Communication protocols and network security

Multicast

Multicast

- **Ways of addressing:**
  - **unicast** (traditional): transmission to a single destination IP address (unique on the Internet / local network)
  - **broadcast**: addressing "all receivers" in a subnetwork (e.g. looking for a router or server, urgent message); doesn't deliver packets outside the network

- How to transmit only to a selected group of addresses, even outside the local network?
  - **multicast** addressing allows delivery to groups of receivers regardless of the borders of the subnetworks
  - IGMP (Internet Group Management Protocol) is used for managing those groups

Multicast

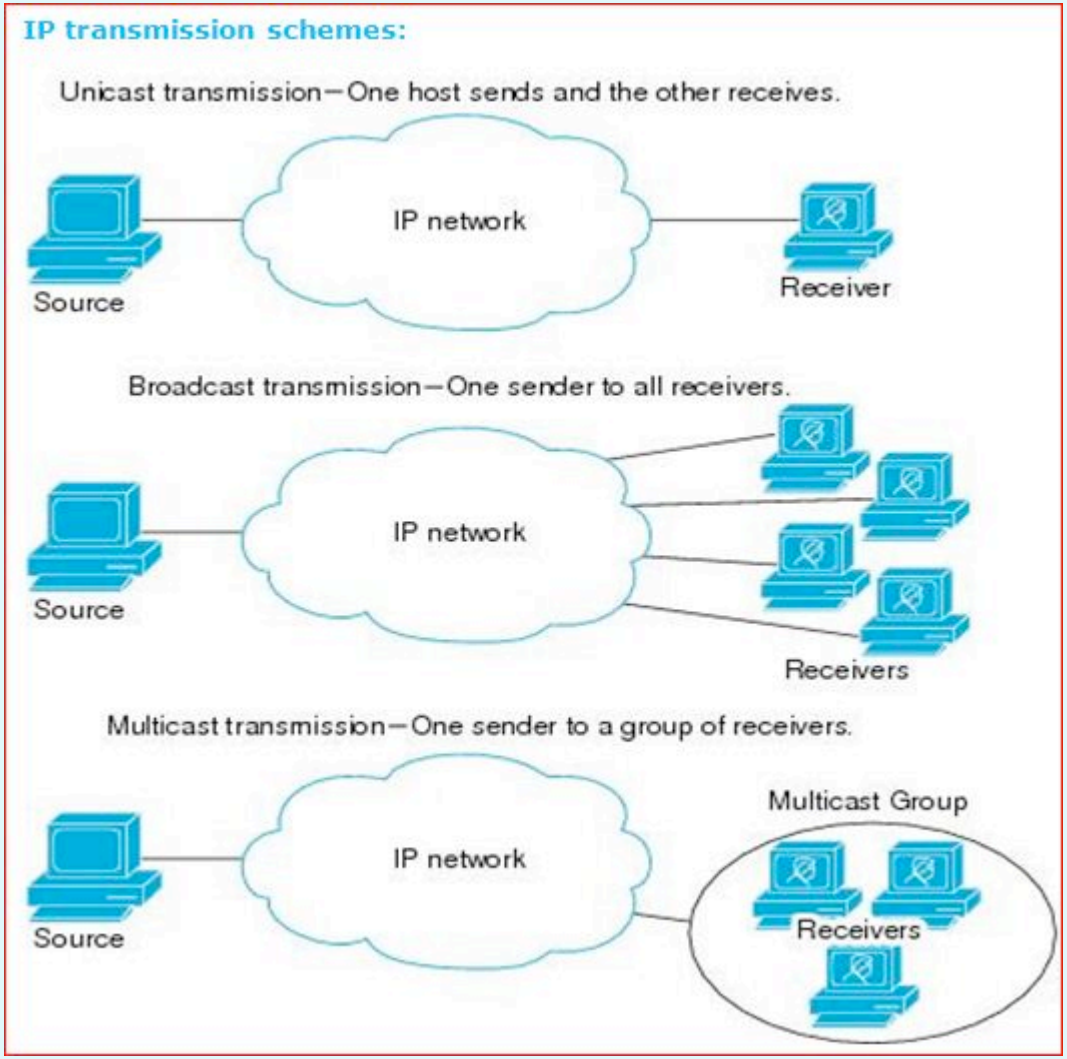
IP transmission schemes:
Unicast transmission—One host sends and the other receives.
IP network
Source
Receiver
Broadcast transmission—One sender to all receivers.
IP network
Source
Receivers
Multicast transmission—One sender to a group of receivers.
IP network
Source
Multicast Group
Receivers

Multicast - example

We want to transmit to 4 of 6 computers in a network. How?

1. ***unicast***: we need 6 copies of the same packet; multiple transmissions can overload the medium.

2. ***broadcast***: address all computers; filtering the right receivers is left up to higher layer protocols.

3. ***multicast***: we transmit to a "special" address representing a GROUP of receivers that listen to the packets targeted at that address
   - similar to *broadcast*: everyone receives the packet
   - but: filtering occurs at the network level - IP (sometimes even at the data-link layer)

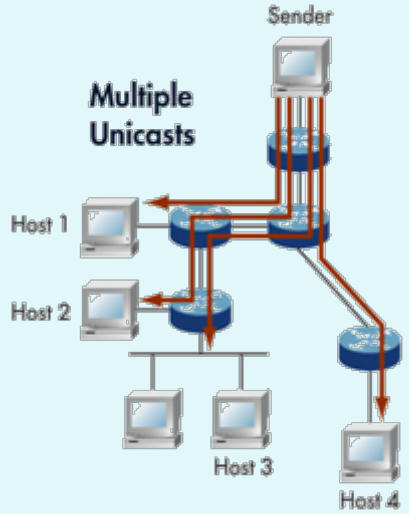

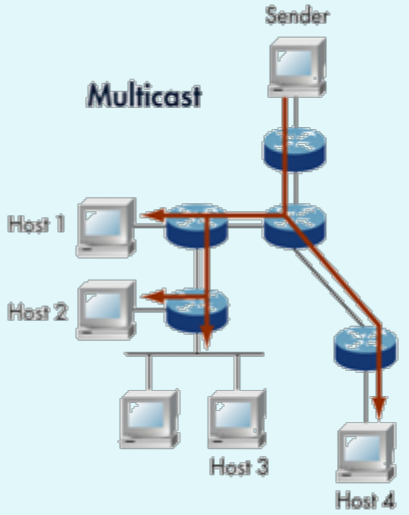

Multicast: packet routing

- *broadcast* packets are not forwarded by routers (everyone would receive them!), meaning they stay inside the local network

- **multicast routing** is practical: a single packet is replicated by the router and forwarded only through those interfaces where there are listeners to that packet. Group names are 32 bit numbers (almost).
- Challenges for the protocol:
  - finding out where the packet receivers are,
  - multicast requires additional work: routing protocols, forwarding information about the listeners,
  - multicast addresses don't form a (sub)network -> the mask has 32 bits. Therefore they require special input in routing tables
    - *challenge: they can also have more special inputs. Why?*
  - security: an eavesdropper (illicit listener) can subscribe to listening the packets and thus become a legitimate receiver
  - what to do when only one receiver signalizes it didn't receive the packet

Multicast

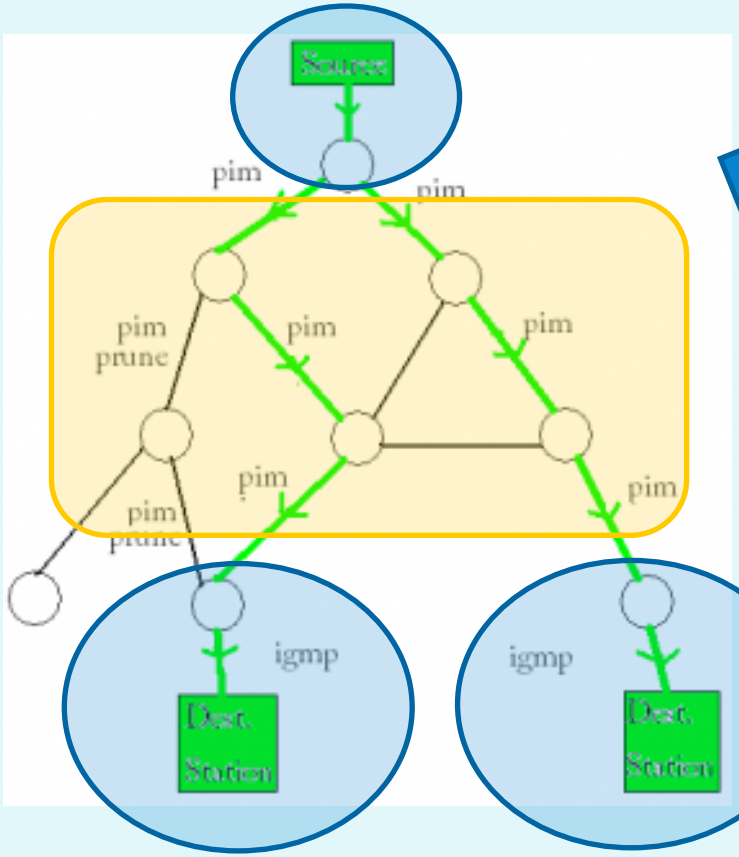

subscribing to multicast traffic (IGMP)

multicast routing (PIM)

Multicast applications

- sending large files over a network (central office to affiliates) - reliable transmission
- software updates in a large network
- data streaming (e.g. sending information about shares to all financial companies
- audio/video streaming
- video on demand (watching a TV program)
- conference holding (consideration: better use of a conference center that decides who can speak and whose packets are to be forwarded to others)
  - *challenge: think about how a conference is held using a multicast approach*
- real-time applications with RTP, which is used for ensuring continuous and quality deliveries within environments that use multicast

IPv4 and IPv6 addressing

IPv4 addressing

- multicast group names are actually specially reserved IPv4 addresses:
  `224.0.0.0 - 239.255.255.255` (class D)
- Special addresses inside that range:

Address range	Description
`224.0.0.0 - 224.0.0.255`	Reserved for **well-known** multicast addresses
`224.0.0.1`	All **systems** (interfaces and routers)
`224.0.0.2`	All **routers**
`224.0.1.0 - 238.255.255.255`	**Globally-scoped** multicast addresses (Internet)
`239.0.0.0 - 239.255.255.255`	**Locally-scoped** multicast addresses (local network)

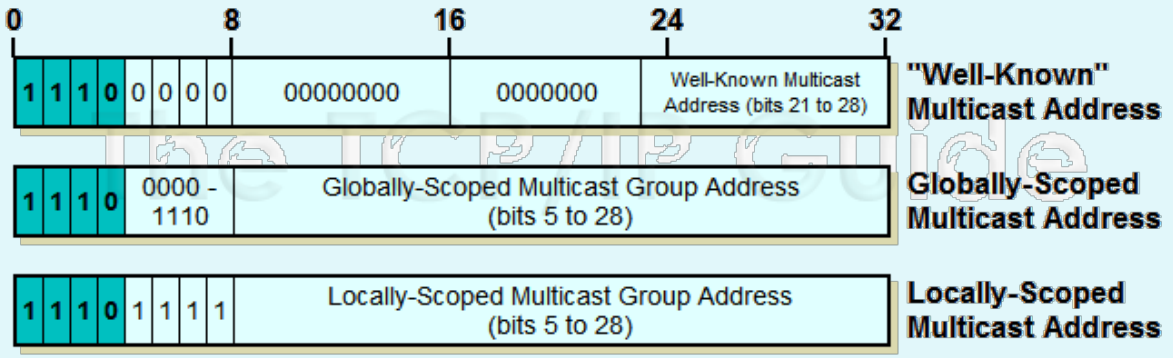

IPv6 addressing

1.) multicast group names are 128 bit numbers - IPv6 address, starting with **FF**

2.) **FF02::1** (link local: all INTERFACES)

3.) **FF02::2** (link local: all ROUTERS)

4.) IPv6 address structure :

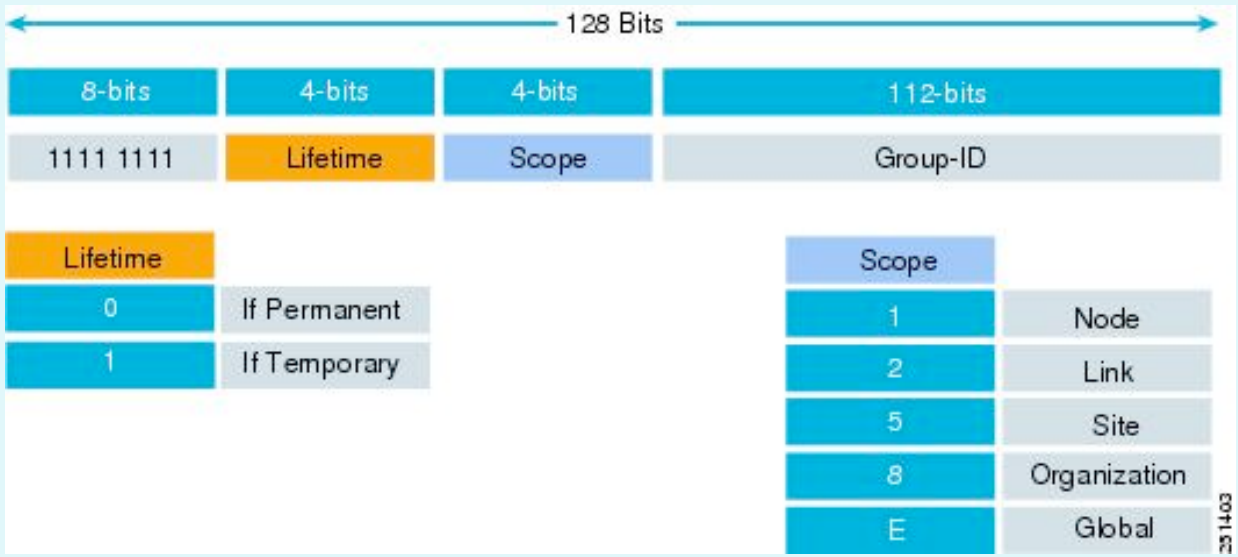

Address mapping

- Ethernet and FDDI frameworks use 48 bit addresses. Multicast group addresses range from 01-00-5e-00-00-00 to 01-00-5e-ff-ff-ff.
- The 01-00-5e prefix represents the multicast frame, the next bit is 0, and the rest 23 bits constitute the name of the multicast group.
- since multicast IP addresses are comprised of 28 variable bits, the mapping isn't unique! Only the 23 low-order bits are mapped to the frame. That means that 32 (25) addresses are mapped to the same address on the second layer.
  - *challenge: what does the router have to do then?*
- The network layer decides whether the datagrams are important for receiving or not.

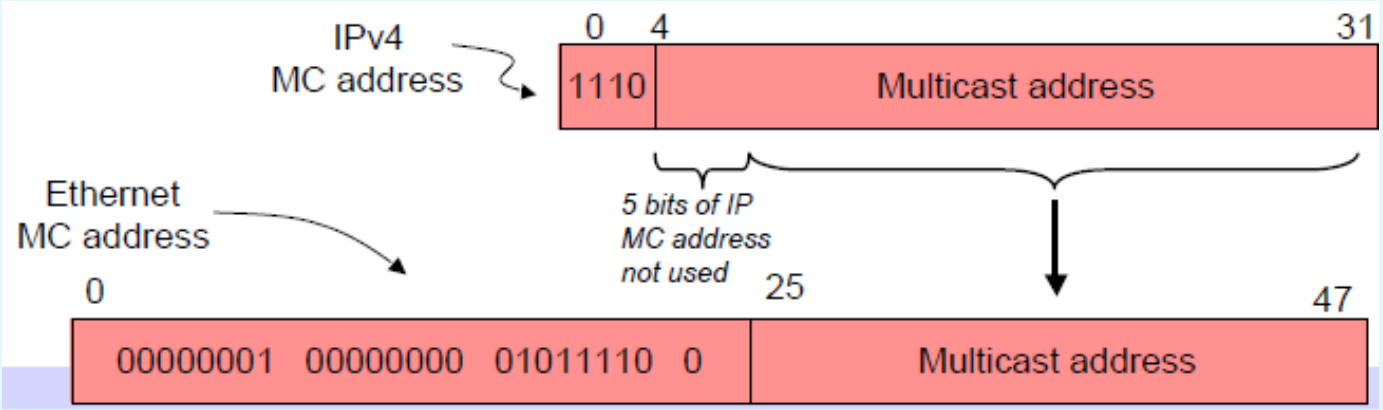

Subscription to multicast traffic: IGMP and MLD protocols

Click here to Enroll Now

IGMP protocol

- network layer protocol IPv4, protocol number 2
- RFC 2236, *Internet Group Management Protocol, Version 2*, RFC 3376, *Internet Group Management Protocol, Version 3*
  - *required: find it on the Internet and read it - further reading!*
  - *challenge: find the other RFC documents related to IGMP*
- IGMP takes care of managing who the multicast receivers are. It allows:
  - establishing group memberships
  - leaving a group
  - detecting other interfaces in the group

IGMP protocol

- the IGMP communication occurs between a host and an immediately-neighboring multicast router

- routers get the task of connecting to the multicast tree structure based on the IGMP protocol

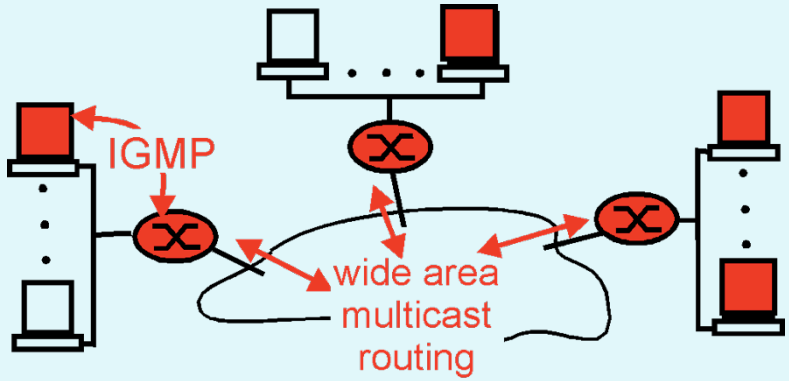

IGMP versions

There are 3 versions: IGMP v1, v2 and v3.

- IGMPv1: Interfaces can connect to groups. There are no messages for leaving a group. Routers use timeouts to detect groups of no concern for the interface.
- IGMPv2: Messages for leaving a group are added. That allows for faster notification about unnecessary traffic termination.
- IGMPv3: Bigger changes in the protocol. Interfaces can determine a LIST of other interfaces from which they wish to receive traffic. The network blocks all traffic from other interfaces.

IGMP protocol

- IGMP messages are 8 bytes long

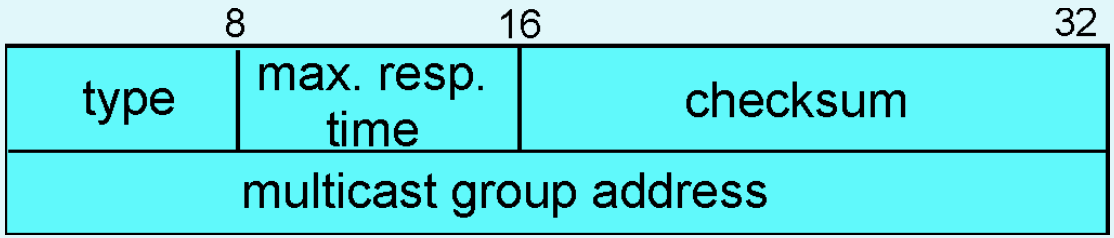

- **type** – type of the message:
  - `17 (0x11): Group Membership Query`
  - `18 (0x12): Group Membership Report IGMP v1`
  - `22 (0x16): Group Membership Report IGMP v2`
  - `34 (0x22): Group Membership Report IGMP v3`
  - `23 (0x17): Leave Group Report IGMP v2`
- **response time** - maximum time allowed for an IGMP Group Membership Query recipient to respond
- **checksum** - (doesn't cover the IP header)
- **multicast group address** - IPv4 address of the multicast group

IGMP protocol

- How to accomplish group management with IGMP

Action	IGMP message	IP Destination Address	IGMP Multicast Group Address
join a group	Group Membership Report	group address	group address
list of group members	Group Membership Query	group address	group address
list of existing groups	Group Membership Query	all interfaces (224.0.0.1)	0.0.0.0
acknowledge being a member of the queried group	Group Membership Report	group address	group address
leave the group	Group Leave Report	all routers (224.0.0.2)	group address

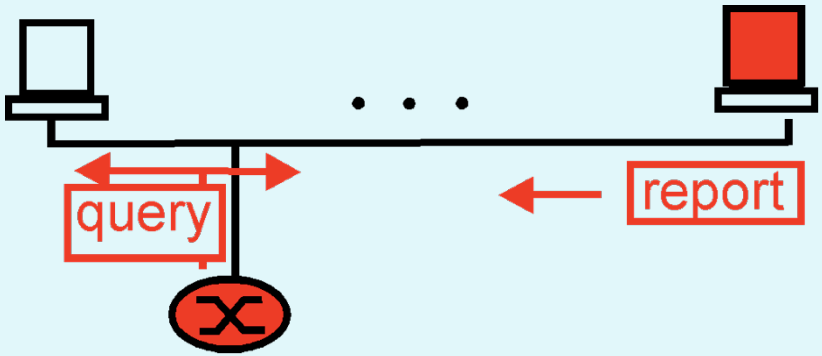

IGMP Protocol

- Special message: IGMPv3 Group Membership report

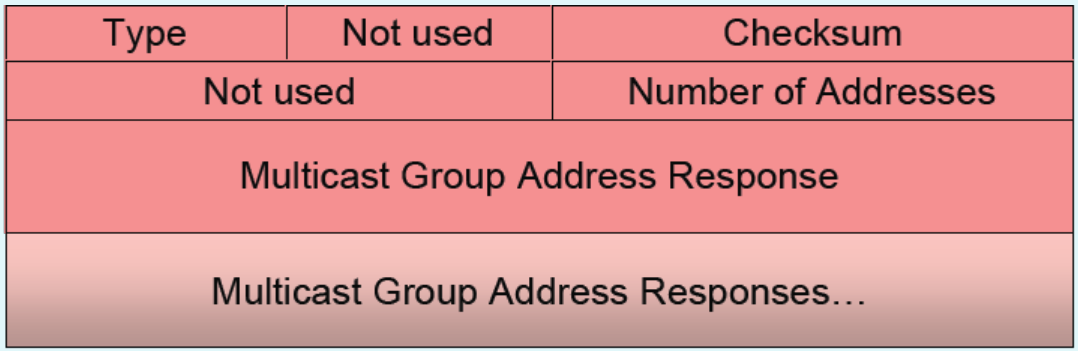

- Type= 0x22
- the responses from all interfaces in the group are in the same packet
- the interface waits for the responses of the other recipients in the group before it answers
  - the special format of the package ensures the avoidance of multiplied multicast traffic

IGMP protocol: subscription to a source

- for joining a group, a GMR message is sent with value TTL=1 (delivered only to the nearest router)
- the router recognizes that it must forward the group packets to the new subscriber (how? mapped multicast address / datagram copies to the IP address)

- the router informs the neighboring routers that it has a new subscriber. If every router were to pass the same message onwards, there would be a problem - the packets would be cross-forwarded across all connections in the network. Solutions:
  - **use of the RPL algorithm** (Reverse Path Lookup): we reject all multicast packets coming from routers that don't connect to the source of the packet through the shortest path
  - **routers have special routing protocols** for multicast traffic: e.g. the PIM-SM protocol (Protocol Independent Multicast - Sparse Mode)

Reverse path lookup: example

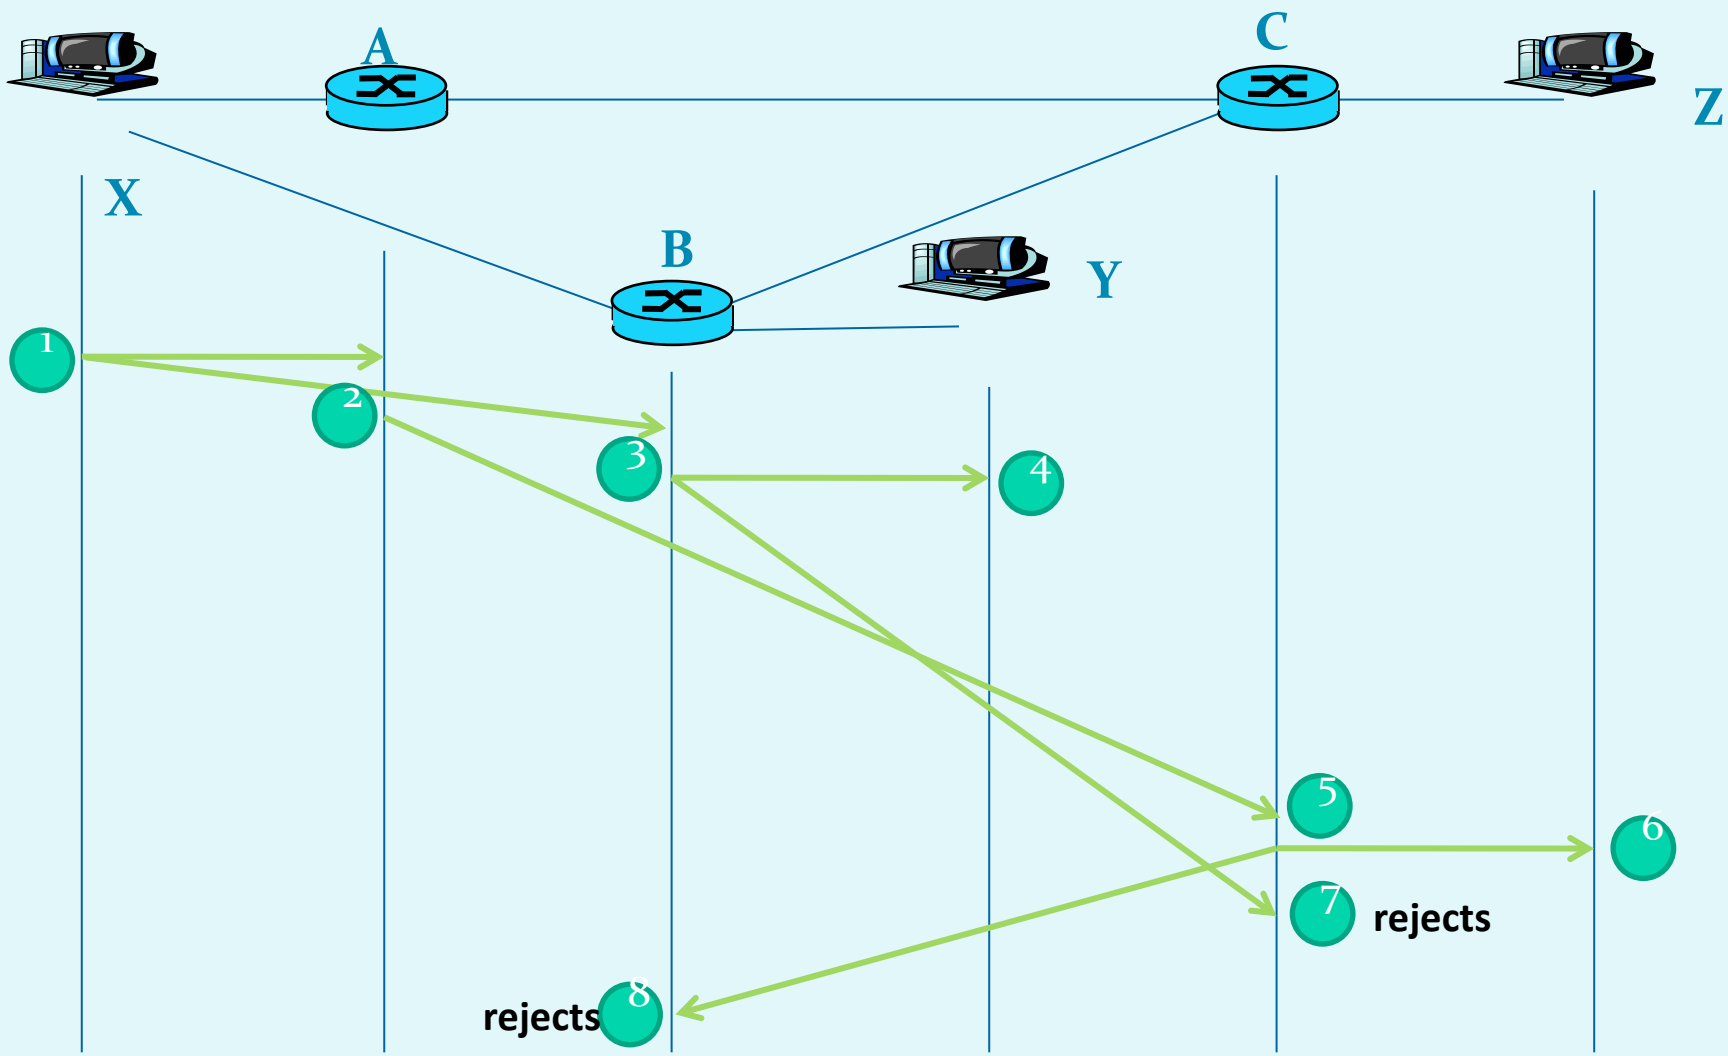

MLD protocol

- Multicast Listener Discovery, RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
  - *required: find it on the Internet and read it - further reading!*
  - *challenge: look for the differences between MLD and IGMP*
  - *challenge: what about the coexistence of IGMP (IPv4) and MLD (IPv6)?*

MLD protocol

- Basically it's a multicast protocol for IPv6 and has the same functionality as IGMP

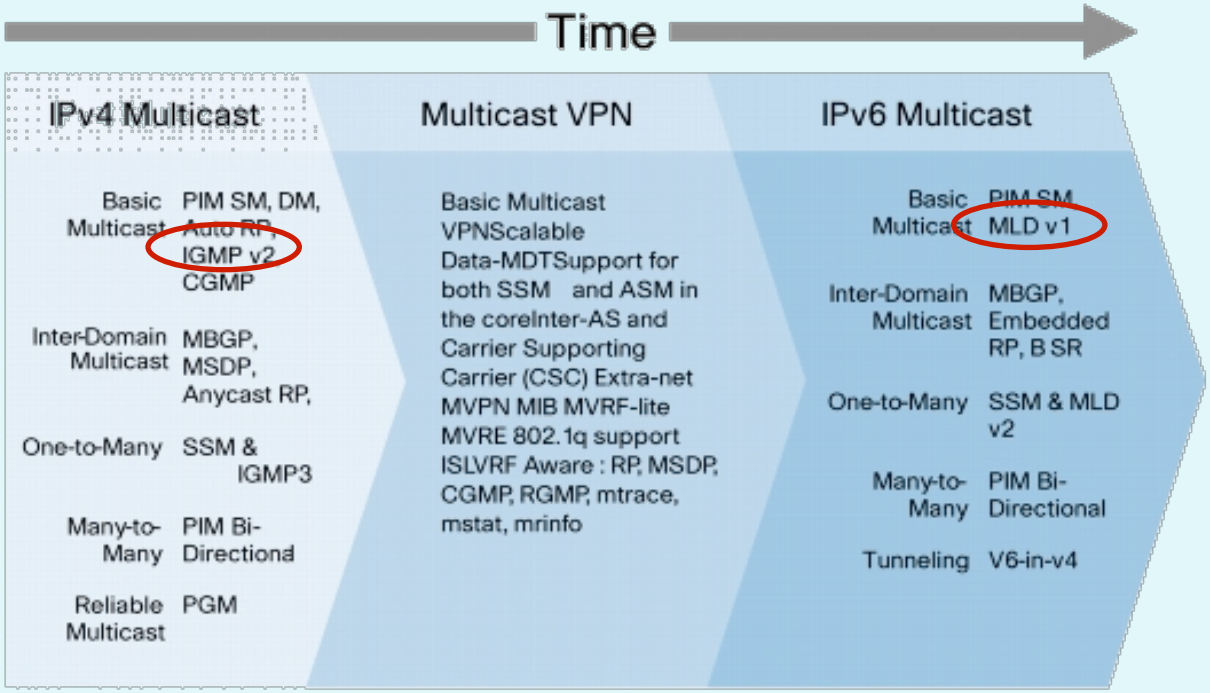

IGMP and MLD protocol

- MLD:

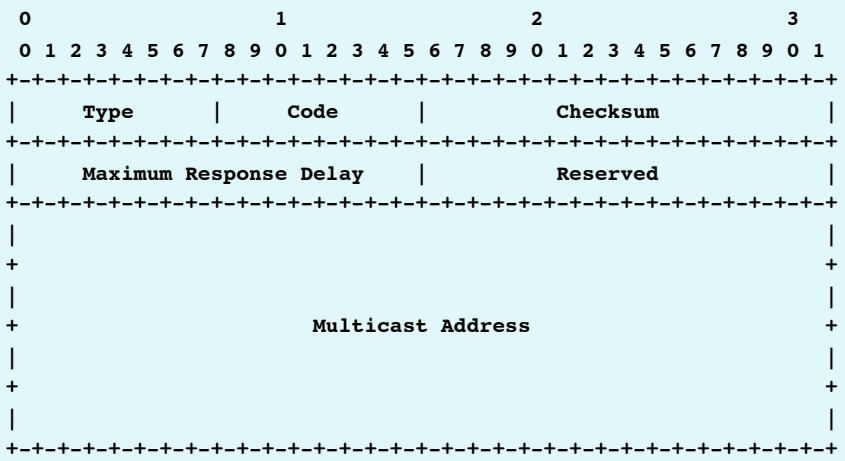

- IGMP:

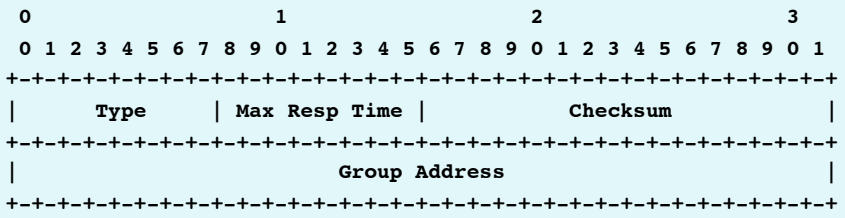

Multicast trees

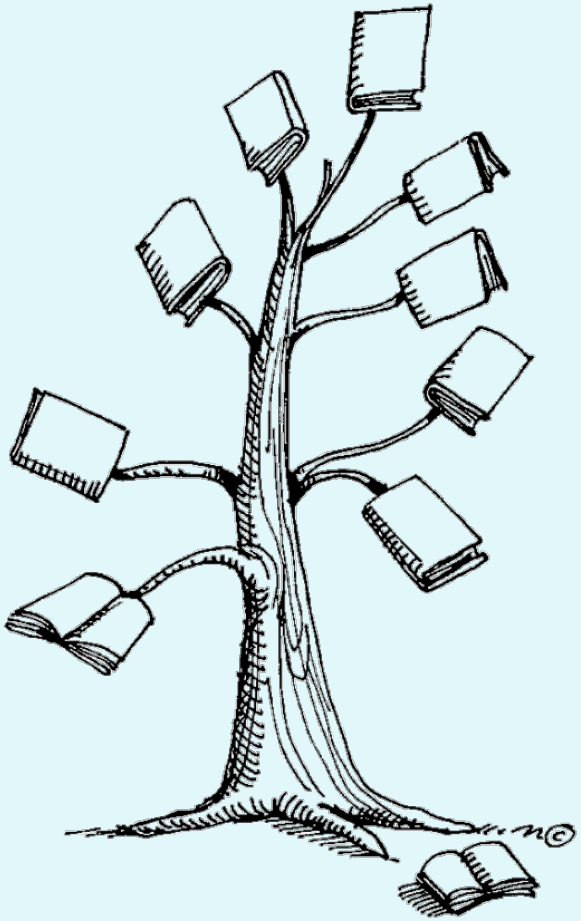

Traffic multicast

- packets move in the form of multicast trees
- a tree can optimize different criteria :
  - figure 1: total path length (number of hops) of all datagrams
  - figure 2: shortest path for every datagram separately (minimum spanning tree)

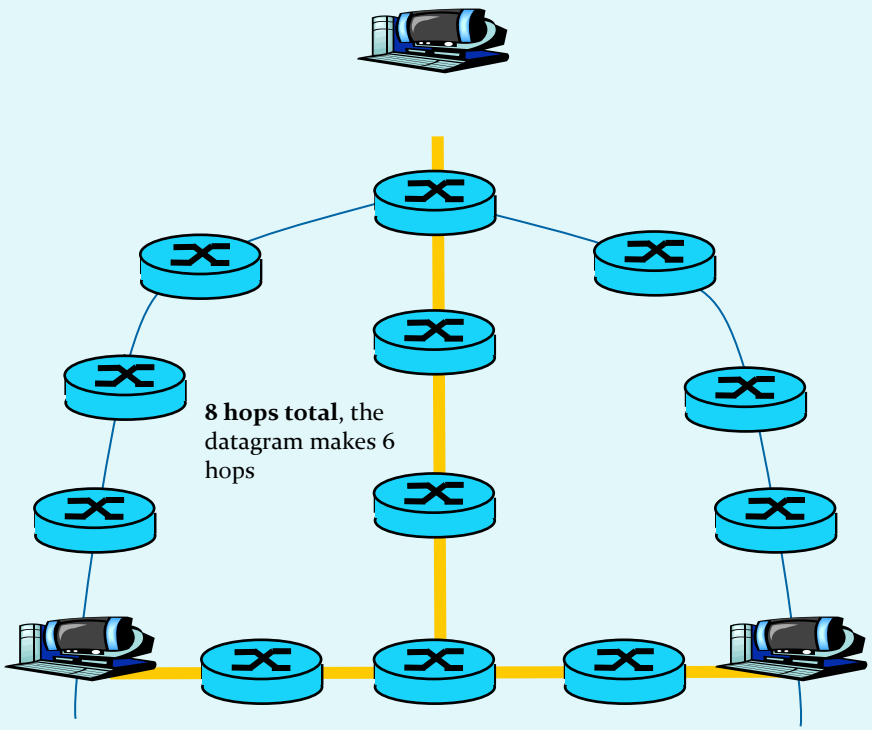

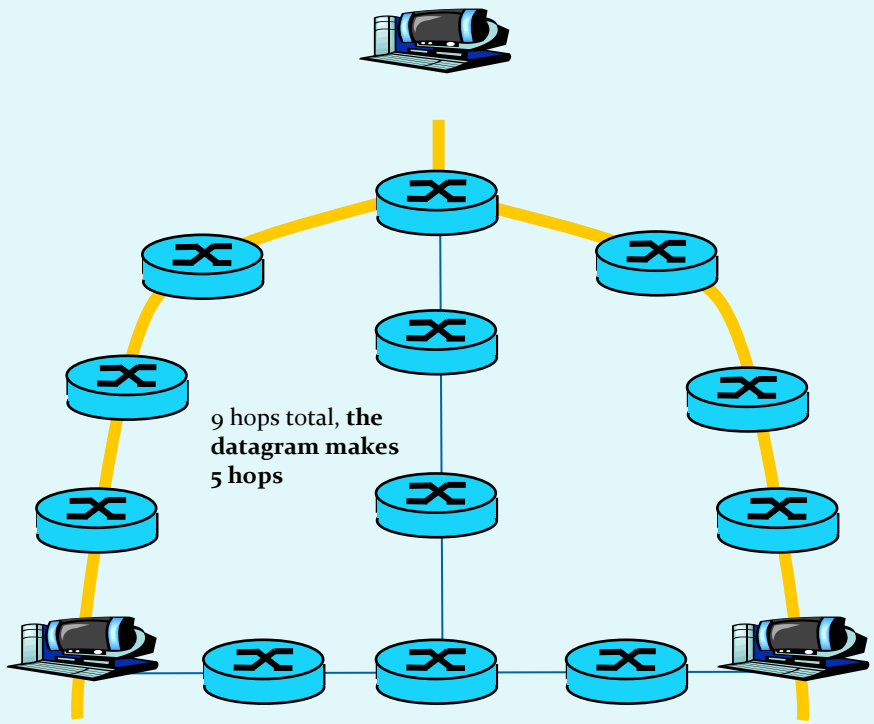

Multicast routing

- Duty of the routing: find a tree of connections that connects all routers in the same multicast group
- For communication between routers we need multicast routing algorithms (working at the network layer), like: PIM, DVMRP, MOSFP and BGP.

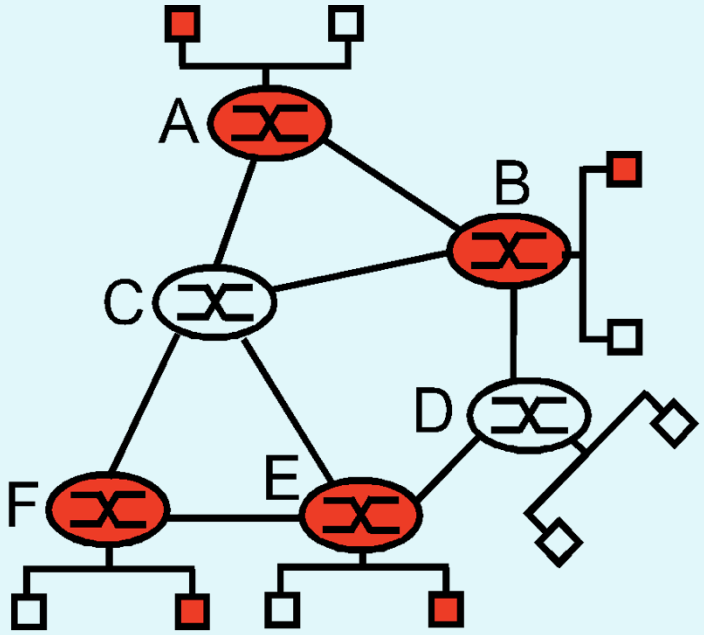

how to connect the red routers into a common tree?

Two solutions for finding a multicast tree

- using a single tree for all routers for routing multicast traffic, we find a single tree (***group-shared tree***) - left figure
- determining a separate tree for every member in the group (***source-based tree***); for N members of the group we have N trees (for every multicast group) - right figure

common tree

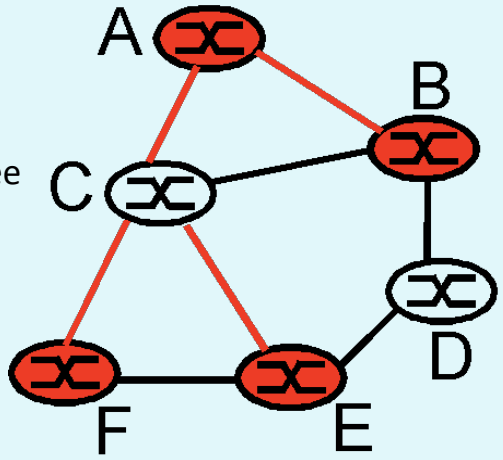

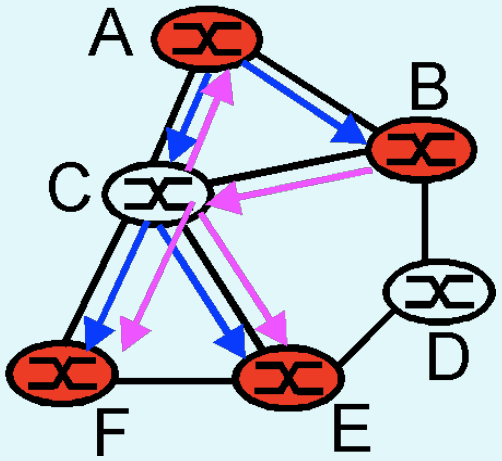

separate tree for A (blue) and tree for B (pink)

Determining a common tree (*group-shared*)

1. finding a tree with the **minimal total cost** (Steiner's algorithm is used for spanning trees; the problem is NP-hard), *left figure*
   or
2. determining the **central node** ("*rendez-vous point*") (the unicast routing rules define how to route to it); the router joins the tree when, on its way to the central node, it encounters the first node already in the tree, *right figure*

minimal total cost

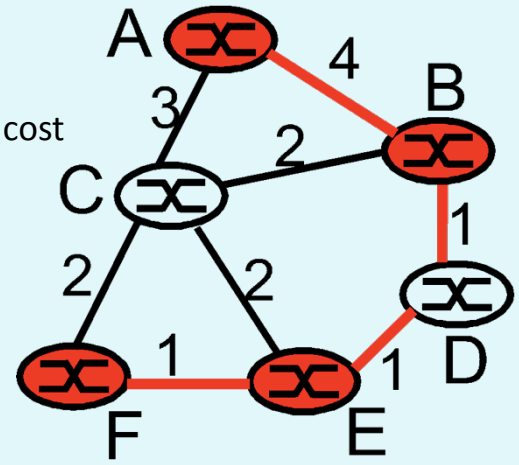

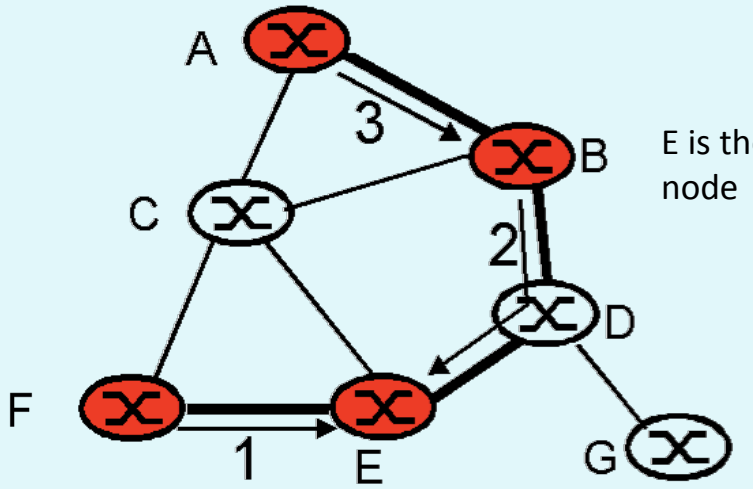

E is the central node

Determining trees for separate senders (source-based)

1. Finding the **shortest path tree** in a graph (using Dijkstra's algorithm which constructs a tree of the shortest connections (edges) relative to a given starting node), *left figure*
   - the routers have to know the states of all connections (edges) (*link-state*)

   or
2. Using **RPL** (***Reverse Path Lookup***): doesn't accept messages from routers that aren't on the shortest path to the source of the message, *figure right*

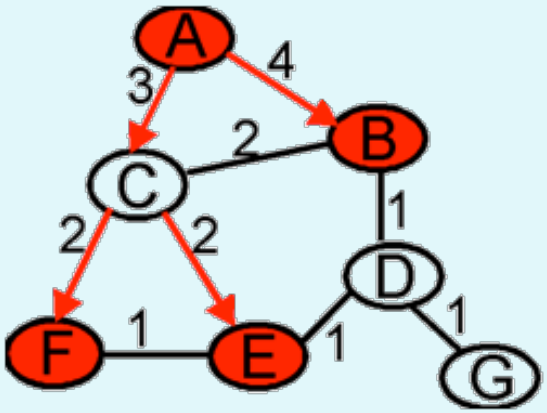

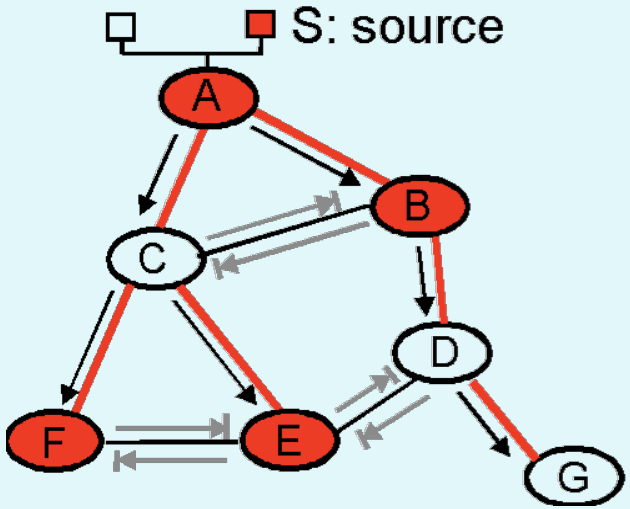

Multicast routing

Routing protocols

- they handle the communication between routers in a network
- divided by 2 criteria (2x2=4 groups)
  1. *sparse-mode* / *dense-mode*
     - *sparse-mode*: certain nodes require inclusion to the tree (*pull* principle)
     - *dense-mode*: we multicast the multicast packets around the entire network, and routers are cut off if they're unneeded (push principle). Two ways:
       - *broadcast and prune* (uses *prune* and *graft* messages): the structure is periodically reinitialized
       - *domain-wide* reporting (routers register clients on the traffic with broadcasting)
  2. intra (within a domain) / inter-domain (between domains)

Routing protocols

- there is a connection between the operation mode and the type of tree built by the protocol

Protocol	Mode	Tree type	Type
PIM-SM	sparse	common	intra and inter-domain
PIM-DM	dense	separate	intra-domain
CBT	sparse	common	intra and inter-domain
MOSPF	dense	separate	intra-domain
BGMP	dense	separate	intra-domain
DVMRP	dense	separate	intra and inter-domain

PIM-SM (*Protocol Independent Multicast - Sparse Mode*)

- PIM-DM: dense-mode, separate tree
- PIM-SM: sparse-mode, common tree, occasionally separate
  - *challenge: read RFC 4601 and study it*

- protocols PIM-SM and PIM-DM are suitable for routers that are already running unicast routing. They are independent from the unicast protocol
- messages use the IP network protocol with protocol number 103
- messages between routers are *unicast* or *multicast* to the address 224.0.0.13 (all PIM routers)

PIM-SM operation

architecture establishment

- candidate bootstrap routers (c-BSR) announce their presence (**message type BOOTSTRAP**) and designate the main bootstrap router BSR
- candidate central (*rendezvous*) routers (c-RP) announce their presence to the BSR router (**message type CANDIDATE-RP-ADVERTISEMENT**)
- The BSR designates the final rendezvous router (RP) for every group and announces them with **messages of type BOOTSTRAP**

data transmission

- PIM-SM routers detect each other and maintain communication with **HELLO messages**
- the interface that sends information to the group address multicasts a datagram to the local segment of the network
- the designated router on the network encapsulates the datagram in a **REGISTER type message** and sends it to the RP
- the RP decapsulates the datagram and multicasts it across the multicast tree

subscription maintenance

- when the RP detects there are no more recipients in the group, it sends a **REGISTER-STOP message** to all designated routers
- when a new user wants to join a group, it sends a **JOIN/PRUNE message** with a list of all desired groups and allowed recipients prejemnikov

Packet format - header content

- The header is 32 bits long
- version = 2
- type:

value	meaning
0	hello
1	register
2	register stop
3	join/prune
4	bootstrap
5	assert
6	candidate-rp-advertisement

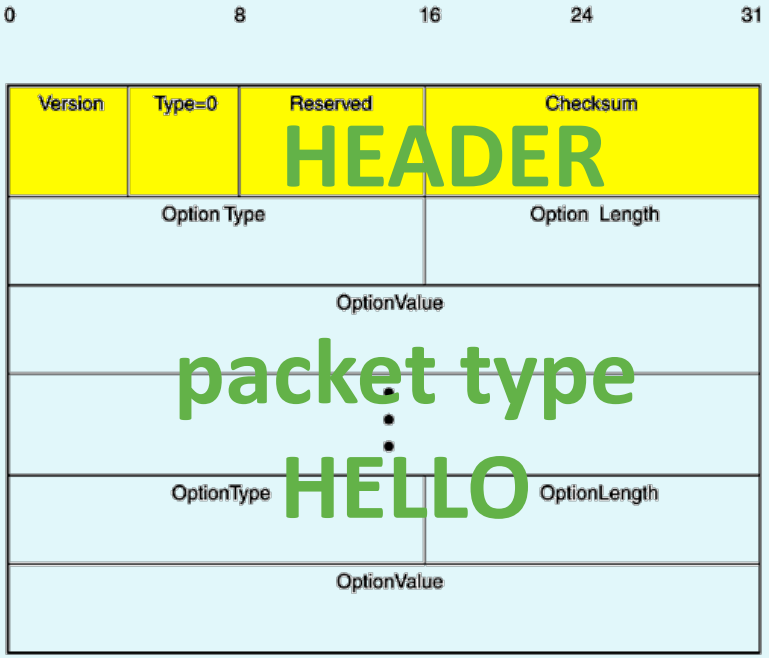

PIM-SM packet format - HELLO packet

- intended for connection maintenance between routers
- if the router designated for transmitting multicast traffic doesn't respond, another one is designated
- the packet contains a set of TLV values, such as e.g. timeout

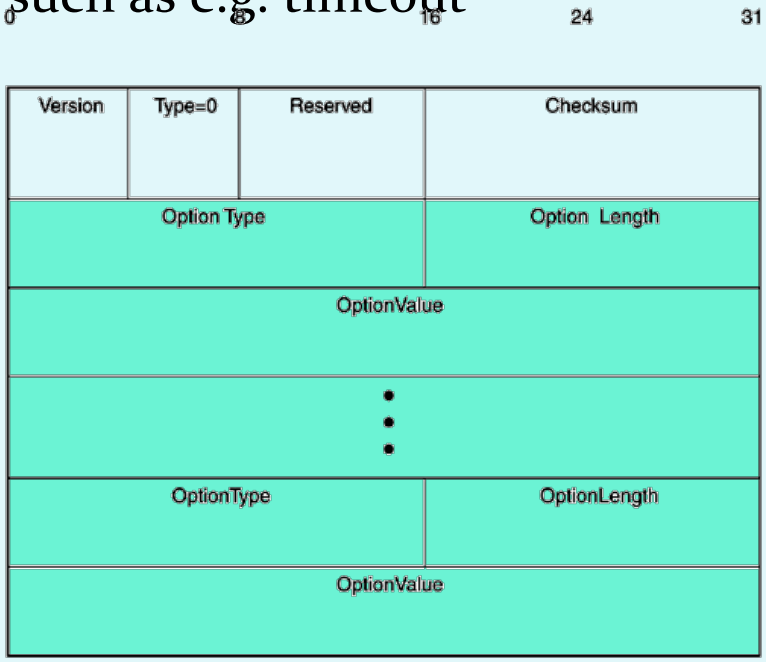

PIM-SM packet format - REGISTER and REGISTER-STOP

- the REGISTER message carries the contents of the multicast message to the central route (unicast)
  - B (border router) - the message reached the router directly connected to the interface,
  - N (null) - the packet is empty, for establishing a link

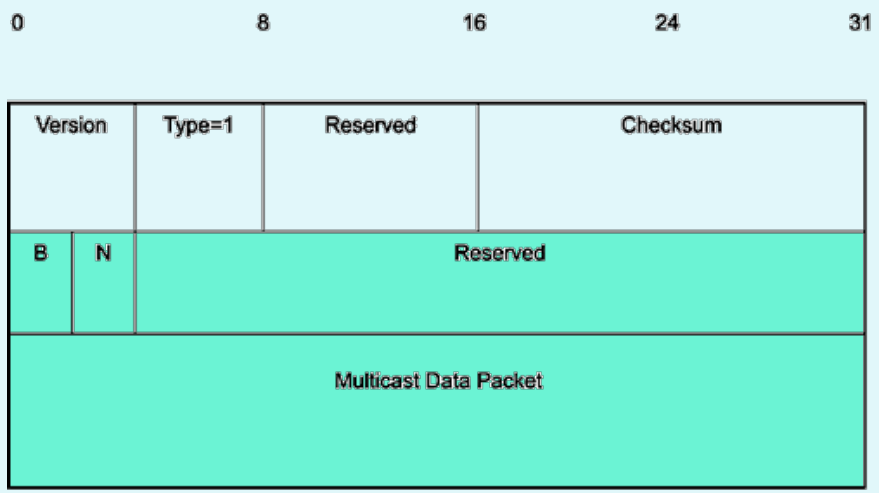

- the REGISTER STOP message is sent by the rendezvous router to the designated router to signal not to send any messages (no recipients / already receiving messages from elsewhere)

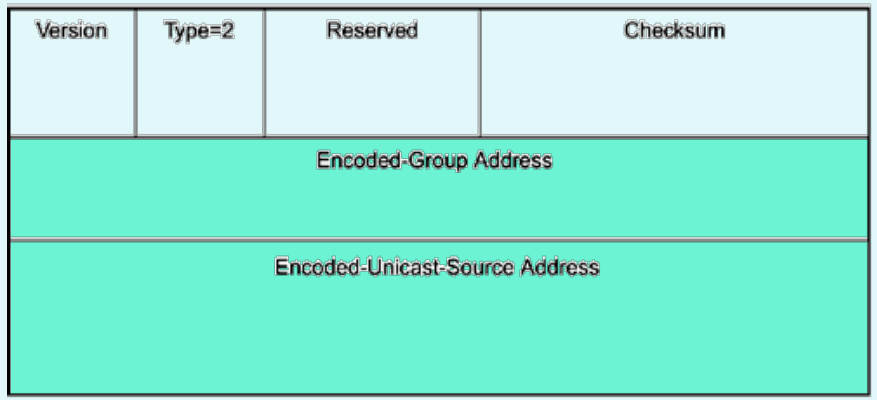

PIM-SM packet format - JOIN/PRUNE

- allows the host to (un)subscribe to receiving multicast traffic
- PIM-SM's Number of Pruned Sources is 0 (because it uses a common tree)
- (Un)subscription fields:
  - Encoded Join Source Address
  - Encoded Pruned Source Address

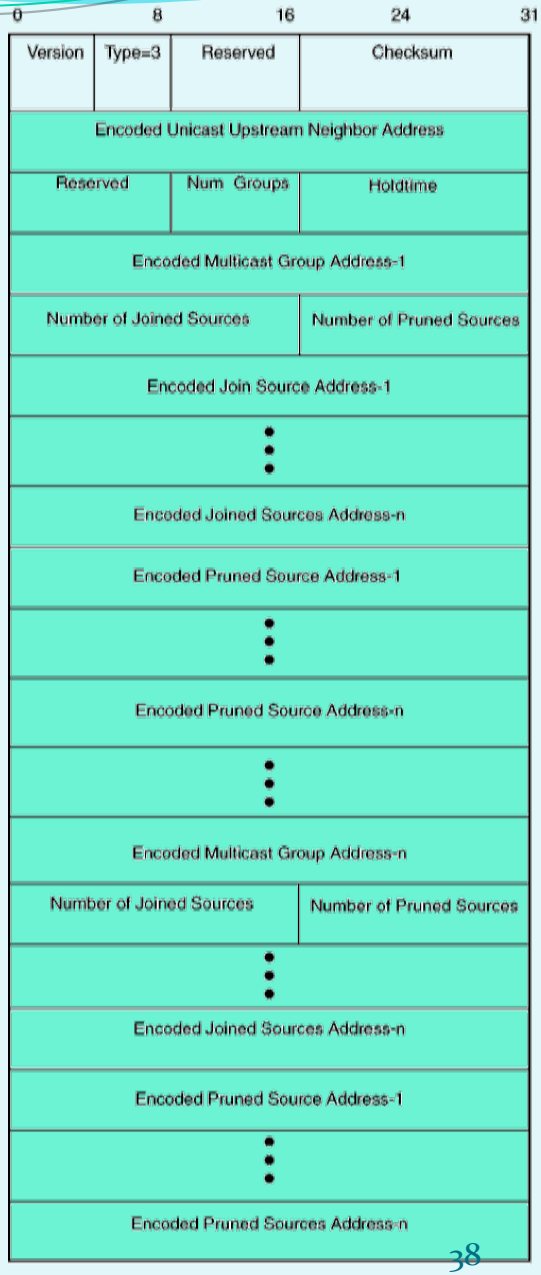

Other routing protocols

- MOSPF
  - Multicast OSPF
  - added: a special packet format that shares information about multicast traffic
    - *challenge: find the RFC documents that describe MOSPF and read them!*
- DVMRP
  - Distance Vector Multicast Routing Protocol
  - transmitted through IGMP packets (type 13)
    - *challenge: read RFC 1075 and study how the protocol works*

MBONE

- connections between networks capable of multicast traffic
  - at first inside the Internet, used by workstations with virtual connections
    - *challenge: read RFC 2715*
  - 1995: MBONE contains 901 routers (DVMRP is used) and it's present in 20 countries
  - 1999: 4178 routers, increased use of RTP, service providers become overloaded
  - IETF sets up the MBONE task force with the task of establishing multicast routing across the entire Internet (development of the MSDP protocol: Multicast Source Discovery Protocol)
    - *challenge: read RFC 1112, what is Any Source Multicast architecture (ASM)?*

We continue next time!

- authentication, authorization and accounting - AAA!