Multicast

FEUP
MPR

Overview

- Multicast
  Service model, concept, examples
- Multicast addresses
  IPv4, IPv6, L2, mapping
- Group Management
  IGMP v1, v2, v3, MLD
  Switch broadcast and multicast, IGMP snooping, CGMP
- Multicast routing
  Source and shared trees
  Multicast forwarding
  PIM-SM, PIM-DM
  MBGP, MSDP

Multicast

Service model, concept, examples

Multicast Service Models

- Principle
  - » Same data reaches multiple receivers
  - » Source avoids transmitting it for each receiver
- Implemented at

  link, network, application layer
- xCast
  - » Broadcast

    To all nodes on a (small) network
  - » Multicast

    To a group
  - » Anycast

    Packet to **1**of **n** hosts

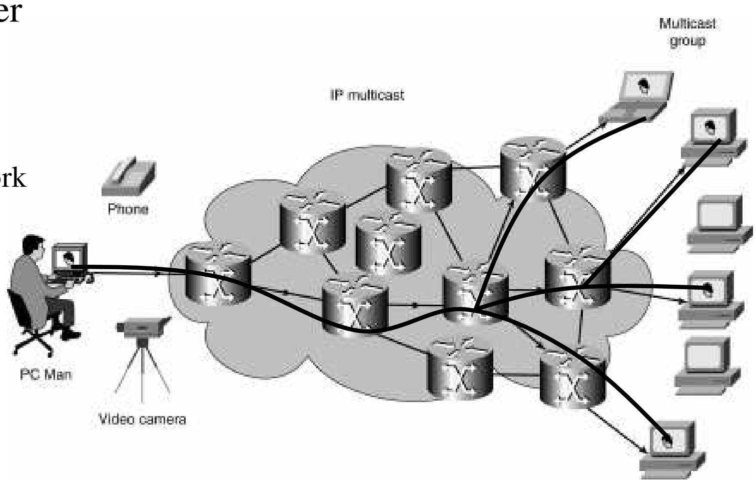

Multicast Group Concept

- Group of receivers express
  - » interest in receiving a data stream
- The group
  - » does not have physical / geographical boundaries
  - » members can be located anywhere on the Internet
- An host interested in receiving the data
  - » must join the group
  - » then, it receives the data stream

Examples of Multicast Applications

- Audio-video distribution

  1 N

- Audio-video conferences

  N N

- File distribution

  stock market quotes, software releases

Multicast addresses

IPv4, IPv6, L2, mapping

IPv4 Multicast Addresses

- Model
  - network level: data packets remain the same; only address changes
  - routers required to replicate and route packets
- Multicast address arbitrary group of IP hosts that
  - have joined the group
  - want to receive traffic sent to this group
- IP Multicast group addresses old class D
  - range 224.0.0.0 - 239.255.255.255
  - used only as destination addresses of IP packets
  - source address is always the unicast source address

Reserved IPv4 Multicast Addresses

- Reserved Link Local Addresses, 224.0.0.0 - 224.0.0.255
  - » used by network protocols on a local network segment
  - » packets transmitted with time-to-live (TTL) of 1

    224.0.0.1 - all systems on this subnet

    224.0.0.2 - all routers on this subnet

    224.0.0.5 - OSPF routers

    224.0.0.12- DHCP server/relay agent
- Globally Scoped Addresses, 224.0.1.0 - 238.255.255.255
  - » Multicast data across the Internet
  - » Some are reserved (e.g. 224.0.1.1 is used by Network Time Protocol)
- Limited Scope Addresses, 239.0.0.0 - 239.255.255.255
  - » Constrained to a local group or organization
  - » Enables address reuse

IPv6 Multicast Addresses

```
|    8   |  4 |  4 |               112 bits                      |
+--------+----+----+---------------------------------------------+
|11111111|flgs|scop|             group ID                        |
+--------+----+----+---------------------------------------------+
```

- » RFC 3513 - IPv6 Addressing Architecture (see also RFC 3956)
- » Flags 0|0|0|T

  T = 0 permanently-assigned ("well-known") multicast address; defined by IANA

  T = 1 non-permanently-assigned ("transient") multicast address
- » Scop: limits the scope of the multicast group

  1 - interface-local scope
  - > single interface on a node; loopback transmission of multicast

  2 - link-local scope; 5 - site-local scope; E - global scope
  - > Same scope as unicast

  4 - admin-local scope; 8 - organization-local scope
  - > administratively configured; organization-local: spans multiple sites

  6, 7, 9, A, B, C, D – unassigned
  - > available for administrators to define additional multicast regions

Layer 2 Multicast Addresses

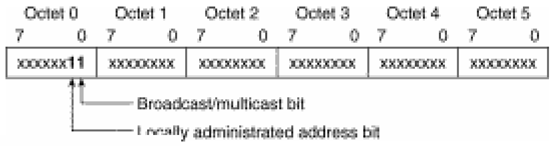


```
+----------------+----------------+----------------+
|cccccc0gcccccccc|ccccccccmmmmmmmm|mmmmmmmmmmmmmmmm|
+----------------+----------------+----------------+  00:C0:DF:08:D5:99

where "c" are the bits of the assigned company_id, "0" is the value of the
universal/local bit to indicate global scope, "g" is individual/group bit,
and "m" are the bits of the manufacturer-selected extension identifier.
```

- A network interface receives packets destined for
  its MAC address(es) and broadcast address (0xFFFF.FFFF.FFFF)
- IEEE 802.3 standard
  bit 0, first octet, indicates group/broadcast/multicast frame
- How can multiple hosts
  receive the same packet and differentiate among multicast groups?

Ethernet MAC Address Mapping

» IANA owns block of Ethernet MAC addresses
  - Starting with 0x 01:00:5E
  - Half of this block is allocated for multicast addresses
    23 bits
    0x 01:00:5e:00:00:00 - 0x 01:00:5e:7f:ff:ff.

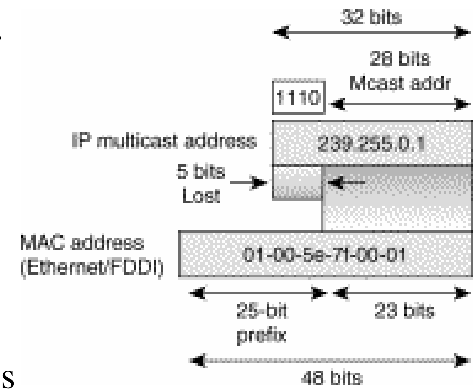

» Lower 23 bits of the IP multicast group address
  are placed into these 23 Ethernet bits

» 5 bits of the IP multicast address are dropped
  an L2 address is not unique; it serves 32 different multicast groups

IPv6 Multicast and Layer 2

- » IANA got new block of MAC addresses for IPv6 Multicast
- » Leading two bytes = 0x 33:33
- » Following 4 bytes/32bits available for address mapping
  from the last 32 bits of the 128 bit Multicast address

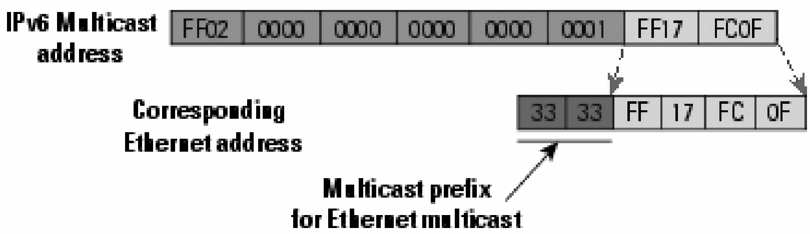

Group Management

IGMP v1, v2, v3, MLD

Bridge broadcast and multicast, IGMP snooping, CGMP

IGMP - Internet Group Management Protocol

- IGMPv1, RFC 1112
- IGMP used to

  dynamically register hosts in a multicast group on a particular LAN
- 2 type of IGMP v1 messages
  - » <u>Membership report</u>
    - – Host sends IGMP membership report to indicate that it is interested in joining the group
  - » <u>Membership query</u>
    - – Router periodically sends IGMP membership query
    - – To verify that at least one host on the subnet is interested in receiving traffic
    - – When there is no reply to 3 consecutive IGMP membership queries

      router stops forwarding traffic directed toward that group

IGMP v1 – Host State Machine, RFC1112

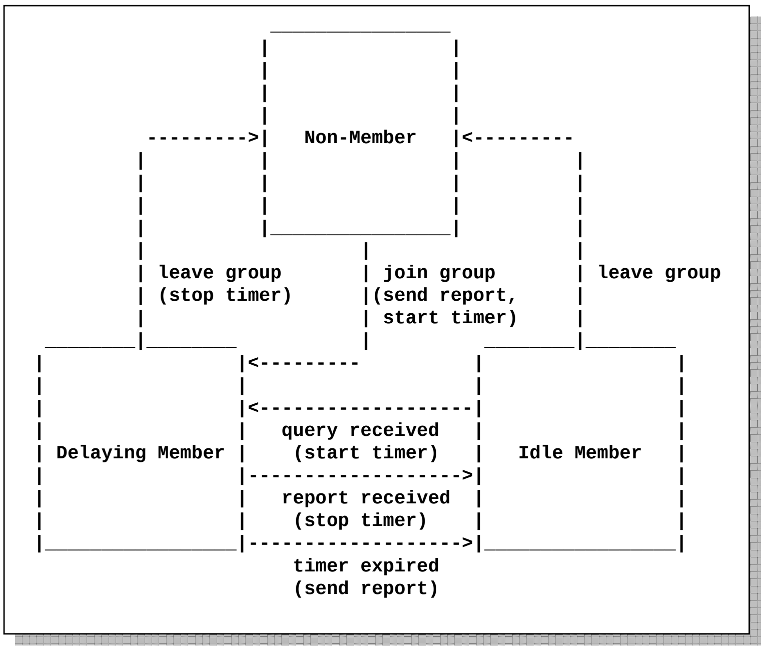

IGMP v2 e v3

- IGMPv2
  - » RFC 2236
  - » 4 type of messages
    - > Membership query
    - > V1 membership report, V2 membership report
    - > Leave group
  - » IGMPv2 works basically as IGMPv1. Main difference:
    - – there is a leave group message
      - > Host can communicate to the local multicast router the intention to leave the group
      - > Reduces the leave latency compared to IGMP Version 1
- IGMPv3
  - » RFC 3376
  - » Enables host to receive only from specified sources
  - » New messages
    - – Group-Source Report   host describes which sources it wants (not) to receive data from
    - – Group-Source Leave   host defines what (source, group) it wants to leave

IGMP v2 – Host State Machine, RFC2236

```
                         ________________
                        |                |
                        |                |
                        |                |
                        |                |
              --------->|   Non-Member   |<---------
             |          |                |          |
             |          |                |          |
             |          |                |          |
             |          |________________|          |
             |                   |                  |
             | leave group       | join group       | leave group
             | (stop timer,      |(send report,     | (send leave
             |  send leave if    | set flag,        |  if flag set)
             |  flag set)        | start timer)     |
     ________|________           |          ________|________
    |                 |<---------          |                 |
    |                 |                    |                 |
    |                 |<-------------------|                 |
    |                 |   query received   |                 |
    | Delaying Member |    (start timer)   |   Idle Member   |
 ---->|               |------------------->|                 |
|   |                 |   report received  |                 |
|   |                 |    (stop timer,    |                 |
|   |                 |     clear flag)    |                 |
|   |_________________|------------------->|_________________|
| query received    |        timer expired
| (reset timer if   |        (send report,
|  Max Resp Time    |         set flag)
|  < current timer) |
 -------------------
```

Group Management – IPv6

- Multicast Listener Discovery – MLD
- MLD v1
  - RFC 2710
  - Similar to IGMPv2
- MLD v2
  - RFC 3810
  - Similar to IGMPv3

Bridges and LAN Switches – Broadcast and Multicast

- Broadcast
  - frame forwarded towards the active ports
    - other than the port from which the frame was received
- Multicast
  - » Same way host decides if it receives the frame
  - » Alternative 1
    - – Take advantage of learning bridge property
    - – Host belonging to a multicast group sends periodically a frame whose L2 source address is the L2 multicast address
  - » Alternative 2 switch snoops IGMP messages
    - – Host sends Join/Leave messages to multicast routers through switch
    - – Switch detects them and adds entry to the corresponding ports
  - » Alternative 3 Cisco Group Management Protocol (CGMP)
    - – Layer 3 signalling between the router and the switch
    - – Host informs de router (IGMP); router informs the switch (CGMP)

Multicast routing

Source and shared trees

Multicast forwarding

PIM-SM, PIM-DM

MBGP, MSDP

Multicast Routing

- Internet multicast
  implemented over networks supporting hardware multicast/broadcast
  by extending routing and forwarding functions

- Based on the concept of distribution tree
  controls the paths of IP multicast traffic

- 2 types of multicast distribution trees
  source trees, shared trees

Source Tree

- The simplest multicast distribution tree
  - root source of the tree
  - leaves group receivers
  - Shortest Path Tree (SPT)
- Notation (S,G)
  - S is the IP address of the source
  - G is the multicast group address
- One tree
  - for each source sending to the group
- n sources Ł n trees

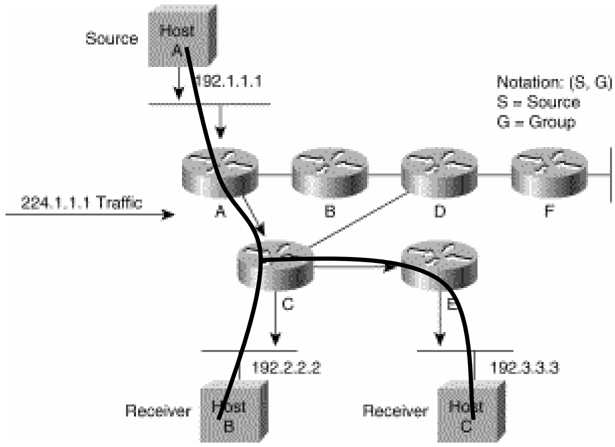

Host A Shortest Path Tree
(192.1.1.1, 224.1.1.1)

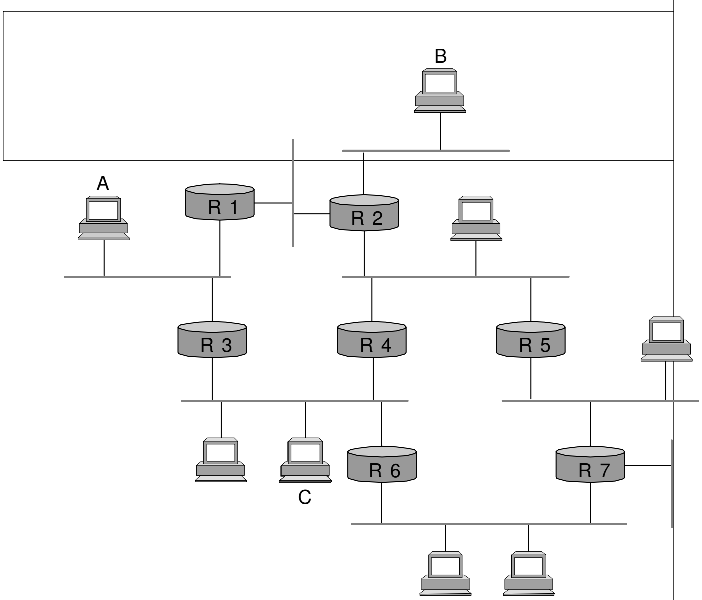

A,B,C – sources
- receivers

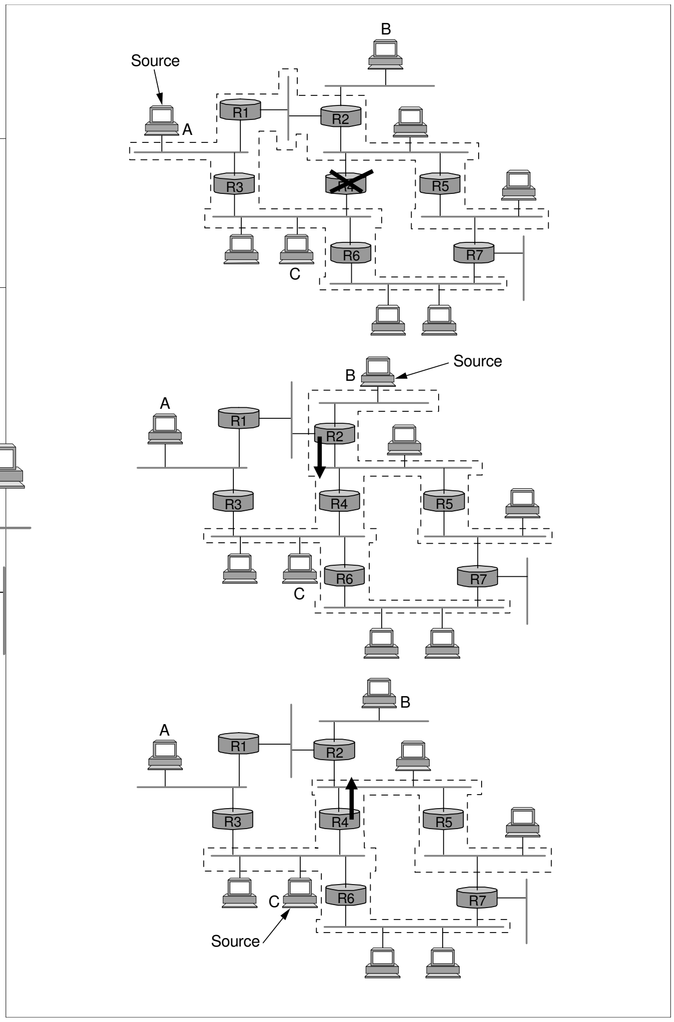

Shared Tree

- Single common root
  - the *rendezvous point (RP)*
- Sources send traffic to *RP*
- Then, traffic
  - is forwarded down shared tree,
  - and reaches all receivers
- Notation (*, G),
  - *, all sources
  - G, the multicast group

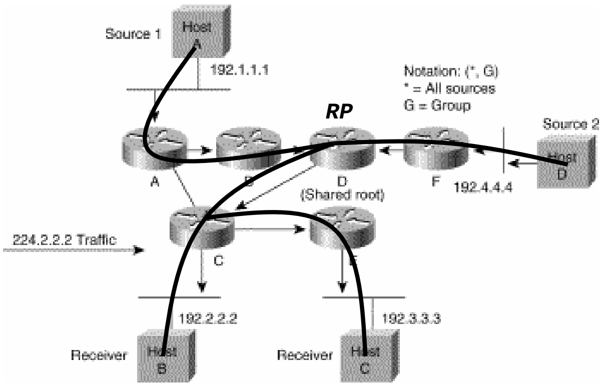

Shared Distribution Tree
(*, 224.2.2.2)

Multicast Trees

- Loop-free
- Messages replicated at routers, where the tree branches
- Members of multicast groups can join or leave at any time
- Distribution trees are dynamically updated
- Source tree
  - » optimal paths between source and receivers
  - » minimal network latency
  - » Routers must maintain path information for each source
    - thousands of sources * thousands of groups   scalability problem
- Shared tree
  - » minimum amount of state in each router
  - » paths between source and receivers might not be the optimal

Multicast Forwarding

- In unicast routing
  - » router sends data towards the destination address
  - » router scans through its routing table, and
    - – forwards a single copy of the unicast packet
    - – in the direction of the destination

- In multicast routing
  - » Router sends traffic away from the source
  - » If there are multiple downstream paths
    - – the router replicates the packet
    - – forwards the traffic down the appropriate downstream paths
  - » Reverse path forwarding
    - – Unicast routing tables determine upstream and downstream neighbors
    - – Router forwards a multicast packet only if it is received on the upstream interface
      - > distribution tree becomes loop-free

Reverse Path Forwarding

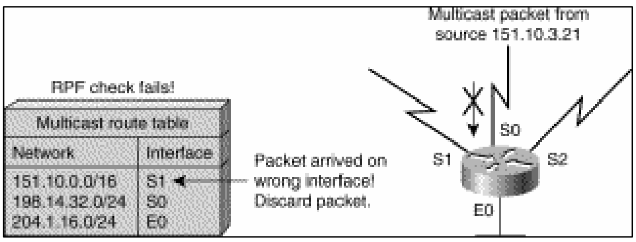

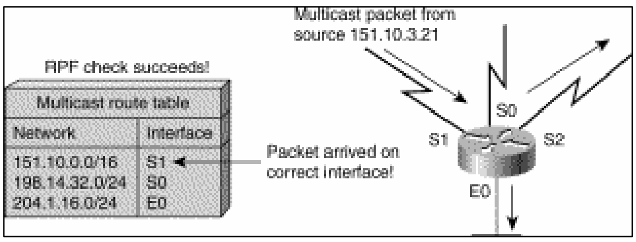

- Router looks <u>source address</u> in the unicast routing table

  to determine if packet arrived through the best interface to the source

  - Correct interface packet is forwarded
  - Uncorrect interface packet is dropped

IP Multicast Protocols

- Tree construction protocol
  - Protocol Independent Multicast, Sparse Mode (PIM-SM)
  - Protocol Independent Multicast, Dense Mode (PIM-DM)
- Advertise reverse paths towards sources
  - Multiprotocol Border Gateway Protocol (MBGP)
- Disseminate information about sources
  - » Multicast Source Discovery Protocol (MSDP)

Protocol-Independent Multicast (PIM)

IP routing protocol independent

- But uses information from unicast routing protocols
  - » to populate the unicast routing table
  - » including OSPF, BGP, static routes
- 2 modes
  - » PIM Dense Mode PIM (PIM-DM)
  - » PIM Sparse Mode (PIM-SM)

Protocol-Independent Multicast (PIM)

- PIM Dense Mode PIM (PIM-DM)
  - » Based on push model Ł flood multicast traffic to every corner of the network
  - » PIM-DM initially floods multicast traffic throughout the network
    - – Routers that do not have any downstream neighbors prune back the tree
  - » Supports only source trees
  - » Efficient if there are active receivers on every subnet
- PIM Sparse Mode (PIM-SM), RFC 2362
  - » Based on pull model Ł only networks towards active receivers forward traffic
  - » Builds shared tree to distribute information from active sources
    - – then, traffic can remain on the shared tree or switch to a source tree
  - » PIM-SM uses an *Rendez-Vous Point* (RP)
    - – Administratively configured; root of the shared tree
    - – Sources register with the RP; data is forwarded down the shared tree
    - – Under high traffic conditions, routers may dynamically create source trees
  - » PIM-SM scales well

PIM-SM

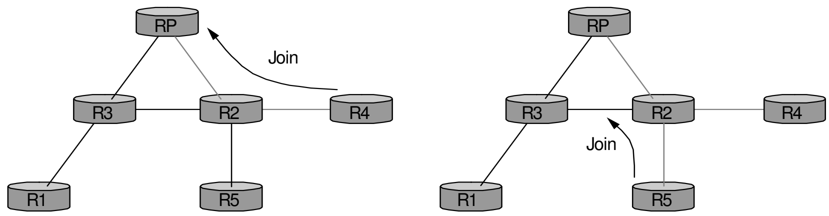

(a) (b)

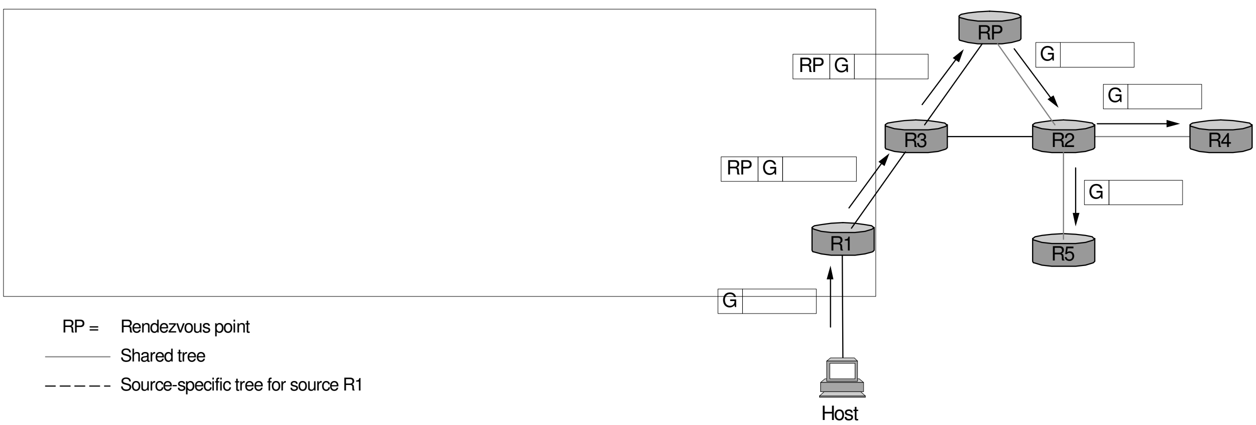

PIM-SM: Route Optimization

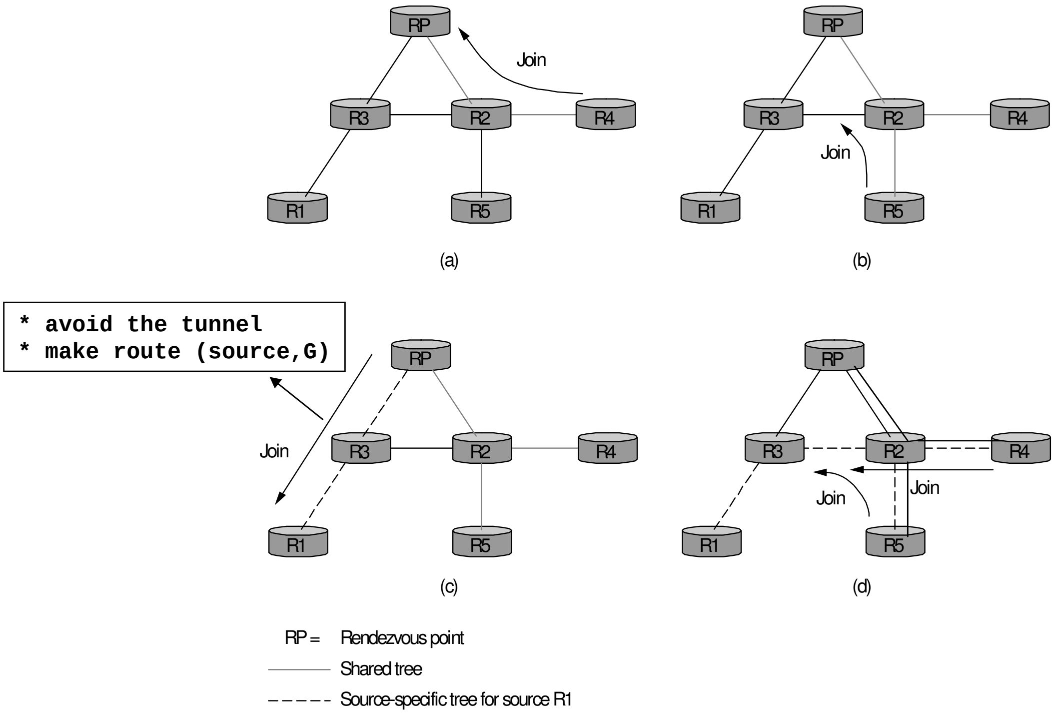

Related Multicast Protocols

- Multiprotocol Border Gateway Protocol (MBGP)
  - RFC 2283
  - Multicast routing between providers; interdomain routing management
- Multicast Source Discovery Protocol (MSDP)
  - Enables each domain to have an RP
  - RPs of different domains exchange information
    - about themselves and the sources they know
  - Inter ISP (RP) protocol