

4.3 Multicast

4.3.1 Multicast Principles

4.3.2 Multicast in LANs

4.3.3 Multicast in WANs

4.3.1 Multicast Principles

Definition of Multicast

Multicast is the transmission of a data stream from one sender to many receivers with packet duplication and forwarding inside the network.

Why is Multicast Important for Multimedia?

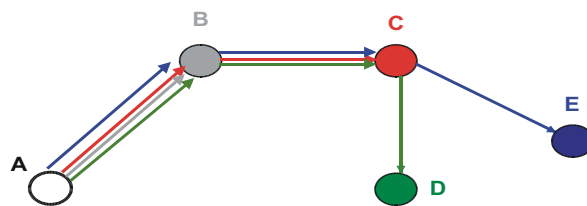
Multicast applications often require 1:n communication:

- Videoconferencing
- TeleCooperation (CSCW) with a shared workspace
- near-Video-on-Demand
- TeleTeaching, TeleSeminars

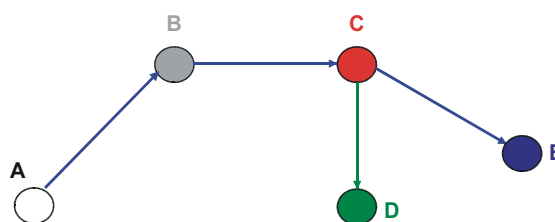
Digital video streams have high data rates (>1 MBit/s). n point-to-point connections would overload the network.

Multicast in Our Example

Assumption: All nodes are group members (receivers)



(a) Four end-to-end connections



(b) One multicast connection

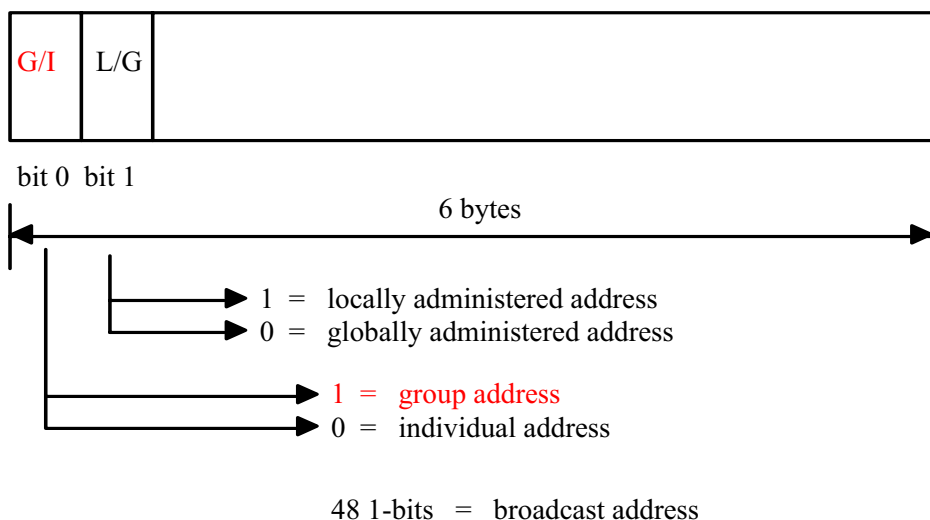
4.3.2 Multicast in LANs

Ethernet, Token Ring, FDDI

- The topology has broadcast property.
- Layer-2 frame addresses according to IEEE 802.2 allow group addressing for multicast. If a station is a member of the group it will forward the incoming packet to the higher layers.
- But: For a long time layers 3 and higher of the Internet did not support group addresses! Thus, if the Internet protocol stack was used on a LAN, this effectively prevented multicasting in layer 2!

Frame Address Format in LANs

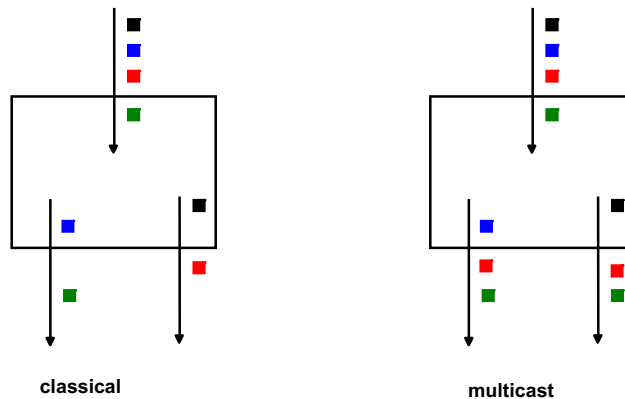
According to IEEE 802 and ISO 8802



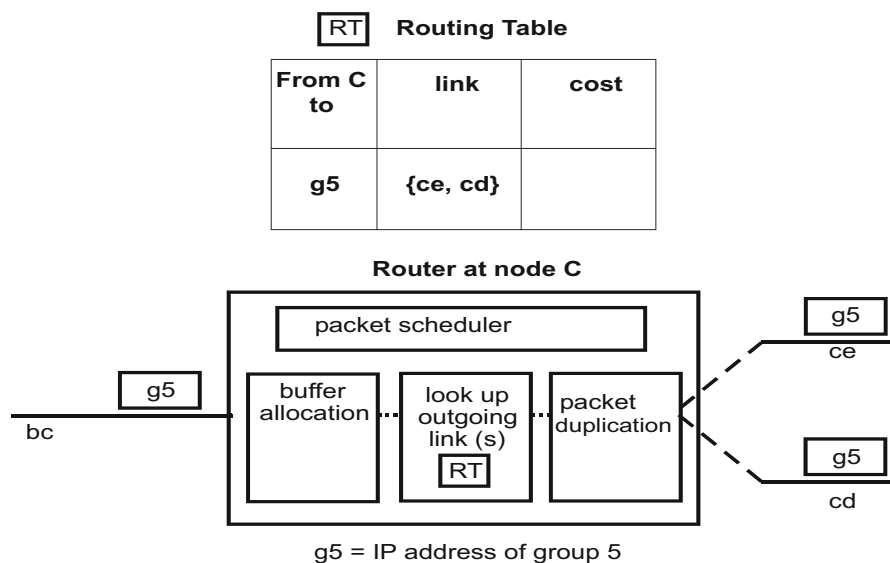
4.3.3 Multicast in WANs

Packets are duplicated in the routers.

- Copy incoming packets/cells and enqueue them into several output buffers. Can be done efficiently.
- Where and how to duplicate is controlled by entries in the routing table
- Should multicast be connection-oriented or connectionless?



Router with Multicast Support



Multicast IP in the Internet

Multicast in the IP layer of the Internet dates back to the late 1980's. At that time IP Version 4 was extended to include multicast addressing and multi-cast routing („multicast IP“).

The MBone is a multicast backbone on the Internet; it is an overlay network of multicast-enabled routers and end systems. The MBone is mainly used for research purposes; ISPs still hesitate to open it to the general public because it is difficult to manage.

Principles of Multicast IP

IP packets are sent to a group address (an IP address of type D).

The protocol is connectionless (a datagram protocol).

The service is “best effort“:

- No error control
- No flow control

The service is receiver-oriented:

- The sender sends multicast packets to the group.
- The sender does not know the set of receivers, has no control over it.
- Each host on the Internet can join the group.

Limiting the scope of a transmission is done by setting the Time-To-Live parameter of the IP packets. There are general conventions on TTL thresholds for multicast IP, such as:

1: this LAN only

<= 16: this company only

<= 32: this country only

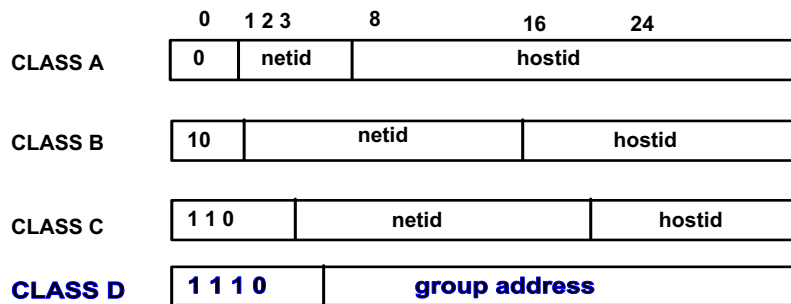
etc.

Group Addressing

Internet Group Addressing

- Group members are IP hosts
- Can be anywhere in the Internet
- The group address controls routing and packet duplication

But: Group addresses must be assigned dynamically! We need an efficient, scalable group address management protocol.

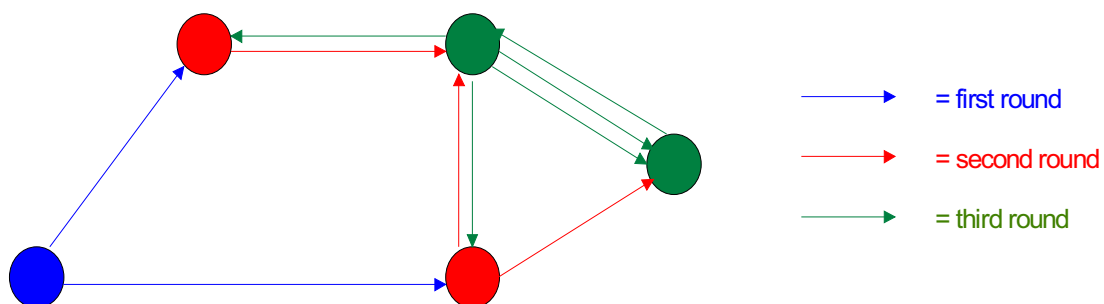


Multicast Routing: Flooding

The simplest way to get a packet stream to **every** node in the network is flooding (broadcasting).

Algorithm Flooding

When a packet arrives it is forwarded on **each** outgoing link except the one on which it has arrived.



Problem: Infinite Number of Packets

Solution

Place a “hop counter“ in the packet header.

- Initialize it with the diameter of the network
- Decrement it by 1 on each hop
- Duplicate packets get the hop counter of the original packet.
- When the hop counter is 0 the packet is not forwarded by the router.

In the example shown on the previous page this algorithm would eliminate the third round of messages since the diameter of the graph is 2.

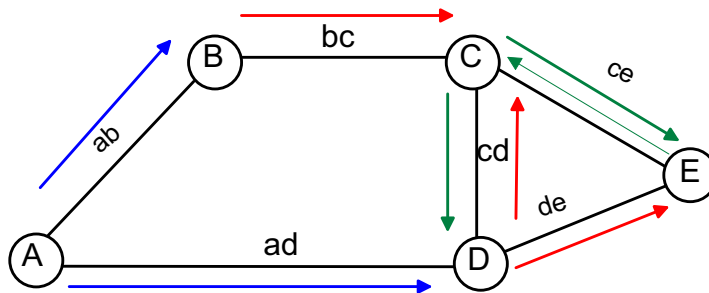
Reverse Path Broadcasting (RPB)

Reverse Path Broadcasting (RPB) is more efficient than flooding. It takes advantage of the fact that each node knows a packet's shortest path from the sender from the classical routing table! This path is called the reverse path.

The basic idea is now that a node forwards only those packets on all outgoing links that have arrived on the shortest path from the sender.

This algorithm generates much fewer packets than pure flooding. But it is still possible that the same packet arrives twice at a node.

Example for Reverse Path Broadcasting (still incomplete)

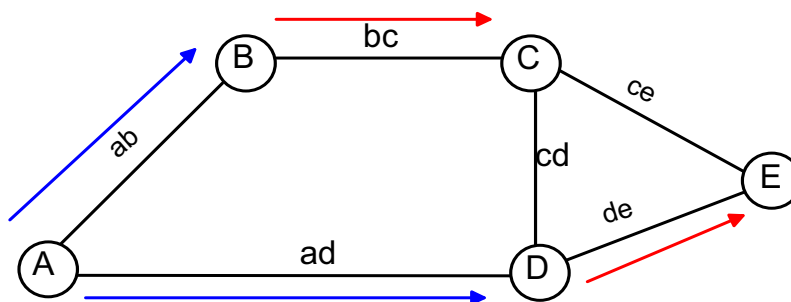


As we can see there are still unnecessary packets: D and E receive the packet twice, C even three times.

Example for Reverse Path Broadcasting (complete algorithm)

When each node sends some extra information to its neighbors RPB can avoid the unnecessary packets. In the extra information each node explicitly informs his neighbor who is on his shortest path to the sender.

The complete RPB algorithm then works as shown below.



Truncated Reverse Path Broadcasting (TRPB)

RPB still delivers all packets to **all** nodes in the network. TRPB allows subnetworks (typically LANs) to only participate in a multicast if they contain at least one interested host. A simple protocol called **IGMP** (Internet Group Management Protocol) was defined for this purpose. Via IGMP the hosts in a LAN tell their router whether they are interested in a particular multicast stream. If no local host is interested in a multicast group the router will stop forwarding packets onto the LAN.

But this only solves the problem for **local hosts** in a LAN. The multicast tree still includes all **routers** in the network.

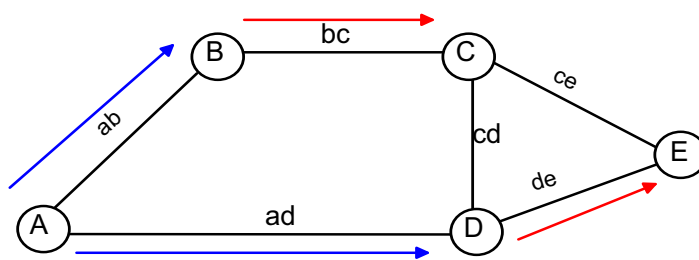
Reverse Path Multicasting (RPM)

The RPM algorithm adds prune messages to the protocol. A prune message travels from a leaf towards the root of the multicast tree. It tells a higher-level node that there are no interested receivers below. This reduces the broadcast tree to a true multicast tree.

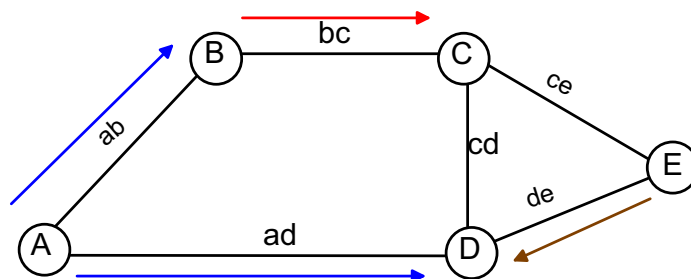
Algorithm “Pruning“

If a router's sons are all tree leaves that are not interested in the multicast stream the router will send a Non-Membership Report (NMR) message to his father. If an inner router receives NMRs from all its sons it will in turn send a NMR to its father. NMR messages contain a time-stamp. When the timer expires the pruning is cancelled, and multicast packets are again forwarded.

Example for Reverse Path Multicasting (1)

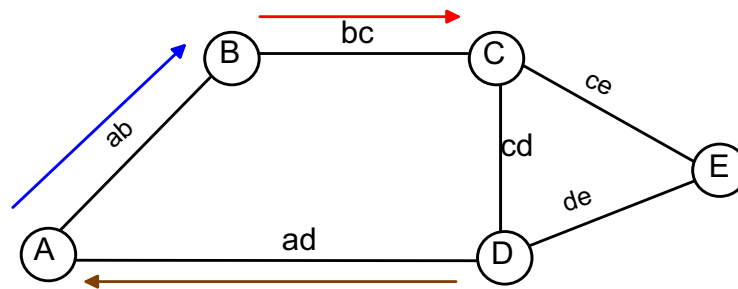


(a) Tree in the initial RPM phase



(b) E sends a “prune“ message

Example for Reverse Path Multicasting (2)



(c) D sends a “prune” message

Multicast in IP Version 6 (IPv6)

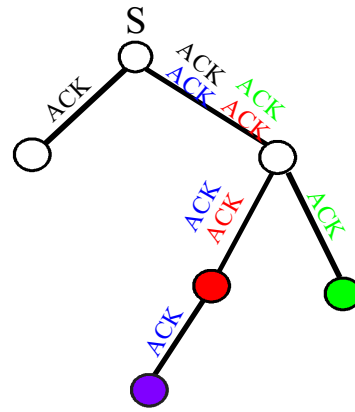
The multicast capability is fully integrated into Version 6 of the IP protocol.

- All IPv6 routers understand group addresses, and are able to route multicast packets.
- The IGMP protocol is integrated into the ICMP protocol.

Reliable Multicast

Two basic schemes

- ACK with retransmission: causes the ACK implosion problem. Many „ACK saving protocols“ suggested in the literature.
- FEC: Not very good for burst errors, no 100% guarantee



ACK implosion

Multicast With QoS Guarantees

Guaranteed Quality-of-Service (QoS):

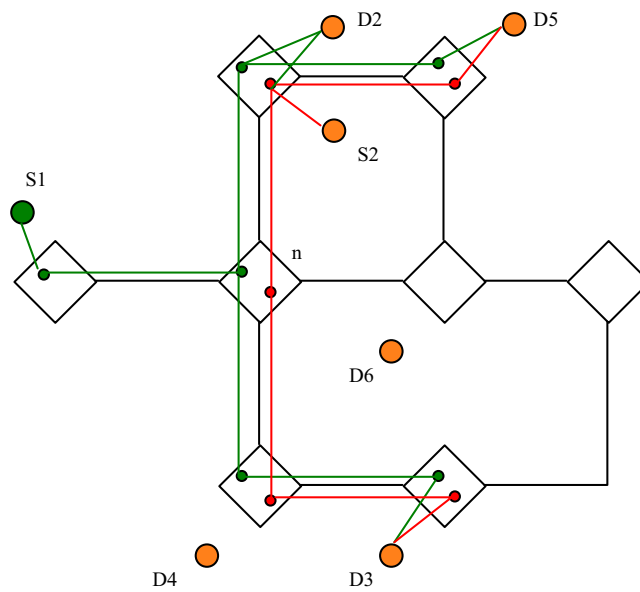
- maximum end-to-end delay
- maximum delay jitter (variance in delay)
- minimum reliability

for a given traffic load described by a flow specification.

This would require resource reservation on all links and nodes in the network:

- Bandwidth
- CPU cycles
- buffer space
- "schedulability"

Dynamic Join and Leave With QoS Guarantees



Open Problems

- Should reservation protocols be connection-oriented or connectionless (or “soft-state”)?
- How can dynamic join-and-leave for group members be combined with QoS guarantees?
- What are the best, scalable tree routing algorithms for multicast? Should they be sender-oriented or receiver-oriented?
- What are the best, scalable reliable multicast protocols?
- Should we allow applications to “inject” filters into the network nodes?

Conclusions on Multicast

Multicast in worldwide interconnected networks is very important for the efficient support of multimedia applications.

Today, there is no multicast protocol at all supporting all the requirements.

Multicast IP and the MBone protocols are useable, but do not support QoS for continuous media. They are not very scalable.

The next generation of Internet protocols will support group addressing and multicast routing, perhaps differentiated services, but no resource reservation and no QoS-based routing.