IP Multicast

Jean-Yves Le Boudec

2019

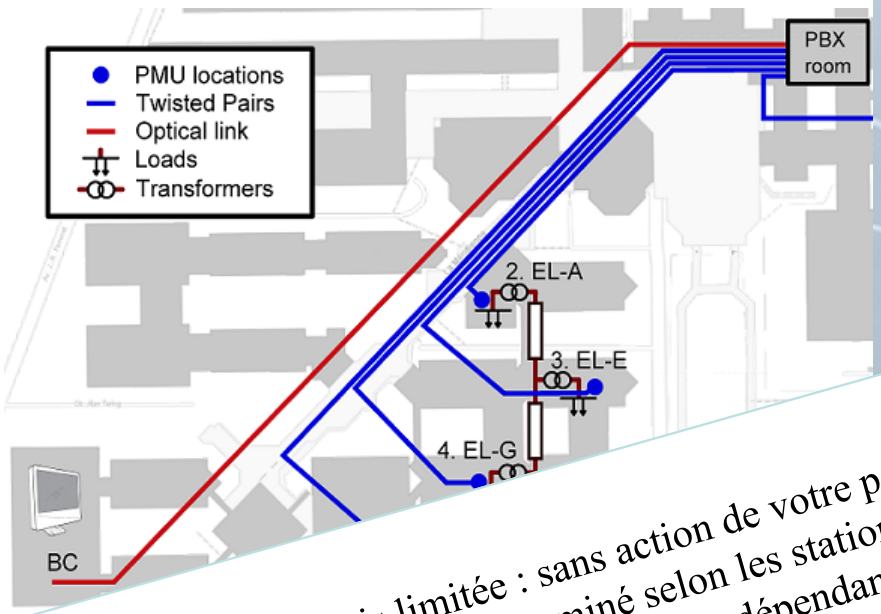

EPFL

La durée d'écoute est désormais limitée : sans action de votre part (un simple clic), la diffusion s'arrête au bout d'un temps déterminé selon les stations. En effet, pour nous, diffuseurs, les technologies actuelles imposent un coût dépendant de la durée et du nombre d'auditeurs. Plusieurs éléments nous indiquent que les internautes ayant accès à l'internet illimité ne coupent pas l'écoute, lorsqu'ils quittent leur ordinateur allumé. Radio France ne peut continuer à financer pour celui qui n'écoute pas. C'est pourquoi nous avons mis en place ce système de confirmation, un peu contraignant, mais qui nous permet de mieux contrôler les coûts de diffusion.

http://viphttp.yacast.net/V4/radiofrance/fip_bd.m3u

IP Multicast

Unicast = send to one destination

Multicast = send to a *group* of destinations

IP has multicast addresses:

224.0.0.0/4 (i.e. 224.0.0.0 to 239.255.255.255) and ff00::/8

For IPv6, bits 13-16 = scope, e.g. ff02/16 = link local, ff05 = site local

An IP multicast address is used to identify a group:

Any Source Multicast (ASM): the group is identified by the multicast address. Any source can send to this group.

Source Specific Multicast (SSM): the group is identified by (s, m) where m is a multicast address and s is a (unicast) source address. Only s can send to this group.

By default 232.0.0.0/8 and ff3x::/96 are SSM addresses (x=scope bits. e.g. ff35::/96 = site-local). See RFC7371.

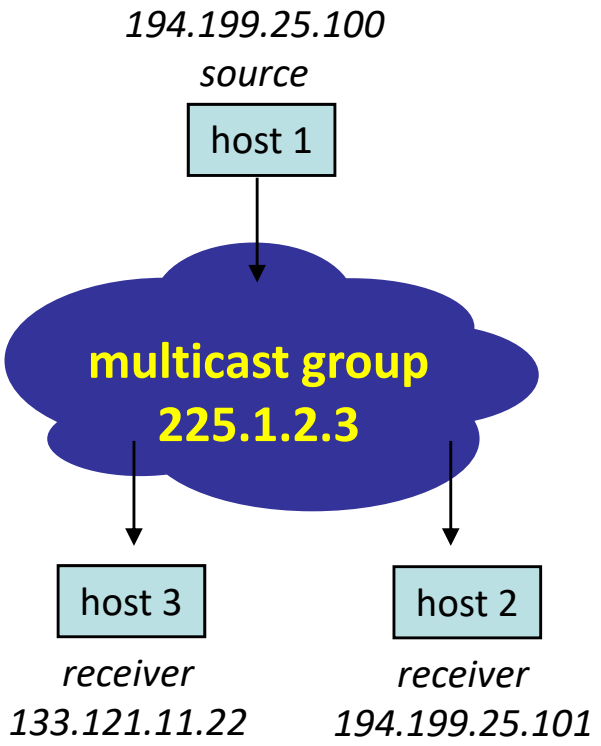

Operation of IP Multicast: destinations need to explicitly join multicast group

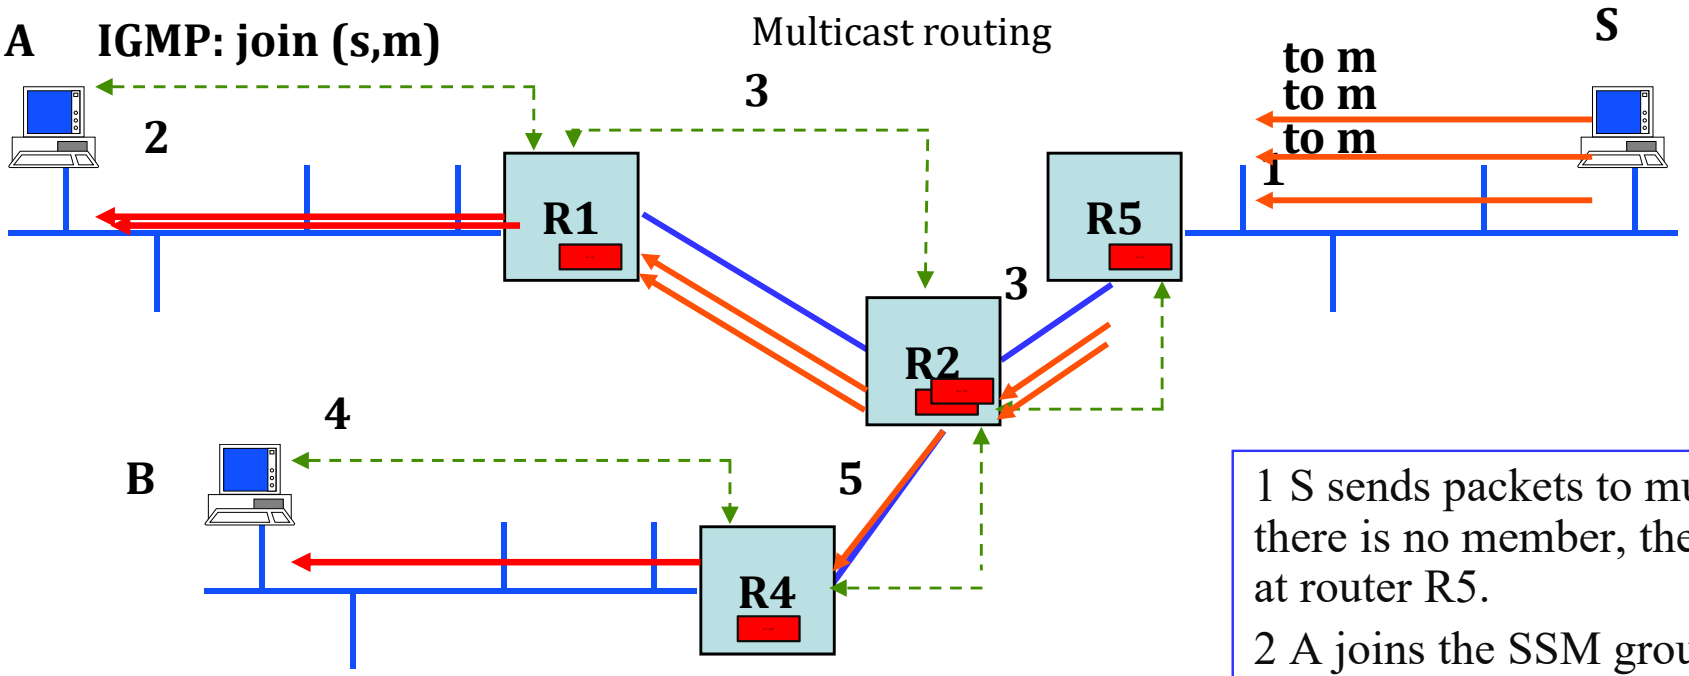

source simply sends one single packet for *n* destination

Destinations subscribe via IGMP(Internet Group Management Protocol, IPv4) or MLD (Multicast Listener Discovery --IPv6); join messages sent to router

routers build distribution tree via a multicast routing protocol (PIM-SM) or by other method

packet multiplication is done by routers

1 S sends packets to multicast address m; there is no member, the data is simply lost at router R5.

2 A joins the SSM group (s,m).

3 R1 informs the rest of the network that (s,m) has a member at R1 using a multicast routing protocol e.g. PIM-SM; this results in a tree being built. Data sent by S now reach A.

4 B joins the multicast address m.

5 R4 informs the rest of the network that m has a member at R4; the multicast routing protocol adds branches to the tree. Data sent by S now reach both A and B.

Multicast enabled Routers Must Keep Additional State Information

In addition to IP principles #1 and #2, an IP router does exact match for multicast groups.

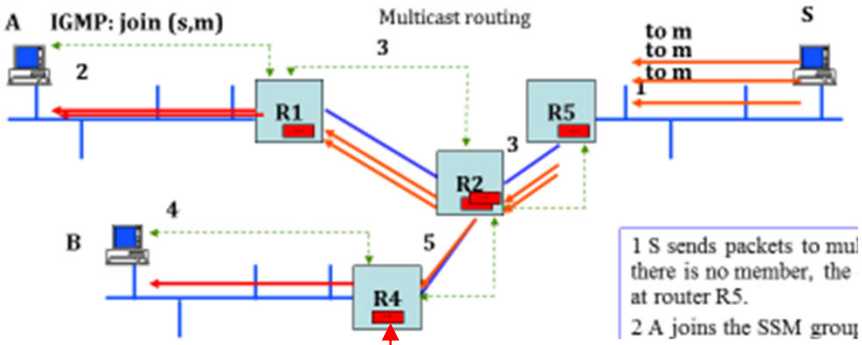

Multicast state information is kept in router for every known multicast group:

(s, m) or (*,m)	// id of group
valid incoming interfaces	// for security
outgoing interfaces	// this is the routing info
other information required by multicast routing protocol	

Multicast forwarding table entries cannot be aggregated

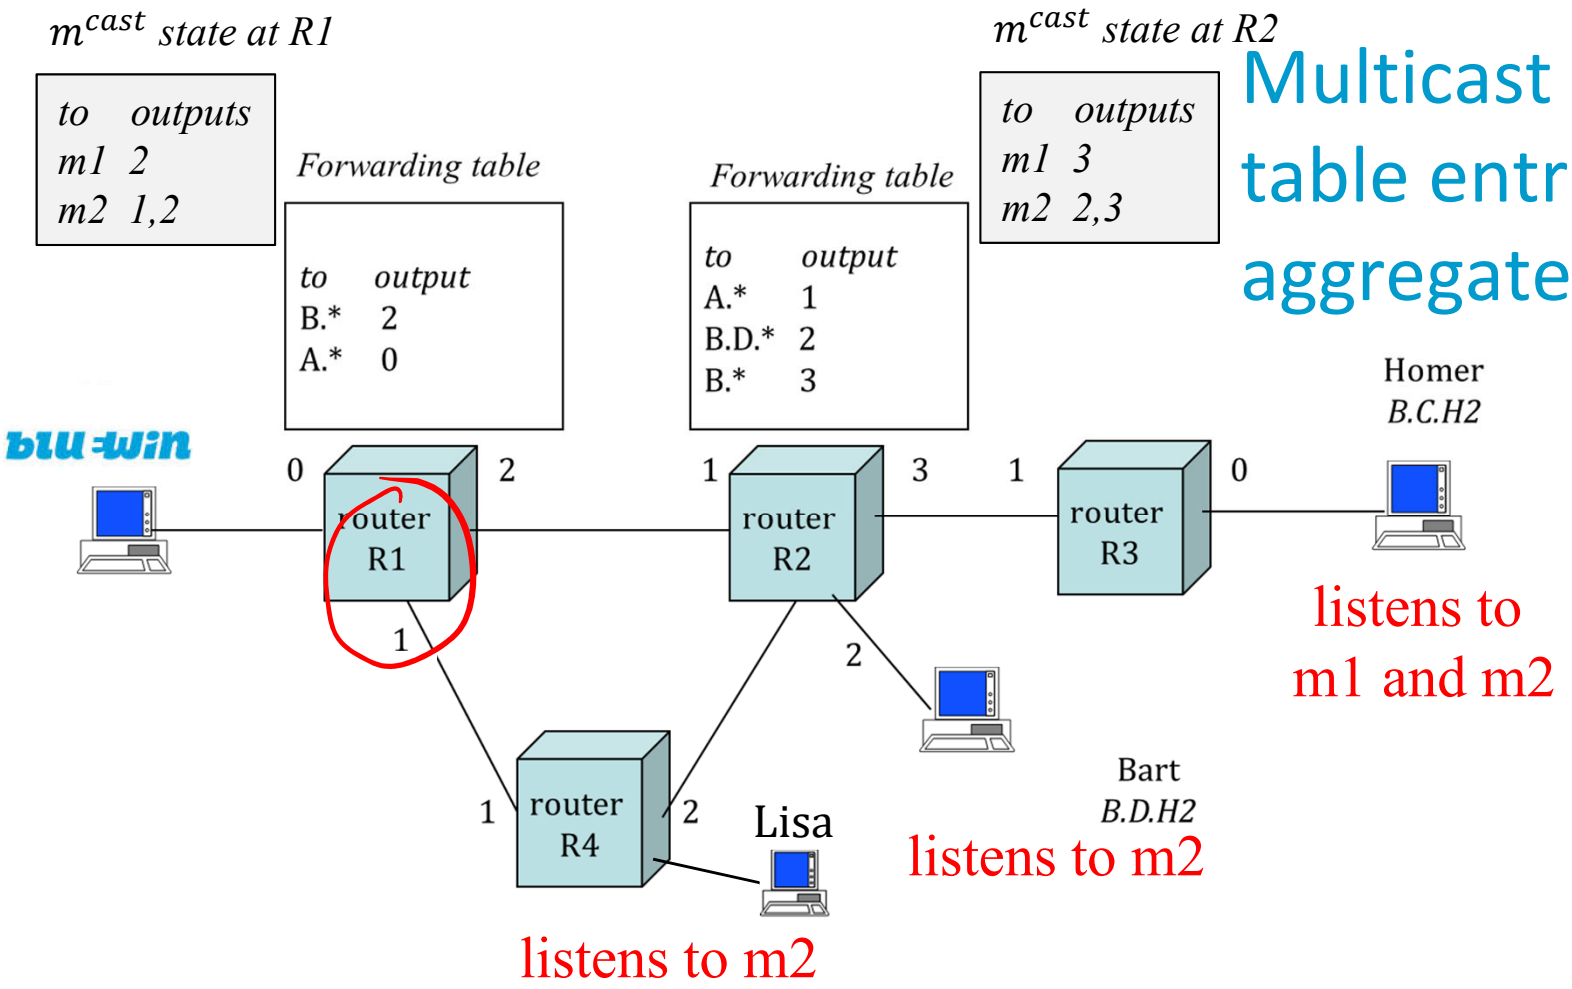

Unicast: R1 has one single entry for all addresses starting with *B*

Multicast: R1 needs the explicit list of all interfaces on which there is a listener, for every multicast address – since the location of listeners depends on applications and users, not on the network topology.

Multicast addresses are purely logical – no topological information

Is there Multicast ARP ?

Recall ARP = find MAC address that corresponds to an IP address;
here the target MAC address is a multicast MAC address.

There is no ARP for multicast. IP multicast address is algorithmically mapped to a multicast MAC address.

- Last 23 bits of IPv4 multicast address are used in MAC address
- Last 32 bits of IPv6 multicast address are used in MAC address

MAC multicast addr.	*Used for*
01-00-5e-XX-XX-XX	IPv4 multicast
33-33-XX-XX-XX-XX	IPv6 multicast

IP dest address	229.130.54.207
IP dest address (hexa)	e5-82-36-cf
IP dest address (bin)	...-10000010-...
Keep last 23 bits (bin)	...-00000010-...
Keep last 23 bits (hexa)	02-36-cf
MAC address	01-00-5e-03-36-cf

Several multicast addresses may correspond to same MAC address

- if needed, operating system removes packets received unnecessarily; it is hoped that this rarely happens

All multicast is handled by MAC layer as ASM

(i.e. MAC multicast address depends only on IP multicast IP address m not on source address s, even if m is an SSM address)

MAC Multicast

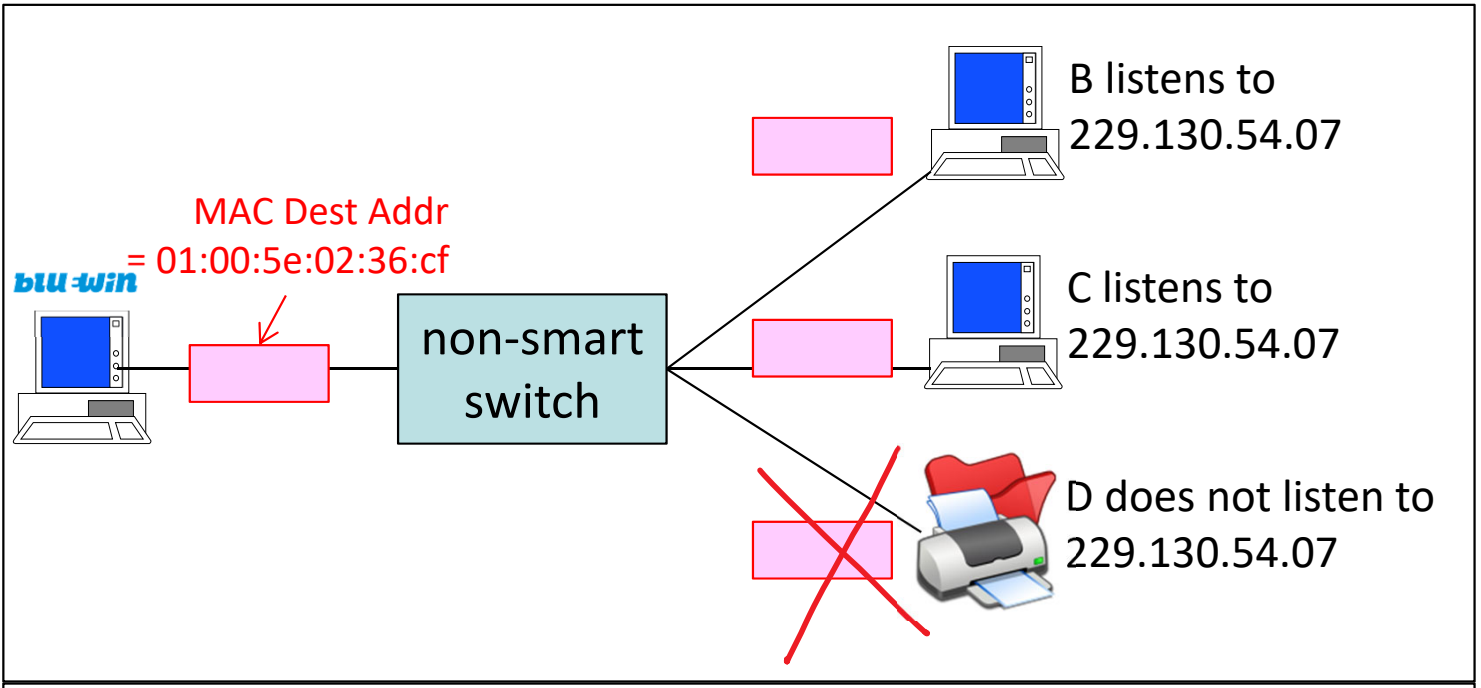

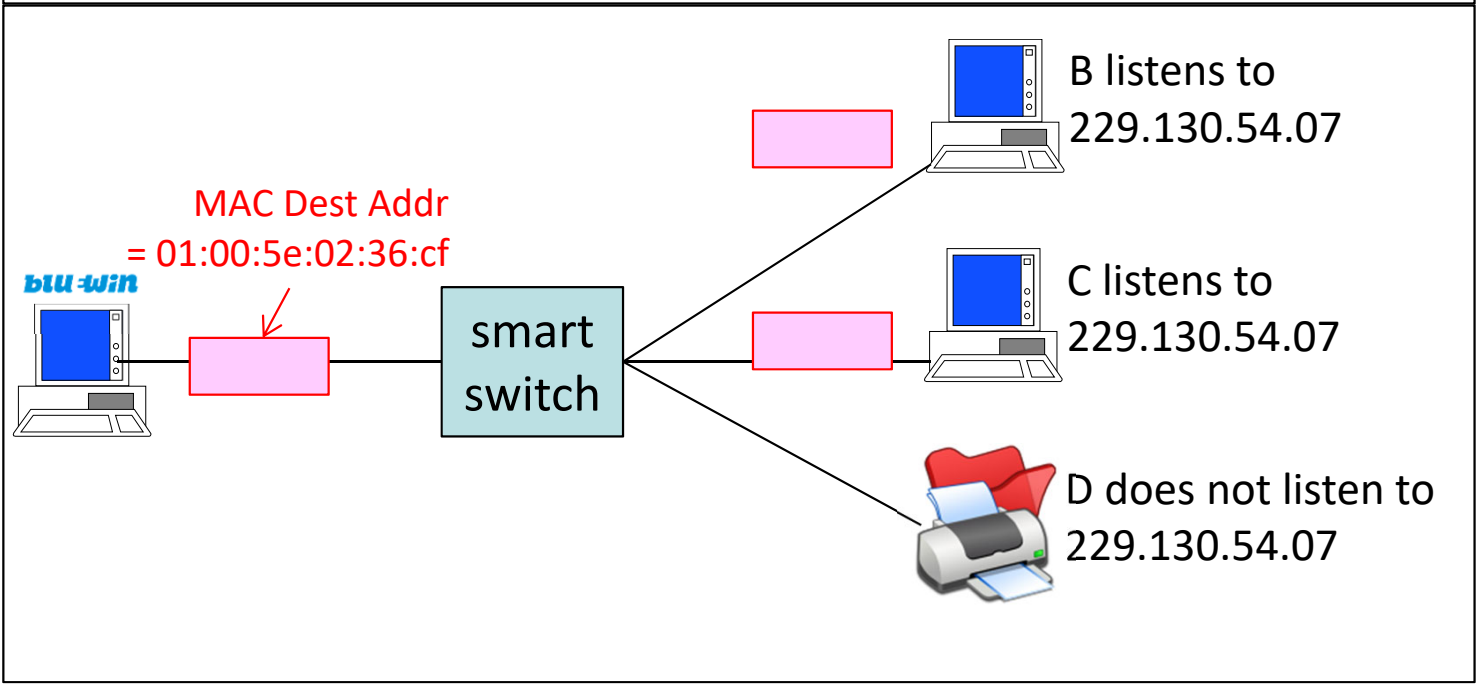

Some (non smart) switches simply treat multicast frames as broadcast.

Some smarter switches simply listen to IGMP/MLD and overhear who listens – deliver only to intended recipients – but do not distinguish SSM from ASM.

Multicast Routing

There are many multicast routing protocols. In practice, widespread is PIM: Protocol Independent Multicast. It supports ASM and SSM and exists in two versions: sparse and dense.

PIM-DM (Dense Mode) makes heavy use of broadcast and can be used only in small, tightly controlled networks.

PIM-SM (Sparse Mode) is more reasonable and is used e.g. for TV distribution.

When used with SSM, PIM-SM is very simple: it uses Reverse Path Forwarding: when a router (such as R1) needs to add a receiver, it sends a PIM/JOIN towards the source, using unicast routing. This creates the distribution tree on the fly.

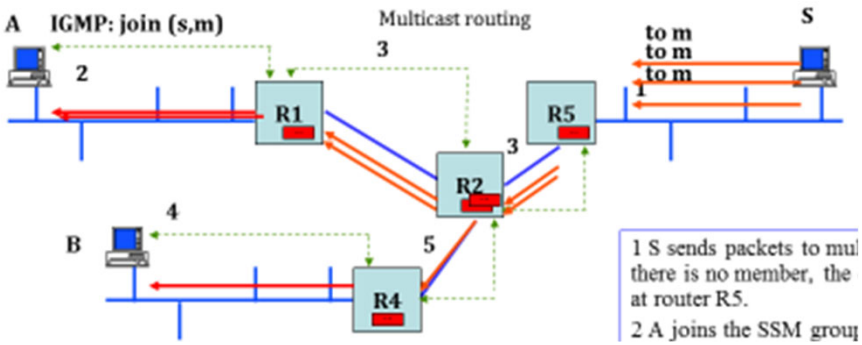

PIM-SM for ASM is more complicated; it uses one multicast router as Rendez-vous Point (RP): destination routers create a tree from RP, using RPF; router closest to source sends source packets to RP; if there exists an interested receiver in the domain, RP creates a tree from source (using RPF) otherwise drops; destinations create trees from sources, using RPF.

Security of IP Multicast

IP multicast makes life easier for attackers (e.g. Denial of Service, witty worm)

- mitigations: limit multicast rate and number of groups; control which multicast group is allowed (access lists)

SSM is safer as routers and destination can reject unwanted sources

IGMP/MLD is not secured and has the same problems as ARP/NDP

- mitigated by same mechanisms: sniffing switches observe all traffic and implement access-lists

multicast capable networks must deploy exhaustive filtering and monitoring tools to limit potential damage

Multicast in Practice

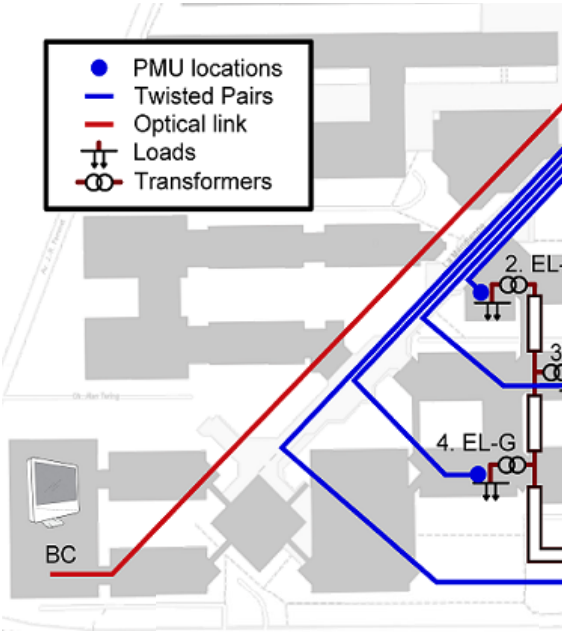

Multicast is good for sources : one packet sent for n destinations -- multiplication is done repeatedly, $O(\log(n))$ times

Multicast suffers from per-flow state in routers

Multicast is not supported everywhere, but is (with PIM-SM):

- At EPFL and other academic networks
- Internet TV distribution
- In some corporate networks for news, sensor streaming, time synchronization, large videoconferences etc…
- In industrial networks (smart grids, factory automation)

Works only with UDP, not with TCP

a) Say what is true

A. A
B. B
C. C
D. A and B
E. A and C
F. B and C
G. All
H. None
I. I don't know

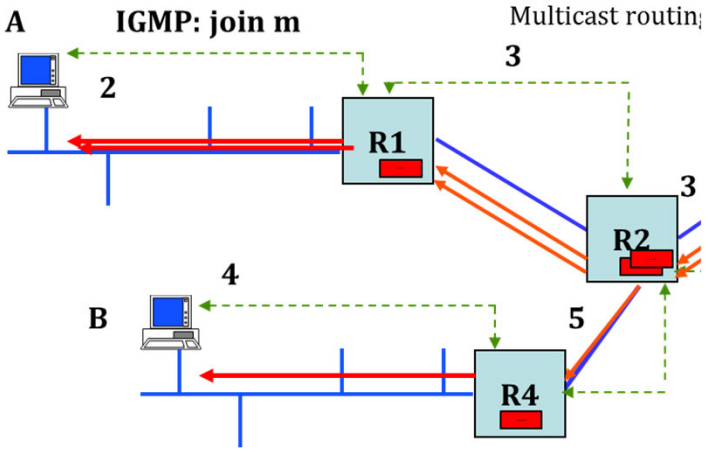

A. In order to send to a multicast group a system must first join the group with IGMP or MLD
B. In order to receive from a multicast group a system must first join the group with IGMP or MLD
C. A system can know whether a packet is multicast by analyzing the IP destination address.

b) The destination MAC address is…

A. A group address derived from the last 23 bits of the IPv6 destination address
B. A group address derived from the last 24 bits of the IPv6 destination address
C. A group address derived from the last 32 bits of the IPv6 destination address
D. A broadcast address
E. The MAC address of an ARP server
F. I don't know

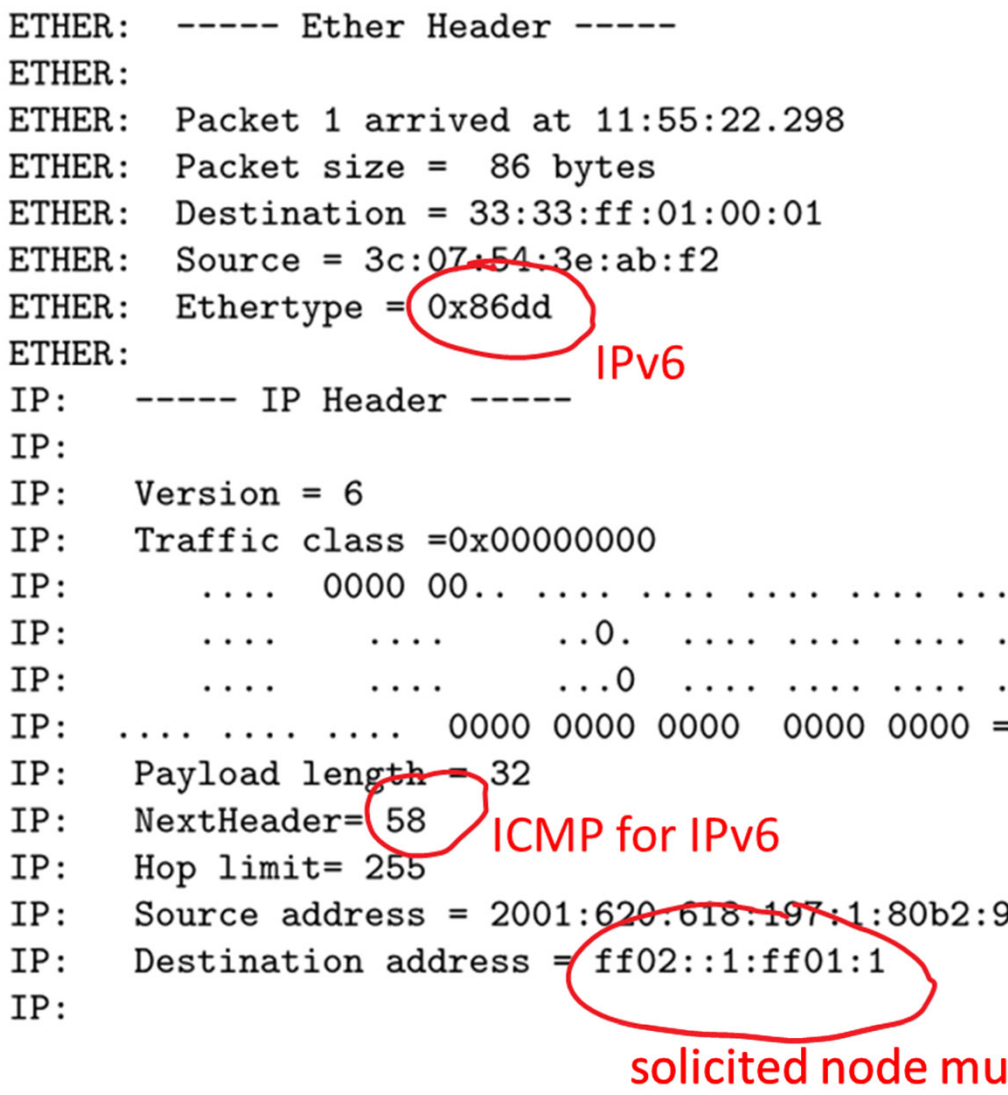

c) Switches handle all multicast as ASM. What is the implication ?

A. an SSM receiver may receive unwanted traffic at the MAC layer
B. SSM traffic is not supported
C. A and B
D. None
E. I don’t know

Solution

a) F

b) C

c) A; a receiver that subscribes to group (s,g) will also receive (at the MAC layer) the traffic destined to group (s’,g) if there is a subscriber to this group on the same LAN. The IP layer will filter out this traffic. In practice, you allocate different SSM IP addresses to different groups in order to avoid this problem.

Conclusion

IP multicast came as an after-thought and uses a different principle than IP unicast (exact match versus longest prefix match) – is not widely deployed

IP multicast addresses cannot be aggregated

IP multicast require the deployment of a solution to compute the multicast trees between routers (with a multicast routing protocol such as PIM-SM or with a network management application, SDN)