Overview

1. IP Multicast
   - DVMRP
   - PIM
   - MBone routing
2. Real-time Traffic on the Internet
   - Jitter and Adaptive Playout
   - RTP and RTCP
3. Multimedia Conferencing
   - Audio tools
   - Video Tools
   - Shared Media Spaces
   - Session Directories

1 IP Multicast

1.1 Requirements

- Support many to many conferencing
- Scale to large groups
  - Data should only traverse each link once
- Robustness to link failures
  - No fate sharing
  - End systems don't know topology

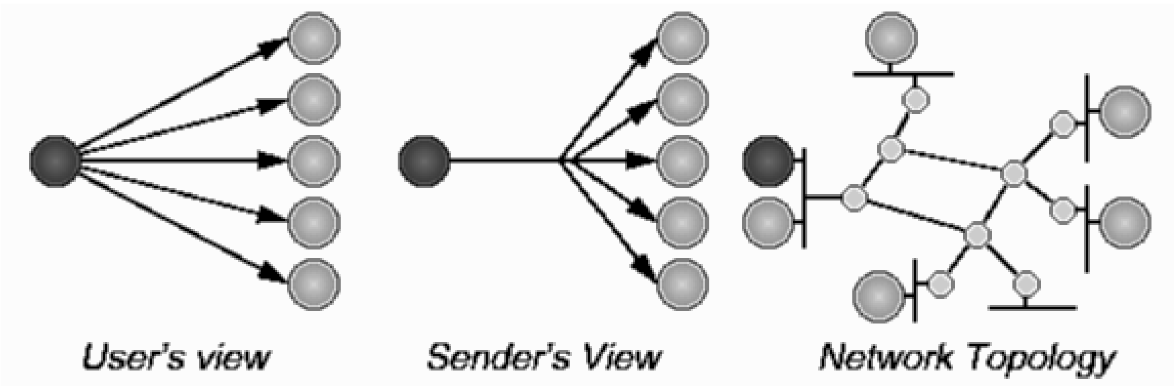

1.2 What is Multicast?

- Dynamically Constructs efficient, shortest path distribution trees from senders to receivers
- Service Model
  - Receivers announce interest
  - Senders just send
  - Routers conspire to deliver sender's data to all interested receivers

1.2.1 Class D addresses

- Traditionally, IP addresses were divided into classes:
  - Class A for companies with millions of hosts
  - Class B for companies with thousands of hosts
  - Class C for companies with hundreds of hosts.
- Class D is multicast group addresses.
  - A multicast sender just sends to a class D multicast address
  - Multicast receivers express an interest in a class D multicast address.
  - If they choose the same address, the network delivers traffic from the senders to the receivers.
  - Class D addresses are in the range 224.0.0.0 to 239.255.255.255, with 224.2.X.Y allocated for conferencing.

1.2.2 Multicast routing models

- DVMRP
  - Distance Vector Multicast Routing Protocol
  - forms the basis of most of the current Mbone
- PIM
  - Protocol Independent Multicast
  - offered in Cisco router
- MOSPF
  - Multicast OSPF
- CBT
  - Core Based Trees
  - from UCL
  - better scaling properties at the expense of non-shortest path trees.

1.3 Distance Vector Multicast Routing Protocol (and DM-PIM)

- Basic Model is "Flood and Prune"

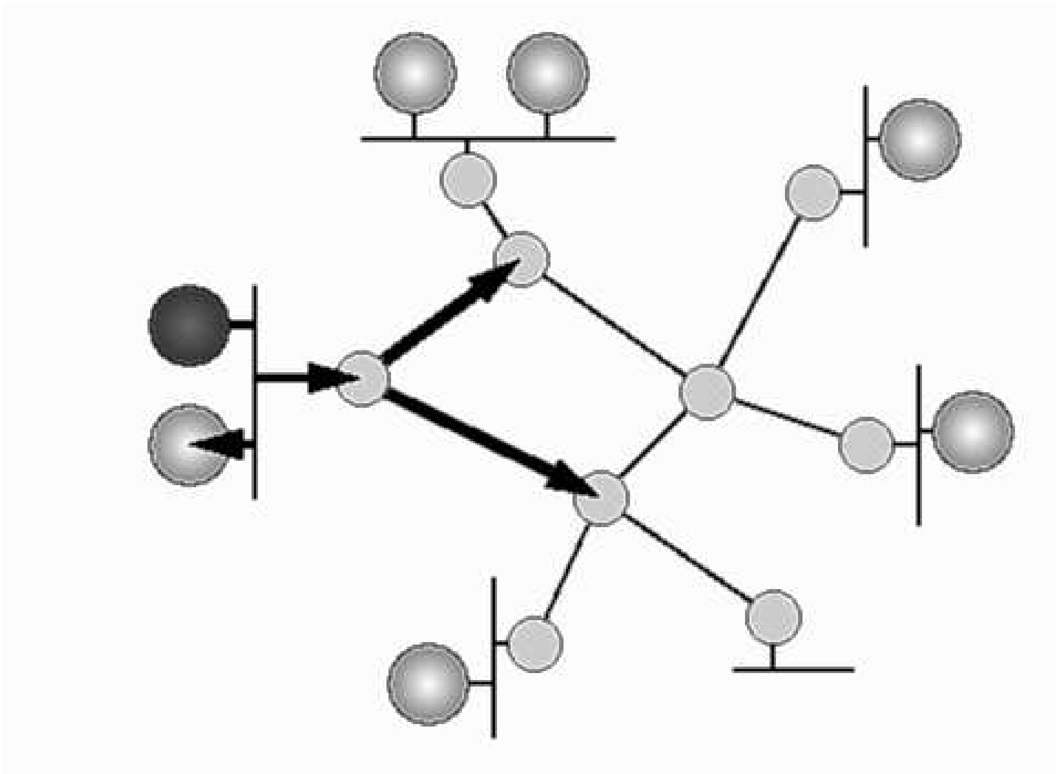

1.3.1 Flood and Prune

Branches not on the shortest path tree to a receiver are pruned off:

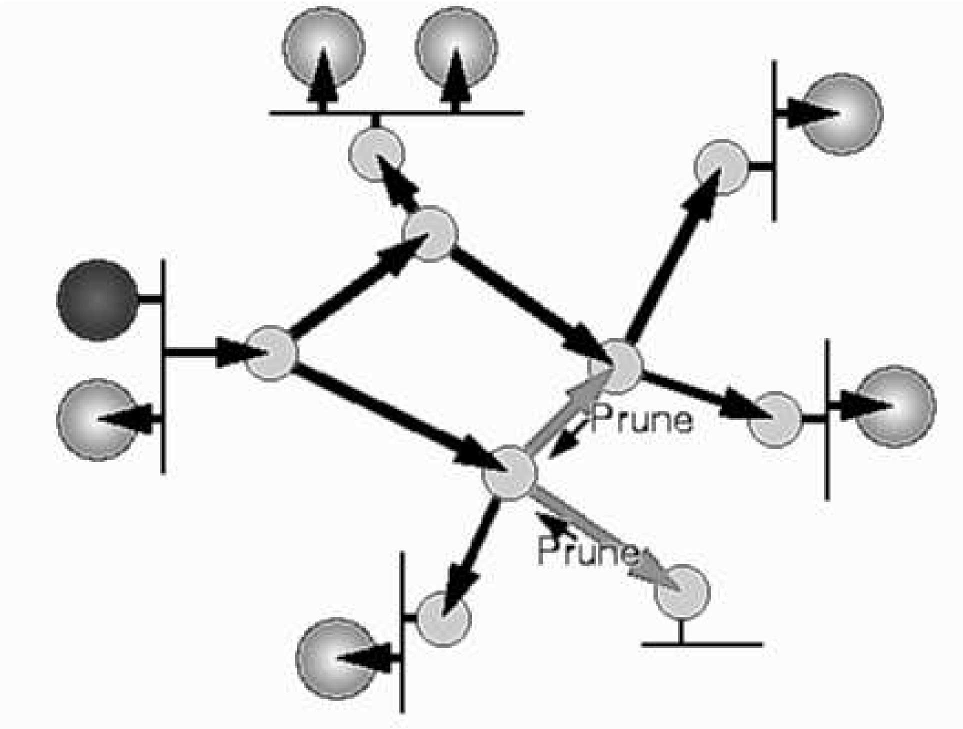

1.3.2 Shortest Path Tree

Resulting in a shortest path distribution tree rooted at the sender:

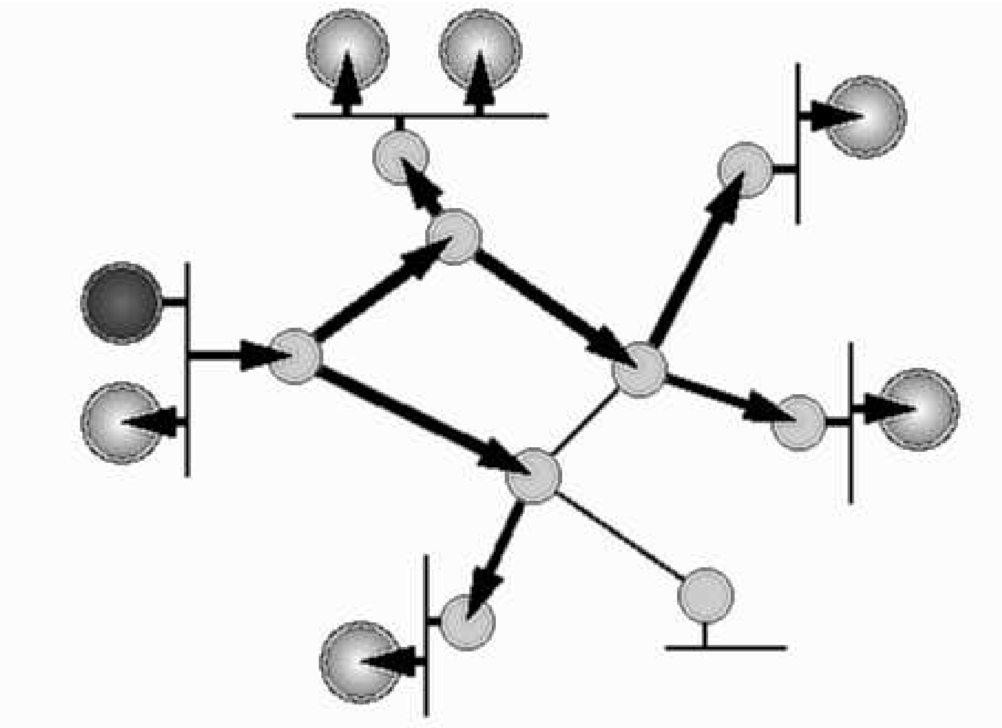

1.3.3 Local Tree Maintenance

Changes in membership are handled efficiently and locally via grafting and pruning:

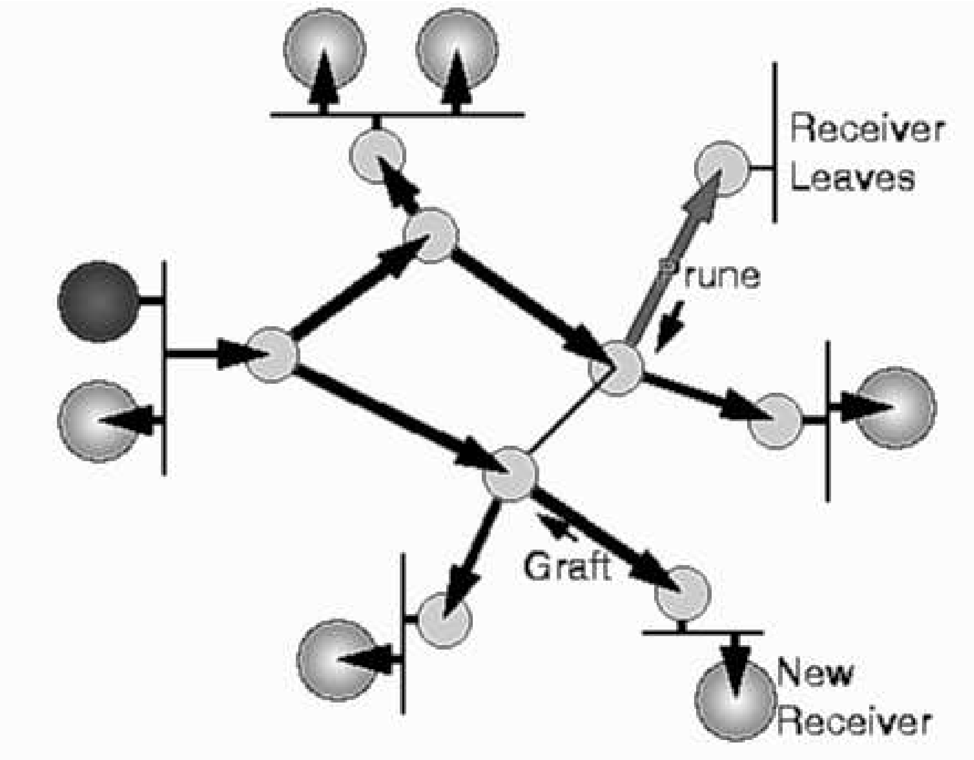

1.3.4 A Shortest Path Tree per Sender

Distribution trees are per-{sender,group}, and are triggered by data packets

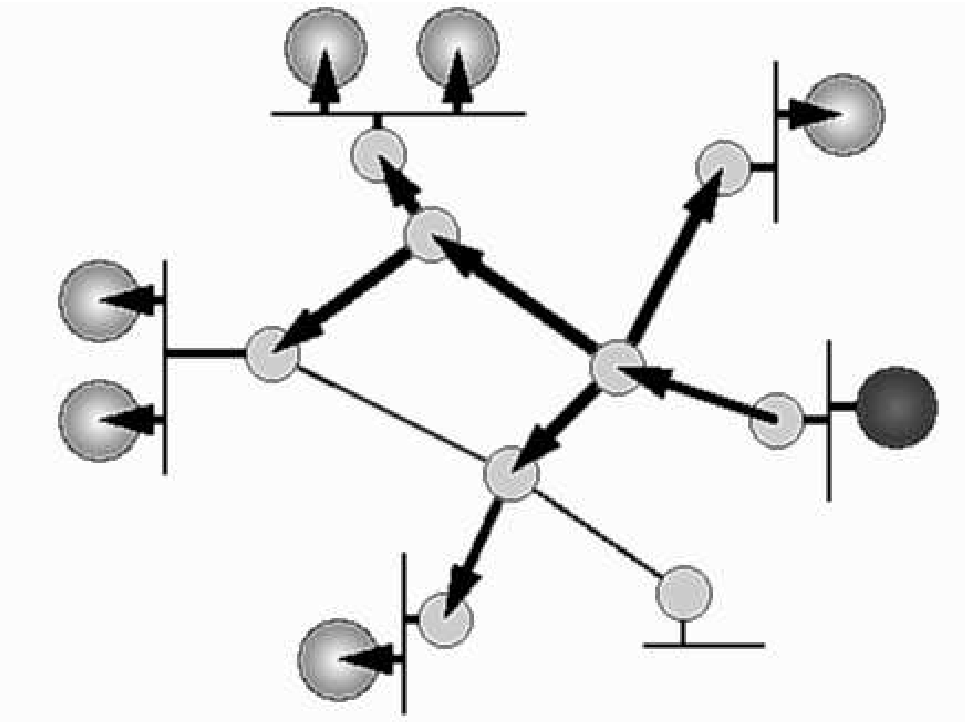

1.4 Protocol Independent Multicast

- Instead of flooding and pruning to directly build a shortest path tree, Sparse-mode PIM initially builds a shared tree.
- The shared tree is built by sending join messages towards a Rendezvous Point (RP)
- Once data is flowing, the shared tree can be converted to a shortest path tree if required.

1.4.1 Building a Shared Tree

- Source S sends traffic to group G
- S's local router forwards group G traffic to the Rendezvous Point (RP) for group G
- R's local router sends a join message for G to RP

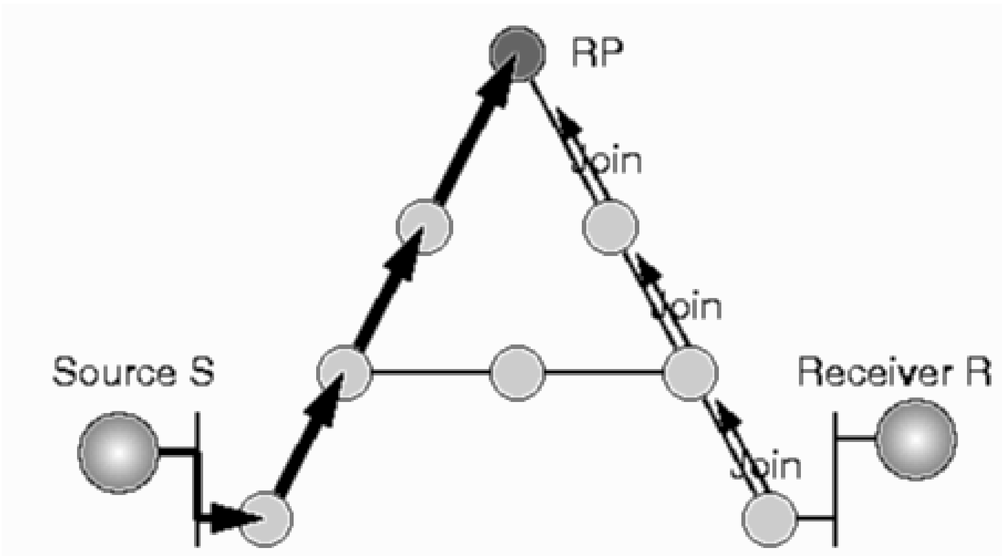

1.4.2 SM-PIM Shared Tree

- Initially data flows over RP tree

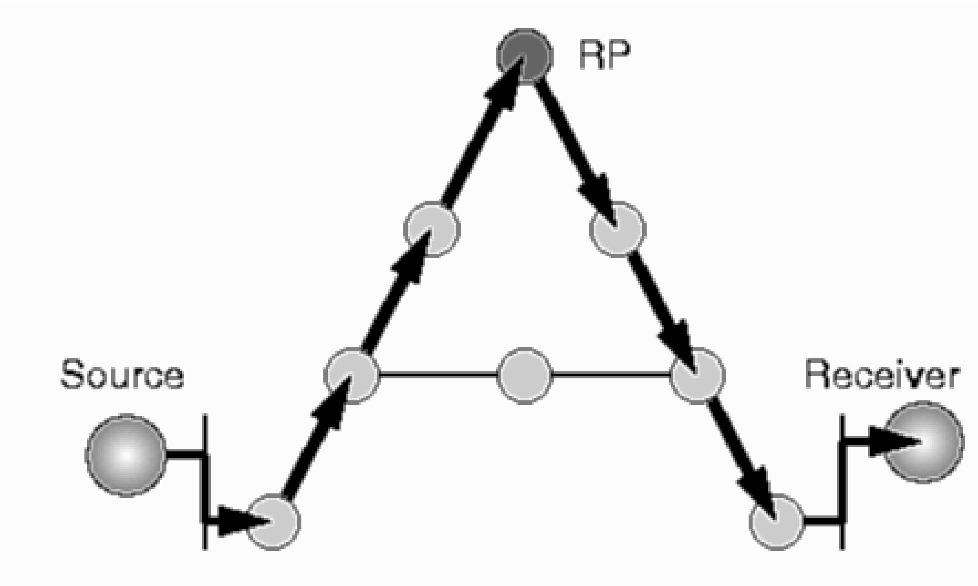

1.4.3 Transferring to SP tree

- Router A sees traffic from S arriving over its B interface, when its unicast route to S is over its C interface.
- If traffic warrants it, A sends a join towards S for {S,G}.

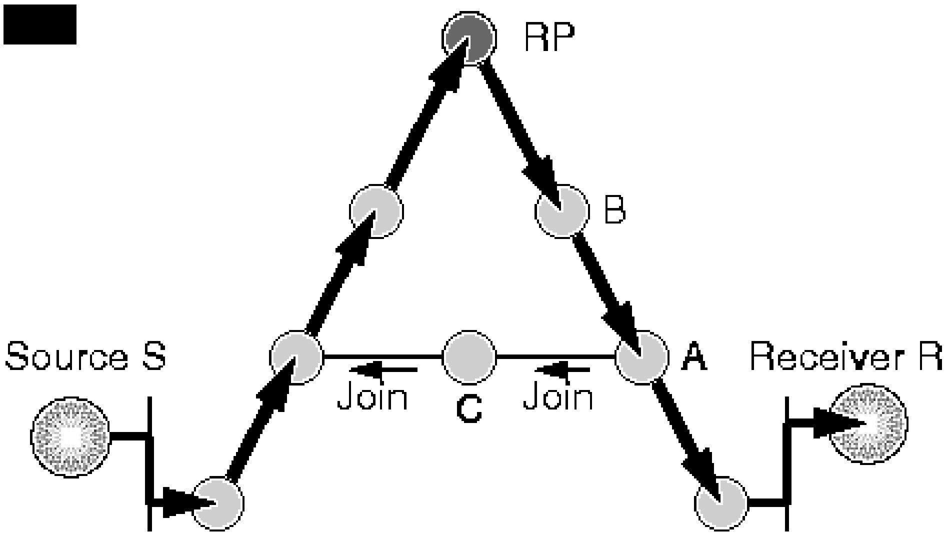

- When the join towards S is successful, group G traffic from S starts to flow along the shortest path tree. A now receives the traffic twice.

- A then sends a prune for $\{S,G\}$ towards the RP for group G, and the transfer to SP tree is complete

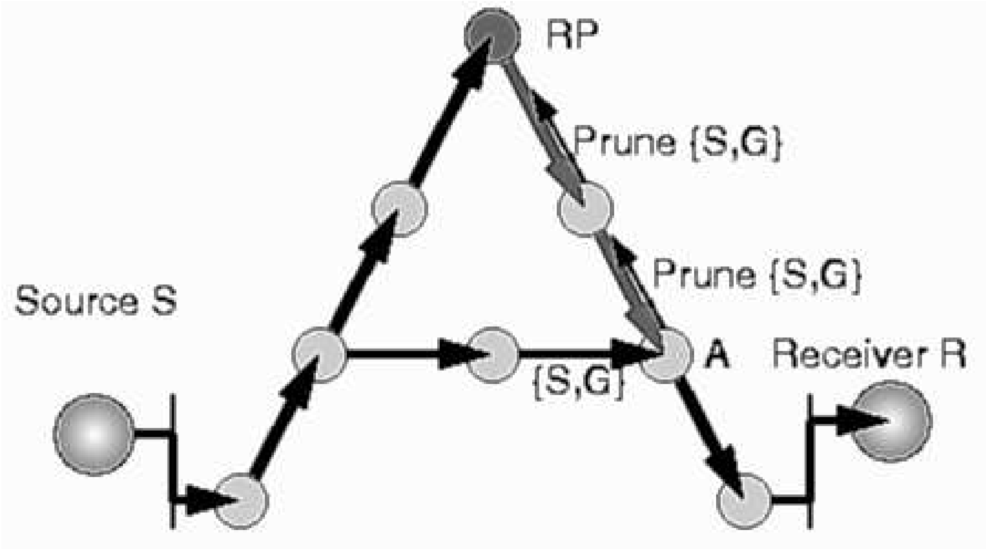

1.4.4 Sparse Mode PIM

- Receiver R receives S,G on the shortest path tree.
- Receiver T continues to receive {S,G} on the RP tree, although it too can switch to a shortest path tree if required

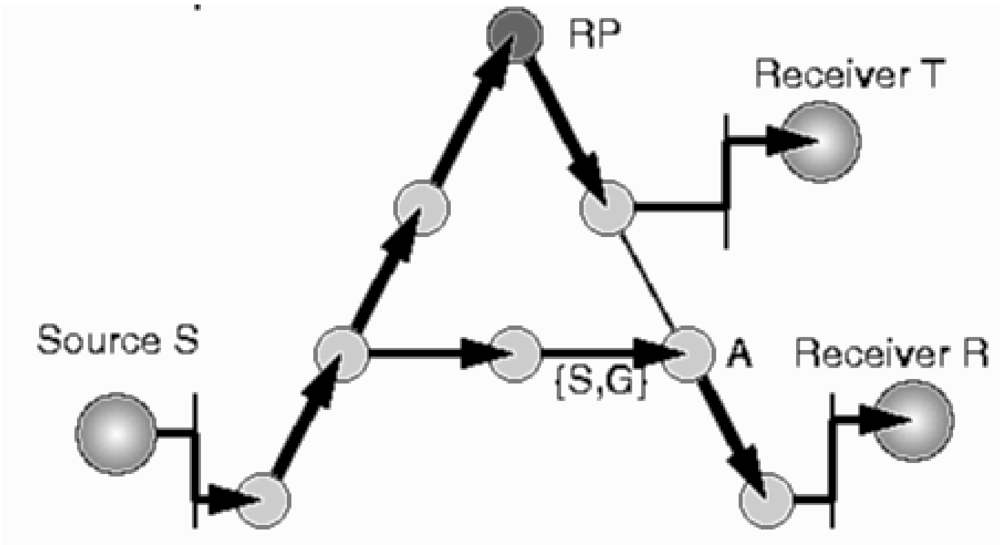

1.5 Routing, Tunnelling and the Mbone

- PIM normally uses the underlying unicast routing to determine its tree formation.
- DVMRP implements its own routing protocol for multicast traffic independently of unicast routing.
- In the current Mbone, DVMRP is usually implemented as an overlay of multicast *tunnels* because most unicast routers do not yet support multicast.
- In a multicast tunnel, an IP multicast packet is encapsulated in a unicast packet to carry it across multicast unaware links and routers.

1.5.1 Tunnels

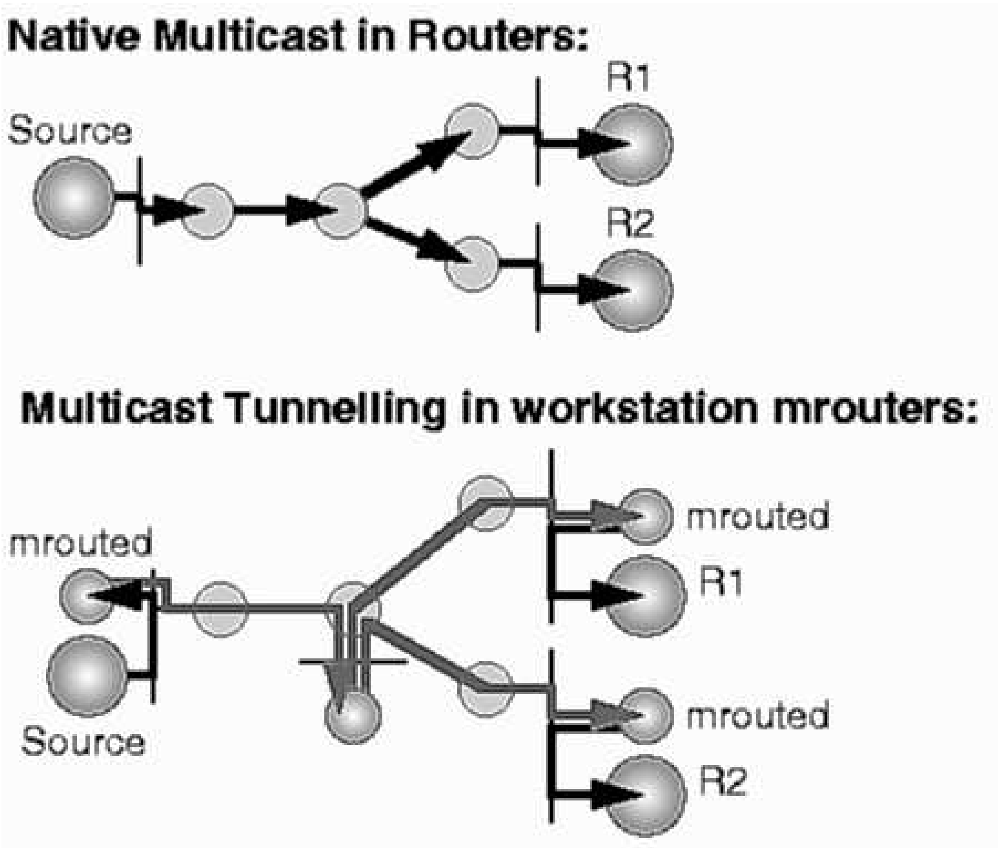

Native Multicast in Routers:
R1
Source
R2
Multicast Tunnelling in workstation mrouters:
mrouted
mrouted
R1
Source
mrouted
R2

- Multicast can be tunnelled across multicast unaware links, but this can limit performance.
- Workstations are not ideal routers
- At best, multicast traverses your LAN twice
- At worst, many tunnels traverse your LAN
- Tunnels require manual configuration

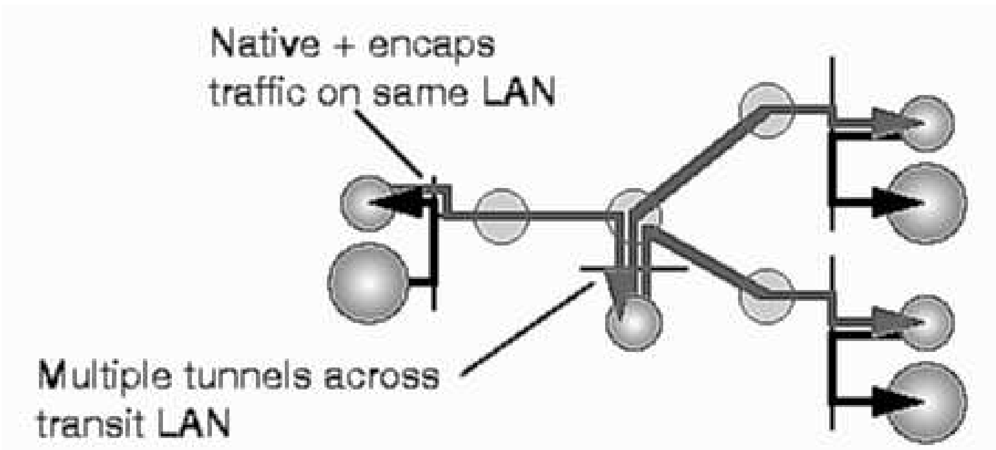

1.5.2 DVMRP Routing

- Each link is given a routing metric (or cost)
- Given two alternative paths from source to a router, the higher cost path is pruned, and the lower cost path used.
- E.g. there are two paths from SB to F. The path via E (cost 3) is pruned and the route via C (cost 2) is retained

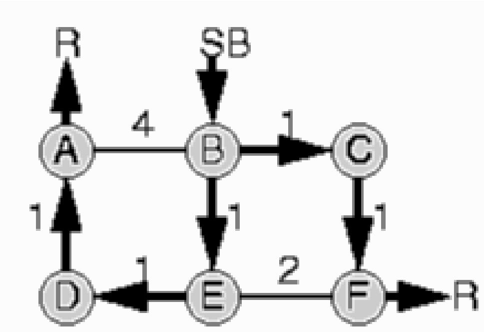

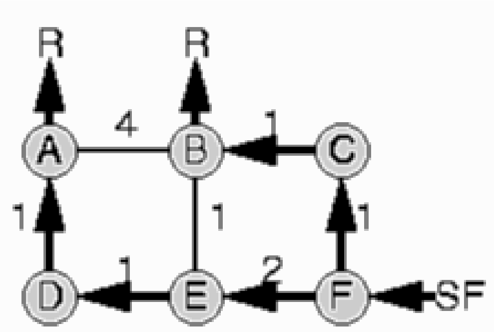

- DVMRP is basically multicast RIP.
  - It was not really intended to be used as one global routing domain (as it is currently in the Mbone)
- DVMRP considers a route with metric 32 to be *infinite*
  - This helps it prevent loops and achieve convergence.
  - It also is uncomfortably small, so it is advisable to use as small a metric on your links as possible.

1.5.3 DVMRP Scoping

- mrouted supports two mechanisms to limit the scope of multicast traffic:
  - TTL scoping
  - Administrative Scoping

1.5.4 Time To Live

- When any IP packet (unicast or multicast) is sent, a field called *time to live* (TTL) is set to some positive value
- Each router the packet passes through decrements the packet by one.
- When the TTL field reaches zero, the packet is dropped. This is supposed to prevent packets looping in the internet forever if a routing loop occurs (and to provide lifetime guarantees for transport mechanisms).
- mrouted leverages the IP TTL field to limit the scope of multicast traffic to particular regions.

1.5.5 TTL Scoping

- When an application is started, it is configured to send its data with a particular TTL.
- Each multicast interface or tunnel is configured with a TTL threshold.
- Every time a multicast packet passes through a multicast router, its TTL is decremented.
- If a packet has a TTL less that the threshold of an interface or tunnel, the packet is not forwarded out of that interface or tunnel.
- If all the tunnels out of a region have the threshold x, then sending data with TTL x-1 will confine the traffic to that region.