4.3 Multicast

4.3.1 Multicast Principles

Definition of Multicast

Multicast is defined as 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)

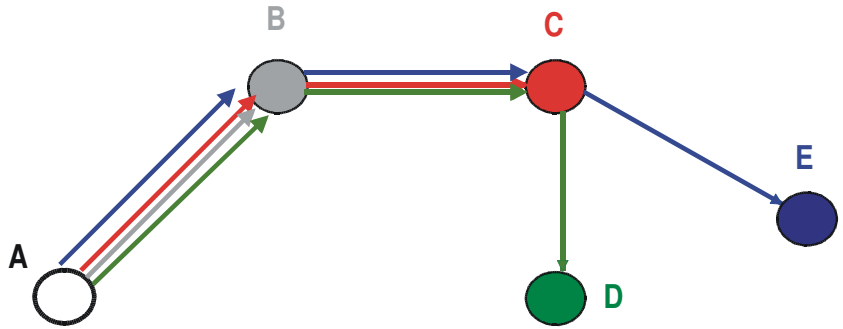

(a) Four end-to-end connections

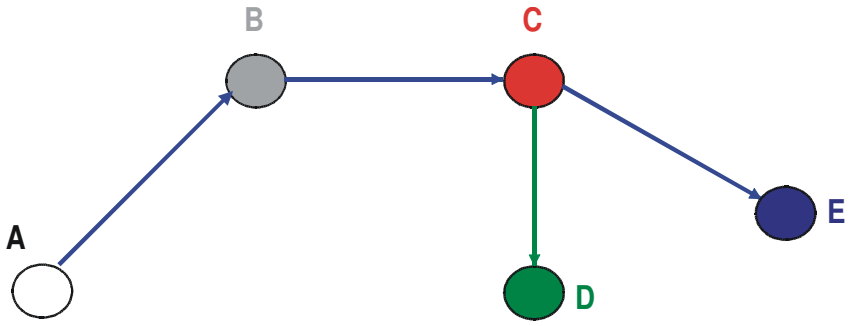

(b) One multicast connection

4.3.2 Multicast in LANs

Ethernet, Token Ring, FDDI

- Die Topologie hat Broadcast-Eigenschaft.
- Die Schicht-2-Adressen nach IEEE 802.2 erlauben die Verwendung von Gruppenadressen für Multicast. Dadurch lässt sich Multicast in einem LAN-Segment leicht und effizient realisieren.
- **Aber:** Ab Schicht 3 wurden in der Internet-Protokollarchitektur lange Zeit nur Peer-to-Peer - Adressen unterstützt! Und im weltweiten Verbund (insbesondere im Internet) muss Multicast auch WAN-Strecken überbrücken.

Frame Address Format in LANs

According to IEEE 802 and ISO 8802

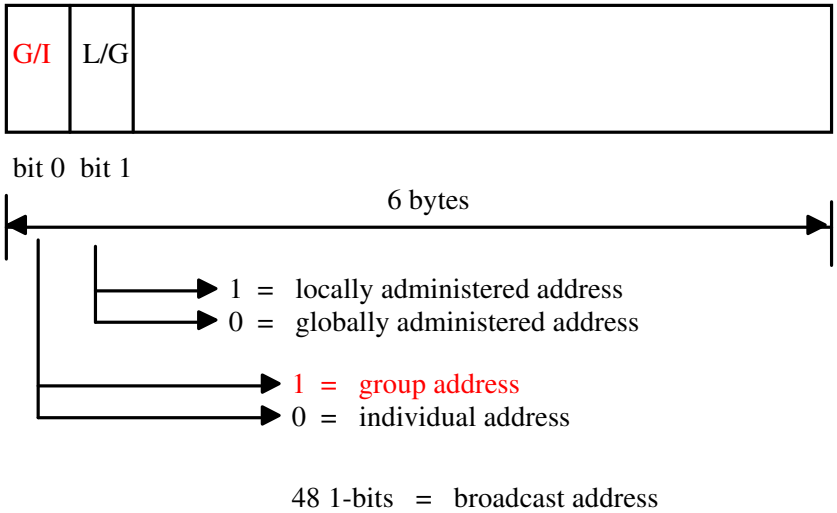

4.3.3 Multicast in WANs

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?

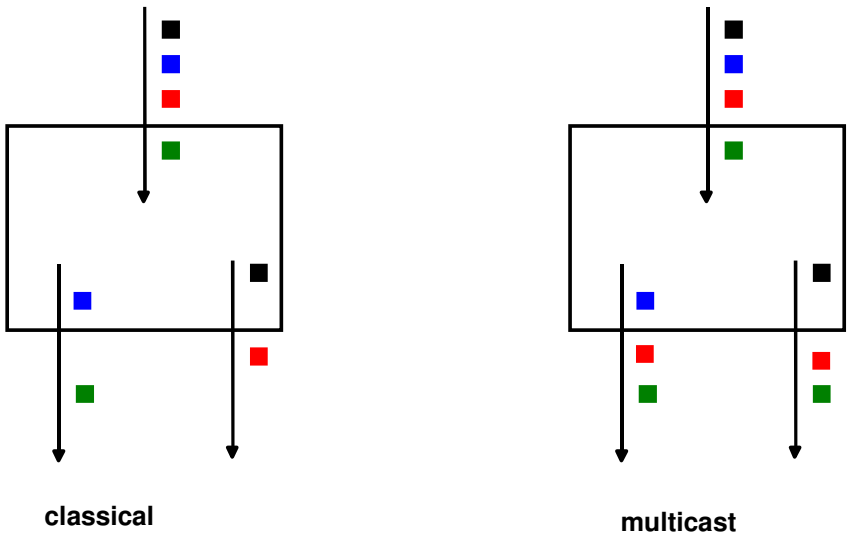

Router with Multicast Support

RT **Routing Table**

From C to	link	cost
g5	{ce, cd}	

Router at node C

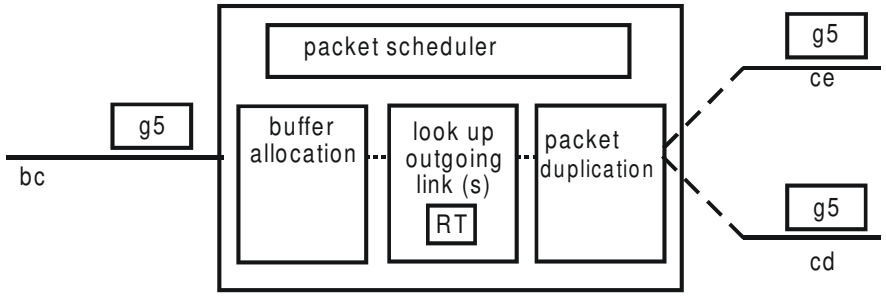

g5 = IP address of group 5

Multicast IP in the Internet

Schon heute gibt es Multicast im Internet. Dies wurde durch eine Erweiterung des IP-Protokolls der Version 4 erreicht (“Multicast IP“). Der Multicast-Backbone **MBone** im Internet ist ein Overlay-Netz multicast-fähiger Router.

Die Protokolle und der MBone sind allerdings als experimentell anzusehen, sie haben noch nicht Produkt-Qualität. Viele Internet Service Provider verzichten auf die Bereitstellung von Multicast, da es schwer zu administrieren ist (Fehlersuche, Accounting).

Prinzipien des Multicast-IP-Protokolls

- Übertragung von IP-Datenpaketen an eine **Gruppenadresse** (IP-Adresse vom Typ D)
- verbindungslos (Datagrammdienst)
- Best-Effort-Prinzip (keine Dienstgütegarantien):
  - keine Fehlerkontrolle
  - keine Flusskontrolle
- **empfängerorientiert**:
  - Der Sender sendet Multicast-Pakete an die Gruppe.
  - Der Sender kennt die Empfänger nicht, hat auch keine Kontrolle über diese.
  - Jeder Host im Internet kann einer Gruppe beitreten.
- Eine Beschränkung des Sendebereiches ist nur durch den Time-To-Live-Parameter möglich (TTL = hop counter im Header des IP-Pakets)

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, scaleable group address management protocol.

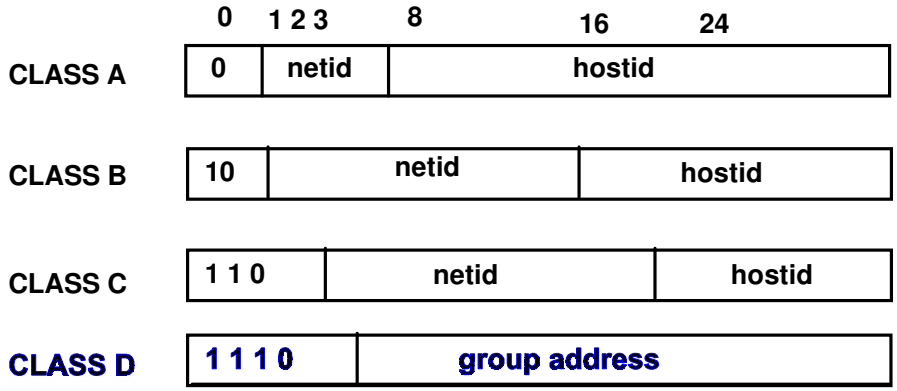

Multicast Routing: Flooding

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 fowarded on **each** outgoing link except the one on which it has arrived.

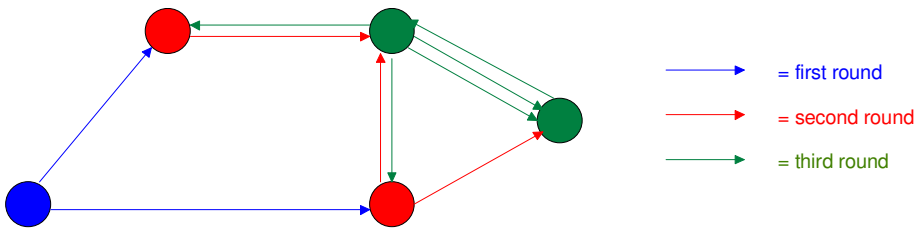

Problem: Infinite Number of Packets

Solution

“hop counter“ in the packet header

- Initialized with the diameter of the network
- Decremented by 1 on each hop
- Duplicate packets get the hop counter of the original packet
- When the hop counter reaches 0 the packet is discarded by the router

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

Example for Reverse Path Broadcasting (still incomplete)

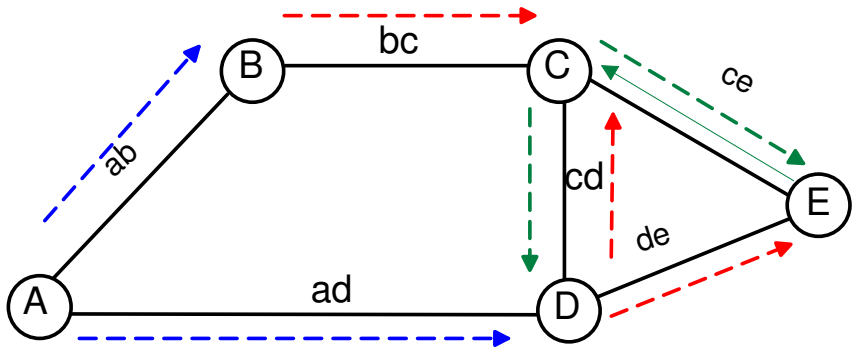

As we can see there are still unnecessary packets: D and E receive the packet twice, C even three times.

Example for Reverse Path Broadcasting (complete algorithm)

When each node sends some extra information to its neighbors RPB can avoid the unnecessary packets. The extra information is the fact whether a link between two neighbors is on the shortest path to the sender.

The complete RPB algorithms then works as shown below.

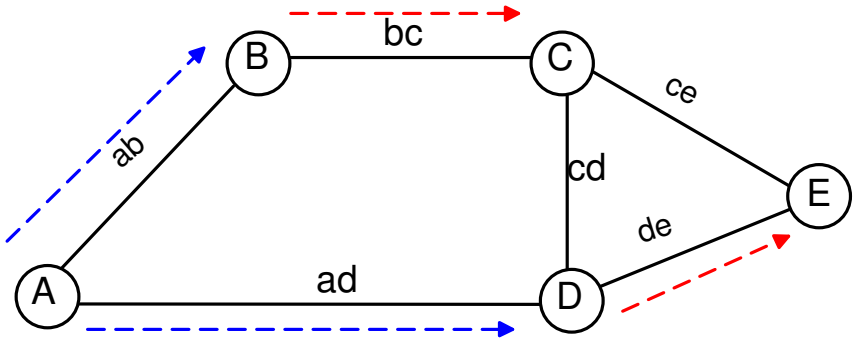

Truncated Reverse Path Broadcasting (TRPB)

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 leafs 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)

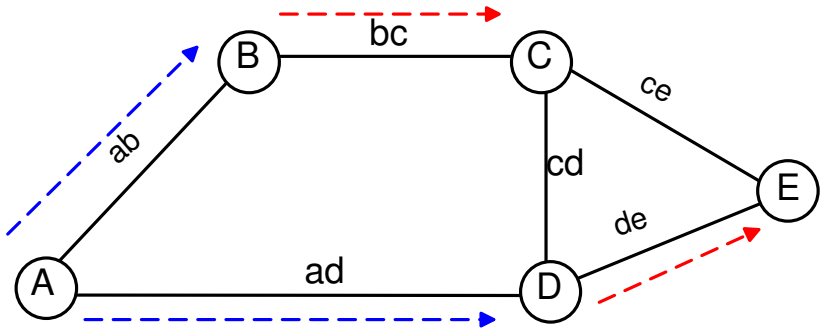

(a) Tree in the initial RPM phase

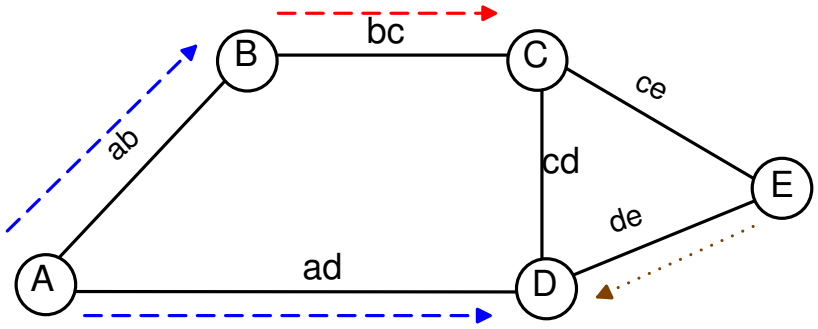

(b) E sends a “prune“ message

Example for Reverse Path Multicasting (2)

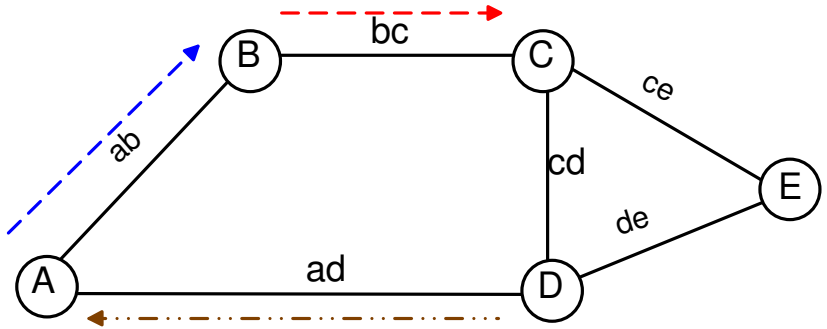

(c) D sends a “prune“ message

Multicast in IP Version 6 (IPv6)

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

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

There is still no consensus on QoS-based routing or resource reservation in routers for the Internet.

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

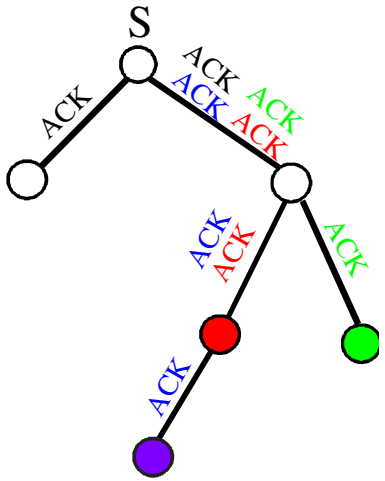

ACK implosion

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"

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, scaleable tree routing algorithms for multicast? Should they be sender-oriented or receiver-oriented?
- What are the best, scaleable 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 scaleable.
- The next generation of Internet protocols will support group addressing and multicast routing, perhaps resource reservation, but not QoS-based routing.