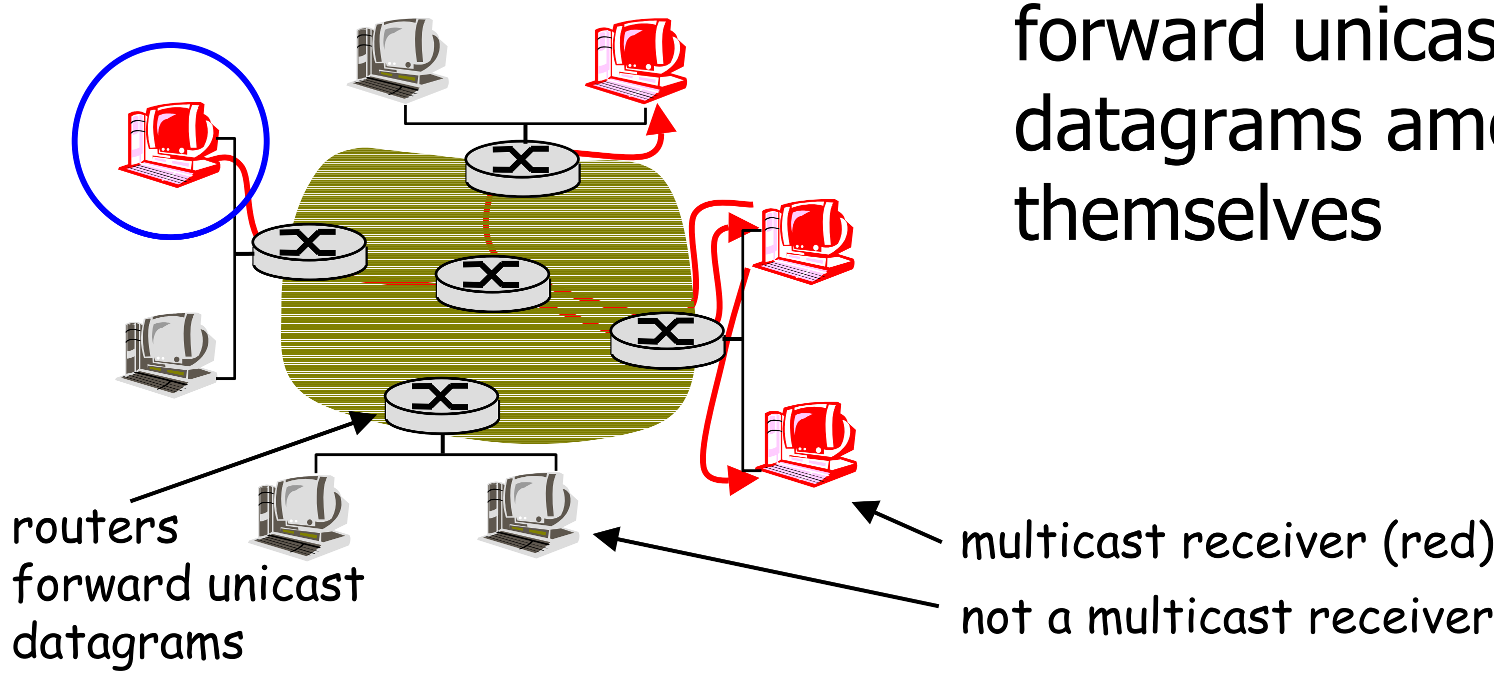

- end systems involved in multicast copy and forward unicast datagrams among themselves

Network Multicast

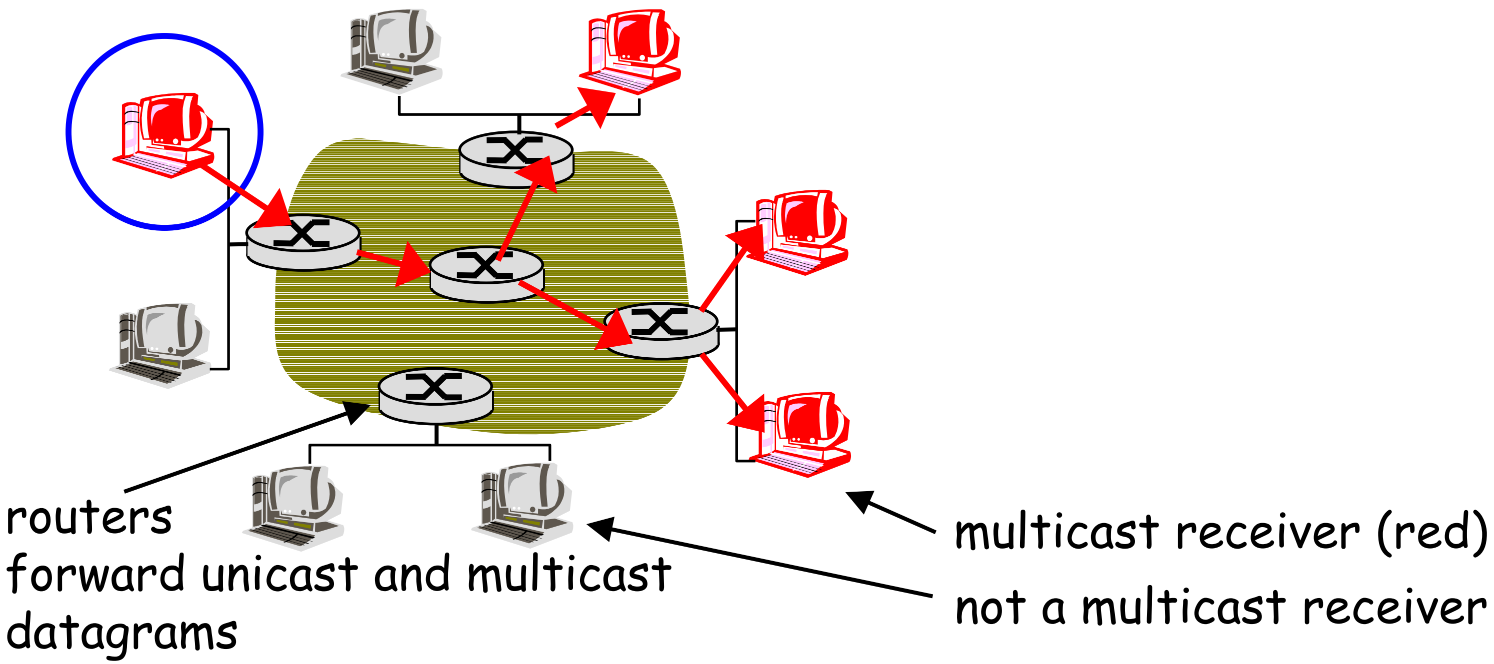

Tunneling

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

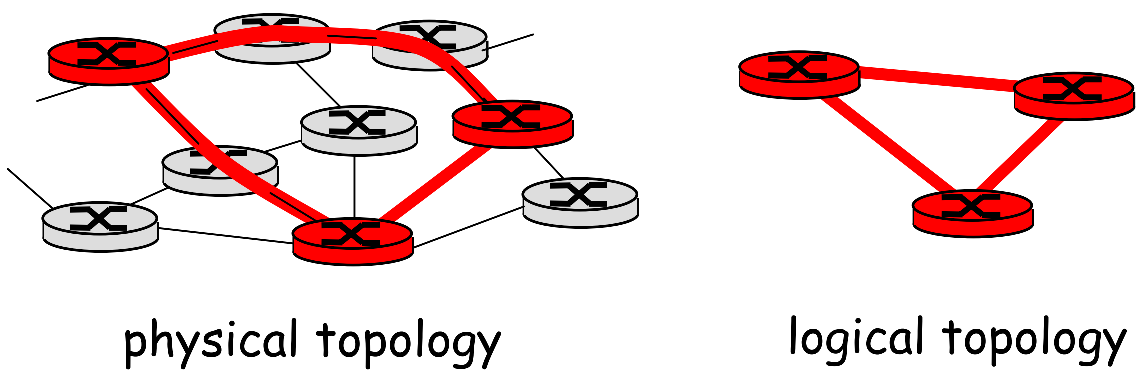

- mcast datagram encapsulated inside "normal" (non-multicast-addressed) datagram
- normal IP datagram sent thru "tunnel" via regular IP unicast to receiving mcast router
- receiving mcast router unencapsulates to get mcast datagram

IP Multicast Architecture

- Based on Steve Deering's original proposal (SIGCOMM 88 Proceedings; his PhD thesis)
- It consists of three basic components:
  - Group addressing based on **globally unique identifiers** (IP multicast addresses)
  - Separation of senders and receivers with **anonymous receiver affiliation**
  - A **tree-based routing and group structure**

IP Multicast Architecture (contd..)

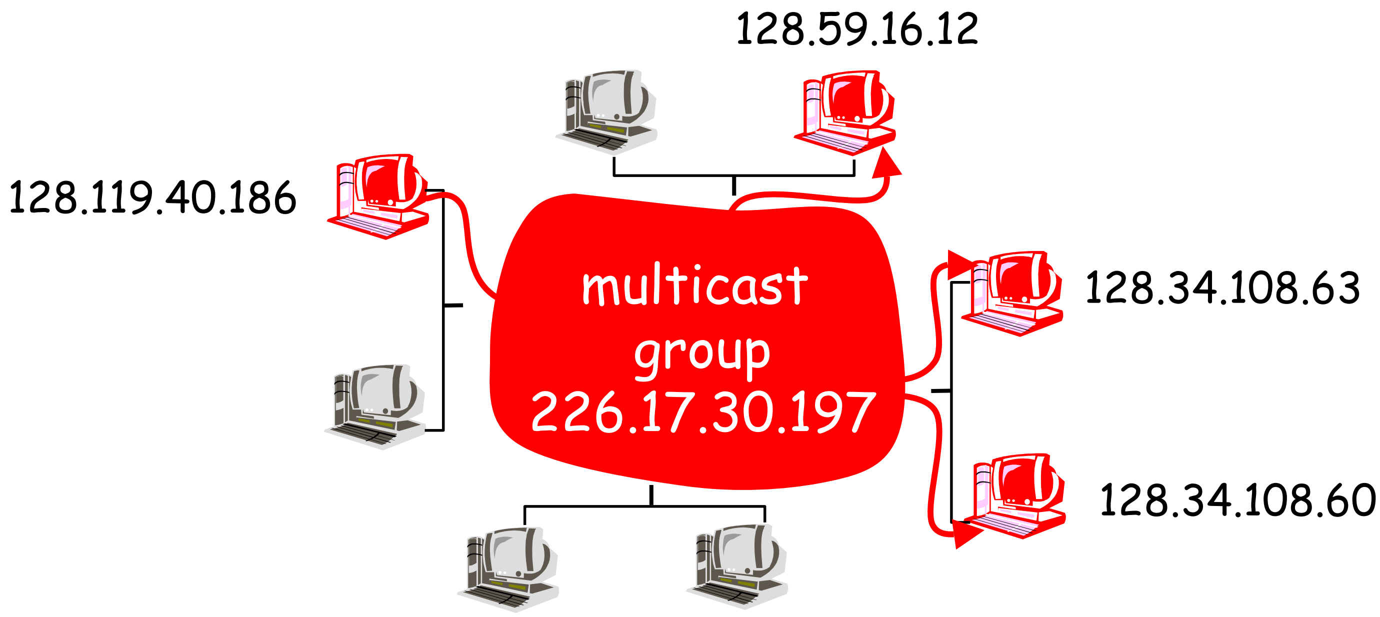

multicast group concept: use of indirection

- hosts addresses IP datagram to multicast group
- routers forward multicast datagrams to hosts that have "joined" that multicast group

IP Multicast Addresses

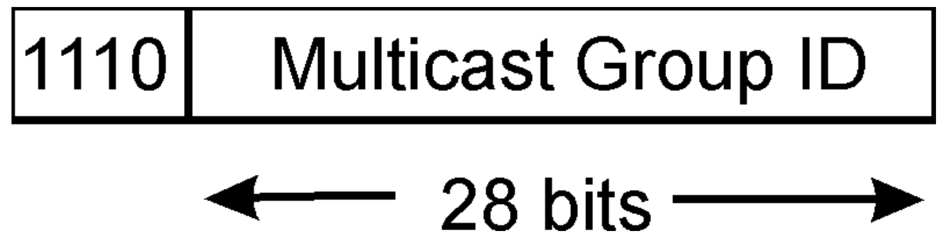

- 28 bits for groups ~ 250 million groups
- **Two kinds of supported addresses:**
  - **Permanent,** e.g.,
    - 224.0.0.1 all systems on LAN
    - 224.0.0.2 all routers on LAN
  - **Temporary - must be created before they can be addressed (join and leave)**

IP Multicast Addressing Problem

- There are plenty of IPv4 and IPv6 addresses for globally unique multicast addresses!
- **The problem is how to assign these addresses permanently or temporarily to multicast groups.**
- **Using globally unique multicast addresses requires Internet-wide coordination in the assignment of identifiers to groups!**
- **Evolving solution: Make multicast identifiers locally unique; examples are now evolving**
  - (e.g., EXPRESS, Simple Multicast use the IP address of a core and a multicast address that the core uses for only one group as the identifier of a group).

IGMP: Internet Group Management Protocol

- *host:* sends IGMP report when application joins mcast group
- *router:* sends IGMP query at regular intervals
  - host belonging to a mcast group must reply to query

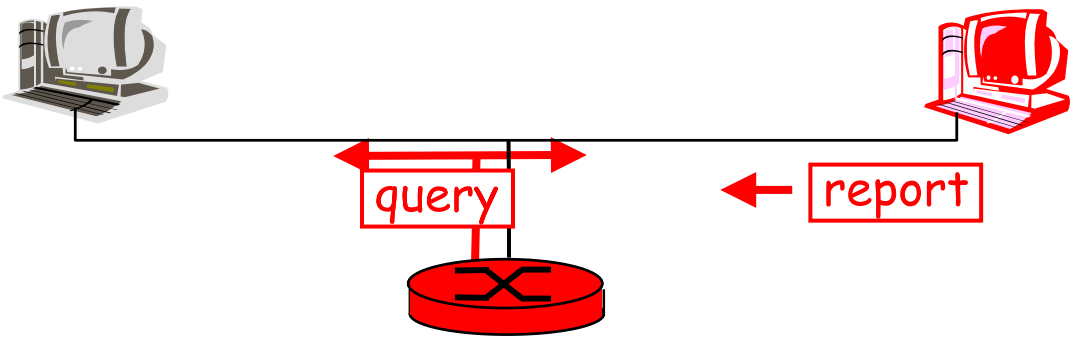

Joining a mcast group: two-step process

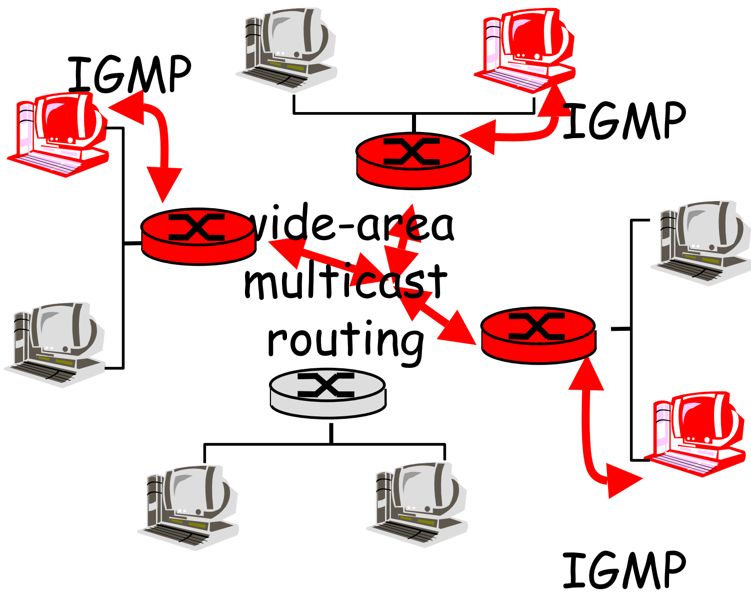

Per Source Per Group Tree

Shortest Path Tree

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

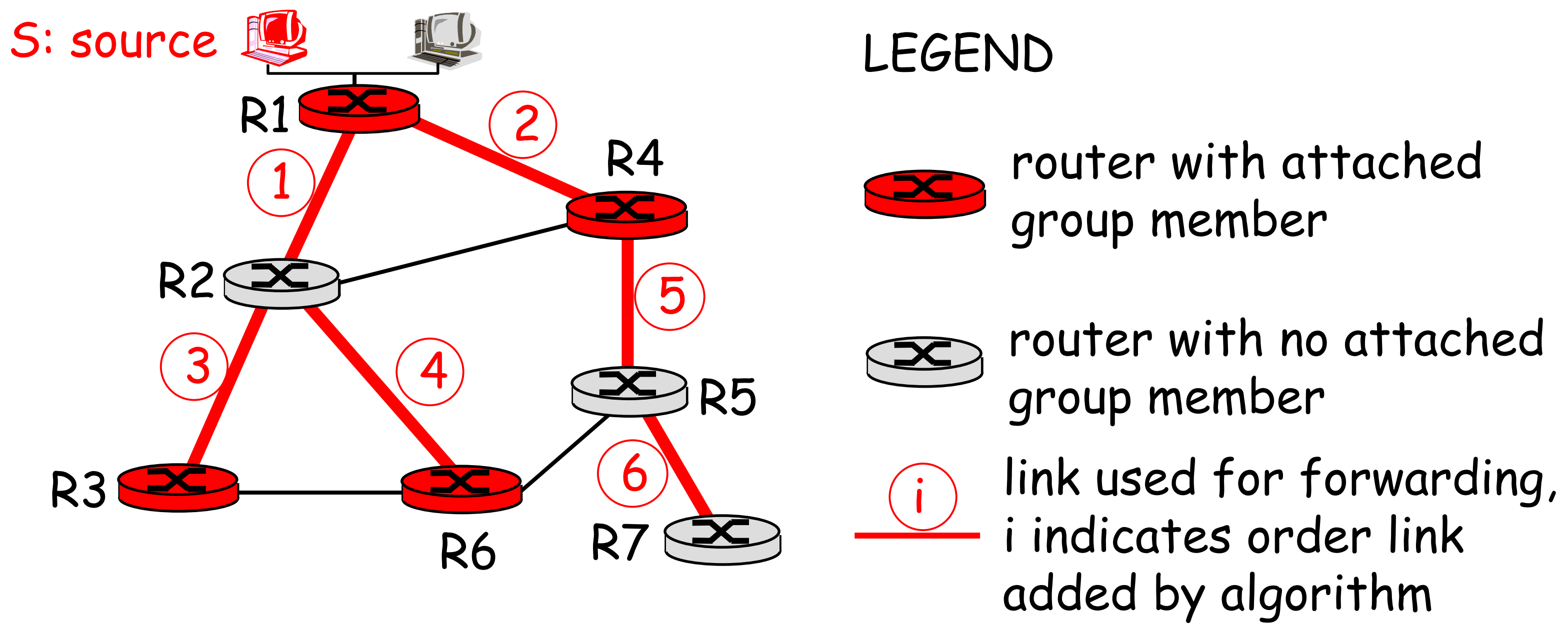

Reverse Path Forwarding

- rely on router's knowledge of unicast shortest path from it to sender
- each router has simple forwarding behavior:

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

then flood datagram onto all outgoing links

else ignore datagram

Reverse Path Forwarding: example

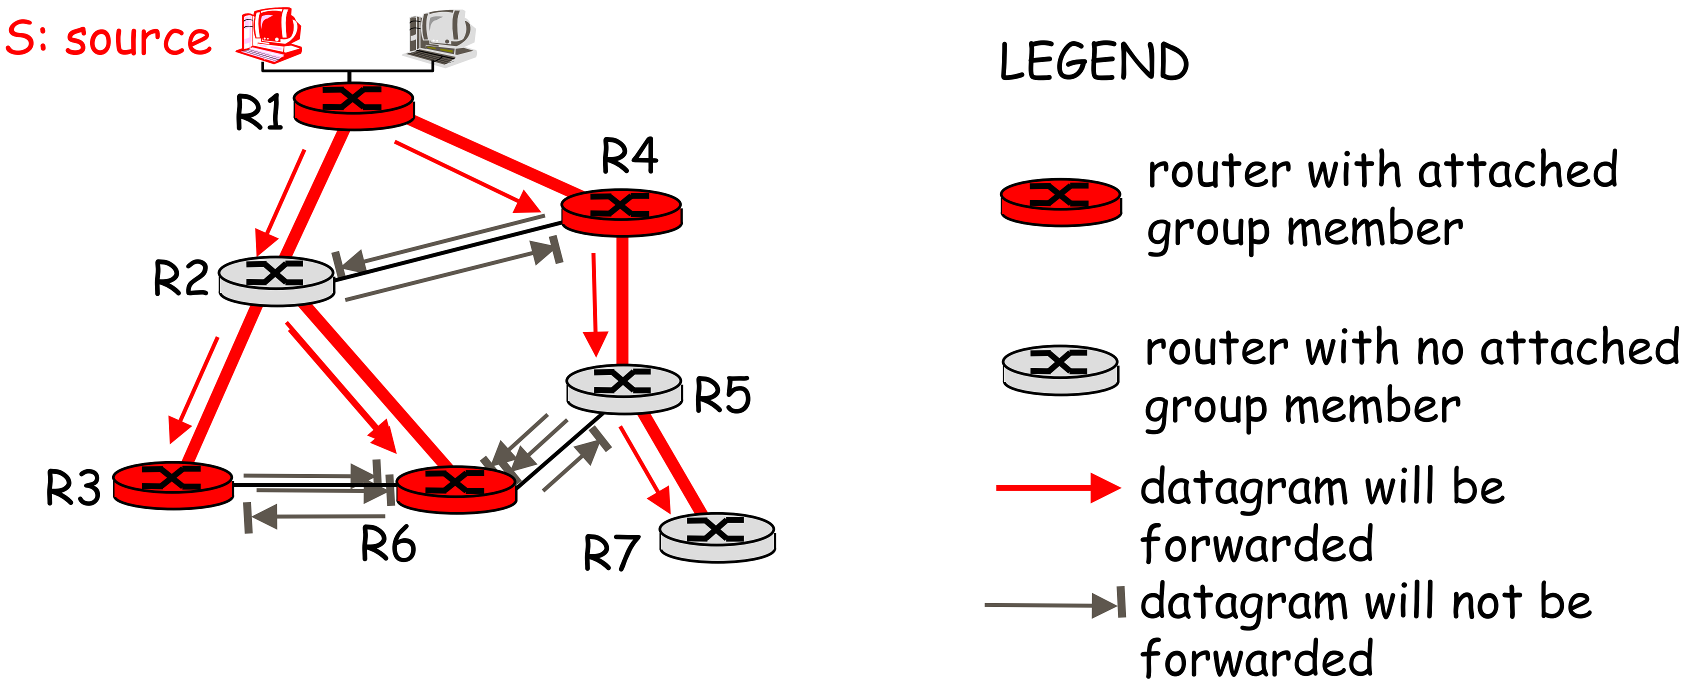

- result is a source-specific *reverse* SPT
  - may be a bad choice with asymmetric links

Reverse Path Forwarding: pruning

- "prune" msgs sent upstream by router with no downstream group members

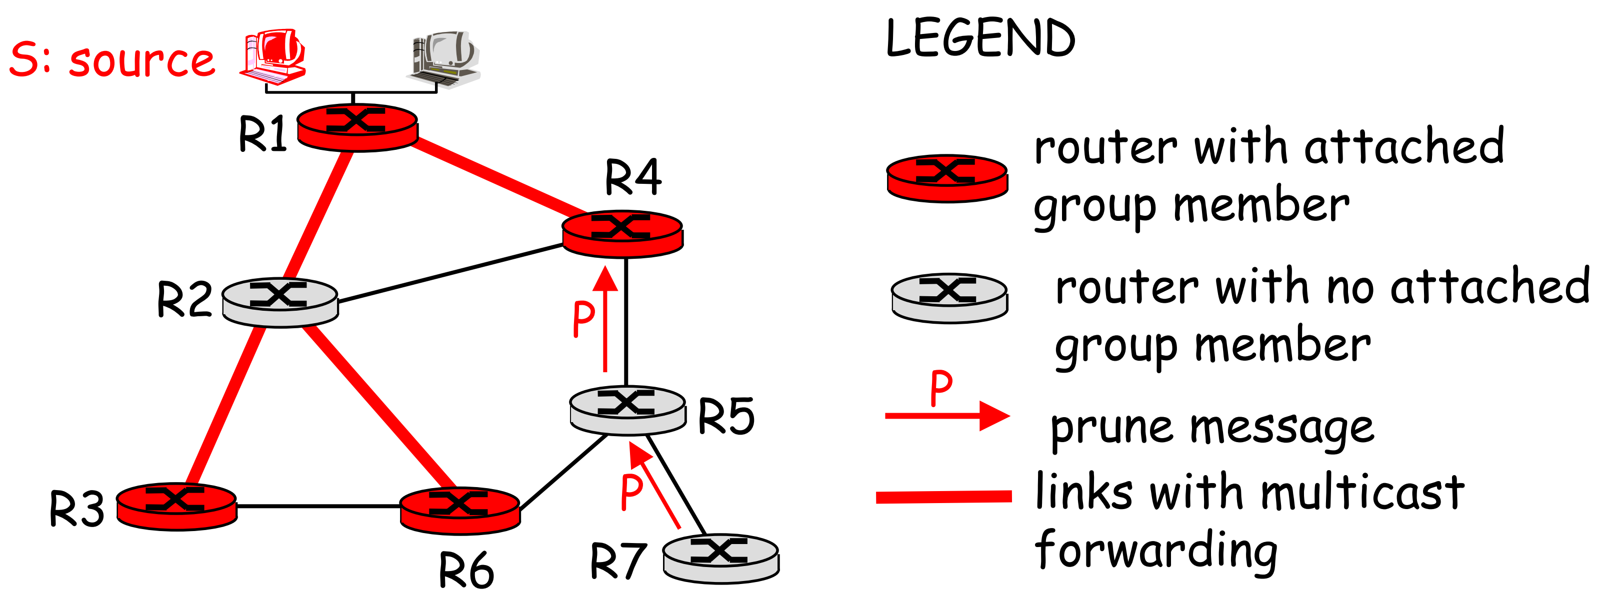

Distance Vector Multicast Routing Protocol (DVMRP)

- 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
- odds and ends
  - commonly implemented in commercial routers
  - Mbone routing done using DVMRP

Problems with Basic Approach

- Flooding the Internet to establish a multicast routing tree is not an option!
- This is the key reason why multicasting (based on DVMRP) cannot be deployed on a large scale.
- The problem stems from the lack of addressing information in an IP multicast address!
- **Where is a multicast group on the Internet?**

- **The solution to this problem is to use special nodes as the "address" of a group**

Per Group Shared Tree

Shared Tree

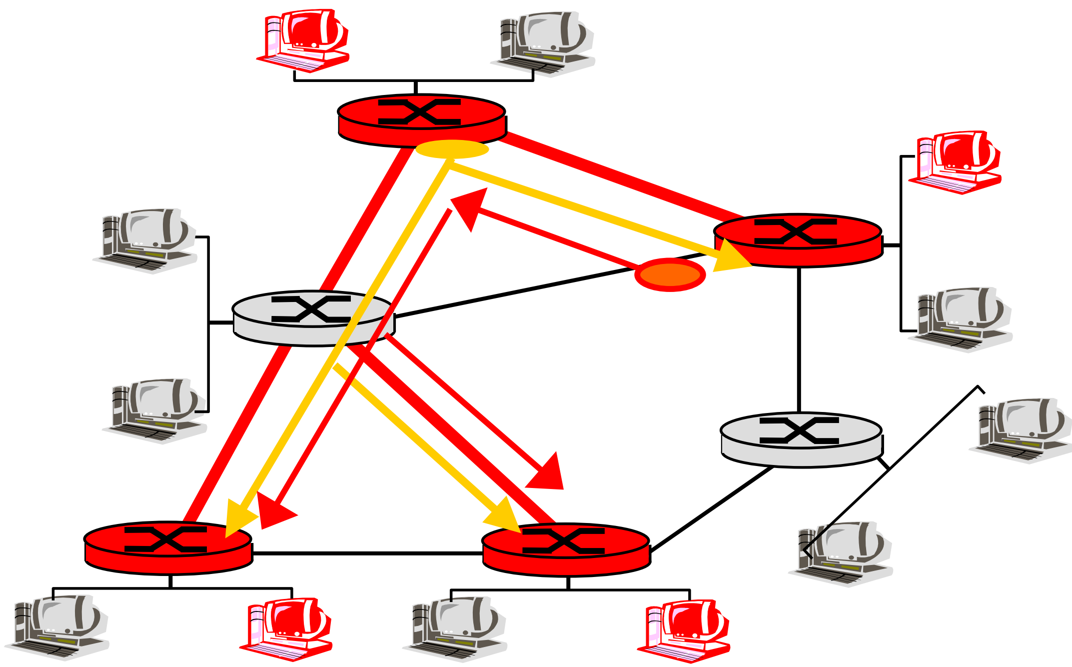

Shared tree

Shared-Tree: Steiner Tree

- Steiner Tree: minimum cost tree connecting all routers with attached group members
- Problem is NP-complete, 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-msg* addressed to center router
  - *join-msg* "processed" by intermediate routers and forwarded towards center
  - *join-msg* either hits existing tree branch for this center, or arrives at center
  - path taken by *join-msg* becomes new branch of tree for this router

Center-based trees: an example

Suppose R6 chosen as center:

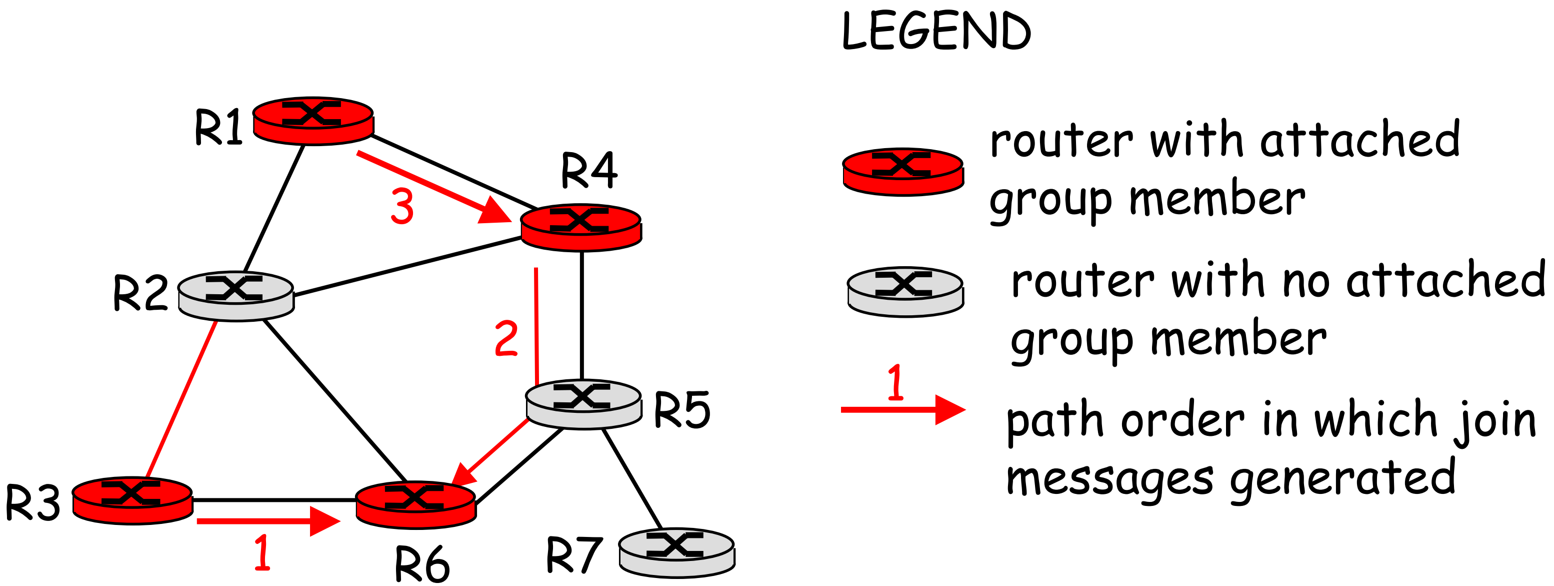

PIM: Protocol Independent Multicast

- 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:

- # networks with group members small wrt # interconnected networks
- group members "widely dispersed"
- bandwidth not plentiful

Consequences of Sparse-Dense Dichotomy:

Dense

- group membership by routers *assumed* until routers explicitly prune
- *data-driven* construction on mcast tree (e.g., RPF)
- bandwidth and non-group-router processing *profligate*

Sparse:

- no membership until routers explicitly join
- *receiver- driven* construction of mcast 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 (core)-based approach
- router sends *join* msg 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

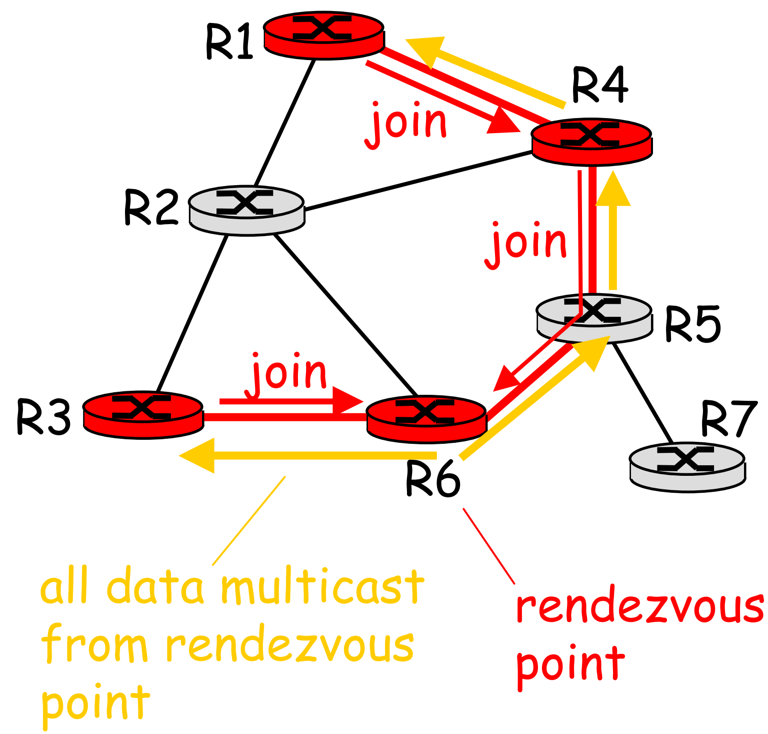

PIM - Sparse Mode

<u>sender(s):</u>

- unicast data to RP, which distributes down RP-rooted tree
- RP can extend mcast tree upstream to source
- RP can send *stop* msg if no attached receivers
  - "no one is listening!"