Multicast

- **Multicast = Sending messages to multiple receivers.**
- **Unicast = Sending messages to a single receiver.**
- **Broadcast = Sending messages to ALL receivers in a network. (sometimes also called *flooding,* if the network is multihop).**

Layering of Multicast

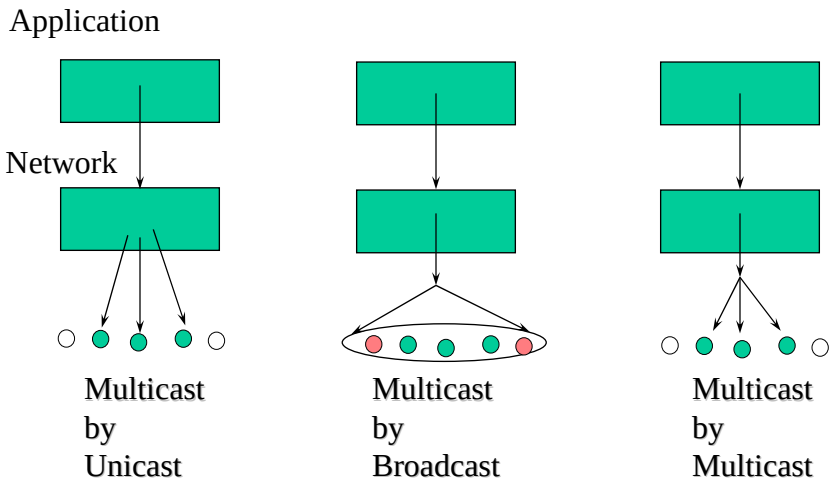

Example

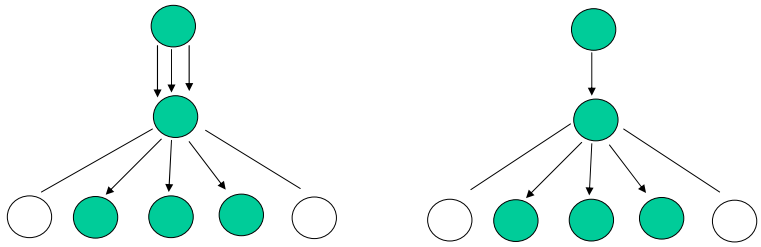

Multiple unicast **Multicast**

Applications of Multicast

- **Tele/Video conferencing.**
- **Distributed interactive games, whiteboards, content delivery (push applications).**
- **Video/media distribution.**
- **Various protocol needs,**
  - Replicated database updates.
  - Router updates, etc.
- **Major application models**
  - Point-to-multipoint (single source, multiple receivers).
  - Multipoint-to multipoint (multiple sources, multi receivers).
  - Source(s) may or may not be receiver(s).
- **Note, many hosts and most routers on current Internet are not multicast capable!!**

Multicast Using Multiple Unicast

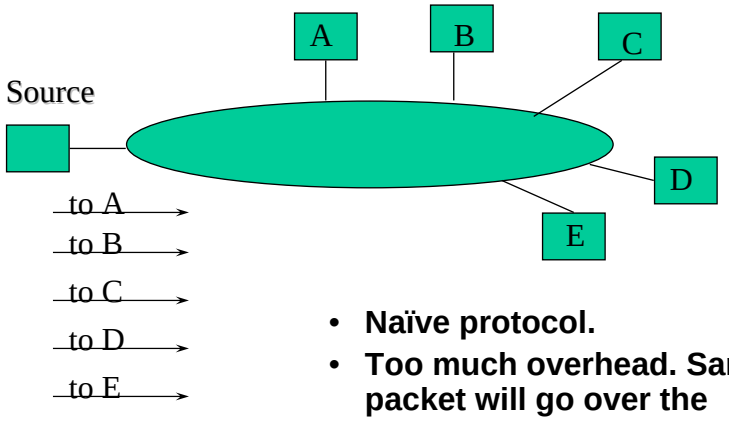

- **Naïve protocol.**
- **Too much overhead. Same packet will go over the same link many times for different destinations.**

Mulicast Using Multi-destination Addressing

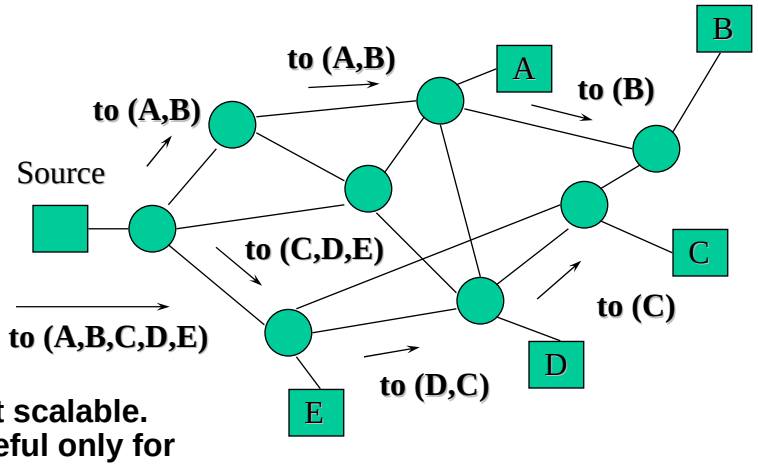

- **Not scalable. Useful only for small no. of receivers.**

Multicast Using Multicast

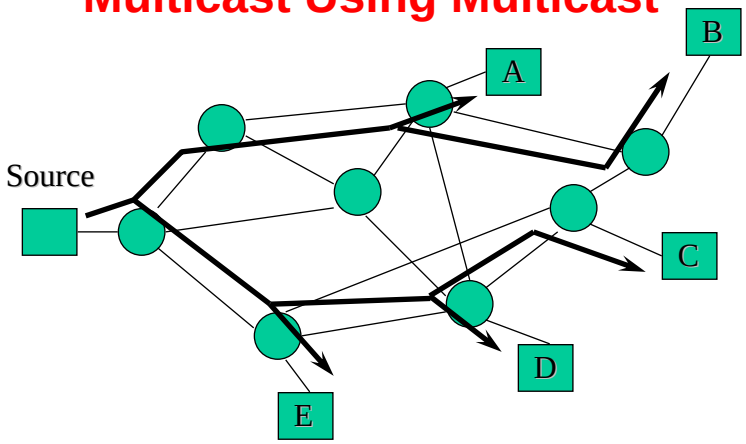

- **Forward via the shortest-path tree rooted at source. (Source-based tree protocol).**

Theoretical Basis

- **Ideally, use a Steiner Tree.**
- **Minimum cost Steiner Tree problem is NP complete.**
  - Graph G = (V, E)
  - Positive edge weights W(e)
  - R a subset of V. R is the set of receivers.
  - Find a minimum-cost subtree of G that includes all of R.
  - Polynomial time heuristic exists. Less than twice optimum. But hard to use in a distributed fashion. Not scalable. Non-incremental changes.
  - Shortest path tree – easy approach. But bound unknown! Incremental changes.

Design Choices for Multicast Routing

- **Addressing – how the mulitcast destinations (receivers) are to be indicated?**
  - List approach – not scalable
  - Group approach – no control; who maintains the group??
- **Routing protocol**
  - Re-use the unicast routing infrastructure?
  - Desirable. But is it too constraining?

Multicast on the Internet (IP Multicast)

- **Concept of multicast group**
  - Set of nodes wishing to listen to a particular “connection.”
  - Sender(s) need not be group member(s). Sender does not know group members.
  - Each group has a “group address.”
  - Nodes may join/leave group at will. Group size unrestricted.
- **Multicast addressing**
  - Packets sent to a group address, not to individual node address.
  - Group address is a valid IP address.
  - A router will multicast the packet on a LAN (link-layer multicast) if there is a host on the LAN that is a member of the destination group.

Multicast Addressing

- **Class D IP address:**
  - 224.0.0.0 thru' 239.255.255.255. High order bits are 1110.
  - Flat address space. No subnetting.
- **Link-layer multicast over IEEE 802 style LANs**
  - E.g., Ethernet, wireless LANs
  - Lower 23 bits of mcast IP address is copied to the lower 23 bits of MAC address and a specific mcast bit is set.
  - Node picks up the packet if there is an address match with the lower 23 bits for a mcast group the node is listening to.

Group Membership Protocol and Link-layer Multicast: Big Picture

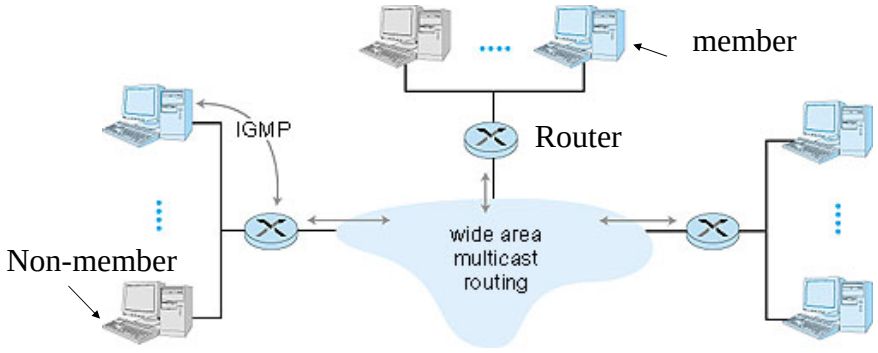

- **Local router knows that a host on the LAN is a group member via IGMP (Internet Group Management Protocol).**
- **Local router forward the packet to host(s) on LAN via link-layer multicast.**

IGMP (Internet Group Management Protocol)

- **Used by a host to declare membership to a multicast group to the nearest router.**

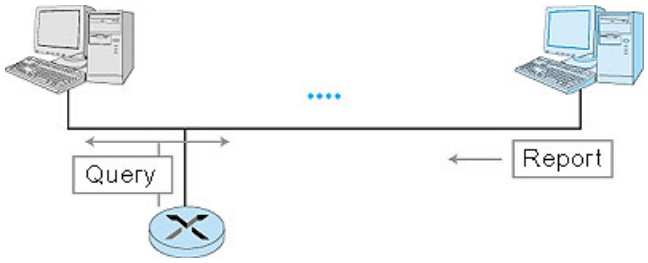

- **Router sends IGMP query. Any member host responds with a IGMP report. Randomized delay for response to discourage reply storm.**

IGMP (contd.)

- **All the router cares about if there is <u>any host at all</u> on the LAN that is a member of a specific group**
  - Feedback suppression: If a host hears a report for the same group by another host, it suppresses its own report.
- **IGMP v1: Periodic querying by router. Query goes on a reserved mcast address. Explicit Join messages. But no explicit leave.**
- **IGMP v2:**
  - Adds group-specific queries.
  - Adds explicit group leave messages by hosts. Lower "leave latency."
- **Note, nobody explicitly keeps track all hosts on the Internet that constitute a group.**

IP Multicast Components

- **Group membership protocol**
  - Router “knows” whether an attached LAN has a member host for a group.
  - Then it multicasts the packets addressed to that group on the LAN.
- **Link-layer multicast**
  - Hosts listening to that address pick up the packets.
- **Network-layer (IP layer) multicast**
  - Source must somehow route the packet to each router that have an attached LAN with at least one host listening to that mcast address.

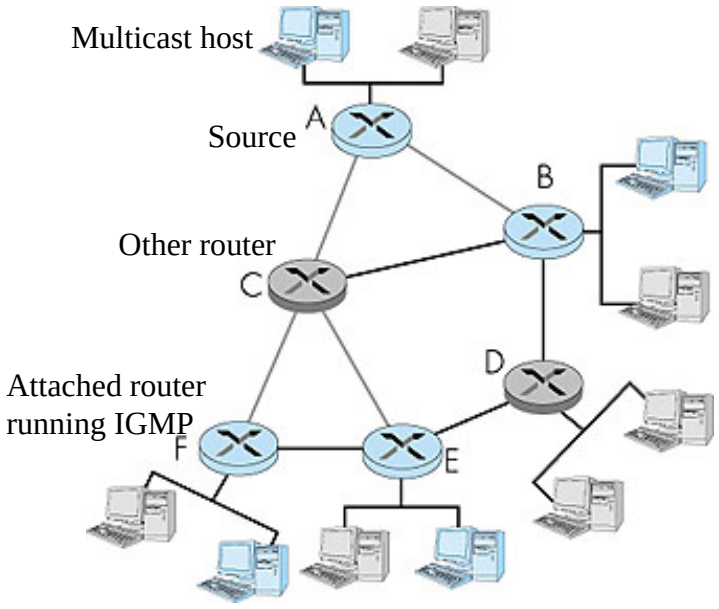

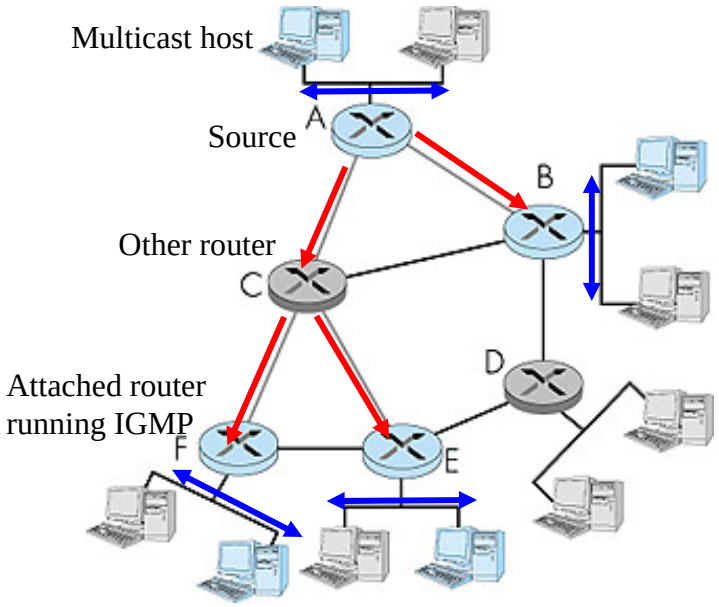

Basic Multicast Routing: Reverse Path Forwarding

- **One of the original mcast routing protocols on the Internet. Basis of MBone.**
- **Packet arrives on a link L from Source. Check if L is the link the router would use to send packets to Source.**
  - If true, Transmit packet on all links except L.
  - If false, Discard packet.
- **Why this rule?**
  - Loop freedom.
  - Routing via Shortest Path Tree from destinations to source (reverse tree).

Reverse Path Forwarding

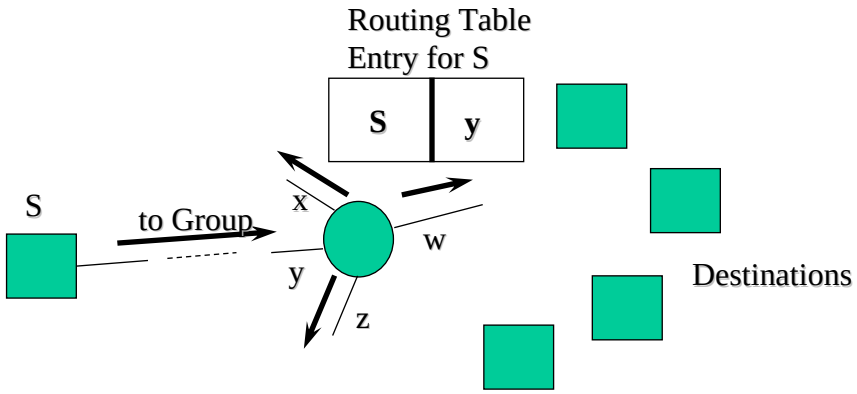

Reverse Path Forwarding

Broadcast via reverse path forwarding.
Need to prune tree for multicast.

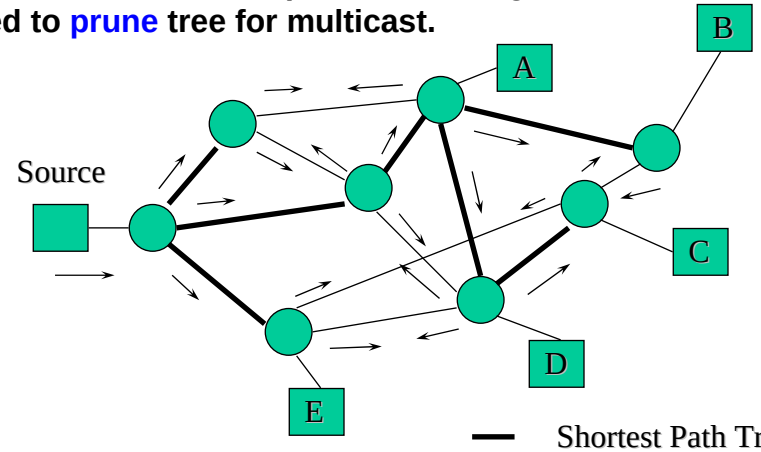

Flood + Prune + Graft = Multicast Forwarding

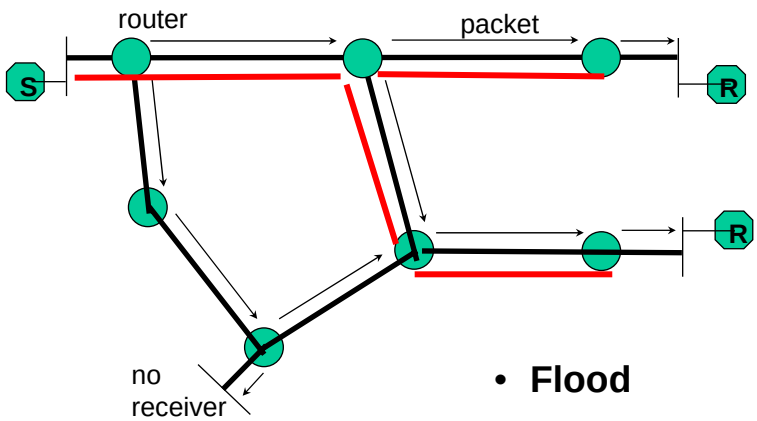

- **Flood**

Flood + Prune + Graft = Multicast Forwarding

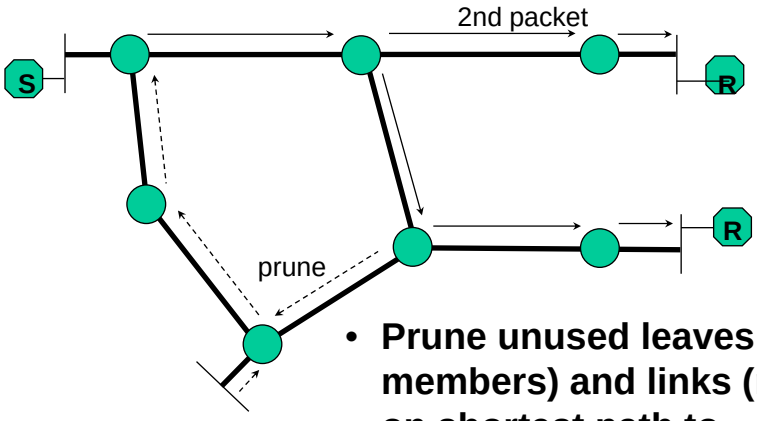

- **Prune unused leaves (no members) and links (not on shortest path to source).**

Flood + Prune + Graft = Multicast Forwarding

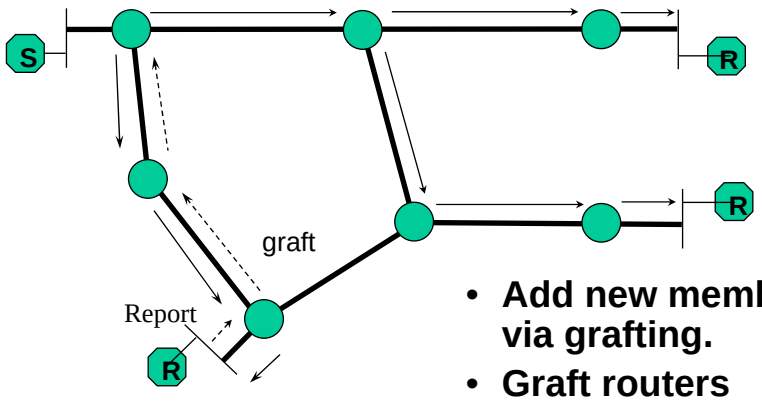

- **Add new members via grafting.**
- **Graft routers towards shortest path to source.**

DVMRP: Distance Vector Multicast Routing Protocol

- **Protocol based on reverse path tree based multicast forwarding.**
- **Uses a distance vector protocol similar to RIP**
  - Computes the next hop on the shortest path to every source.
  - Floods the first packet for a <source,group> pair in the entire internetwork based on TTL.
  - All edge routers get this packet.
  - Prune the broadcast tree as described before.

Mbone: Internet Multicast Backbone

- **Today, most routers on Internet are not multicast capable.**
- **Mbone = Virtual network overlaying Internet**
  - All nodes on the overlay network are multicast capable. Use DVMRP.
  - Talk to neighboring Mbone routers using IP-in-IP encapsulation (tunneling).

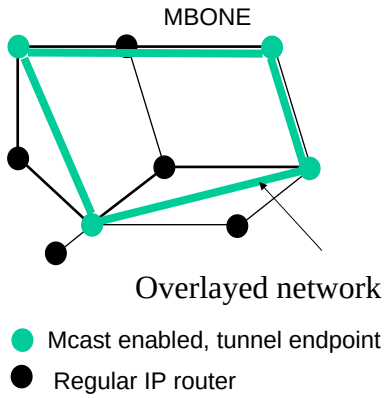

IP-in-IP Encapsulation (Tunneling)

- Put a whole IP packet with a multicast destination inside a “regular” IP packet with a unicast destination
- The unicast destination is the next hop Mbone router (tunnel endpoint).
- A specific header in the “wrapper” packet indicates that it is using tunneling (protocol type = 4).
- Tunnel endpoint decapsulates the packet for processing

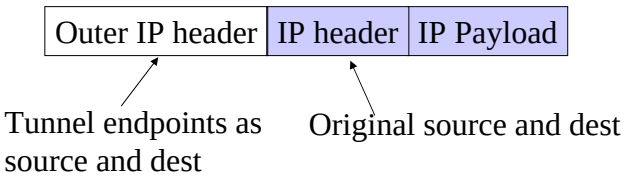

DVMRP Performance

- **Tree is source-based. One tree per <source,group>.**
- **Uses broadcast and prune to build tree. Sometimes called implicit join.**
  - Works better for dense mode (DM) multicast meaning a large group size compared to network size.
- **Must broadcast+prune periodically to accommodate link cost changes.**
- **Must maintain per <source, group> state on ALL routers regardless of their own membership**
  - Non-tree routers may still need to "graft" somebody in future.
  - They need to know which direction to forward graft messages.
- **=> DVMRP not good for**
  - Scattered group membership.
  - Large no, of sources / groups.