CS 640: Introduction to Computer Networks

Aditya Akella

Lecture 12 - Multicast

Multicast

- Unicast: one source to one destination
  - Web, telnet, FTP, ssh
- Broadcast: one source to all destinations
  - Never used over the Internet
  - LAN applications
- *Multicast*: one source to many destinations
  - Several important applications
- Multicast goal: efficient data distribution

2

Multicast – Efficient Data Distribution

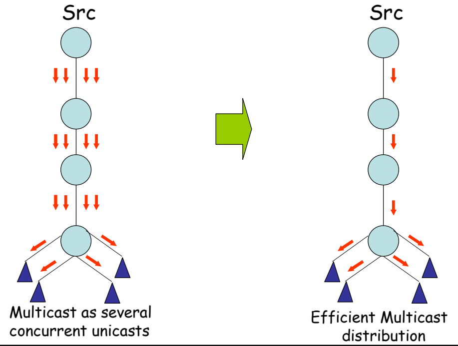

3

Multicast Example Applications

- Broadcast audio/video
- Push-based systems
- Software distribution
- Teleconferencing (audio, video, shared whiteboard, text editor)
- Multi-player games
- Server/service location
- Other distributed applications

4

IP Multicast Architecture

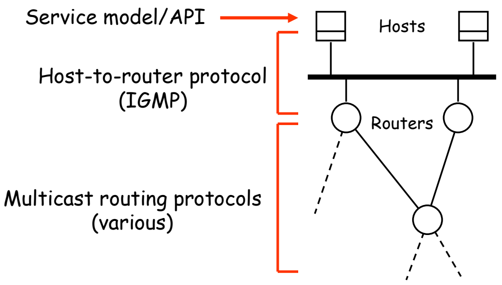

5

IP Multicast Service Model (rfc1112)

- Each group identified by a single IP address
- Groups may be of any size
- Members of groups may be located anywhere in the Internet
  - We will focus on an internetwork
- Members of groups can join and leave at will
- Senders need not be members
- Group membership not known explicitly

6

IP Multicast Addresses

- Class D IP addresses
  - 224.0.0.0 - 239.255.255.255

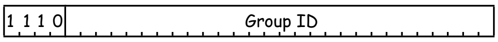

- How to allocate these addresses?
  - Well-known multicast addresses, assigned by IANA
  - Transient multicast addresses, assigned and reclaimed dynamically
    - e.g., by "sdr" program
- Interested recipients must *join* a group by selecting the appropriate multicast group address

7

IP Multicast Architecture

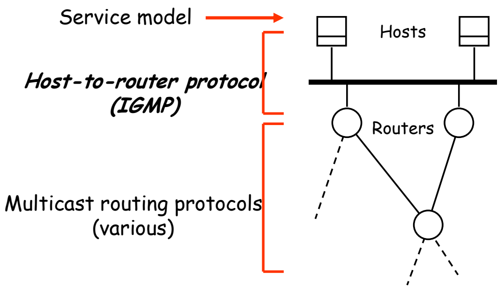

8

Internet Group Management Protocol

- End system to router protocol is IGMP
- Each host keeps track of which mcast groups it has subscribed to
  - Socket API informs IGMP process of all joins
- Objective is to keep router up-to-date with group membership of entire LAN
  - Routers need not know who all the members are, only that *members exist*

9

How IGMP Works

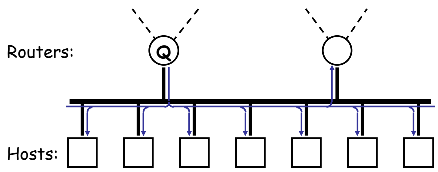

- On each link, one router is elected the "querier"
- Querier periodically sends a Membership Query message to the all-systems group (224.0.0.1), with TTL = 1
- On receipt, hosts start random timers (between 0 and 10 seconds) for each multicast group to which they belong

10

How IGMP Works (cont.)

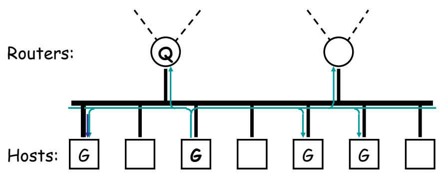

- When a host's timer for group G expires, it sends a Membership Report <u>to group G</u>, with TTL = 1
- Other members of G hear the report and stop their timers
- Routers hear <u>all</u> reports, and time out non-responding groups
  - "Soft state" again

11

How IGMP Works (cont.)

- Note that, in normal case, only one report message per group present is sent in response to a query
- Query interval is typically 60-90 seconds
- When a host first joins a group, it sends one or two immediate reports, instead of waiting for a query

12

IP Multicast Architecture

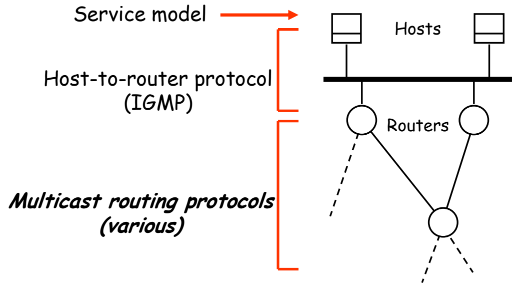

13

Routing Techniques

- Basic objective – routers must collectively build distribution tree for multicast packets
- Flood and prune based approach for DV-networks
  - Begin by flooding traffic to entire network
  - Prune branches with no receivers
  - Examples: DVMRP
- Link-state based networks use a different approach
  - Routers advertise groups for which they have receivers to entire network
  - Compute trees on demand
  - Example: MOSPF
- There are several others: PIM-SM, PIM-DM, CBT…
  - These are "rendezvous-based" approaches
  - Independent of underlying routing protocol

14

MOSPF: Example

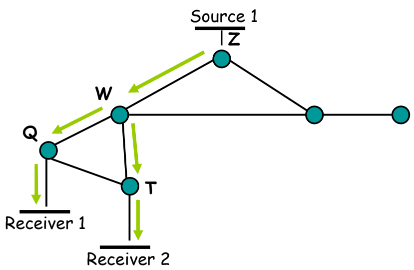

15

Link Failure/Topology Change

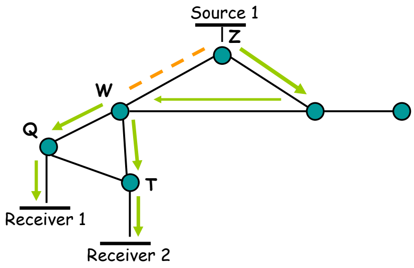

16

Impact on Route Computation

- Hard to pre-compute multicast trees for all possible sources and all possible groups
  - Otherwise, may end up with a lot of unwanted state where there are no senders
- Compute on demand when first packet from a source S to a group G arrives
- New link-state advertisement
  - May lead to addition or deletion of outgoing interfaces if it contains different group addresses
  - May lead to re-computation of entire tree if links are changed

17

Distance-Vector Multicast Routing

- DVMRP consists of two major components:
  - A conventional distance-vector routing protocol (like RIP)
  - A protocol for determining how to forward multicast packets, based on the routing table
- DVMRP router forwards a packet if
  - The packet arrived from the link used to reach the source of the packet (reverse path forwarding check – RPF)
  - If downstream links have *not pruned* the tree

18

Example Topology

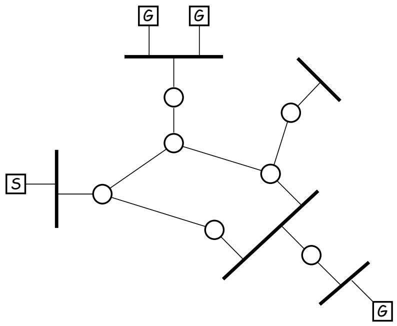

19

Broadcast with Truncation

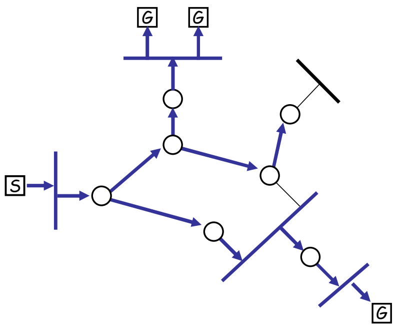

20

Prune

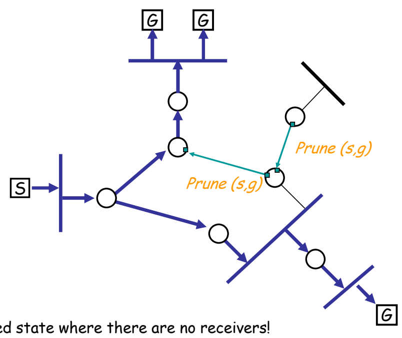

Unwanted state where there are no receivers!

21

Graft

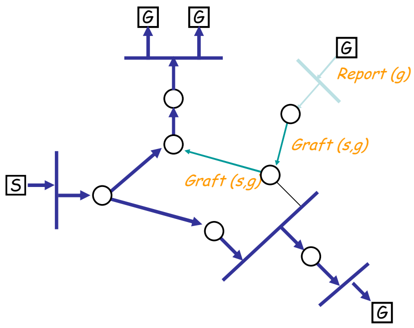

Source-based Trees

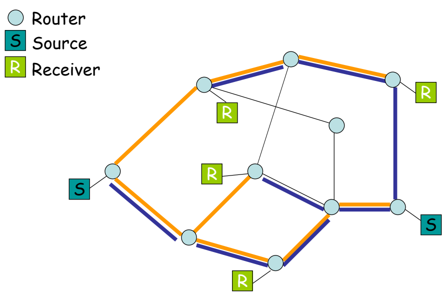

Both protocols discussed today use this approach

Shared Tree

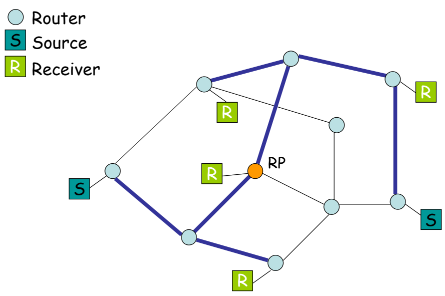

Shared vs. Source-Based Trees

- Source-based trees
  - Shortest path trees – low delay, better load distribution
  - More state at routers (per-source state)
  - Efficient for *dense-area multicast*
- Shared trees
  - Higher delay (bounded by factor of 2), traffic concentration
  - Choice of core affects efficiency
  - Per-group state at routers
  - Efficient for *sparse-area multicast:* PIM-SM

25