Lecture 10 (Feb 19, 2004)

Outline

- Network-layer Multicast

One to many communication

- Application level one to many communication
- multiple unicasts

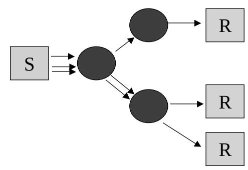

- IP multicast

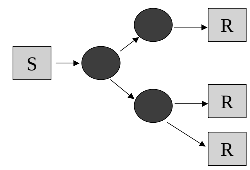

Types of Multicast

- At network-layer
  - Topic of this lecture
- Sequence of unicasts
  - Separate streams of unicast traffic for each destination from the source
  - Does not require support at network-layer
- Application-layer multicast
  - Based on unicasts
  - Constructs an overlay structure
  - Source unicasts to a subset of receives, these receivers unicast to another subset, which unicast to another subset and so on to reach the whole multicast group

Why Multicast

- When sending same data to multiple receivers
  - better bandwidth utilization
  - less host/router processing
  - quicker participation
- Application
  - Video/Audio broadcast (One sender)
  - Video conferencing (Many senders)
  - Real time news distribution
  - Interactive gaming

IP multicast service model

- Invented by Steve Deering (PhD. 1991)
  - It's a different way of routing datagrams
- RFC1112 : Host Extensions for IP Multicasting - 1989
- Senders transmit IP datagrams to a "host group"
- "Host group" identified by a class D IP address
- Members of host group could be present anywhere in the Internet
- Members join and leave the group and indicate this to the routers
- Senders and receivers are distinct: i.e., a sender need not be a member
- Routers listen to all multicast addresses and use multicast routing protocols to manage groups

IP multicast group address

- Things are a little tricky in multicast since receivers can be *anywhere*
- Class D address space
  - high-order three 3bits are set
  - 224.0.0.0 ~ 239.255.255.255
- Allocation is essentially random – any class D can be used
  - Nothing prevents an app from sending to any multicast address
  - Customers end hosts and ISPs are the ones who suffer
- Some well-known address have been designated
  - RFC1700
  - 224.0.0.0 ~ 224.0.0.25
- Standard are evolving

Getting Packets to End Hosts

- Packets from remote sources will only be forwarded by IP routers onto a local network only if they know there is at least one recipient for that group on that network
- Internet Group Management Protocol (IGMP, RFC2236)
  - Used by end hosts to signal that they want to join a specific multicast group
  - Used by *routers* to discover what groups have have interested member hosts on each network to which they are attached.
  - Implemented directly over IP

IGMP – Joining a group

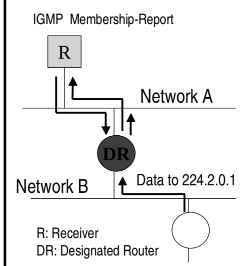

Example : R joins to Group 224.2.0.1

- R sends **IGMP Membership-Report to 224.2.0.1**
- DR receives it. DR will start forwarding **packets for 224.2.0.1** to Network A
- DR periodically sends **IGMP Membership-Query to 224.0.0.1** (ALL-SYSTEMS.MCAST.NET)
- R answers **IGMP Membership-Report to 224.2.0.1**

IGMP – Leaving a group

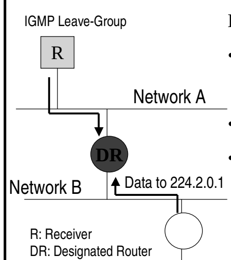

Example : R leaves from a Group 224.2.0.1

- R sends **IGMP Leave-Group**
- DR receives it.
- DR stops forwarding **packets for 224.2.0.1** to Network A if no more 224.2.0.1 group members on Network A.

Challenges in the multicast model

- How can a sender restrict who can receive?
  - need authentication, authorization
  - encryption of data
  - key distribution
  - still an active area of research

IP multicast routing

- Purpose: share Group information among routers, to implement better routing for data distribution
- Distribution tree structure
  - Source tree vs shared tree
- Data distribution policy
  - Opt in (ACK) type vs opt out (NACK) type
- Routing protocols are used in conjunction with IGMP

Source distribution tree

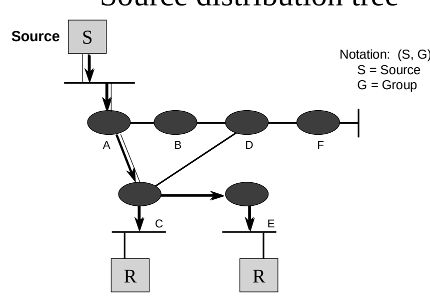

Shared distribution tree

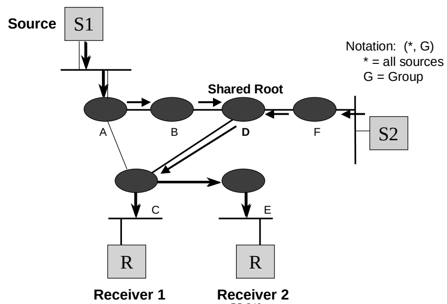

Source tree characteristics

- Source tree
  - More memory O (G x S) in routers
  - optimal path from source to receiver, minimizes delay
- good for
  - small number of senders, many receivers such as Radio broadcasting application

Shared tree characteristics

- Shared tree
  - Less memory O (G) in routers
  - Sub-optimal path from source to receiver, may introduce extra delay (source to root)
  - May have duplicate data transfer (possible duplication of a path from source to root and a path from root to receivers)
- good for
  - Environments where most of the shared tree is the same as the source tree
  - Many senders with low bandwidth (e.g. shared whiteboard)

Data distribution policy

- Opt out (NACK) type
  - Start with “broadcasting” then prune brunches with no receivers, to create a distribution tree
  - Lots of wasted traffic when there are only a few receivers and they are spread over wide area
- Opt in (ACK) type
  - Forward only to the hosts which explicitly joined to the group
  - Latency of join propagation

Protocol types

- Dense mode protocols
  - assumes dense group membership
  - Source distribution tree and NACK type
  - **DVMRP** (Distance Vector Multicast Routing Protocol)
  - **PIM-DM** (Protocol Independent Multicast, Dense Mode)
  - Example: Company-wide announcement
- Sparse mode protocol
  - assumes sparse group membership
  - Shared distribution tree and ACK type
  - **PIM-SM** (Protocol Independent Multicast, Sparse Mode)
  - Examples: Futurama or a Shuttle Launch

DVMRP

exchange distance vectors

- Each router maintains a ‘multicast routing table’ by exchanging distance vector information among routers
  - First multicast routing protocol ever deployed in the Internet
    - Similar to RIP
  - Constructs a source tree for each group using reverse path forwarding
    - Tree provides a shortest path between source and each receiver
- There is a “designated forwarder” in each subnet
  - Multiple routers on the same LAN select designated forwarder by lower metric or lower IP address (discover when exchanging metric info.)
- Once tree is created, it is used to forward messages from source to receivers
- If all routers in the network do not support DVMRP then unicast tunnels are used to connect multicast enabled networks

DVMRP
broadcast & prune

- Flood multicast packets based on RPF (Reverse path forwarding) rule to all routers.
- Leaf routers check and sends prune message to upstream router when no group member is on their network
- Upstream router prune the interface with no dependent downstream router.
- *Graft* message to create a new branch for late participants
- Restart forwarding after prune lifetime (standard : 720 minutes)
- draft-ietf-idmr-dvmrp-v3-09.txt (September 1999)

RPF(reverse path forwarding)

- Simple algorithm developed to avoid duplicate packets on multi-access links
- RPF algorithm takes advantage of the IP routing table to compute a multicast tree for each source.
- RPF check
  1. When a multicast packet is received, note its source (*S*) and interface (*I*)
  2. If *I* belongs to the shortest path from *S*, forward to all interfaces except *I*
  3. If test in step 2 is false, drop the packet
- Packet is never forwarded back out the RPF interface!

DVMRP (1)
form a source tree by exchanging metric

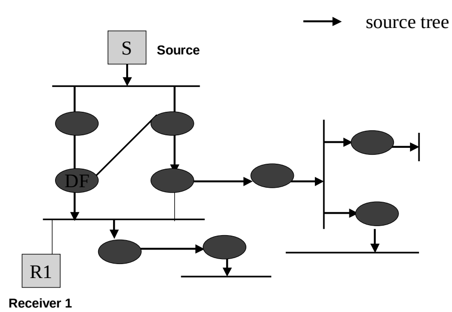

DVMRP (2)
broadcast

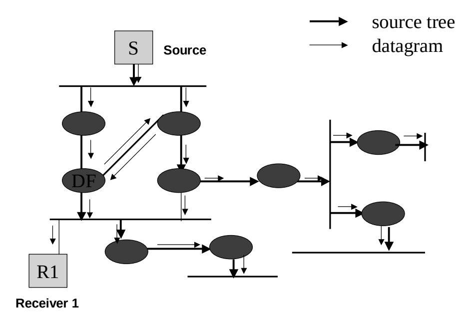

DVMRP (3)
prune

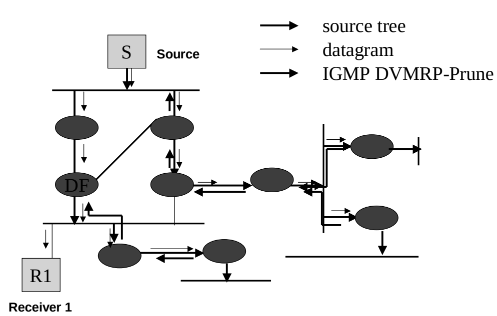

DVMRP (4)
X and Y pruned

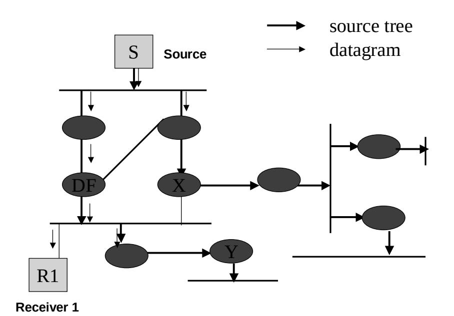

DVMRP (4)

New member

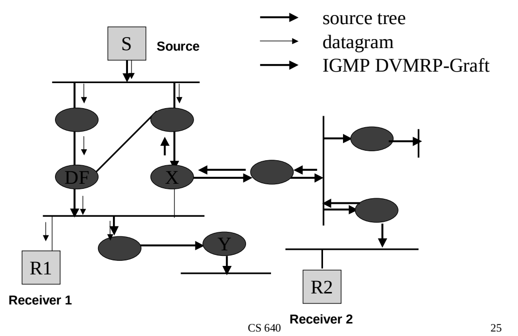

DVMRP (4)

New branch